Belief Records and the Price of Disaster Fear

July 22, 2026

Abstract

In a complete-market economy in which investors hold a disaster model until its since-adoption record rejects it, the equilibrium risk-neutral disaster intensity factors exactly into a representative-agent benchmark times a wealth-weighted power mean of current beliefs. Inversion recovers the market's priced disaster belief, not the physical law. However, that belief aggregate is not the complete pricing state: continuation values depend on the law governing belief revision, and that law depends on the records under which current beliefs have survived. Therefore, conditional on the current option state, proxies for those survival records forecast the slope of the disaster-tail term structure—a residual that belief-sufficient models set to zero.

1. Introduction

Index options price rare disasters, and the disaster distribution they imply is not the one the macroeconomic record measures. The physical disaster sizes of Barro (2006), inverted through a frictionless representative-agent kernel, overstate the tail options price on average; yet that same risk-neutral tail spikes far above any physical benchmark after a crisis (Backus et al., 2011; Bollerslev and Todorov, 2011; Andersen et al., 2015b). We show the gap has an exact equilibrium interpretation: the surface prices beliefs about a fixed physical law, and inversion recovers the priced belief aggregate exactly, not the physical law. But that aggregate alone is not enough to price continuation claims, whose value depends on the law governing belief revision.

Investors in our complete-market Lucas economy each hold a single model of the disaster rate and revise it *lazily*—held while the record accrued since adoption stays unsurprising under it, abandoned for the best alternative at a surprise (Ortoleva, 2012). From common priors and common data, heterogeneous only in their revision tolerances, they come to disagree endogenously. Current beliefs thus determine current subjective disaster laws, while survival records determine when those laws change.[1] In equilibrium the risk-neutral disaster intensity factors exactly,

$$U_t \;=\; U^{\mathrm{RA}}\,\kappa_t,$$

into a representative-agent benchmark U^{RA}—the intensity a correct-belief agent prices off the macro-measured physical law—times a *belief multiplier* κ_t: a wealth-weighted *power* mean, of order the common risk tolerance, of current disaster likelihood ratios (Jouini and Napp, 2007; Cvitanić and Malamud, 2011). Representative-agent inversion therefore recovers a priced belief aggregate, $\lambda^\star\kappa_t$—the object the disaster-option literature measures (Baron et al., 2024; Chen et al., 2012)—not the physical disaster law. However, extracting a spot intensity, estimating its transition law as a function of that intensity, and using the fitted law to price longer maturities imposes precisely the sufficiency restriction the theorem refutes.

The reason is the state. Prices of continuation claims capitalize how the cross-section will evolve, and its evolution is not a functional of the cross-section: revision timing turns on the survival records and tolerances that the record-resolved state Ψ_t carries and the current belief distribution S_t does not (Figure 1 summarizes the mechanism). Under CRRA the split is exact. *First*, the fear factor U_t, the multiplier κ_t, the short rate, and the risk-neutral size shape are functionals of S_t alone. *Second*, the price–dividend ratio, the level of the disaster premium,

[1] A lazily held belief is not a sufficient statistic for its own revision—the individual-state result of the companion paper (Teguia, 2026), whose equilibrium incidence this paper derives.

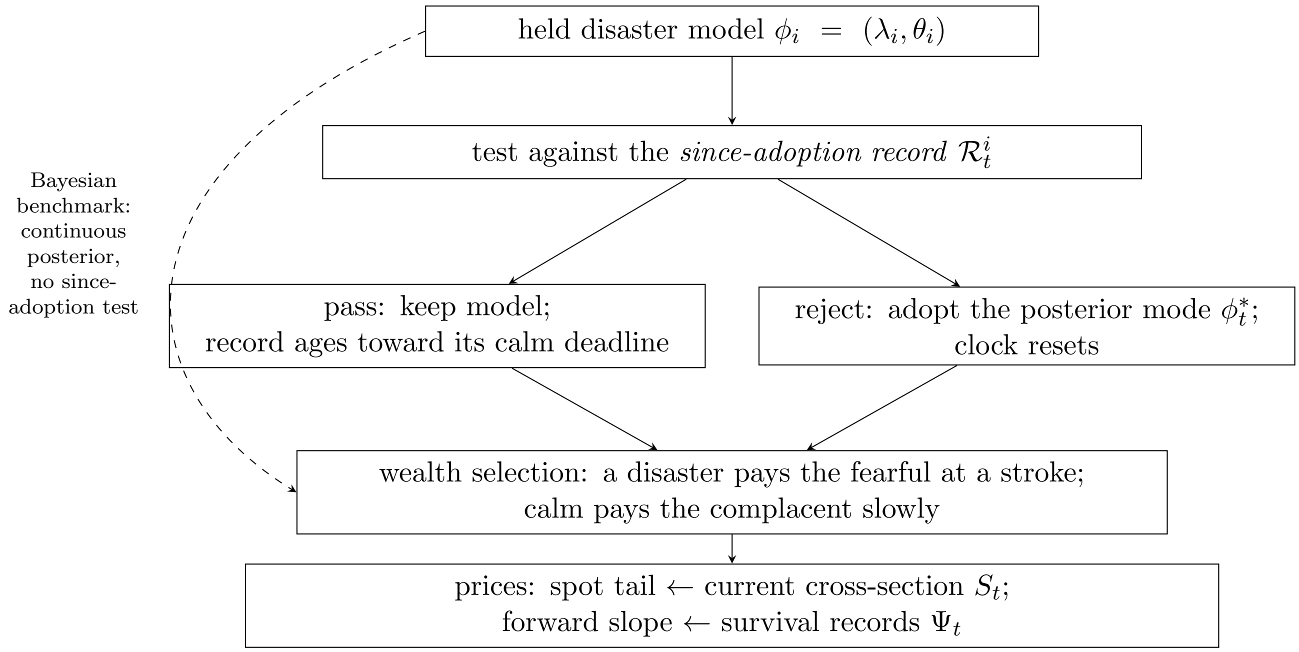

Figure 1: The mechanism. An investor holds one disaster model and tests it only on the record accrued since adoption. An unsurprising record keeps the belief and ages toward a scheduled calm deadline; a surprising one triggers revision to the common posterior mode. Wealth meanwhile reallocates toward whichever views the path vindicates. The spot tail prices the current cross-section of held beliefs; the term-structure slope prices the records under which they survived—the coordinate no belief-sufficient model carries. A Bayesian benchmark bypasses the record loop (dashed): her posterior is her state, continuously updated with no since-adoption test, so her forward surface reads off the current cross-section.

and the term structure of the tail load on the full Ψ_t. Consequently, economies with identical current beliefs, wealth shares, short rates, and spot disaster tails carry different forward tails, disaster premia, and price–dividend ratios (Figure 2). The spot surface reads *what the market fears*; the term structure reads *how long it has feared it*.

Call a model *belief-sufficient* if its pricing state is the current cross-section of beliefs, however generated—the field's workhorses, from latent intensities (Wachter, 2013; Gabaix, 2012; Seo and Wachter, 2019) to Bayesian and filtered posteriors (Collin-Dufresne et al., 2016; Dew-Becker et al., 2025), heterogeneous cross-sections (Chen et al., 2012), and diagnostic distortions (d'Arienzo, 2020; Wang, 2021). The record's imprint is absent at the short end and enters at first order beyond it—two dates with the same spot tail share the front of the curve and part only further out—so the test is a residual one: after controlling for the current option state, survival-record proxies—time since the last disaster (observable), accumulated calm surprise (a model-guided reconstruction), distance to a rejection boundary (structural)—should forecast the slope of the disaster-tail term structure, while any belief-sufficient model sets that residual to zero. The raw path-dependence documented in the surface runs the other way—strongest at the short end and *weakening* with maturity (Guyon and Lekeufack, 2023)—so ours is the orthogonal residual that survives the control and *enters* where that raw effect fades. A nonzero residual rejects belief-sufficiency *given* the maintained spanning of the current option state. Attributing the residual to *testing* rests on two signs no exogenous clock produces—calm deadlines that shorten in the held intensity, and a disaster-triggered upward re-anchoring arriving with the wealth reweight. The rejection does not rest on the attribution—any history-bearing failure of belief-sufficiency is caught; the two signs mark *testing* as the source. Time since the last disaster is recovered from the path alone; the remaining proxies are model-based reconstructions (Remark 1)—testable where a latent disaster intensity, the literature's dark matter (Campbell, 2018; Chen et al., 2024), is not.

The record-dependent transition law also produces asymmetric dynamics: a disaster reweights wealth toward the fearful and can trigger a coordinated *upward* re-anchoring, while calm erodes fear through gradual selection and, under monotone re-selection, staggered rejection deadlines, and demographic turnover sustains the cycle in a stationary economy. The level implication is a *countercyclical* priced tail: the power mean sits below the arithmetic mean, so dispersion alone prices the tail beneath the representative-agent benchmark, and the inverted intensity $\lambda^\star \kappa_t$ falls through a calm and jumps at impact. The term-structure implications are the record's: a disaster inverts the disaster-tail curve, the inversion unwinds on a schedule timed by the crisis's age, and the calm-conditional calendar spread—in traded form, deep out-of-the-money put calendar spreads—is priced by the same record channel that every belief-sufficient model sets to zero. Conditional on the spot state, the record wedge is usually a few percent of the twelve-year forward tail and organizes a large share of the calm-conditional forward variation; and because two investors can share the current predictive belief while their beliefs survived different records, the channel moves prices without moving fundamental forecasts. The slope prediction (Theorem 1) is untested and untargeted—new to the data and to the calibration alike.

Two caveats are in order. First, the calibration is illustrative rather than estimated: absent data on disagreement about the option-relevant disaster intensity, the belief block is pinned to a few stylized option-tail targets, and we make no out-of-sample claims.[2] Second, we use CRRA preferences because they keep the factorization and the spot/forward separation exact; the cost is the familiar safe-rate problem of power utility. Still, the model is falsifiable: wherever it pins a sign, the opposite cyclicality in the data rejects it.

Related literature. The paper adds no disaster process to rare-disaster asset pricing (Barro, 2006; Wachter, 2013; Gabaix, 2012; Seo and Wachter, 2019): the physical law is fixed, and a heterogeneous, lazily revising cross-section moves the priced tail. The level it moves is shared territory. A power mean of order below one sits under the arithmetic mean, so a cross-section right on average prices the tail below the representative-agent benchmark—dispersion alone is complacency (Jouini and Napp, 2007; Cvitanić and Malamud, 2011)—and the far-tail gap of Backus et al. (2011), the investor-fears wedge of Bollerslev and Todorov (2011), and the unspanned tail factor of Andersen et al. (2015b) all sit at this front of the curve; the selection engine—disasters paying the fearful at a stroke, calm paying the complacent slowly—is likewise the market-selection mechanism of Blume and Easley (2006) and Kogan et al. (2006), a functional of the current cross-section. Chen et al. (2012) price a given cross-section of disaster beliefs; we ask why identical current cross-sections carry different forward prices. What the field has measured sits on the shared side of the line (Internet Appendix Table 1); the level's one directional prediction is a countercyclical inverted intensity $\lambda^{\star}\kappa_t$, and the procyclical *physical* tail of Gormsen and Jensen (2025) is reconciled with it in Internet Appendix **??**.

A large literature lets rare events and slow learning generate *persistent*, belief-driven dynamics

[2] Turnover enters the stationary economy of Section 6 through a standard overlapping-generations structure; the model dispersion of entering cohorts adds an exogenous source of disagreement to the endogenous one of the homogeneous-start baseline.

(Collin-Dufresne et al., 2016; Wachter and Zhu, 2025; Ghaderi et al., 2022; Kozlowski et al., 2020)—post-disaster fear that persists for years under Bayesian latent-regime learning, option surfaces generated by filtered disaster states (Seo and Wachter, 2019), tail beliefs scarred by observed draws—but in every case the persistence is a property of the *belief*: the posterior is the pricing state, and two markets that share it share their forward surfaces. The closest neighbor, Dew-Becker et al. (2025), filters a latent value smoothly into a perceived intensity that is its own Markov pricing state, delivering the leverage effect and volatility dynamics we concede; ours needs Ψ_t, so the pricing state, not the shared aggregator, separates the two; a record coordinate is one structural source of the excess volatility of long-maturity claims relative to short-end dynamics (Giglio and Kelly, 2018).

Against non-Bayesian revision the division of labor is exact: the rule is Ortoleva (2012)'s (see also Menzies and Zizzo 2009; Cho and Kasa 2015), the enlarged individual state is the companion paper's (Teguia, 2026), and this paper owns the price decomposition the enlargement creates in equilibrium. Models of costly attention and infrequent adjustment place a time-since-last-action coordinate in the agent's state (Abel et al., 2013; Bacchetta and van Wincoop, 2010) and are likewise outside the belief-sufficient class; what testing adds are the two signs above—a surprise clock, shorter for more fearful models, fired upward by a disaster—content an attention cost does not supply and the data can check. And where memory determines beliefs and prices (Malmendier and Nagel, 2011; Nagel and Xu, 2022; Wachter and Kahana, 2024; Bordalo et al., 2020), the remembered past moves the *current predictive belief*, so memory and the forecast move together; here two investors can share the current predictive belief and differ in their expected revision times, because their beliefs survived different records—market memory that prices without predicting fundamentals. Price-extrapolation (Bastianello and Fontanier, 2026), robust learning (Hansen and Sargent, 2010), and eroding confidence (Lanzani, 2025) each move a current belief object that remains sufficient; our investors commit to a single model and test it.

Section 2 establishes the result in a two-economy example. Sections 3 and 4 derive its two ingredients—what the spot tail recovers, and why the recovered aggregate does not determine its own evolution—and Section 5 combines them in the full equilibrium. Section 6 disciplines the mechanism against the option-tail evidence. Proofs are in Appendices A–B and the Internet Appendix.

2. The main result in a two-economy example

This section constructs the main result in miniature: in two dates and one formula, a pair of economies with identical current pricing states for the spot tail and different continuation states—the difference in a coordinate no belief-sufficient state can carry.

The priced object. Disasters arrive at the Poisson rate $\lambda^\star = 3\%$ per year with sizes $J \sim F_J$; investors have CRRA utility with $\alpha = 3.5$ and entertain two rival models of the arrival rate, a complacent $\lambda_{lo} = 1\%$ and a fearful $\lambda_{hi} = 6\%$—the physical and preference blocks of Table I. The simplest disaster claim pays one dollar at the first disaster inside a window; its price per unit of physical probability is the *risk-neutral* tail, which factors (Proposition 2) into a fixed risk-aversion markup $\mathbb{E}_f[e^{\alpha J}] = 2.82$—common to every comparison below, so we drop it—times a belief multiplier. When a share s of wealth holds the fearful model, the multiplier is the intensity-channel form of Internet Appendix Corollary **??**,

$$\kappa \;=\; \Big[s \left(\lambda_{hi}/\lambda^\star\right)^{1/\alpha} + (1-s) \left(\lambda_{lo}/\lambda^\star\right)^{1/\alpha} \Big]^{\alpha},$$

the wealth-weighted *power* mean of the held rates, not their average. With a quarter of wealth fearful, $\kappa_0 = 0.573$: the market prices the tail at 57 cents per benchmark dollar.

Two updaters. Hand the data to a Bayesian with an even prior over the two models—the section's public prior. Each calm year multiplies her posterior odds on the fearful model by $e^{-0.05}$; five quiet years carry her predictive rate from 3.5% to 3.19%. Her belief *glides*—and, the operative property, her posterior is a sufficient statistic for her own future: tell me her belief today and I can price every claim at every horizon.

The lazy investor of Definition 1 does not glide. She *holds* the fearful model and tracks one object, the record since she adopted it: through calm, the quiet spell Δ, whose surprise under her own model—essentially $e^{-\lambda_{hi}\Delta}$—drains deterministically and crosses her tolerance $\varepsilon = 0.650$ at the *deadline*

$$\Delta^{\star} \;\approx\; \frac{\log(1/\varepsilon)}{\lambda_{hi}}\,,$$

the zero-count term of the exact criterion (14), whose full two-sided form sets the deadline at $\Delta^{\star} = 11.9$ years (the zero-count term alone would say 7.2); the exact date is the example's clock, deterministic and finite by Internet Appendix Proposition **??**. At the deadline she abandons the model and re-anchors at the posterior mode, after any quiet spell the *complacent* model (Internet Appendix Lemma **??**).

Where the wealth goes. In calm the fearful investor is long insurance that never pays, and her wealth bleeds to the complacent: the fearful share obeys the replicator of Internet Appendix Proposition **??**,

$$d\log\frac{s_t}{1-s_t} \;=\; -\frac{\lambda_{hi}-\lambda_{lo}}{\alpha}\,dt \;=\; -1.43\% \text{ per year},$$

so five quiet years carry s from 0.250 to 0.237 and the multiplier from 0.573 to 0.558—a smooth cent-and-a-half decline. At a disaster the insurance pays: wealth is reweighted by each model's likelihood ratio, $s \mapsto s\,(\lambda_{hi}/\lambda^{\star})^{1/\alpha}/G$, jumping the fearful share from 0.250 to 0.357. The channel is smooth in calm and proportional at crashes—wealth never lurches on a schedule—and it is a functional of the current cross-section alone: hand me today's s and I know today's drift and today's reweight exactly. A Bayesian economy—formally, any belief-sufficient benchmark sharing the date-t cross-section—has the *same* wealth dynamics; selection acts on realized likelihoods, not on the updating rule, so it is the engine the benchmark shares, not the one that separates.

Up by surprise, down on schedule. At a crash the two engines fire together: the reweight lifts the multiplier from 0.573 to 0.706 through wealth alone, while one disaster against five quiet years flips the public mode, and every investor whose updated record leaves the typical set re-anchors *upward* at once—through the purse and through the mind, the purse in proportion, the mind in lumps.

Two histories, one cross-section. Now build two tester economies identical today: in each, a quarter of wealth is fearful, so both price the spot tail at $\kappa_0 = 0.573$. They differ only in what today's cross-section cannot see, the records. In *Aged*, the fearful cohort adopted its model eight years earlier, so it carries eight quiet years of record and its deadline arrives in year four ($\Delta^\star - a \approx 3.9$). In *Fresh*, the cohort adopted at date 0 with a blank record; its deadline lies beyond the horizon. Price the five-year-forward tail claim—the tail over years five to six, conditional on calm: a calendar spread of deep out-of-the-money puts.

	spot κ_0	forward κ_5 (calm)	forward / spot
Aged	0.573	0.333	**0.58**
Fresh	0.573	0.558	**0.97**

Against the Fresh counterfactual, the shared selection over the full five years accounts for 1.5 cents and the deadline for the 22.4-cent gap between the rows: selection is the channel the two economies share while both run it; what separates the rows is beliefs re-anchoring, and nothing else. (The table reports the left-endpoint forward multiplier κ_5; the year-five-to-six tranche integrates the forward multiplier over the window, the two differing only through the within-window drift.)

Example Proposition. In the two-date binary economy, let Aged and Fresh share the cross-section s_0 and differ only in the fearful cohort's adoption date. Then (i) the spot tails coincide, $\kappa_0 = \kappa_0'$; (ii) the calm-conditional forward tails differ, $\kappa_5 = \lambda_{lo}/\lambda^\star < \kappa_5'$; and (iii) in any economy whose date-0 pricing state is the date-0 cross-section—the Bayesian of this section included—the two forwards coincide.

(Proof in Internet Appendix **??**.) Item (iii) covers the Bayesian outright—her posterior *is* her state—and the wedge is not a wealth effect in disguise: the wealth channel is common to the class, so conditioning on the spot surface nets it out.

The example assumes what the model must deliver: deterministic calm deadlines and a common re-selection destination; the split surviving a continuum of tolerances, general records, disasters, and turnover; and no finite statistic of the cross-section recovering the records. The next two sections derive the ingredients separately—Section 3 what any current cross-section prices, Section 4 why that cross-section does not determine its own evolution—and Section 5 combines them in Theorem 1, whose Remark 1 conditions on the spot surface to expose the residual in data. Every number is the calibration of Table I.

3. The economy, and what the spot tail recovers

We build the economy on one modeling choice: investors agree about everything except the law of disasters. The setting is a continuum-of-agents Lucas tree whose dividend carries Brownian risk and rare downward jumps drawn from a continuous size distribution. Investors share the diffusion—which the continuous record pins down on any interval—and dispute the disaster *law*. In the baseline they dispute its arrival rate λ, the size law held common; disagreement over the size index θ is an extension. We carry the joint law $\phi = (\lambda, \theta)$ as the general object and state every result in the *intensity channel* (every *held* model carries the correct size law, $f_{\theta(\phi)} = f$; the family $\{F_\theta\}$ itself remains nondegenerate and available for the size-disagreement extension) unless the size extension is flagged. Every quantitative result is intensity-channel; size disagreement is what makes the pricing state infinite-dimensional (Proposition 3) and is otherwise developed in the Internet Appendix. The object to track is the cross-section of held ϕ and how it moves. Why the disaster law, and not the diffusion, is the admissible thing to disagree about—the diffusion and the realized jumps are pinned down by the consumption path itself, whereas the intensity and the size law are learned only from the rare disasters as they accrue—is established in Internet Appendix **??**.

A. The aggregate endowment

Fix a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ satisfying the usual conditions, with $\mathbb{P}$ the physical law. On it live a standard Brownian motion W, a Poisson process N of intensity $\lambda^\star$ independent of W, and an i.i.d. sequence of disaster sizes $\{J_k\}_{k\geq 1}$ independent of (W, N). Write $\{T_k\}_{k\geq 1}$ for the jump times of N and collect the disasters in the integer-valued random measure

$$\mu(dt, dx) = \sum_{k\geq 1} \delta_{(T_k, J_k)}(dt, dx)$$

on $[0, \infty) \times \mathcal{J}$, a marked Poisson process in the sense of Brémaud (1981).

Assumption 1 (Primitives of the endowment): *The drift $\mu_0 \in \mathbb{R}$, the diffusion coefficient $\sigma > 0$, and the disaster intensity $\lambda^\star > 0$ are constants known to the modeler (the fixed physical law; μ_0 and σ are common knowledge among investors, while $\lambda^\star$ and F_J are known to the modeler alone—the objects investors dispute), and the size support $\mathcal{J} = [\underline{J}, \overline{J}] \subset (0, \infty)$ is a known compact interval. The disaster-size law F_J is a Borel probability measure on $\mathcal{J}$ with continuous, strictly positive Lebesgue density f. Aggregate consumption (the Lucas-tree dividend) $C_t > 0$ solves*

$$\frac{dC_t}{C_{t-}} = \mu_0\, dt + \sigma\, dW_t - \int_{\mathcal{J}} \left(1 - e^{-x}\right) \mu(dt, dx), \qquad C_0 > 0, \tag{1}$$

so that, in logs,

$$d \log C_t = \left(\mu_0 - \tfrac{1}{2}\sigma^2\right) dt + \sigma\, dW_t - \int_{\mathcal{J}} x\, \mu(dt, dx) :$$

between disasters C is geometric Brownian, and at T_k it falls by $C_{T_k} = C_{T_k-} e^{-J_k}$, i.e. $\Delta \log C_{T_k} = -J_k$. The filtration $(\mathcal{F}_t)$ is the $\mathbb{P}$-augmentation of the natural filtration of the path C.

Assumption 2 (The belief family): *The size family $\{F_\theta : \theta \in \Theta\}$ of Section 3.B, with Lebesgue* *densities f_θ on $\mathcal{J}$ ($\Theta \subset \mathbb{R}$ a compact interval, and $\Lambda = [\underline{\lambda}, \overline{\lambda}] \subset (0, \infty)$ containing $\lambda^\star$ in its interior), satisfies:*

(i) Full support: *$(x, \theta) \mapsto f_\theta(x)$ is jointly continuous with $f_\theta(x) > 0$ on $\mathcal{J} \times \Theta$, hence bounded and bounded away from zero on that compact set.*

(ii) Regularity: *$\theta \mapsto \log f_\theta(x)$ is continuously differentiable with score $\dot{\ell}_\theta(x) := \partial_\theta \log f_\theta(x)$, and the Fisher information $I(\theta) := \int_{\mathcal{J}} \dot{\ell}_\theta(x)^2 \, F_\theta(dx)$ is finite and strictly positive on Θ.*

(iii) Identifiability: *$\theta \neq \theta' \Rightarrow F_\theta \neq F_{\theta'}$.*

(iv) Strict KL identifiability: *$\theta \mapsto \mathrm{KL}(F_J \| F_\theta)$ has a unique minimizer $\theta^\star$ on Θ.*

(v) Spanning: *for every $\alpha > 0$, the family $\{(f_\theta/f)^{1/\alpha} : \theta \in \Theta\}$ is linearly independent in $L^1(\mathcal{J}, f)$—no nontrivial finite linear combination vanishes f-a.e.*

(vi) Smooth pseudo-true point: *$\theta^\star \in \operatorname{int}\Theta$; for each $x \in \mathcal{J}$, $\theta \mapsto \log f_\theta(x)$ is twice continuously differentiable on a neighbourhood of $\theta^\star$, with $|\dot{\ell}_\theta(x)|$ and $|\partial^2_\theta \log f_\theta(x)|$ dominated there by an envelope $\Psi(x)$ satisfying $\mathbb{E}_{F_J}[\Psi] < \infty$ and $\mathbb{E}_{F_J}[\Psi^2] < \infty$ (the scalar envelope $\Psi(\cdot)$ on $\mathcal{J}$ bears no relation to the record-resolved state Ψ_t of Definition 4); and, writing $M(\theta) := \mathbb{E}_{F_J}[\log f_\theta(J)]$, the pseudo-Hessian $H := -M''(\theta^\star)$ is strictly positive and $B := \mathbb{E}_{F_J}[\dot{\ell}_{\theta^\star}(J)^2] < \infty$.*

Assumption 1 draws a deliberate line between what investors agree on and what they may dispute. Under $\mathbb{P}$ the disaster measure μ has compensator

$$\nu(dt, dx) = \lambda^\star \, dt \, F_J(dx), \tag{2}$$

the arrival rate $\lambda^\star \, dt$ times the size law $F_J(dx)$. Investors share the drift μ_0, the diffusion σ, and the support $\mathcal{J}$, and agree that disasters arrive as a marked Poisson process; what they dispute is the disaster *law* itself—both how often disasters strike and how large they tend to be. The reason the line falls here is informational: the diffusion is common knowledge because the continuous record pins it down on any interval in the high-frequency limit (Internet Appendix **??**), whereas the disaster law is learned only from the rare disasters as they accrue.

B. The disputed object and the belief rule

An investor holds one dogmatic model of the disaster process: a pair $\phi = (\lambda, \theta)$ consisting of an arrival intensity λ and an index θ of the size law F_θ. Both are parameters of a distribution—the intensity of an arrival law, the index of a size law—and both are disputed. The true process has intensity $\lambda^\star$ and size law F_J; neither is known to the investor, and her held pair is tested jointly—counts and sizes together—against its since-adoption record; the full public history enters only through the re-selection posterior.

The held model is a point $\phi = (\lambda, \theta)$ in the compact parameter set $\Phi := \Lambda \times \Theta$, with $\Lambda = [\underline{\lambda}, \overline{\lambda}] \subset (0, \infty)$ containing $\lambda^\star$ in its interior and $\Theta \subset \mathbb{R}$ a compact interval. Under ϕ the disaster process is a marked Poisson process with intensity λ and i.i.d. marks drawn from F_θ on $\mathcal{J}$, where $\{F_\theta : \theta \in \Theta\}$ has Lebesgue densities f_θ satisfying Assumption 2 (Internet Appendix **??**).

The belief rule is the dynamic hypothesis-testing rule of Ortoleva (2012), applied to the joint model $\phi = (\lambda, \theta)$; the dynamics it generates are derived in Section 4.

Definition 1 (Dynamic Ortoleva rule on the joint model): *All objects are adapted to the public filtration $(\mathcal{F}_t)$, the $\mathbb{P}$-augmentation of the natural filtration of the path C (Assumption 1). Investor i has a permanent surprise threshold $\varepsilon^i \in (0, 1)$ and a* prior over models*: a second-order posterior ρ^i_t on the menu Φ, the Bayesian update of a prior ρ_0 over models by the likelihood of the* entire *marked-Poisson record on $[0, t]$,*

$$\rho^i_t(\phi) \;\propto\; \rho_0(\phi)\, e^{-\lambda t}\, \lambda^{N_t} \prod_{k:\, T_k \le t} f_\theta(J_k), \qquad \phi = (\lambda, \theta), \tag{3}$$

which is $\mathcal{F}_t$-measurable and accumulates on the full history. Her held model $\phi_i(t) \in \Phi$ is piecewise constant, equal to the mode of her posterior ρ_t^i over models at her most recent adoption date and held until the next rejection. The rule runs as follows.

Step 0 (initial model—the best given the prior over models). *At $t = 0$ the investor adopts the maximum-a-posteriori model of the prior ρ_0,*

$$\phi_i(0) \;=\; \phi_0 \;\in\; \arg\max_{\phi\in\Phi} \rho_0(\phi), \qquad a_0^i = 0,$$

the common starting model under the homogeneous start of §3.C. Write a_t^i for the most recent adoption date up to t and $\tau_t^i = t - a_t^i$ for the regime age.

Step 1 (the test—the since-adoption history against the held model). *The held model is tested* only *on the history accrued since it was adopted: the record on the window $(a_t^i, t]$,*

$$\mathcal{R}_t^i \;=\; \big(n_t^i,\; (J_k : a_t^i < T_k \le t)\big), \qquad n_t^i \;=\; N_t - N_{a_t^i}, \tag{4}$$

its count and realized sizes, both $\mathcal{F}_t$-measurable.[3] *Under the held model $\phi = (\lambda, \theta)$—which makes the disaster process a marked Poisson process of intensity λ with i.i.d. marks F_θ (Assumption 2)—a window of length Δ has predictive law $Q_{\phi,\Delta}$, the law of the window record* under the held model's own measure $\mathbb{P}^\phi$, *with density (against counting measure on the count and Lebesgue measure on the sizes)*

$$q_{\phi,\Delta}\big(n, (x_1, \ldots, x_n)\big) \;=\; \underbrace{e^{-\lambda\Delta}\,\frac{(\lambda\Delta)^n}{n!}}_{\textit{count: }\text{Poisson}(\lambda\Delta)} \;\cdot\; \underbrace{\textstyle\prod_{j=1}^n f_\theta(x_j)}_{\textit{sizes: i.i.d. } F_\theta}.$$

Departure of the realized record from the held model is measured by the deviance*—the saturated count deviance at the empirical rate n/Δ plus the size log-likelihood ratio of the record against*

[3] The record is recovered from the path by Lemma IA.1(ii), which rests only on Assumption 1 and is independent of the rule, so the forward reference is non-circular, and the since-birth count N_t enters only through n_t^i. Because the held model is time-homogeneous, the predictive law $Q_{\phi,\Delta}$ below depends on the window only through its length $\Delta = \tau_t^i$, so the test window is never empty once $\tau_t^i > 0$.

its best in-family fit $\hat\theta$ (the size MLE on the realized sizes),

$$D_{\phi,\Delta}(\mathcal{R}) \;=\; \underbrace{\Delta\, d_P\Big(\tfrac{n}{\Delta}\,\Big\|\,\lambda\Big)}_{count} \;+\; \underbrace{\sum_k \log\frac{f_{\hat\theta}(J_k)}{f_\theta(J_k)}}_{size}, \qquad d_P(a\|\lambda) = a\log\tfrac{a}{\lambda} - a + \lambda,$$

the Poisson count deviance (with $d_P(0\|\lambda) = \lambda$) plus the parametric size deviance, each non-negative and zero only at a record the held model fits exactly. The surprise *of a record is its predictive deviance tail under $Q_{\phi,\Delta}$,*

$$\bar\varepsilon(\mathcal{R};\phi,\Delta) \;=\; Q_{\phi,\Delta}\big(\{\mathcal{R}' : D_{\phi,\Delta}(\mathcal{R}') \geq D_{\phi,\Delta}(\mathcal{R})\}\big).$$

The held model passes *the test at t iff the realized record is unsurprising, $\bar\varepsilon(\mathcal{R}^i_t;\phi_i,\tau^i_t) > \varepsilon^i$—equivalently, iff $\mathcal{R}^i_t$ lies in the mass-$(1-\varepsilon^i)$ typical set, the acceptance region $A_{\tau^i_t}(\phi_i,\varepsilon^i) = \{\mathcal{R} : D_{\phi_i,\tau^i_t}(\mathcal{R}) < r_{\phi_i,\tau^i_t}(\varepsilon^i)\}$ with $r_{\phi,\Delta}(\varepsilon)$ the $(1-\varepsilon)$-quantile of $D_{\phi,\Delta}$ under $Q_{\phi,\Delta}$ (at atoms of D under $Q_{\phi,\Delta}$—the count atoms, the zero-count record included—the p-value rule $\bar\varepsilon > \varepsilon$ governs and the closed-threshold form is its shorthand; the formal conventions precede Internet Appendix Definition* **??***).*

Step 2 (keep, or update and re-select). *As long as the held model passes, it is kept and only the regime age advances ($\dot\tau^i_t = 1$). It is* rejected *at the first instant the record leaves the typical set, $\tau^{\text{rev}} = \inf\{t > a^i : \bar\varepsilon(\mathcal{R}^i_t;\phi_i,\tau^i_t) \leq \varepsilon^i\}$, an $(\mathcal{F}_t)$-stopping time. At a rejection the investor (i)* updates the prior over models on the full history*—the posterior ρ^i_t of (3) evaluated at that instant—and (ii)* re-selects *the new model as the tie-broken mode of that posterior. Fix once and for all the lexicographic total order $\preceq$ on Φ: $(\lambda,\theta) \preceq (\lambda',\theta')$ iff $\lambda < \lambda'$, or $\lambda = \lambda'$ and $\theta \leq \theta'$. Let $\mathfrak{m}(\rho) := \min_{\preceq} \arg\max_{\phi\in\Phi} \rho(\phi)$ denote the $\preceq$-smallest maximiser; it exists and is unique because* $\arg\max_{\phi\in\Phi}\rho$ *is a non-empty compact subset of the compact Φ (the posterior is continuous on Φ), on which the lexicographic minimum is attained. The investor adopts*

$$\phi_i' \;=\; \mathfrak{m}(\rho^i_t) \;\in\; \arg\max_{\phi\in\Phi} \rho^i_t(\phi), \tag{5}$$

resets the adoption date to that instant (so τ^i restarts at 0), and returns to Step 1.[4]

Because the tolerance is held *fixed*, even a correct-on-average model is rejected and re-adopted without end (Lemma IA.2): revision is recurrent, not terminal. The second-order posterior, updated on the full record and never reset, nonetheless concentrates the re-selected mode at the pseudo-true $\phi^\star$ (Berk, 1966), so recurrent rejection and a consolidating destination coexist—the trigger frequentist, the re-selection Bayesian; Internet Appendix **??** places the rule in the sequential-testing and learning literatures. The since-adoption window gives a held model a fresh sample on which to fail, and the surviving record—not the current belief—organizes the dynamics of Section 4.

C. The population, and a homogeneous start

A continuum of investors, indexed by $i \in [0,1]$, watch the same path (1) and share everything in it but one thing. They share the undisputed primitives of Assumption 1—the drift μ_0, the diffusion σ, the support $\mathcal{J}$—and the menu of Assumption 2, but not the physical disaster law $(\lambda^\star, F_J)$, which is known to the modeler and is precisely what they learn and may dispute. The one thing that distinguishes them is a *tolerance for surprise*.

Assumption 3 (One-dimensional heterogeneity; homogeneous start): *Each investor i has a surprise tolerance $\varepsilon^i \in (0,1)$; the cross-section of tolerances is the primitive (non-random) continuum law G on $(0,1)$, with continuous and strictly increasing G, so ε^i is investor i's coordinate under G rather than an independent draw. This tolerance is her one permanent type. At $t = 0$ all investors share a common second-order prior $\rho_0 = \rho_0^\Lambda \otimes \rho_0^\Theta$ on $\Phi = \Lambda \times \Theta$, a product of a continuous prior on intensity and one on the size index, with full support and a continuously differentiable log-density; all hold the common model $\phi_0 = (\lambda_0, \theta_0) \in \arg\max_\phi \rho_0(\phi)$, differing only in ε^i.*

[4]The menu Φ permits re-selection in either direction in each coordinate, so no model is absorbing; since re-selection may return a model near the incumbent, the belief-moving event is the *switch* $\phi_i' \neq \phi_i$, not the rejection itself.

The environment and the disputable object are now fixed: a Lucas disaster economy in which investors share the diffusion and the occurrence of disasters but hold rival, identifiable laws $\phi = (\lambda, \theta)$ for the *intensity and size* of disasters, revised only when the rare-event record makes the held law surprising. We now ask what such a cross-section prices, deferring to Section 4 how otherwise-identical investors—differing only in tolerance—come to hold it and how it moves. The object to track is the *risk-neutral disaster-size law*—the distribution of disaster sizes priced into options, the empirical "investor fears" of Bollerslev and Todorov (2011). We build it in three steps: the change-of-measure density each held model induces and the equilibrium pricing kernel, a power-mean aggregate of those densities formed with the *fixed* Negishi multipliers (§3.D); and the risk-neutral disaster-size law it implies. That law, we show, is the physical law tilted by risk aversion and by a size-belief factor in which the *stochastic* consumption shares weight the held models' likelihood ratios—a tilt that moves with the cross-section while the physical disaster process stays fixed (§3.E). Preferences are common CRRA and disaster risk is spanned, as in the heterogeneous-belief disaster economy of Chen et al. (2012); the departure is that the beliefs are lazy-Ortoleva held models, not Bayesian, and the dispersion is endogenous (Section 4).

D. The market and the pricing kernel

Each investor prices under her current held model $\phi_i(t-)$ (Definition 1); her belief is a change of measure from the truth with density M_t^i, the point-process likelihood ratio of Brémaud (1981, Ch. VI) (Definition A.1).[5]

Assumption 4 (Market and preferences): *A continuum of investors $i \in [0, 1]$ have common CRRA utility $u(c) = c^{1-\alpha}/(1-\alpha)$, $\alpha > 0$ with $\alpha \neq 1$ (the log case $\alpha = 1$ is the continuous*

[5]Between revisions the held model is fixed and the disaster process is marked Poisson with intensity λ_i and size law F_{θ_i}; globally the subjective law is a marked point process with predictable compensator $\lambda_i(t-)\, F_{\theta_i(t-)}(dx)\, dt$ (Brémaud, 1981). Because the intensity is bounded and the per-disaster ratio is continuous and strictly positive on the compact $\mathcal{J} \times \Phi$, M^i is a strictly positive $(\mathcal{F}_t, \mathbb{P})$-martingale with $\mathbb{E}[M_t^i] = 1$ (Appendix A). M^i is built from the *realized* held-model process—the density that makes M^i a $\mathbb{P}$-martingale and feeds the Negishi construction of Internet Appendix **??**.

extension $u(c) = \log c$*; the disaster results use* $\alpha > 1$*), and common discount rate* $\delta > 0$*. Markets are complete: investors trade a riskless bond, the Lucas tree (a claim to the aggregate endowment* C *of Assumption 1), and a complete set of disaster-contingent (Arrow) claims spanning the jump risk, so there is a unique common state-price density* ξ_t *(Cox and Huang, 1989; Karatzas et al., 1990). Budget (Negishi) multipliers* $\zeta_i > 0$ *satisfy* $\int_0^1 \zeta_i^{-1/\alpha}\, di \in (0, \infty)$*.*
(iv) Symmetric allocation. *We select the symmetric, equal-Pareto-weight equilibrium:* $\zeta_i \equiv \zeta$ *(normalized* $\zeta = 1$*) for every date-0 investor. The allocation is selected in consumption-share space; we do* not *impose equal initial financial wealth, which is inconsistent with a common* ζ *under heterogeneous beliefs (Remark A.1). The constant* ζ *is fixed once, independently of the economy's realization: the same number in any two economies compared below. This selection is without loss for what the wedge delivers: its form, the sign and range of its magnitude, and its entire stationary law are invariant or robust to the Pareto weights (shown in the Internet Appendix); the weights fix only which cross-section enters the kernel at a date.*

The pricing weights here are *stochastic* consumption shares, so the share process is part of the state: the consumption-share-weighted cross-section S_t of held models, the wealth-weighted distribution whose emergence from otherwise-identical investors and law of motion are derived in Section 4.[6]

Proposition 1 (Equilibrium kernel): *Under Assumptions 1–4, and the maintained integrability condition that lifetime utilities and present-value budgets are finite (implied by the transversality condition* (T) *of Section 5,* $\delta > \sup_S m(S)$*, where* $m(S)$ *is the equilibrium value-growth rate defined there), a competitive equilibrium (Definition A.2) exists and is unique; its state-price density is*

$$\xi_t = e^{-\delta t}\, C_t^{-\alpha}\, \Xi_t^{\alpha}, \qquad \Xi_t := \int_0^1 \zeta_i^{-1/\alpha} \big(M_t^i\big)^{1/\alpha}\, di \ \in (0, \infty), \tag{6}$$

[6] This is the complete-market, heterogeneous-belief Lucas economy of Basak (2005); Jouini and Napp (2007) and, with disasters, Chen et al. (2012), with one difference: the pricing weights are stochastic consumption shares, which the fixed-weight consensus of Jouini and Napp (2007)—built on exogenous diffusion beliefs—cannot absorb.

and the equilibrium consumption shares are

$$s_t^i := \frac{c_t^i}{C_t} = \frac{\zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}}{\Xi_t}, \qquad \int_0^1 s_t^i \, di = 1. \tag{7}$$

The kernel's belief factor is the multiplier-weighted $\frac{1}{\alpha}$-power mean Ξ_t^α of the densities, which is not *the share-weighted arithmetic mean $\langle M\rangle_t := \int_0^1 s_t^i M_t^i \, di$ (the two coincide, up to a price-irrelevant constant, only when the cross-section is degenerate—all M_t^i equal). Under the symmetric allocation (Assumption 4(iv)) $\zeta_i \equiv \zeta$ cancels from (7) and rescales (6) by a constant, so $s_t^i = (M_t^i)^{1/\alpha} / \int_0^1 (M_t^j)^{1/\alpha} dj$.*

Sophisticated in the rule, priced by naive conditions. The valuation primitive is sophistication with respect to the rule: each investor evaluates plans under the density built from her realized revision process, her own future revisions included, so the martingale method applies without qualification and her wealth at every revision is the equilibrium continuation value by definition. None of that sophistication is load-bearing for prices. On every inter-revision interval the realized path satisfies the naive Euler equation of the currently held model treated as permanent, and those conditions, concatenated along the realized re-adoptions, support the state-price density for every claim; what naivety alone does not pin is the consumption *level* at each re-solve, which sophistication supplies—and at logarithmic utility not even that gap opens (Lemma A.1, Appendix A). No result below rests on an investor hedging her own future revisions.

The kernel splits into an endowment term $e^{-\delta t} C_t^{-\alpha}$—the representative-agent kernel of the true Lucas economy—times a belief factor Ξ_t^α that the lazy beliefs contribute. The level Ξ_t is a *fixed-weight* (Negishi) power mean of the densities M_t^i; it moves not because the weights move but because the densities do, which is what makes the priced belief object time-varying even though every primitive is fixed. The state-dependent consumption shares s_t^i are the marginal pricing weights of its *local dynamics*—they govern the disaster jump ratio and the calm drift, not the level. Three statistics must be kept distinct: the head-count mean held model $\int \phi_i \, m_t$, the wealth-weighted mean $\int \phi_i s_t^i \, di$, and the priced power mean Ξ_t^α; the kernel prices the last, and the first two are descriptive shadows. The proof (Appendix A) is the complete-market construction: the individual static optimum (martingale method, Cox and Huang, 1989), market clearing, and strict concavity for uniqueness.

E. The disaster-fear factor

Markets are complete, so there is a unique equivalent martingale measure $\mathbb{Q}$, with density process $Z_t = \xi_t \exp(\int_0^t r_s\, ds)/\xi_0$ (r the equilibrium short rate). Since the bank account is continuous, Z and ξ share the same jump ratio, and the risk-neutral disaster law is fixed by that jump ratio alone—the diffusive change of measure does not touch the size law.

Definition 2 (Risk-neutral disaster law, fear factor, and the investor-fears wedge): *Let, at a disaster of size x occurring at t,*

$$g_t(x) := \frac{\Xi_t}{\Xi_{t-}}\bigg|_{J=x} = \int_0^1 s^i_{t-}\, r_i(x;t)^{1/\alpha}\, di = \int_0^1 s^i_{t-} \Big(\tfrac{\lambda_i(t-) f_{\theta_i(t-)}(x)}{\lambda^\star f(x)}\Big)^{1/\alpha} di, \tag{8}$$

the share-weighted aggregator whose α-th power $g_t(x)^\alpha$ is the $(1/\alpha)$-power mean of the held models' disaster likelihood ratios (the formulas below use g_t^α). The risk-neutral disaster compensator *is $\nu^{\mathbb{Q}}(dt,dx) = Y_t(x)\,\nu(dt,dx)$ with $Y_t(x) = e^{\alpha x} g_t(x)^\alpha$, i.e.*

$$\nu^{\mathbb{Q}}(dt,dx) = \lambda^\star\, e^{\alpha x}\, g_t(x)^\alpha\, dt\, F_J(dx). \tag{9}$$

As a predictable compensator $\nu^{\mathbb{Q}}$ is evaluated at the left limit: g_t is built from the pre-jump shares s^i_{t-} above, so $g_t, Y_t, U_t, f^{\mathbb{Q}}_t$ enter jump compensators through their predictable $t-$ versions, while S_t (and g_t, U_t) without qualification denote the current post-jump state. At a disaster epoch the post-jump values are the reweighted-and-relocated ones; (10)–(11) hold at every t under this convention, compensators taking the $t-$ versions. For instance, at a disaster time t, $g_t(x)$ inside a compensator means the predictable $g_{t-}(x)$, while g_t in the spot law (10) is the post-jump aggregate. Writing $\mathbb{E}_f[h] := \int_{\mathcal{J}} h(y)\, f(y)\, dy$ for expectation under the physical size law, its size *density, intensity, and tail are*

$$f_t^{\mathbb{Q}}(x) = \frac{e^{\alpha x} g_t(x)^\alpha f(x)}{\mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]}, \qquad \lambda_t^{\mathbb{Q}} = \lambda^\star\, \mathbb{E}_f\big[e^{\alpha J} g_t(J)^\alpha\big], \qquad \bar{U}_t(x) := \int_x^{\bar{J}} f_t^{\mathbb{Q}}(y)\, dy. \tag{10}$$

The disaster-fear factor *is the risk-neutral disaster intensity* $U_t := \lambda_t^{\mathbb{Q}}$*; the* frictionless benchmark *is its correct-belief value* $U^{\mathrm{RA}} := \lambda^\star\, \mathbb{E}_f[e^{\alpha J}]$ *(the value at* $g_t \equiv 1$*); and the* investor-fears wedge *is the gap from the physical rate,*

$$U_t - \lambda^\star = \lambda^\star\big(\mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha] - 1\big).$$

Finally the belief multiplier *is the normalized fear factor,*

$$\kappa_t := \frac{U_t}{U^{\mathrm{RA}}} = \frac{\mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]}{\mathbb{E}_f[e^{\alpha J}]} \in (0, \infty),$$

the $e^{\alpha J}$*-weighted mean of* g_t^α *relative to its correct-belief value.*

Built from the pre-jump shares s_{t-}^i (8), the tilt $g_t(\cdot)$ is $\mathcal{F}_{t-}$-measurable—the predictable integrand of every disaster compensator below—and carries no left-limit subscript.

Proposition 2 (The disaster-fear bridge): *Under Assumptions 1–4, in the equilibrium of Proposition 1:*

(i) (tilted law) *the risk-neutral disaster law* (9)–(10) *is the physical size density* f *tilted by the risk-aversion factor* $e^{\alpha x}$ *and the belief factor* $g_t(x)^\alpha$*, with fear factor* $U_t = \lambda^\star\, \mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]$*;*

(ii) (the bridge) *the fear factor factorizes into a preference benchmark* U^{RA} *and a belief multiplier* κ_t*,*

$$U_t \;=\; U^{\mathrm{RA}}\, \kappa_t, \qquad U^{\mathrm{RA}} = \lambda^\star\, \mathbb{E}_f[e^{\alpha J}], \qquad \kappa_t = \frac{\mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]}{\mathbb{E}_f[e^{\alpha J}]}; \tag{11}$$

(iii) (belief-driven variation) *holding* $(\alpha, \lambda^\star, F_J)$ *fixed,* U^{RA} *is constant, so*

$$d \log U_t = d \log \kappa_t, \tag{12}$$

and $f_t^{\mathbb{Q}}$, $\bar{U}_t$, *and* U_t *are functionals of the belief aggregate* g_t (8) *alone.*

U_t loads on the belief state g_t and on no diffusive quantity—the exact zero volatility loading the AFT separation reports; the mapping to the measured BT and AFT objects is in Appendix B.A.

Index options trade across *maturities*, and a disaster-tail claim of maturity τ prices the expected *path* of the risk-neutral disaster law over $[t, t+\tau]$, not the instant alone—the *forward fear factor*

$$F_t(\tau) := \mathbb{E}\big[U_{t+\tau} \mid \mathcal{F}_t\big] = U^{\mathrm{RA}}\, \mathbb{E}\big[\kappa_{t+\tau} \mid \mathcal{F}_t\big], \qquad \tau \geq 0, \tag{13}$$

the expected risk-neutral disaster intensity at horizon τ, whose origin is the bridge $F_t(0) = U_t = U^{\mathrm{RA}}\kappa_t$. (A maturity-$\tau$ option prices the *risk-neutral* forward intensity $\mathbb{E}^{\mathbb{Q}}[U_{t+\tau} \mid \mathcal{F}_t]$, the option-implied forward tail risk of the affine surface literature (Andersen et al., 2015b); the physical $F_t(\tau)$ equals it at the origin and departs for $\tau > 0$ only by the forward-intensity risk premium, and the survival records enter both, so the term-structure slope's dependence on Ψ_t below is invariant to the measure.) Away from the origin, with U^{RA} fixed, the curve moves with the expected forward multiplier $\mathbb{E}[\kappa_{t+\tau} \mid \mathcal{F}_t]$—the path the cross-section is expected to trace, the belief *dynamics*—so the surface relates to beliefs at every maturity, not only the shortest. Its slope carries the survival *records* Ψ_t of the standing models, which S_t does not encode; the precise statement, and the proof that this record-dependence is absent for every Markov benchmark, is Theorem 1, once the dynamics of Section 4 are in place.

We work throughout in the *intensity channel*: the size law is common ($f_{\theta_i} = f$), so $r_i = \lambda_i/\lambda^\star$ is x-free, the aggregate is the scalar $G_t = \mathbb{E}_{S_t}[(\lambda/\lambda^\star)^{1/\alpha}]$ (Internet Appendix Corollary **??**), and the multiplier is $U_t/U^{\mathrm{RA}} = G_t^\alpha$. This is the object the disaster-option literature inverts, and the setting of every quantitative result. Disagreement over the size index θ is a transparent extension—the aggregator g_t then carries the full joint object and the whole shape $f_t^{\mathbb{Q}}$ moves with the cross-section—which we develop in closed form in the Internet Appendix; the few statements that invoke it are flagged.

Two scope limits hold throughout: the deep disaster tail, not the frequent small jumps of the near-the-money surface; and a multi-year contraction idealized, as in Wachter (2013); Gabaix (2012), by an instantaneous jump.

The factorization's invariance to the Pareto weights, and the robustness of the wedge's sign, range, and stationary magnitude to that allocation, are collected in the Internet Appendix; Remark IA.2 delimits which belief processes move the factor and which make it flat.

The factorization prices *any* cross-section of held beliefs: the spot option-implied tail identifies $\lambda^\star \kappa_t$, a wealth-weighted aggregate of current beliefs. Whether that aggregate prices claims across maturities depends on whether it determines its own transition law. Section 4 shows that it does not.

4. Belief dynamics and the continuation state

We turn the revision rule into a law of motion for the wealth-weighted cross-section of held models. The section states three objects—the model each investor holds and tests, the hazard at which she abandons it, and the law of motion they induce on the cross-section—and then develops their implications. We introduce only what those statements require; the single-agent foundation of the hazard is the companion paper's result (Teguia, 2026), and the deadline structure, the common-destination lemma, and the proofs are in Internet Appendix **??**.

A. The test object, the revision hazard, and the law of motion

An investor holds one disaster law $\phi = (\lambda, \theta)$ and tests it against her *since-adoption record* $\mathcal{R}_t^i$ (4)—the count n_t^i and realized sizes accumulated since she last revised, taken at regime age τ_t^i. She keeps ϕ while that record stays unsurprising under it and abandons it for the menu's best alternative at a surprise; the re-selection target is the mode (5) of the public second-order posterior ρ_t (3), common across investors. Two facts about this object organize what follows, stated here and established in Internet Appendix **??**. First, the disaster count is informative about the held *intensity*, so the test fires not only at a disaster but also in a long calm, when too few disasters have arrived to be consistent with the held λ_i. Second, the destination factorizes: because both the likelihood and the common prior separate across λ and θ, the joint re-selection mode splits into an intensity mode and a size mode taken separately (Lemma IA.1(iii)).

The kernel (6) and the bridge (11) aggregate held models by consumption share, so the object the law of motion governs is their wealth-weighted distribution. Three objects now appear together and are worth keeping distinct: the pricing state S_t, the wealth-weighted cross-section any belief model would carry; the public posterior ρ_t, the common destination a reviser re-selects to; and the record-resolved state Ψ_t, the cross-section together with the survival records—the coordinate that governs the forward surface and that S_t does not encode.

Definition 3 (The pricing state): *The* pricing state S_t *is the wealth-weighted law of held models: the probability measure on $\Phi = \Lambda \times \Theta$ with $\int_\Phi h\, dS_t = \int_0^1 h\big(\phi_i(t)\big)\, s_t^i\, di$ for every bounded h. Write $\bar{\lambda}_t^S := \int_\Phi \lambda\, dS_t = \int_0^1 s_t^i \lambda_i(t)\, di$ for the wealth-weighted mean held intensity.*

Wealth here is type-attached: each share s_t^i evolves along investor i's own path, and S_t is the pushforward of those shares—never a weight field on Φ. The projection is well posed because relative returns are functionals of the held model alone: the calm replicator drift $(\bar{\lambda}_t^S - \lambda_i)/\alpha$ and the disaster reweight $r(\phi_i, x)^{1/\alpha}$ (Internet Appendix Proposition **??**), so the tilt and every selection covariance below are defined on the held-model coordinate without reference to a state field.

Definition 4 (The record-resolved state and the post-disaster map): *The* record-resolved state Ψ_t *is the wealth-weighted law of the full per-investor state: the probability measure on* $\Phi \times \mathcal{R} \times [0, \infty) \times (0, 1)$ *with*

$$\int h \, d\Psi_t = \int_0^1 h\big(\phi_i(t),\, \mathcal{R}_t^i,\, \tau_t^i,\, \varepsilon^i\big)\, s_t^i \, di$$

for every bounded h*, where* $(\mathcal{R}_t^i, \tau_t^i)$ *are the since-adoption record and regime age of Definition 1, completed by the public posterior* ρ_t*—a deterministic functional of the full public record, common across investors—whose mode is the re-selection destination* $\phi_t^* = \mathfrak{m}(\rho_t)$*. Its held-model marginal is the pricing state,*

$$S_t = \pi_\phi(\Psi_t), \qquad \pi_\phi : (\phi, \mathcal{R}, \tau, \varepsilon) \mapsto \phi.$$

At a disaster of size x*, the* post-disaster state Ψ_t^x *is obtained from* Ψ_{t-} *by, in order, (a) the share reweight* $s_{t-}^i \mapsto s_{t-}^i r_i(x;t)^{1/\alpha}/g_t(x)$ *(the tilt of* (8)*); (b) the record update* $\mathcal{R}^i \mapsto \mathcal{R}^i \oplus x$*; (c) the test re-evaluation on the augmented record (Definition 1, Step 1; propensity* (15) *below); and (d) for every triggered rejector the reset of the since-adoption record and regime age,* $(\mathcal{R}^i, \tau^i) \mapsto (\varnothing, 0)$*, with the held model relocated to the common destination* ϕ_t^* (5) *when it differs—a self-re-adoption resets the window and moves nothing.*

Formally, Ψ_t is a finite Borel measure on the per-investor state space $\Sigma := \Phi \times \mathsf{R} \times \mathbb{R}_+ \times (0, 1)$ of tuples $(\phi, \mathcal{R}, \tau, \varepsilon)$, with $\mathsf{R} := \bigsqcup_{n \geq 0} \mathcal{J}^n$ the space of finite size records—a standard Borel space as a countable disjoint union of compacts; Σ carries the product σ-algebra, $S_t = \pi_\phi \Psi_t$ is the pushforward under the first coordinate, and every integral against Ψ_t runs over Σ; S_t and Ψ_t are thus random elements of standard Borel spaces of measures, so conditioning on them (Proposition 5) is well posed.

The public posterior is carried alongside Ψ_t because it accumulates on the entire public history and is therefore *not* recovered from the since-adoption records inside Ψ_t. The instantaneous risk-neutral objects $\bar{\lambda}_t^S$, $g_t(\cdot)$, U_t, and $f_t^{\mathbb{Q}}$ see only the marginal S_t; the forward-looking valuation of Section 5 loads on the full pair, and we write $v(\Psi_t)$ and $\mathcal{A}_\Psi$ for $v(\Psi_t, \rho_t)$ and $\mathcal{A}_{\Psi,\rho}$, displaying ρ_t only where it does analytic work—the calm boundary map of Lemma B.1 and the post-disaster relocation (d) above. Markov, valuation, and residual statements in Ψ_t are read on this pair throughout.

The rate at which an investor abandons her model is the *revision hazard*, the first passage of the since-adoption deviance to the tolerance boundary; its calm expression: with a clean record, rejection arrives on a schedule, the *calm deadline*

$$\Delta^\star(\phi, \varepsilon) \;=\; \inf\bigl\{\Delta > 0 :\; \bar{\varepsilon}\bigl((0, \varnothing); \phi, \Delta\bigr) \le \varepsilon\bigr\} \;<\; \infty, \tag{14}$$

the first calm length at which the zero-count record itself becomes surprising, finite for every (ϕ, ε) (proved with the hazard in Internet Appendix **??**). At a disaster the record gains a point and the test re-evaluates: an investor at $(\phi, \tau, \mathcal{R}, \varepsilon)$ is *triggered* by a size-x disaster iff $D_{\phi,\tau}(\mathcal{R} \oplus x) \ge r_{\phi,\tau}(\varepsilon)$, so her trigger propensity—the physical arrival rate of her rejection—is

$$\lambda^\star\, \pi(\phi, \tau, \mathcal{R}; \varepsilon), \qquad \pi(\phi, \tau, \mathcal{R}; \varepsilon) = \int_{\mathcal{J}} \mathbf{1}\bigl\{D_{\phi,\tau}(\mathcal{R} \oplus x) \ge r_{\phi,\tau}(\varepsilon)\bigr\}\, f(x)\, dx. \tag{15}$$

Internet Appendix **??** develops the full hazard, and the deadline of Definition 1 is its calm expression.

Between revisions the cross-section drifts by *selection* (Internet Appendix Proposition **??**); the economics is in Internet Appendix **??**.

Proposition 3 (Law of motion of the pricing state): *In the equilibrium of Proposition 1, for every bounded h on Φ the pricing state evolves by:*

(i) (selection drift — continuous, calm)

$$\frac{d}{dt}\langle h, S_t\rangle\Big|_{\text{sel}} = \tfrac{1}{\alpha} \int_{\Phi} \bigl(\bar{\lambda}^S_t - \lambda\bigr)\, h(\phi)\, S_t(d\phi);$$

(ii) (calm switch flow — calm) *the staggered calm-deadline revisers leave their held models and arrive at the common destination* $\phi_t^* = \mathfrak{m}(\rho_t)$*, the same for every investor revising at* t *(Internet Appendix Lemma* **??***); aggregated, they contribute*

$$\frac{d}{dt}\langle h, S_t\rangle\Big|_{\mathrm{sw}} = \int_\Phi \big[h(\phi_t^*) - h(\phi)\big]\,\Phi_t^{\mathrm{sw}}(d\phi),$$

with Φ_t^{sw} *the calm switch flow* (**??**) *of Internet Appendix Lemma* **??***—finite-variation, and continuous where the calm-deadline law is atomless; on a flat deadline level set of positive* G*-mass it carries a predictable finite-variation jump—the deadline dates are announced, adoption dates and types being* $\mathcal{F}_t$*-measurable coordinates of the state—, which by Lemma B.1 leaves every traded price continuous; disaster-triggered switches are* not *part of this term and enter only at (iii);*

(iii) (disaster jump — at disaster epochs) *at a disaster of size* x *the state jumps in two steps—the marginal of the post-disaster map* Ψ_t^x *of Definition 4. First the shares reweight (Internet Appendix Proposition* **??***(ii)),*

$$\widetilde{S}_t^x(d\phi) = \frac{r(\phi,x)^{1/\alpha}}{g_t(x)}\,S_{t-}(d\phi), \qquad r(\phi,x) = \frac{\lambda f_\theta(x)}{\lambda^\star f(x)}; \tag{16}$$

then every triggered *switcher relocates to the common destination* ϕ_t^* *with record and age reset,* $(\mathcal{R},\tau) \to (\varnothing, 0)$*; the final* S_t *is the held-model marginal of the relocated record-resolved state* Ψ_t^x*—the relocation acts on* Ψ*, not on* Φ*. The triggered set is evaluated on the* record-resolved *state* Ψ_{t-}*, not on* Φ *so* S_t *is not a partition of* S_{t-} *over* Φ*, and* (16) *equals* S_t *only when the triggered set is empty. Both steps are a single jump at the disaster epoch, not a continuous rate; the continuous calm drift is the sum of (i) and (ii).*

Under size disagreement the disaster tilt loads on S_t *through the aggregator*

$$x \longmapsto g_t(x) = \int_\Phi \Big(\tfrac{\lambda}{\lambda^\star}\Big)^{1/\alpha}\Big(\tfrac{f_\theta(x)}{f(x)}\Big)^{1/\alpha} S_t(d\phi), \tag{17}$$

which by the spanning condition (Assumption 2(v)) no finite collection of linear moments of S_t determines; the instantaneous risk-neutral law $(U_t, f_t^{\mathbb{Q}})$ is therefore infinite-dimensional in S_t.[7]

The triggered set is genuinely record-level: investors sharing a held model ϕ may carry different $(\mathcal{R}, \tau, \varepsilon)$ and so need not trigger together.

The derivation—the deadline structure, the finiteness of the switch flow, and the price-continuity of its predictable jumps—is in Internet Appendix **??**.

B. What the law of motion implies

The three forces give the wedge its *fast-in, slow-out* shape. Only the disaster jump is synchronized—all triggered switchers land on the single common mode ϕ_t^* (Internet Appendix Lemma **??**), so the repricing is a coordinated wave rather than idiosyncratic noise—and it moves $\bar{\lambda}_t^S$ up at one instant through the reweight (Internet Appendix (**??**))—signed in the intensity channel, and whenever λ is positively associated with $r(\cdot, x)^{1/\alpha}$; the relocation adds under the complacent-switcher condition stated in Proposition 6. Selection and the calm switch flow are its staggered relaxation, both pushing $\bar{\lambda}_t^S$ down in calm (the switch flow under monotone-reselection)—a finite-history regime: the condition holds once the public mode has followed the calm evidence down, while in the locked stationary limit with a fearful projection the switch flow turns non-negative and the relaxation is selection and turnover—, only as fast as the surviving disagreement that selection consumes, spread across the population by the deadline dispersion. These signs are the wave priced in Section 5.C; the equations fix only directions, while the wedge's magnitude waits on the stationary cross-section of Section 6.

The law of motion is not autonomous in the cross-section. The instantaneous tail reduces to the scalar G_t in the baseline (Internet Appendix Corollary **??**), but the switch flow and its

[7] This rules out any finite collection of *linear* moments; the reduction is in fact tighter—no *continuous* finite-dimensional statistic of S_t is sufficient for $(U_t, f_t^{\mathbb{Q}})$ under size disagreement (Proposition IA.1, Internet Appendix **??**, by Brouwer's invariance of domain). The continuity qualifier is essential: in the weak topology $\mathcal{P}(\Phi)$ is Polish, so a measurable bijection furnishes a discontinuous one-dimensional sufficient statistic.

destination draw on the record-resolved Ψ_t, so the *forward* surface depends on Ψ_t even where the instantaneous tail does not—the enlargement the rest of the paper prices, and the coordinate the term-structure test isolates.

A last implication is that the dispersion the cross-section carries is manufactured, not assumed: investors identical in priors, data, preferences, and initial model—heterogeneous only in their surprise tolerances—watching one path *become* a non-degenerate cross-section.

Proposition 4 (Endogenous emergence of disagreement): *Start from the homogeneous configuration of Assumption 3: at $t = 0$ every investor holds the common model ϕ_0, so $S_0 = \delta_{\phi_0}$ and $\mathrm{Var}_{S_0}(\lambda) = 0$. Disagreement is then generated by the rule, not assumed: there is a finite $t > 0$ with $\mathrm{Var}_{S_t}(\lambda) > 0$ in the consumption-share measure S_t of Definition 3 (a positive-G-mass split transfers to positive S_t-mass by the absolute continuity established in the proof) whenever either*

(a) (disaster split) *the first disaster T_k precedes the earliest calm deadline—so the population is still homogeneous at T_k^- ($S_{T_k^-} = \delta_{\phi_0}$) and the post-disaster surprise of the updated one-count record is the operative trigger—the rejecting cohort $\{\varepsilon^i \geq \bar{\varepsilon}(\mathcal{R}_{T_k}; \phi_0, T_k)\}$—the post-disaster surprise of the one-count record over $(0, T_k]$— is a strict subset of positive G-mass of the population, and the common destination differs from the incumbent* in intensity, $\lambda(\phi^{(k)}) \neq \lambda(\phi_0)$ *(a size-only move, λ unchanged, leaves* $\mathrm{Var}_S(\lambda) = 0$*); or*

(b) (staggered calm) *the calm-deadline map $\varepsilon \mapsto \Delta^\star(\phi_0, \varepsilon)$ is non-constant on a positive-G-mass set, the calm intensity mode $\tau \mapsto \lambda^{\mathrm{m}}_\tau$ is strictly decreasing (Internet Appendix Remark* **??***), and the realized path stays disaster-free long enough for two positive-G-mass threshold groups to reach their distinct deadlines (so the staggered revisions fire before a disaster resets the count).*

The section brackets the pricing question. The cross-section S_t prices the spot tail (Section 3); its transition law is not a functional of S_t, because the switch flow and the disaster relocation consult the records and tolerances that Ψ_t carries and S_t does not (Proposition 3). Section 5 derives where those additional coordinates enter equilibrium prices.

5. The pricing state across maturities

Sections 3 and 4 separate the state that prices the spot tail from the state that governs its evolution. This section derives how that separation appears across maturities: the theorem first, then the decomposition that organizes its empirical content, then the time series of the wedge. Every result is state-contingent—valid for any cross-section and free of demographic assumptions; only unconditional averages await the stationary primitive of Internet Appendix Remark **??**.

A. The forward tail

Proposition B.4 says the instantaneous tail is pinned by S_t and the forward-looking objects are not. The next result makes that gap operational. It locates exactly where the record-resolved state enters the *term structure* of the risk-neutral tail, gives the channels their signs, and shows the dependence is a property of the revision rule rather than of the disaster law—absent for any Markov benchmark.

Theorem 1 (Spot sufficiency and forward insufficiency): *Suppose the age structure of Ψ_t admits a density, so the calm flow of held models reaching their deadline $\Delta^\star(\phi, \varepsilon)$ (14) at t is a finite measure Φ_t^{sw} (Proposition 6). The forward fear factor $F_t(s)$ of (13) decomposes as follows. In the intensity channel:*

(i) (Level on S_t.) $F_t(0) = U_t = U^{\mathrm{RA}}\kappa_t$ *is a functional of the pricing state S_t alone (Proposition 2, Proposition B.4).*

(ii) (Slope splits, into two history channels.) *The right-hand slope at the origin decomposes*

into a part measurable in S_t and a part measurable in Ψ_t but not in S_t,

$$\partial_s^+ F_t(0) = \underbrace{U_t\big[-b_{\rm sel}(S_t)\big] + \lambda^\star\,\Theta_{\rm rw}(S_t)}_{\textit{selection drift and disaster reweight: } \mathrm{b}_S(S_t)} + \underbrace{U_t\,c_{\rm sw}\big(\Phi_t^{\rm sw}\big)}_{\textit{calm re-leveling}} + \underbrace{\lambda^\star\,\Theta_{\rm sw}(\Psi_t)}_{\textit{disaster-triggered re-leveling}},$$

where $b_{\rm sel}(S_t) \geq 0$ is the calm selection drift of (20) and $\Theta_{\rm rw}(S_t) := \int_{\mathcal{J}}\big(U(\widetilde{S}_t^x) - U_t\big)f(x)\,dx$ the compensator of the reweight-only *fear jumpboth S_t-measurable—through the relevant selection and reweight moments in general, through the scalar g_t alone in the binary intensity channel—. The two Ψ_t-channels are the staggered revisions: the calm flow*

$$c_{\rm sw}\big(\Phi_t^{\rm sw}\big) = \frac{\alpha\,\mathbb{E}_f\big[e^{\alpha J} g_t(J)^{\alpha-1}\int_\Phi\big(r(\phi_t^*, J)^{1/\alpha} - r(\phi, J)^{1/\alpha}\big)\Phi_t^{\rm sw}(d\phi)\big]}{\mathbb{E}_f\big[e^{\alpha J} g_t(J)^\alpha\big]},$$

non-positive under monotone-reselection; and $\lambda^\star\Theta_{\rm sw}(\Psi_t)$, where $\Theta_{\rm sw}$ is the per-disaster, mark-averaged relocation increment of U—so $\lambda^\star\Theta_{\rm sw}$ is the compensator of the relocations fired at the record-dependent rate $\lambda^\star\,\pi(\phi, \tau, \mathcal{R}; \varepsilon)$ (15), whose destination intensity is non-decreasing in the disaster count (Proposition 6(ii)). Under monotone re-selection $c_{\rm sw} \leq 0$, and under the complacent-switcher condition of Proposition 6 (for $\alpha \geq 1$, or its likelihood-ratio form) also $\Theta_{\rm sw} \geq 0$; the theorem leaves both unsigned. Both $\Phi_t^{\rm sw}$ and the trigger propensity π are functionals of the record-resolved state Ψ_t and not of its marginal S_t.

(iii) (Two histories, one cross-section.) *For record-resolved states Ψ_t, Ψ_t' with a common held-model marginal $\pi_\phi\Psi_t = \pi_\phi\Psi_t' = S_t$, the forward curves share their level and their $\mathrm{b}_S(S_t)$-slope, $F_t(0) = F_t'(0) = U_t$; they differ only through the two staggered-revision channels.* Conditional on no disaster over the horizon—*write $F_t^0(s) := \mathbb{E}[U_{t+s} \mid \mathcal{F}_t,\ N_{(t,t+s]} = 0]$ for the calm-conditional curve—only the calm channel operates and the state whose calm channel is larger in the weighted sense $c_{\rm sw}(\Phi_t'^{\rm sw}) \leq c_{\rm sw}(\Phi_t^{\rm sw})$—for instance a dominating deadline flow with the common destination—carries, under monotone re-selection, the weakly steeper forward-tail decline, $\partial_s^+ F_t^{0\prime}(0) \leq \partial_s^+ F_t^0(0)$; under the theorem's two sign conditions the channels are oppositely signed, and the net is a cross-sectional condition*

on the balance of calm-deadline proximity against disaster-trigger propensity. The wedge $F_t(s) - F_t'(s)$ is zero at $s = 0$ and first order in s; its full-horizon law is Markov in Ψ_t, not autonomous in S_t (Proposition B.2).

(iv) (No Markov-state counterpart.) *Replace lazy revision by any Markov benchmark—a model whose pricing state is Markov in the current cross-section. The slope's dependence on survival history is a property of the revision rule, present for any disaster law and absent across the* belief-sufficient *Markov class—models whose pricing state is Markov in a current belief state;*

Three remarks on the statement. The forward fear factor is the expected future risk-neutral disaster intensity at horizon s, the term structure that the deep out-of-the-money tail traces across maturities. The compensator Θ_{rw} is the change in U produced by the size-x reweight (16) of the aggregator; the sign of c_{sw} under monotone re-selection is scheduled relief (Internet Appendix Proposition **??**: the calm destination intensity is non-increasing), the condition asking that the calm destination sit weakly below every calm switcher's held intensity. The theorem holds in the sophisticated equilibrium and rests on none of its sophistication: the price system is supported interval-by-interval by naive first-order conditions (Lemma A.1)—anticipation of one's own revisions enters consumption levels, never prices—and the record reaches the forward curve through the realized-density representation, not through any investor's foresight. In part (iv), the worked instance is heterogeneous Bayesian learning of the disaster law: each posterior is a sufficient statistic for its own future, so the wealth-weighted cross-section is itself Markov; there is no deadline and no trigger threshold, hence neither Φ_t^{sw} nor Θ_{sw}, and $\partial_s^+ F_t(0)$ is measurable in S_t. The forward curve is then pinned by the current cross-section at every horizon (Internet Appendix IA.9); an affine latent-state model and a representative filtering economy give the same conclusion through their own Markov states. And a benchmark that enlarges its state with ages, clocks, or records is history-bearing by construction and outside the class the comparison targets.

B. The benchmark decomposition and the observable signature

Write S_t for the current pricing cross-section (Definition 3) and Ψ_t for the record-resolved state that also carries the since-adoption records, ages, and revision hazards (Definition 4), with $S_t = \pi_\phi \Psi_t$ its held-model marginal.

Proposition 5 (The benchmark decomposition: current cross-section versus survival history): *In the intensity channel, every integrable equilibrium price functional F_t admits an orthogonal decomposition into a part carried by the current cross-section and a part carried by the survival history,*

$$F_t = \underbrace{B_t(S_t)}_{\text{current cross-section}} + \underbrace{\mathcal{H}_t(\Psi_t)}_{\text{survival-history residual}}, \qquad B_t := \mathbb{E}[F_t \mid S_t], \quad \mathbb{E}[\mathcal{H}_t \mid S_t] = 0, \tag{18}$$

with B_t a functional of S_t and the residual $\mathcal{H}_t = F_t - \mathbb{E}[F_t \mid S_t]$ a functional of Ψ_t orthogonal to S_t.

(i) (Instantaneous prices carry no residual.) *The fear factor U_t and multiplier κ_t, the risk-neutral disaster law, and the short rate r_t are functionals of S_t alone (Proposition B.4); for each, $B_t = F_t$ and $\mathcal{H}_t \equiv 0$.*

(ii) (Forward-looking prices carry a residual.) *The price–dividend ratio $v(\Psi_t)$, the* level *of the disaster premium, and the forward fear factor $F_t(s) = \mathbb{E}[U_{t+s} \mid \mathcal{F}_t]$ at horizons $s > 0$ have $\mathcal{H}_t \not\equiv 0$: the price–dividend ratio and the premium level carry it at date t itself, while the forward fear factor's residual is zero at horizon zero and first order in s, its slope loading on the survival-history coordinates of Ψ_t—the calm deadline-proximity flow Φ_t^{sw} and the disaster-trigger propensity π—and it is absent at horizon zero, entering at first order (Propositions B.2 and 1).*

(iii) (The residual is the lazy-revision signature.) *For any* Markov benchmark*—a model whose pricing state is Markov in the current cross-section, the* belief-sufficient *class of the Introduction—*

every forward object is a functional of that state, so $\mathcal{H}_t \equiv 0$. The residual is therefore present for any *disaster law under lazy revision and* absent *across the belief-sufficient Markov class.*

The projection itself is generic—any conditioning state supports one; the content is where the residual vanishes. A heterogeneous Bayesian economy is one such benchmark, because each investor's posterior is a sufficient statistic for her own future belief and the wealth-weighted cross-section is therefore itself Markov; an affine latent-state specification and a representative filtering economy are others (single-agent case: Internet Appendix IA.9). It is exactly what the term-structure test isolates (Remark 1): conditioning on the current option state removes $B_t(S_t)$ and leaves $\mathcal{H}_t(\Psi_t)$—the records carrying how, and for how long, current beliefs have survived.

The instantaneous price objects—the short rate, the price–dividend ratio, and the disaster premium—are placed in this decomposition and derived in Appendix B: the short rate and the risk-neutral tail carry no residual, while the price–dividend ratio and the premium level inherit the survival-history residual through $v(\Psi_t)$.

Remark 1 (The observable signature, and what it is not). Theorem 1 yields an over-identifying restriction on the left-tail surface, and fixes the maturity at which it bites. The level of the surface—the short-dated risk-neutral tail—is pinned by the current cross-section S_t, so once one conditions on it the age structure is irrelevant at the short end; the dependence on survival history enters the *slope*, zero at the short end and first order beyond it. The test is therefore a residual one: regress a forward-tail slope (a long-minus-short tail measure) on variables proxying the two staggered-revision channels—time since the last disaster, accumulated calm surprise, or an estimate of the distance-to-rejection $\Delta^\star - \tau$—*after* controlling for the current option state (the short-dated tail block: level, volatility, skew), realized volatility, and current skew. The rule predicts a non-zero residual loading beyond the short end—zero at horizon zero, first order in maturity—that any Markov benchmark sets to zero. The record-resolved coordinate is a distribution over records, ages, and tolerances, and the three proxies span its two named channels—calm-deadline proximity and disaster-trigger propensity—not the state; a nonzero residual loading is evidence for the mechanism, while a null rejects the proxied channels rather than the class distinction itself. The class the residual rejects is belief-sufficiency, not every history-bearing alternative: staggered adjusters on exogenous clocks (Abel et al., 2013; Bacchetta and van Wincoop, 2010) also escape it and would also carry a slope residual. Attributing the residual to *testing* rests on two signs no exogenous clock produces: calm deadlines that *shorten in the held intensity*, so the fearful revise down sooner, and a disaster-fired *upward* re-anchoring wave arriving together with the wealth reweight. Internet Appendix Proposition **??** makes the comparison exact: under any record-independent clock the clock's enlarged state (held model, regime age, cohort age) is itself sufficient—record proxies price zero conditional on the clock's own full state—while under testing they price beyond even that state. What is over-identifying is the restriction that the residual exists—absent at the short end, entering at first order beyond it—not a claim that three regressors exhaust it. The combined unconditional slope is unsigned—the calm-deadline and disaster-trigger channels oppose—so the signed form of the test is calm-conditional: conditional on no disaster over the window, greater proximity of the fearful mass to its rejection deadlines lowers the forward slope. An implementation: an international index-options panel—disasters are too rare in any one market—regressing a fixed-moneyness left-tail calendar slope on a flexible short-dated surface block, the record proxies, and their interaction with the short-dated fear level, the interaction licensed by the deadline's shortening in the held intensity; market and time effects absorb common variation, and the deep-out-of-the-money calendar spreads of Remark 3 are the traded object. The traded time-t spread prices the *unconditional* forward object, disaster-trigger states included; the calm-conditional signed prediction is carried by the no-disaster subsample and the interaction terms, not by the raw spread level.

This residual is a finer object than the raw path-dependence of the implied-volatility surface (Guyon and Lekeufack, 2023; Andrès et al., 2024), and the two must not be conflated.

That literature regresses implied volatility on the past return path *without* conditioning on the current pricing state, and finds the explanatory power strongest at short maturities and weakening as maturity grows. In the present model that raw pattern is reproduced through S_t, which itself encodes the recent path—the short-dated tail is the tightest function of S_t, while longer maturities average over future selection and turnover that revert toward the stationary cross-section—and it is *generic*: any Markov benchmark also encodes the past. The rule's signature is the orthogonal residual, history that still moves the forward-tail slope after the current state is removed, which Theorem 1 shows is present under lazy revision and absent across the Markov class. The conditioning set carries an identification assumption: in the binary-menu intensity channel the spot tail level identifies S_t exactly, so short-dated tail controls span the S_t-block; with richer menus those controls are structural filters for S_t rather than exact replacements, and the maintained assumption is that they span the S_t-measurable part of the slope—the interpretation under which residual loading on record proxies falsifies belief-sufficiency rather than omitted current-state components.

Remark 2 (The option-implied forward object). The surface prices the risk-neutral forward tail, $F_t^{\mathbb{Q}}(s)$. Under the kernel of Proposition 1 the $\mathbb{Q}$-compensator of disasters is $\lambda^\star e^{\alpha x} g_t(x)^\alpha f(x)\,dx\,dt$—arrival intensity $\lambda_t^{\mathbb{Q}} = \lambda^\star \mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha] = U_{t-}$, size density the normalized (10)—an S_{t-}-measurable tilt of arrival and size (Jacod and Shiryaev, 2003). Applying the tilt to the generator split behind the slope decomposition of Theorem 1(ii) therefore changes only the S_t-measurable loadings: the calm flow Φ_t^{sw} is untouched and the trigger *set* is unchanged pathwise; the compensator of triggered relocations integrates that set against the tilted mark law, so the numerical propensity changes while its record dependence does not—the slope's two staggered-revision channels remain functionals of Ψ_t beyond S_t under $\mathbb{Q}$ exactly as under $\mathbb{P}$. The residual test of Remark 1 thus applies to the option-implied forward surface itself, not only to its physical forecast.

Corollary 1 (The priced calendar spread): *Fix* $0 < s_1 < s_2$ *and let* $P_t(s_1, s_2) := \mathbb{E}_t^{\mathbb{Q}}\big[e^{-\int_t^{t+s_2} r_u\,du} N_{(t+s_1,t+s_2]}\big]$ *be the discounted risk-neutral forward disaster count—the deep-out-of-the-money calendar tranche.*

Let $P^0_t(s_1, s_2)$ be its calm-conditional version, the same expectation conditional on no disaster through $t + s_1$. For record-resolved states Ψ_t, Ψ'_t with common held-model marginal S_t and common public posterior: $P^0_t - P'^0_t$ is first order in $s_2 - s_1$, equal to a fear-path term driven by the calm channel of Theorem 1(iii) plus the short-rate-path term that channel induces through (B1); both are functionals of Ψ_t and not of S_t, and both vanish for any benchmark whose pricing state is Markov in the current belief cross-section. The unconditional difference $P_t - P'_t$ adds the disaster-triggered channel through the state at $t + s_1$.

The short-rate path enters the spread through the same record channel and vanishes with it across the benchmark class; its calibrated magnitude is reported in Internet Appendix **??**.

The theorem's prediction thus attaches to the priced object itself, not only to its physical forecast; the proof, in Appendix B, couples the two economies as in part (iii) of the theorem.

Remark 3 (From counts to puts). The tranche $P_t(s_1, s_2)$ prices the forward disaster *count*; the traded instruments are deep-out-of-the-money put calendar spreads, which filter arrivals through sizes and discounting (Backus et al., 2011; Andersen et al., 2015b). The step between the two adds one further record coordinate: the valuation jump. To the corollary's order—first order in the tranche width, where a single arrival matters and diffusive exercise over the tranche is negligible at deep-out-of-the-money strikes—a put calendar spread at fixed moneyness is the discounted forward integral of a per-arrival value against the mark-resolved compensator. The compensator factorizes at each instant into the arrival scale U_{u-} and the size law $f^{\mathbb{Q}}_{u-}$ of (10), a functional of g_{u-} and hence S_{u-}-measurable (Proposition 2(iii)); for a put written on the index, the arrival's intrinsic payoff is set by the *price* jump $e^{-x}\psi_{u-}(x)$, $\psi_{u-}(x) = v(\Psi^x_{u-})/v(\Psi_{u-})$, so the value flow is $U_{u-}\, h(S_{u-}; \psi_{u-})$, the moneyness loading carrying the valuation-jump ratio—itself record-dependent (Proposition B.2). The record therefore enters a traded put spread through *two* channels: the calm-conditional path of (U_u, S_u, r_u) that Theorem 1 characterizes—zero at the origin, first order beyond it—and the valuation jump ψ, present already in the level (the price–dividend witness) and moving with the record thereafter. The second channel strengthens the record's imprint, at the cost that h is no longer S-measurable for index puts; it is exactly S-measurable for consumption-linked payoffs, and the binary-channel reading—h constant, the spread a constant multiple of the count tranche—is the frozen-ψ approximation. The residual test of Remark 1 applies to traded put calendar spreads with this reading: the short-dated controls absorb the size loading and the frozen-ψ level, the record's imprint rides the arrival-and-state path together with the valuation-jump channel, and the exact-ψ level component is itself record content, picked up by the level diagnostic of Section 5.D.

The mapping from the model's objects to the measured ones—the BT wedge, the AFT tail factor, the calendar tranche—is completed in Appendix B.A.

C. The time series of the wedge

A disaster is an impulse; this subsection traces the response of the priced objects assembled above. The engine is market selection on the realized path (Sandroni, 2000; Blume and Easley, 2006; Yan, 2008; Kogan et al., 2006, 2017): a disaster vindicates the fearful at a stroke, the ensuing calm vindicates the complacent only gradually, and the result is a fast jump and a slow, staggered relaxation (Andersen et al., 2015b). Throughout we keep the descriptive mean belief $\bar{\lambda}_t^S = \int_\Phi \lambda\, S_t(d\phi)$ distinct from the priced fear factor U_t: the kernel prices the power mean, not the mean (Section 3).

Proposition 6 (Disaster jump and calm erosion): *In the equilibrium of Proposition 1, with* Var_{S_t} *and* Cov_{S_t} *the variance and covariance under* S_t*:*

(i) (Calm erosion — exact, both calm channels.) *Between disasters the wealth-weighted mean belief moves continuously away from the predictable deadline atoms (the flat-level-set case of Internet Appendix Lemma* **??***); the displayed law is the atomless-flow component, by*

selection and the calm switch flow,

$$\frac{d\bar{\lambda}_t^S}{dt} = \underbrace{-\tfrac{1}{\alpha}\,\mathrm{Var}_{S_t}(\lambda)}_{\textit{selection}\,\le 0} + \underbrace{\int_\Phi \big(\lambda(\phi_t^*) - \lambda\big)\,\Phi_t^{\mathrm{sw}}(d\phi)}_{\textit{calm switch flow (signed below)}},$$

and, for any bounded h, $\frac{d}{dt}\langle h, S_t\rangle = -\frac{1}{\alpha}\,\mathrm{Cov}_{S_t}(\lambda, h) + \int_\Phi [h(\phi_t^*) - h(\phi)]\,\Phi_t^{\mathrm{sw}}(d\phi)$. On the disaster-free-since-adoption *part of the switch set the destination intensity is non-increasing (scheduled relief, Internet Appendix Proposition* **??***; the calm MAP of Internet Appendix Lemma* **??** *falls within a calm spell), so $\lambda(\phi_t^*) \le \lambda$ and the switch-flow term is non-positive there. Call $\lambda(\phi_t^*) \le \lambda$ on the full switch set the* monotone-reselection condition*.*

(ii) (Disaster jump — intensity channel.) *At a disaster of size x the cross-section reweights by* (16)*; the reweight alone raises the mean belief by*

$$\bar{\lambda}_t^{\tilde{S}} - \bar{\lambda}_{t-}^S = \frac{\mathrm{Cov}_{S_{t-}}\big(\lambda,\, r(\cdot,x)^{1/\alpha}\big)}{\mathbb{E}_{S_{t-}}[r(\cdot,x)^{1/\alpha}]} \;\ge\; 0, \tag{19}$$

≥ 0 in the intensity channel ($f_\theta \equiv f$, $r(\phi,x) = \lambda/\lambda^\star$) and strict if $\mathrm{Var}_{S_{t-}}(\lambda) > 0$*; the triggered* switchers *(Internet Appendix Lemma* **??***) relocate to the common post-disaster mode ϕ_t^*, whose intensity coordinate is non-decreasing in the disaster count (proved by revealed preference in the appendix). Their contribution to the jump is computed under the* post-reweight *switcher measure (the disaster updates in the order reweight-then-relocate of Definition* 4*), $\Delta\bar{\lambda}_t^S|_{\mathrm{switch}} = \widehat{s}_t^{\mathrm{sw}}\big(\lambda(\phi_t^*) - \widehat{\bar{\lambda}}_t^{\mathrm{sw}}\big)$, where $\widehat{s}_t^{\mathrm{sw}}$ and $\widehat{\bar{\lambda}}_t^{\mathrm{sw}}$ are the switchers' mass and mean intensity* after *the $r(\cdot,x)^{1/\alpha}$ reweight; call $\lambda(\phi_t^*) \ge \widehat{\bar{\lambda}}_t^{\mathrm{sw}}$ the* complacent-switcher condition*.*

(iii) (The priced fear factor inherits the pattern — intensity channel.) *Through the bridge*

relation (12)*, in calm* U_t *drifts down,*

$$\frac{d\log U_t}{dt} = -\,\frac{\mathbb{E}_f\big[e^{\alpha J} g_t(J)^{\alpha-1}\ \mathrm{Cov}_{S_t}\big(\lambda, r(\cdot,J)^{1/\alpha}\big)\big]}{\mathbb{E}_f\big[e^{\alpha J} g_t(J)^{\alpha}\big]}\ \le\ 0, \tag{20}$$

the selection *part of the calm drift, strict if* $\mathrm{Var}_{S_t}(\lambda) > 0$*. On a switch-active instant the calm switch flow of (i) adds a further term proportional to* $\mathbb{E}_f\big[e^{\alpha J} g_t(J)^{\alpha-1}\int_\Phi (r(\phi_t^*, J)^{1/\alpha} - r(\phi, J)^{1/\alpha})\ \Phi_t^{\mathrm{sw}}(d\phi)\big]$*, non-positive under monotone re-selection in the* likelihood-ratio *sense—*$r(\phi_t^*,\cdot) \le r(\phi,\cdot)$ *on the switch set, which in the intensity channel is the condition of (i);* U_t *jumps up at a disaster with the reweighted aggregator—the reweight-only component of (ii); the triggered relocation adds under the complacent-switcher condition there, for* $\alpha \ge 1$ *(in its likelihood-ratio form for all* α*). The signs hold in the intensity channel and, generally, whenever* λ *and* $r(\cdot,x)^{1/\alpha}$ *are positively associated across the cross-section.*

(iv) (Staggered recovery.) *The calm decay runs until the next disaster or the calm deadlines* $\Delta^\star(\phi,\varepsilon)$ (14)*, at which revisers re-select a held model that Internet Appendix Lemma* **??** *makes non-increasing on the disaster-free-since-adoption part of the switch set (or under the monotone-reselection condition)—a predictable step, weakly downward there. By Internet Appendix Proposition* **??** *the deadlines are weakly ordered in* ε *and differ wherever the deadline map is non-constant, so the steps are staggered across that part of the population over a horizon set by* $\Delta^\star$*,*

The response is fast in (the instantaneous jump of (ii)) and slow out (the drift of (i),(iii) plus staggered deadline steps). In the decomposition of Proposition 5 the fast-in reweight (19) and the selection drift (20) are the cross-section part $B_t(S_t)$, while the calm switch flow Φ_t^{sw}, the disaster-triggered relocation, and the deadline-timed steps $\Delta^\star(\phi,\varepsilon)$ load on the survival-history residual $\mathcal{H}_t(\Psi_t)$—the *centered* channels of Proposition 5, the raw flows net of their S_t-conditional mean. The staggered, record-timed part of the slow-out—switch flow, triggered relocations, deadline steps, net of their S_t-conditional mean—is the residual, while the smooth selection drift is cross-section part, which a Markov-in-cross-section benchmark also carries:

such a benchmark reweights and drifts but has no deadlines and no triggers, so it carries the fast-in and the smooth drift, not the staggered slow-out (Internet Appendix IA.9). The item-by-item discussion is in Internet Appendix ??.

What the record does and does not price—fragility priced while forecasting no fundamental, the dissociation claimed in either channel, the finite-dimensionality claim only under size disagreement—is discussed in Internet Appendix ??.

D. From the theorem to an empirical test

Existing option evidence establishes that the priced tail is time-varying, persistent, economically consequential, and related to the past path (Bollerslev and Todorov, 2011; Andersen et al., 2015a; Guyon and Lekeufack, 2023). None of it distinguishes a persistent Markov fear state from a current state whose transition law also depends on surviving records. The distinguishing test asks whether path-based record proxies forecast the forward-tail slope after a sufficiently rich control for the current option state.

The ingredients are these. The observed object is the risk-neutral forward tail of Remark 2: writing $\widehat{U}_t(\tau)$ for a model-free option-implied tail measure at maturity τ, the outcome is a slope $\Delta_t := \widehat{U}_t(\tau_1) - \widehat{U}_t(\tau_0)$, $\tau_1 > \tau_0$. The control is a rich vector X_t^{spot} of current-surface measurements—the shortest-maturity tail intensity, implied variance, skew, jump variation, short-maturity surface factors—with no claim that a single spot intensity spans S_t; the spanning of Remark 1 is maintained, not primitive. The record proxies Z_t^{rec} enter in a hierarchy: time since the last disaster (observable, the headline proxy), accumulated calm surprise over a window (a model-guided reconstruction from an inferred intensity, not a measurement), and deadline-proximity mass (structural, recoverable only from an estimated belief block—reconstructions that inherit the dark-matter critique the path-read proxy escapes). The restriction is

$$\Delta_t \;=\; f\big(X_t^{\text{spot}}\big) \;+\; \beta^{\top} Z_t^{\text{rec}} \;+\; \varepsilon_t.$$

First, $\beta \neq 0$ rejects belief sufficiency jointly with the maintained spanning of the controls—not, by itself, the mechanism. *Second*, attribution to testing rests on the two signs of Remark 1: at a fixed spot state, greater deadline proximity predicts a more downward slope, and dates shortly after a disaster differ from same-spot-tail dates in a direction set by time since the event. *Finally*, the model's loading is state-dependent rather than globally monotone—the reset wedge changes sign across the fear cycle (Figure 3, panel B)—so a linear coefficient on time since disaster can vanish under the true model, and a null is informative only when the spot controls are demonstrably rich. A within-design falsification disciplines the spanning burden: by Proposition 5(i) the *level* is S_t-measurable, so the same regression run on the spot tail rather than the slope must assign the record proxies no weight—a record loading in the level equation diagnoses control insufficiency, not the mechanism—while an omitted current-state moment biases both equations, and the theory's fingerprint is the contrast, zero at the origin and first order in maturity. In the calibrated intensity channel the omitted-moment exposure is degenerate, the spot tail itself spanning S_t; it lives where the state is infinite-dimensional, the size channel of Proposition 3. Two scope notes close the mapping. The quantitative exhibits of Section 6 isolate the calm-conditional component of the slope, the deadline channel in its cleanest form; traded calendar spreads load on the combined object—deadline and trigger channels under the risk-neutral law—which the theorem covers and the exhibits do not separately quantify. And the testing rule transfers to a crash-frequency clock or an international panel, while the macro-disaster calibration does not automatically transfer with it.

6. The wedge and the option-tail evidence

The event study of Section 5.C and any unconditional moment presuppose a *stationary* cross-section, which the present primitives do not supply: posterior consistency drives an infinitely-lived investor's held model to the Kullback–Leibler point $\phi^\star$ and the dispersion dies out (Internet

Appendix Remark **??**). We add the one primitive that sustains it—demographic turnover—prove stationarity, and cast the quantitative content as comparative statics: physical and preference parameters from the macro-disaster literature, the belief block pinned in Section 6.B to a few stylized option-tail targets.

We restore a nondegenerate stationary cross-section with the device the rare-disaster literature uses for the purpose: a standard *demographic turnover*, perpetual-youth overlapping generations with fair annuities (Yaari, 1965; Blanchard, 1985; Gârleanu and Panageas, 2015). Each investor faces a constant death hazard $\Omega > 0$; newborns arrive at the same rate, drawn from a fixed type law Π over (ε, ϕ_0) with ϕ_0-marginal Π_ϕ, and enter holding their drawn model with a reset test clock while inheriting the common public posterior. Because young cohorts hold ϕ_0 and have not yet revised, turnover continually replenishes the disagreement that selection consumes—not only a stationarity device but a continuing source of disagreement, testing governing persistence and revision while the newborn law supplies the variation. We state the primitive as Assumption **??** (Internet Appendix) and specialize to a finite menu $\Phi = \{\phi_1, \ldots, \phi_K\}$ for the stationary and quantitative exercise; the assumption, the equilibrium it defines, and all proofs are in Internet Appendix **??**.

Two results carry the economy of Sections 3–5 into this stationary world, and the quantitative exercise needs only their conclusions. First, *the kernel keeps its form* (Internet Appendix Proposition **??**): with fair annuities the death hazard cancels from every survivor's Euler equation and enters only through the composition of the living, so $\xi_t = e^{-\delta t} C_t^{-\alpha} \Xi_t^{\alpha}$ holds over the living cross-section, the short rate (B1) is unchanged—level and all—and turnover acts only by adding a mean-reversion of the cross-section toward the newborn law Π_ϕ at rate Ω. Every *instantaneous* pricing relation of Section 5 therefore transfers verbatim, evaluated at the turnover-stationary S_t, while the forward-looking valuation keeps its form with the cross-section dynamics augmented by this mean-reversion. Second, *that cross-section is stationary* (Internet Appendix Proposition **??**): in the limiting regime where the common posterior has concentrated ($\mathrm{MAP}(\rho_t) = \phi^\star$, so the re-selection destination is calendar-invariant), a finite menu and a constant inflow of fresh types make S_t a time-homogeneous functional of a stationary, ergodic disaster clock, with a unique nondegenerate stationary law whose unconditional moments equal long-run time averages.

The quantitative message turns on two summary properties of the stationary cross-section. The stationary *dispersion*—and through Proposition 6 the amplitude of the fear wave—is governed by the newborn joint law Π—birth models *and* tolerances, through the tolerance-dependent revision hazards—and the turnover rate Ω: more dispersed births leave more disagreement, and Ω sets how fast fresh cohorts replace the selected cross-section. Whether the stationary cross-section is *on average complacent*—the level statement we take to the option evidence below—is likewise a parameter condition on Π_ϕ, not a consequence of the primitives.

Sensitivity of the cycle to the belief block—newborn dispersion, demographic turnover, menu coarseness—is developed in Internet Appendix **??**.

A. Parameters and quantitative message

We discipline the physical and preference parameters with the macro-disaster evidence and hold them fixed throughout; the belief block (Π, Φ) is not separately observed, since disaggregated disagreement data on the disaster intensity priced into options do not exist—though firm-level surveys of subjective *macroeconomic* disaster beliefs document exactly the wide dispersion and repeated revision the mechanism produces (Menkhoff, 2025). We use it two complementary ways. The comparative statics below *vary* the belief dispersion to trace the mechanism, using the option numbers only as a reference, not as inputs. The calibration of Section 6.B instead *pins* the belief block to two stylized option-tail targets (Section 6.B); this is an illustrative consistency check—whether one mechanism, its parameters in economically distinct primary roles, can span the option-tail evidence—not a fit to the option surface and not an out-of-sample prediction. It establishes plausibility, not identification: the mechanism's falsifiable content is not the fit but the residual restriction of Theorem 1, which no Markov-in-cross-section benchmark satisfies and which the calibration neither presupposes nor exhausts. No survey data enter.

Block	Parameter	Value	Source
Physical	disaster intensity $\lambda^\star$	0.03/yr	Barro and Ursúa (2008), consumption disasters (10%)
	mean disaster size $\mathbb{E}_f[b]$	0.21	Barro and Ursúa (2008)
	size law f	power law ($\gamma \approx 7.2$, $b_{\min} = 0.10$)	Barro and Jin (2011)
	trend, diffusion (μ_0, σ)	0.025/yr, $0.02/\sqrt{\text{yr}}$	Barro (2006); Barro and Ursúa (2008)
Preferences	risk aversion α	3.5	Barro (2006); Barro and Ursúa (2008) (equity premium)
	discount δ	0.03/yr	standard
Demography	death hazard Ω	0.02/yr (50-yr horizon)	Blanchard (1985); Gârleanu and Panageas (2015)
Beliefs *(free)*	menu $\{\lambda_{lo}, \lambda_{hi}\}$	varied; $\{0.01, 0.06\}$ in the option-evidence calibration	comparative statics vs. Section 6.B
	newborn law Π	varied; 90% complacent in the option-evidence calibration	comparative statics vs. Section 6.B

Table I: Parameters. The physical and preference blocks are disciplined by the macro-disaster literature; the belief block is illustrative and varied. Disaster sizes are consumption contractions for the unlevered tree; for comparison Barro (2006) report $\lambda^\star \approx 0.017$/yr at a 15% GDP threshold with mean size 0.29. At the calibrated physical size law and risk aversion—menu-free objects—the kernel inputs evaluate to $\mathbb{E}_f[e^{\alpha J}] = 2.82$ and the representative-agent fear factor to $U^{\text{RA}} = 0.085$/yr; the newborn belief block draws the model 90/10 complacent/fearful and the tolerance from a grid on $(0.01, 0.90)$ with one third of its mass below 0.12.

Before beliefs enter, risk aversion alone inflates the disaster-state pricing: a single contraction of size b carries the kernel weight $e^{\alpha J}$ with $J = -\log(1-b)$. At $\alpha = 3.5$ this multiplier is 1.4 at $b = 0.10$, 2.3 at the Barro–Ursúa mean $b = 0.21$, 3.3 at $b = 0.29$, and 6.0 at $b = 0.40$; the power-law tail (Barro and Jin, 2011) pushes the f-average $\mathbb{E}_f[e^{\alpha J}]$ above the value at the mean size.[8] The representative-agent fear factor is $U^{\mathrm{RA}} = \lambda^\star\, \mathbb{E}_f[e^{\alpha J}]$: physical disasters of frequency 3% are priced as if several times more frequent. This effect is the standard rare-disaster channel; our contribution is what the belief cross-section does to it.

The belief cross-section moves the risk-neutral tail several-fold. Under the binary menu, the fear factor is $U_t = G_t^\alpha\, U^{\mathrm{RA}}$ with $G_t = S_t^{hi}(\lambda_{hi}/\lambda^\star)^{1/\alpha} + (1 - S_t^{hi})(\lambda_{lo}/\lambda^\star)^{1/\alpha}$ (Internet Appendix (**??**)). With the illustrative $\{\lambda_{lo}, \lambda_{hi}\} = \{0.01, 0.06\}$ around $\lambda^\star = 0.03$,[9] the ratio U_t/U^{RA} runs from 0.33 when the cross-section is fully complacent ($S^{hi} = 0$) to 2.0 when fully fearful ($S^{hi} = 1$), passing through 0.91 at $S^{hi} = \frac{1}{2}$: across the menu's admissible support the belief multiplier ranges six-fold, *above and below* the representative-agent benchmark, with no change in any physical parameter; the stationary law of S^{hi} occupies a narrower endogenous subset, reported below. The long-run re-selection destination under this misspecified menu is the *fearful* point (footnote above): calm complacency in the stationary economy is sustained by the complacent newborn inflow and by selection, not by learning toward a low intensity—the regime discussion of Internet Appendix **??**. The state-insufficiency theorem does not depend on this configuration; the calibrated cycle does, and a menu containing the truth is a different economy, not reported here. Within the cycle the configuration's reach is one place deep: trigger synchronization and deadline staggering are the rule's, operating under any destination—the locked-destination regime of Internet Appendix **??** confirms both—while the level of calm complacency and its replenishment belong to the menu-and-turnover pair.

The same multiplier, written for a general cross-section in the intensity channel (Internet Appendix Corollary **??**), is $U_t/U^{\mathrm{RA}} = G_t^\alpha$ with $G_t = \mathbb{E}_{S_t}[(\lambda/\lambda^\star)^{1/\alpha}]$ (Internet Appendix (**??**)). Replacing the two-point menu by a continuum of held intensities filling $[\lambda_{lo}, \lambda_{hi}] = [0.01, 0.06]$

[8] The power-law parameters describe the Barro and Jin (2011) law of the gross contraction $z = 1/(1-b) = e^J$ ($\gamma \approx 7.2$, $z_{\min} = 1/(1-0.10) = 1.11$), not a Pareto on b. Under it $\mathbb{E}_f[b] = 1 - \gamma/[(\gamma+1)z_{\min}] = 0.21$ (the Barro–Ursúa mean) and $\mathbb{E}_f[e^{\alpha J}] = \mathbb{E}[z^\alpha] = \gamma z_{\min}^\alpha/(\gamma - \alpha) \approx 2.82$ are jointly delivered; a Pareto on b would imply a smaller mean (about 0.12) and is not what is used.

[9] The two-point menu is deliberately misspecified: $\lambda^\star = 0.03$ lies outside $\{0.01, 0.06\}$. No held point is correct; the consistency results, which identify the destination as the menu's Kullback–Leibler projection, apply through that projection: $T(0.01) = 0.0130 > T(0.06) = 0.0092$, so the projection is λ_{hi} and the locked destination is the fearful point. The multiplier passage in the text is the pricing identity evaluated pointwise in the fearful share s; under the stationary law of s it inherits that law, with no learning input beyond it.

and varying the *shape* of the cross-section gives, at $\alpha = 3.5$, $\lambda^\star = 0.03$, the illustrative values:

cross-section of held λ	mean belief $\bar{\lambda}^S$	multiplier U_t/U^{RA}
$\mathrm{Beta}(2,5)$, complacent-skewed	0.024	0.78
$\mathrm{Uniform}[0.01, 0.06]$	0.035	1.09
$\mathrm{Beta}(2,2)$, symmetric	0.035	1.12
$\mathrm{Beta}(5,2)$, fearful-skewed	0.046	1.51
$\mathrm{Beta}(8,2)$, very fearful	0.050	1.66

The continuum fills the interval $[\lambda_{lo}/\lambda^\star, \lambda_{hi}/\lambda^\star] = [0.33, 2.0]$ between the binary extremes; in the intensity channel it does not raise that ceiling (the ceiling-lift of Remark IA.1 is a size-channel statement), but it makes the multiplier a smooth, dispersion-sensitive functional of the cross-section. These Beta cross-sections are illustrative instantaneous evaluations of the general intensity-channel multiplier (Internet Appendix Corollary **??**); they are not stationary laws of the turnover economy, whose stationarity theorem (Internet Appendix Proposition **??**) is stated under the finite menu. That sensitivity is exact: since $\alpha > 1$ and $y \mapsto y^{1/\alpha}$ is strictly concave, Jensen's inequality applied to $\lambda/\lambda^\star$ under S_t gives

$$\frac{U_t}{U^{\mathrm{RA}}} = G_t^\alpha = \left(\mathbb{E}_{S_t}[(\lambda/\lambda^\star)^{1/\alpha}]\right)^\alpha \;\leq\; \frac{\bar{\lambda}_t^S}{\lambda^\star}, \tag{21}$$

with equality iff the held intensity is S_t-a.s. constant. The consequences—calm complacency at a correct-on-average cross-section, the shortfall growing in dispersion and risk aversion—are developed in Internet Appendix **??**.

Selection is slow relative to disasters. The selection part of the calm erosion of Proposition 6(i), $d\bar{\lambda}^S/dt|_{\mathrm{sel}} = -\frac{1}{\alpha}\,\mathrm{Var}_S(\lambda)$, reallocates wealth at the rate $\frac{1}{\alpha}(\lambda_{hi} - \lambda_{lo}) \approx 0.014/\mathrm{yr}$—the decay of the high-to-low wealth *odds*, an e-folding of about 70 years (a heuristic scale; the full relaxation also carries the calm switch flow and demographic turnover). Disasters recur every $1/\lambda^\star \approx 33$ years and cohorts turn over every $1/\Omega \approx 50$ years. Selection alone is too slow to eliminate fearful wealth between typical disasters; the actual stationary relaxation additionally reflects deadline-driven switching and turnover. The stationary cross-section therefore carries a persistently positive, fluctuating fear share, and with it a non-degenerate, time-varying disaster premium—the model counterpart of the chronic equity premium and the slow post-crisis reversion of the tail factor (Andersen et al., 2015b).

The magnitudes reconcile two option facts that look opposed: the same multiplier $U_t/U^{\mathrm{RA}} = G_t^{\alpha}$—below one in calm, above one after disasters—carries the economy between the modest option-implied disasters Backus et al. (2011) document and the large, spiking tail premium of Bollerslev and Todorov (2011) and Andersen et al. (2015b): the literatures measure it at different points of the fear cycle, against different anchors, and are not in conflict. The countercyclical, below-macro prediction is about the inverted intensity $\lambda^{\star}\kappa_t$, not the raw U_t, which carries the risk-aversion premium $\mathbb{E}_f[e^{\alpha J}]$ on top.

B. Confronting the option-tail evidence

The closed-form objects above line up, one for one, with three model-free option-tail literatures, so the model can be taken to that evidence without estimating the belief process. We hold the physical block at its macro-disaster calibration (Table I), which delivers the representative-agent benchmark $U^{\mathrm{RA}} = \lambda^{\star}\mathbb{E}_f[e^{\alpha J}] = 0.085/\mathrm{yr}$ in closed form, and pin the belief block—newborn dispersion and disaster-trigger mass—to two stylized targets drawn from that evidence—the crisis jump, matched inside the documented band, and the calm moderation $\kappa_t < 1$, matched in deep calm though not at the literature's point level (first row of the table); the post-crisis half-life is generated by the deadline and turnover blocks and reported as an *achieved* moment. Matching the level and crisis facts places the calibrated economy in an empirically relevant region; it does not identify lazy revision—only the forward-record residual of Section 5.D does. The parameters have economically distinct primary roles: the newborn model shares chiefly govern the calm composition of κ_t, and the low-threshold mass chiefly governs the immediate disaster-trigger response. The pinned block, in the intensity channel: the binary menu $\{0.01, 0.06\}$, a newborn law 90% complacent, and tolerances on a two-block grid on $(0.01, 0.90)$ with a third of the mass below 0.12—the disaster-trigger dial. Along realized disaster records the resulting deterministic transport reaches crisis spikes of $U_t/\lambda^\star \approx 4.5$–$4.8$, inside the 3.2–4.8 crisis range, with κ_t traversing the benchmark over the cycle and dipping to 0.53–0.63 in deep calm. The intensity rows of Table II are closed-form implications of the fear factor $U_t = U^{\mathrm{RA}}\kappa_t$ (11) and the finite-menu reduction of Internet Appendix Corollary **??**; the premium rows are *structural* consequences—the dominance of the disaster over the diffusive premium and the loading of the premium on U_t rather than σ—that additionally involve the price–dividend jump ratio ψ_t, a Feynman–Kac functional of the full record-resolved state rather than a closed form (Proposition B.2), and follow from its qualitative behavior rather than a solved level, as the row notes below record.

The three anchors map directly (Table II): Backus et al. (2011)'s modest option-implied disaster law is $\kappa_t < 1$ in deep calm, the option-implied intensity falling below U^{RA} as the quiet record accumulates; Bollerslev and Todorov (2011)'s time-varying fear premium is $U_t > \lambda^\star$ on average, with the disaster premium nearly the whole equity premium since the diffusive $\alpha\sigma^2$ is negligible; and Andersen et al. (2015b,a)'s unspanned, persistent, premium-forecasting tail factor is κ_t, which jumps at disasters, erodes slowly in calm (20), and enters the premium through U_t while σ is held fixed.

Two rows are structural rather than calibrated. The diffusive premium is mechanically small and fixed: $\alpha\sigma^2 = 0.14\%/\mathrm{yr}$ at the Table I diffusion, so by (B27) the disaster premium DRP_t accounts for essentially all of the equity premium wherever it is sizable relative to that part, as at the calibration; the *share* itself is not belief-free, since the level of DRP_t moves with ψ_t and the full pricing state. The predict-the-premium-not-volatility pattern is an identity in this model: DRP_t loads on U_t (B28) while σ is constant, so the fear factor forecasts the premium with a nonzero loading and volatility forecasts it with exactly zero—the Andersen et al. (2015b) separation, an exact zero loading in the model (σ constant) rather than a fitted coefficient. Two of the remaining rows—the crisis jump and the calm moderation—are the calibrated targets: the jump matched inside the documented band, the moderation matched in deep calm though not at the literature's point level; the half-life is an *achieved* moment, generated by the deadline and turnover blocks and reported for comparison. The row is measured on the computed record, peak to halfway back to the pre-disaster level. Its channel decomposition at this calibration: the selection drift (20) contributes about 0.5% per year—a log rate of decline of U_t, as are the rates below—(on its own a half-life above a century—the reason no free relaxation-rate parameter exists to tune), the demographic pull roughly 0.7–1.6%, and the staggered deadline flow the remainder—operative once the public mode has returned to the complacent point, about seven years after a disaster on this record: immediately after a disaster the finite-history mode sits at the fearful point, so calm deadlines re-anchor upward or re-adopt, and the monotone-reselection condition holds only once the mode has followed the calm evidence down. Persistence is generated by the revision architecture, not by a free relaxation-rate parameter. All are stylized: published estimates vary with sample and method, so we match orders of magnitude, not point values.

This is a disciplined confrontation, not a structural estimation: we show that the closed-form fear factor is consistent with the published model-free moments, not that the belief process is identified by them. The one feature the belief block cannot repair is the safe rate. Under power utility the safe rate sits at the high CRRA *level* (Section 6.A)—its realized dynamics are modest, but the level is the standard risk-free-rate artifact at $\alpha = 3.5$; that is a preference limitation, removed only by an Epstein–Zin extension, and not something the belief calibration can absorb.

Remark 4 (What power utility buys, and what is specific to it). Power utility does two things at once. It makes the factorization $U_t = U^{\mathrm{RA}}\kappa_t$ (Proposition 2) *exact*, and it makes the level/slope dissociation *clean*: the kernel's disaster jump ratio is $e^{\alpha x} g_t(x)^{\alpha}$, which depends only on the current cross-section, so the instantaneous tail is a functional of S_t alone and within the implied disaster tail the record Ψ_t enters only the forward surface—the price–dividend level, the valuation jump ψ_t, and the disaster-premium level do load on Ψ_t (Proposition B.2); the dissociation is level-versus-slope of the tail— (Proposition B.4, Theorem 1). Both are specific to the class: an Epstein–Zin kernel carries the valuation-ratio jump $\psi_t = v(\Psi_t^x)/v(\Psi_{t-})$, a functional of the full Ψ_t (Proposition B.2), so under it the record would move the instantaneous tail as well and the level would no longer pin S_t. The two implications cut opposite ways. The record's *economic* role is preference-robust—indeed stronger under Epstein–Zin, where it moves the level as well as the slope—so the priced market-memory coordinate (Internet Appendix **??**) is not an artifact of power utility. But the clean *separation* that makes the restriction sharp—level on S_t, residual slope on Ψ_t, isolated by conditioning on the current tail—is CRRA-specific, and a recursive-utility version would demand a different empirical design, since conditioning on the current tail would then partial out part of the record itself. Power utility is thus the tractable window through which the record's imprint is cleanly visible, at the cost of a safe-rate *level* at the high-α artifact of the class, removable only by that same extension.

The procyclical *physical* tail of Gormsen and Jensen (2025) runs the other way and is reconciled with the countercyclical priced tail in Internet Appendix **??**.

C. The forward slope across records

The theorem's forward content has so far been carried by a constructed pair; this subsection reports it across records the calibrated economy generates along realized histories: finite-history economies from the calibrated initial condition, the public-mode path alive—the cycle of this section, not draws from the locked-destination invariant law (that regime is reported in Internet Appendix **??**). We simulate one hundred fifty such sixty-year histories—disasters at the physical rate, the rule, selection, turnover—and at each terminal date evaluate the calm-conditional forward tail $F^0_{t^*}(s)/U_{t^*}$ together with a reset-record counterfactual that holds the cross-section of held models, wealth shares, tolerances, and the public posterior fixed and resets only the since-adoption records. The sharpest near-match is the theorem's construction realized: among pairs within 0.02 of one another in the spot fearful share S^{hi}—which spans the spot state in the binary channel—the largest twelve-year slope gap separates two histories 0.001 apart: 5.8 percentage points, twenty-two percent of the steeper curve's twelve-year decline, at two and twenty years since the last disaster (Figure 3, panel A: the pair in bold against the twenty-nine records of a narrow comparison window, $S^{hi}_{t^*} \in [0.55, 0.70)$, inside the post-disaster band). Across the 103 near-matched pairs the slope gaps have median 2.1, ninetieth percentile 5.0, and maximum 5.8 percentage points.

The typical magnitude is smaller. Across sampled histories the record wedge—$[F^{\text{reset}} - F^{\text{actual}}]/F^{\text{reset}}$ at twelve years, the reset-record counterfactual against the path as realized—concentrates within a few percent: median 1.1 in deep calm, where little fearful mass remains to re-anchor; either sign after recent disasters, median -1.3; and the constructed 16.2 of Figure 2 (starred, panel B) is a calibrated ceiling in a moderate-fear transient that fixed-date sampling of these finite histories rarely delivers, not a typical sampled magnitude. The negative post-disaster wedges have a timing interpretation: reset clocks, all synchronized at t^*, deliver their first revision wave after the public mode has relaxed to the complacent point, while the realized clocks spend part of theirs in the earlier window where re-selection returns the fearful model—the same opposition that leaves the unconditional slope unsigned (Internet Appendix **??**).

What the record organizes is the residual. Within each pre-specified region we regress the twelve-year slope on S^{hi}, linearly; adding time since the last disaster raises the model-implied R^2, computed from the sampled histories, from 0.64 to 0.92 after disasters ($n = 75$), and moving it from its within-band twenty-fifth to seventy-fifth percentile—six to twenty-five years—raises the forward-to-spot ratio by 5.3 percentage points at fixed S^{hi}; adding deadline-proximity mass raises the calm-band R^2 from 0.87 to 0.99 ($n = 73$), the proximity sign negative as Remark 1 predicts (Table III; Internet Appendix **??**). At the surface, and conditional on the short-dated block throughout: long-maturity tail claims contain variation beyond what the short end implies—related to the excess sensitivity of long-duration claims emphasized by Giglio and Kelly (2018)—and the pace at which a post-crisis inversion of the tail term structure unwinds is timed by the crisis's age, a coordinate a belief-sufficient account leaves redundant given the level. The record wedge is usually a few percent; conditional on the spot state, it organizes a large share of the calm-conditional forward slope's variation.

7. Conclusion

The spot disaster tail recovers a wealth-weighted belief aggregate rather than the physical disaster law, and that aggregate is not the complete pricing state. The kernel prices the *power* mean of beliefs, which sits below the arithmetic mean, so a cross-section correct on average prices the option tail below the representative-agent benchmark in calm and reprices it upward in coordinated jumps when a disaster proves the fearful right (Propositions 2, 6); demographic turnover makes the cycle stationary (Internet Appendix Propositions **??**, **??**). The surface therefore measures beliefs over a fixed quantity of disaster risk—reconciling the modest average tails of Backus et al. (2011) with the spiking risk-neutral tails of Bollerslev and Todorov (2011) and Andersen et al. (2015b) as one belief cycle—and leaves the over-identifying mark: because a held belief is not a sufficient statistic for its own revision, the forward surface is priced off the records under which current beliefs have survived, and record proxies forecast it even after the contemporaneous option state is controlled for—a residual any belief-sufficient model sets to zero (Proposition 5), the deadline and crash signs of Remark 1 separating testing from other history-bearing accounts. The spot surface reads what the market fears; the term structure, how long it has feared it.

The limitations are explicit. CRRA preferences fix the safe-rate *level* at the high-α artifact of the class and let the belief cycle move it over a wide range (quantified in Internet Appendix **??**), and the same choice confines the record to the forward surface, so the clean level/slope separation behind the residual test is itself specific to power utility; the disaster-dynamics signs are intensity-channel results under a positive-association condition we flag rather than assume; and the belief block is calibrated against option-tail moments rather than estimated, since disagreement data on the disaster intensity priced into options do not exist—though surveyed disaster beliefs display the dispersion and revision the mechanism produces (Menkhoff, 2025). The empirical implication is accordingly conditional: after a sufficiently rich control for the current option state, path-based proxies for surviving records should forecast the disaster-tail slope; a nonzero residual rejects belief sufficiency jointly with that spanning condition, and the deadline and disaster-trigger signs separate testing from other history-bearing accounts (Section 5.D). The gap between the priced tail and the *physical* return tail (Gormsen and Jensen, 2025) remains the far side of the same wedge, posed here and not settled. Across maturities, options reveal not only the disaster beliefs investors currently price but the revision process attached to them. What the wedge between the option surface and the macroeconomic record measures, throughout, is not a mismeasurement of either, but the price of disaster fear.

Appendices A and B below collect the proofs of the equilibrium kernel, the bridge, and the pricing results. Appendices C–K, constituting the Internet Appendix to this paper (available online), collect the remaining proofs, the demographic and stationary constructions, the quantitative development, and the discussion sections; appendix lettering and result numbering are continuous across the two documents, and cross-references resolve in both directions.

Notation

Primitives (the fixed physical law)

C_t	aggregate consumption, the Lucas-tree dividend
$\mu_0,\ \sigma$	drift and diffusion of log consumption
$\lambda^\star$	physical disaster intensity
$F_J,\ f$	physical disaster-size law and density on $\mathcal{J} = [\underline{J}, \overline{J}]$
$T_k,\ J_k$	time and size of the kth disaster

Beliefs and the revision rule

$\phi = (\lambda, \theta)$	a held model: disaster intensity λ, size index θ
$\Phi = \Lambda \times \Theta$	model space; F_θ the belief size law
$\theta^\star$	pseudo-true size index (minimizes $\mathrm{KL}(F_J \| F_\theta)$)
$\varepsilon^i,\ G$	investor i's surprise tolerance; its cross-section law
$\rho_t,\ \phi_t^*$	common public posterior; re-selection destination $\phi_t^* = \mathfrak{m}(\rho_t)$
$\mathcal{R}_t^i,\ \tau_t^i$	investor i's since-adoption record and regime age

Preferences and the pricing kernel

α	coefficient of relative risk aversion
s_t^i	investor i's wealth (consumption) share
Ξ_t	belief aggregator (enters the kernel $\xi_t = e^{-\delta t} C_t^{-\alpha} \Xi_t^\alpha$)
$g_t(x)$	share-weighted aggregator; g_t^α is the $(1/\alpha)$-power mean of the disaster likelihood ratios

The fear factor and the belief multiplier

$\lambda_t^{\mathbb{Q}},\ f_t^{\mathbb{Q}}$	risk-neutral disaster intensity and size density
U_t	disaster-fear factor, $U_t = \lambda_t^{\mathbb{Q}}$
U^{RA}	frictionless benchmark $\lambda^\star \, \mathbb{E}_f[e^{\alpha J}]$ (correct-belief value, at $g_t \equiv 1$)
κ_t	belief multiplier U_t/U^{RA}; the bridge is $U_t = U^{\mathrm{RA}} \kappa_t$

Appendix A. Proofs for the equilibrium kernel and the disaster-fear bridge

The kernel comes from individual optimization and market clearing. Each investor solves the standard complete-market problem, maximizing $\mathbb{E}_i\int_0^\infty e^{-\delta t}u(c_t^i)\,dt$ subject to the budget $\mathbb{E}\int_0^\infty \xi_t c_t^i\,dt \le W_0^i$—the objective taken under her subjective belief, the budget under the common state-price density ξ. Writing the subjective expectation as $\mathbb{E}_i[\,\cdot\,] = \mathbb{E}[M_t^i\,\cdot\,]$, the first-order condition $e^{-\delta t}u'(c_t^i)M_t^i = \zeta_i\xi_t$ equates discounted marginal utility under her belief to the shadow price of wealth, so with $u'(c) = c^{-\alpha}$ optimal consumption is $c_t^i \propto (M_t^i)^{1/\alpha}$: the reciprocal of risk aversion applied to her belief density. Market clearing $\int_0^1 c_t^i\,di = C_t$ then aggregates these individual demands across the cross-section into the single belief factor $\Xi_t = \int_0^1 \zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}\,di$ that the kernel carries—a *power mean* of the held densities, not their average. The next result records this equilibrium and its uniqueness.

The belief density is a positive martingale. We verify the claim after Definition A.1. Under $\mathbb{P}$ the disaster measure μ has $(\mathcal{F}_t,\mathbb{P})$-compensator $\nu(dt,dx) = \lambda^\star\,dt\,F_J(dx)$ (2). The held model $\phi_i(t-) = (\lambda_i(t-),\theta_i(t-))$ is predictable (left-continuous and piecewise constant, Definition 1) and proposes the compensator $\nu^i(dt,dx) = \lambda_i(t-)\,dt\,F_{\theta_i(t-)}(dx)$. The process (A1) is the Radon–Nikodym density of the marked-point-process law with compensator ν^i against that with ν (Brémaud, 1981, VI, Theorems 2–4), which is a positive $(\mathcal{F}_t,\mathbb{P})$-local martingale solving $dM_t^i = M_{t-}^i\int_{\mathcal{J}}\big(r_i(x;t)-1\big)\,(\mu-\nu)(dt,dx)$. We check the two hypotheses of Brémaud (1981, VI, T4) under which it is a true martingale with $\mathbb{E}[M_t^i] = 1$. *(a) Bounded intensity.* $\Lambda = [\underline{\lambda},\overline{\lambda}] \subset (0,\infty)$ (Assumption 2), so $\int_0^t \lambda_i(s-)\,ds \le \overline{\lambda}\,t < \infty$. *(b) Bounded likelihood ratio.* The map $(x,\phi)\mapsto r(x;\phi) = \lambda f_\theta(x)/(\lambda^\star f(x))$ is continuous and strictly positive on the compact $\mathcal{J}\times\Phi$ (full support, Assumption 2(i)), hence $0 < \underline{r} := \min_{\mathcal{J}\times\Phi} r \le r_i(x;t) \le \overline{r} := \max_{\mathcal{J}\times\Phi} r < \infty$. Conditions (a)–(b) give $\mathbb{E}[\int_0^t\int_{\mathcal{J}}|r_i(x;s)-1|\,\nu(ds,dx)] \le (\overline{r}+1)\,\overline{\lambda}\,t < \infty$, so the stochastic integral is a true martingale and $\mathbb{E}[M_t^i] = 1$ (Brémaud, 1981, VI, T4). Positivity is immediate from $r_i > 0$. □

Independent increments and the realized-density representation. The product form $\prod_k r_i(J_k; T_k)$ is the independent-increment factorization of the Poisson record—each disaster contributing its own likelihood ratio, with no cross-disaster interaction *given the held-model path*; revisions, which change ϕ_i, are the only channel coupling disasters. The same structure makes the since-adoption segment likelihood the *conditional* Radon–Nikodym increment given the pre-adoption past, so testing the held model on its since-adoption record (Definition 1) is exact as conditional testing—a property of the Poisson record's independent increments, not of general locally equivalent laws. The ratio M_T^i / M_t^i is built from the predictable held-model *process* $\phi_i(s-)$ over $(t, T]$, the models adopted at the agent's own future revisions included, which makes M^i a $\mathbb{P}$-martingale and feeds the Negishi construction of Internet Appendix **??**. The valuation primitive is *sophisticated with respect to the rule*: the investor evaluates consumption plans under this realized-rule density, her own future revisions included, so her problem is time-consistent, the martingale method applies without qualification, and her wealth at each revision is the equilibrium continuation value by definition rather than by convention. Behaviorally she remains the paper's tester—at every date she prices the instantaneous objects under her currently held model, dogmatically, and at a rejection she re-solves under the re-selected model. Lemma A.1 establishes that the sophistication is not load-bearing for prices: the realized path satisfies the naive Euler equation on every inter-revision interval, and the concatenation of those first-order conditions along the realized re-adoptions is exactly the realized-density construction—the state-price density a market of naive revisers supports, with no anticipation of one's own revisions entering the price system. Levels are where the sophistication has bite: a fully self-financing *naive* sequence would price its frozen tail rather than the equilibrium one (the two coincide only at logarithmic utility), while the sophisticated investor's wealth at each revision is the equilibrium continuation value—exactly the level the naive re-solve needs. On the naive reading, the departures from dynamic consistency are confined to the rejection epochs, each gated by the tolerance (Weinstein, 2017)—a per-shift control, with no lifetime bound claimed. The two readings agree on every *instantaneous* object—the short rate and the fear factor depend on the current model $\phi_i(t-)$ alone—and the forward-looking objects are equilibrium prices, pinned by ξ, not by any single investor's first-person forward valuation.

Definition A.1 (Belief density): *Investor i's* belief density *is the Radon–Nikodym derivative of her held-model law $\mathbb{P}^{\phi_i}$ against the physical law $\mathbb{P}$ on $\mathcal{F}_t$,*

$$M_t^i = \frac{d\mathbb{P}^{\phi_i}}{d\mathbb{P}}\bigg|_{\mathcal{F}_t} = \exp\Bigl(-\int_0^t \bigl(\lambda_i(s-) - \lambda^\star\bigr)\,ds\Bigr) \prod_{k:\,T_k \le t} r_i(J_k;T_k), \qquad r_i(x;t) := \frac{\lambda_i(t-)\, f_{\theta_i(t-)}(x)}{\lambda^\star\, f(x)}, \tag{A1}$$

with $M_0^i = 1$: the intensity-wedge drift times the per-disaster likelihood ratios r_i. Investor i's conditional expectation is $\mathbb{E}_t^i[X] = \mathbb{E}_t[(M_T^i/M_t^i)X]$.

Remark A.1 (The allocation is selected, not derived). Under complete markets with heterogeneous beliefs, equal initial wealth and a common multiplier cannot both hold. The CRRA optimum is $c_t^i = \zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}(e^{-\delta t}/\xi_t)^{1/\alpha}$, so ζ_i is pinned by the budget $\zeta_i^{-1/\alpha}\,\mathbb{E}\bigl[\int_0^\infty \xi_t^{1-1/\alpha} e^{-\delta t/\alpha}(M_t^i)^{1/\alpha}\,dt\bigr] = W_0^i$, whose value depends on the law of $\{M_t^i\}$, hence on the surprise threshold ε^i. Common wealth W_0^i forces type-dependent ζ_i, while common ζ_i forces type-dependent wealth—a genuine fork. We take the second branch: ε indexes a revision *rule*, not a preference or endowment we wish to price, so equal weights isolate the belief channel from Pareto-weight dispersion. This setup is the standard perpetual-youth reduced form (Yaari, 1965; Blanchard, 1985; Gârleanu and Panageas, 2015); the general heterogeneous-ζ kernel of Proposition 1 remains available.

Definition A.2 (Competitive equilibrium): *A* competitive equilibrium *is a state-price density $(\xi_t)_{t\ge 0}$ and a consumption allocation (c_t^i) such that (i) each c^i maximizes $\mathbb{E}^i\bigl[\int_0^\infty e^{-\delta t}u(c_t^i)\,dt\bigr]$ subject to the static budget $\mathbb{E}\bigl[\int_0^\infty \xi_t c_t^i\,dt\bigr] \le \mathbb{E}\bigl[\int_0^\infty \xi_t e_t^i\,dt\bigr]$, the expectation under the physical law $\mathbb{P}$ and e^i investor i's endowment—the supporting endowment that implements the selected Pareto weights of Remark A.1 (under heterogeneous beliefs a common multiplier implies type-dependent e^i), not a separate primitive of the economy; and (ii) the goods market clears, $\int_0^1 c_t^i\,di = C_t$ for* *every t.*

Proof of Proposition 1. Step 1 (the aggregator is finite and positive). Fix a path with $N_t < \infty$ disasters on $[0, t]$ (a.s., Assumption 1). By (A1) and the bounds $\underline{r} \le r_i \le \overline{r}$ and $\lambda_i \in [\underline{\lambda}, \overline{\lambda}]$,

$$e^{-(\overline{\lambda}-\lambda^\star)t}\,\underline{r}^{\,N_t} \;\le\; M_t^i \;\le\; e^{(\lambda^\star-\underline{\lambda})t}\,\overline{r}^{\,N_t} \qquad \text{for every } i,$$

so M_t^i is bounded above and below by path-dependent constants uniform in i. Hence $(M_t^i)^{1/\alpha}$ is bounded uniformly in i, and with $\int_0^1 \zeta_i^{-1/\alpha}\,di < \infty$ (Assumption 4),

$$0 \;<\; \Xi_t = \int_0^1 \zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}\,di \;\le\; \left(e^{(\lambda^\star-\underline{\lambda})t}\overline{r}^{\,N_t}\right)^{1/\alpha}\int_0^1 \zeta_i^{-1/\alpha}\,di \;<\; \infty.$$

Step 2 (individual optimum). Markets are complete (Assumption 4), so investor i's dynamic problem is the static budget problem under the unique state-price density ξ (Cox and Huang, 1989; Karatzas et al., 1990). Writing the objective under $\mathbb{P}$ via $\mathbb{E}^i[\,\cdot\,] = \mathbb{E}[M_t^i\,\cdot\,]$, the Lagrangian $\mathbb{E}\big[\int_0^\infty \big(e^{-\delta t}u(c_t^i)M_t^i - \zeta_i\xi_t c_t^i\big)dt\big]$ is concave in c^i (as u is concave), so its pointwise (in (t,ω)) stationarity condition is necessary and sufficient:

$$e^{-\delta t}u'(c_t^i)\,M_t^i = \zeta_i\,\xi_t. \tag{A2}$$

With $u'(c) = c^{-\alpha}$, (A2) inverts to

$$c_t^i = \zeta_i^{-1/\alpha}\big(\xi_t e^{\delta t}\big)^{-1/\alpha}\big(M_t^i\big)^{1/\alpha}. \tag{A3}$$

Step 3 (market clearing fixes the kernel). Integrate (A3) over i and impose $\int_0^1 c_t^i\,di = C_t$.

The factor $(\xi_t e^{\delta t})^{-1/\alpha}$ is constant in i, so

$$
\begin{aligned}
\left(\xi_t e^{\delta t}\right)^{-1/\alpha} \int_0^1 \zeta_i^{-1/\alpha} (M_t^i)^{1/\alpha}\, di &= C_t \\
\left(\xi_t e^{\delta t}\right)^{-1/\alpha} \Xi_t &= C_t \\
\left(\xi_t e^{\delta t}\right)^{-1/\alpha} &= C_t/\Xi_t \qquad \text{(A4)} \\
\xi_t e^{\delta t} &= \left(C_t/\Xi_t\right)^{-\alpha} = C_t^{-\alpha}\Xi_t^{\alpha} \\
\xi_t &= e^{-\delta t} C_t^{-\alpha}\Xi_t^{\alpha}, \qquad \text{(A5)}
\end{aligned}
$$

which is (6); $\Xi_t \in (0, \infty)$ by Step 1, so ξ_t is well-defined and positive.

Step 4 (shares). Divide (A3) by C_t and substitute $(\xi_t e^{\delta t})^{-1/\alpha} = C_t/\Xi_t$ from (A4):

$$
s_t^i = \frac{c_t^i}{C_t} = \frac{\zeta_i^{-1/\alpha}(C_t/\Xi_t)(M_t^i)^{1/\alpha}}{C_t} = \frac{\zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}}{\Xi_t}, \qquad \text{(A6)}
$$

which is (7); $\int_0^1 s_t^i\, di = \Xi_t/\Xi_t = 1$ and $s_t^i > 0$, so s_t is a probability density in i.

Step 5 (existence and uniqueness). Equations (A5) and (A6) construct (ξ, c^i); by Step 2 each c^i is the (unique) individual optimum given ξ, and by Step 3 the goods market clears, so (ξ, c^i) is a competitive equilibrium (Definition A.2). Uniqueness: at any equilibrium the FOC (A2) and clearing hold, and Steps 2–3 derive (6)–(7) from them with no freedom left once the ζ_i are fixed. Uniqueness is stated *given* the selected weights ζ_i: the supporting endowments of the equilibrium definition are chosen to implement them, so each budget holds by construction rather than pinning the weights. In a fixed-endowment economy the weights would in turn be determined by the strictly monotone Negishi map from Pareto weights to excess demands (Cox and Huang, 1989; Karatzas et al., 1990); nothing below requires that step. Under the symmetric allocation $\zeta_i \equiv \zeta = 1$ (Assumption 4(iv); a weight normalization, not equal initial wealth, which under heterogeneous beliefs would make ζ_i type-dependent, Remark A.1) ζ cancels from (A6).

Step 6 (the power mean is not the arithmetic mean). From (A6), $\zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha} = s_t^i \Xi_t$, so $M_t^i = \zeta_i\,(s_t^i\Xi_t)^\alpha$. Substituting into the share-weighted arithmetic mean $\langle M\rangle_t = \int_0^1 s_t^i M_t^i\,di$,

$$\langle M\rangle_t = \int_0^1 s_t^i\,\zeta_i\,(s_t^i\Xi_t)^\alpha\,di = \Xi_t^\alpha \int_0^1 \zeta_i\,(s_t^i)^{1+\alpha}\,di.$$

Under the equal-Pareto-weight selection $\zeta_i \equiv \zeta$ (the case carried downstream), the shares are price-irrelevantly ζ-free and

$$\langle M\rangle_t = \zeta\,\Xi_t^\alpha \int_0^1 (s_t^i)^{1+\alpha}\,di \;\geq\; \zeta\,\Xi_t^\alpha,$$

the inequality by Jensen applied to the convex map $x \mapsto x^{1+\alpha}$ $(\alpha > 0)$ and the probability measure di on $[0,1]$, since $\int_0^1 (s_t^i)^{1+\alpha} di \geq (\int_0^1 s_t^i\,di)^{1+\alpha} = 1$, with equality iff s_t^i is constant in i—equivalently iff the M_t^i are equal across the cross-section. Hence the priced power-mean factor Ξ_t^α and the descriptive arithmetic mean $\langle M\rangle_t$ coincide (up to the price-irrelevant constant ζ) only at cross-sectional degeneracy, and otherwise differ, the arithmetic mean strictly larger; the kernel prices the former. □

Lemma A.1 (Naive Euler support): *In the equilibrium of Proposition 1, fix investor i with revision times $\tau_0 = 0 < \tau_1 < \cdots$, and on $[\tau_k, \tau_{k+1})$ write $M_{t\to s}^{(k)}$ for the likelihood ratio of the model $\phi_i(\tau_k)$ held there, treated as permanent.* (i) (Naive Euler equations.) *For every $\tau_k \leq t \leq s < \tau_{k+1}$ the realized path satisfies*

$$e^{-\delta(s-t)}\left(\frac{c_{i,s}}{c_{i,t}}\right)^{-\alpha} M_{t\to s}^{(k)} \;=\; \frac{\xi_s}{\xi_t}, \tag{A7}$$

the first-order condition of the problem posed at τ_k under $\phi_i(\tau_k)$ held forever; and no consumption–portfolio deviation supported in $[\tau_k, \tau_{k+1})$ and self-financing at the market prices ξ raises the objective $\mathbb{E}_{\tau_k}^{(k)} \int e^{-\delta(s-\tau_k)} u(c_s)\,ds$. (ii) (SDF support.) *Consequently ξ is determined interval-by-interval by naive first-order conditions, $\xi_s/\xi_{\tau_k} = e^{-\delta(s-\tau_k)}(c_{i,s}/c_{i,\tau_k})^{-\alpha} M_{\tau_k\to s}^{(k)}$ on $[\tau_k, \tau_{k+1})$, and their concatenation along the realized re-adoptions is the realized-density representation of Appendix A; every claim is priced by this ξ, with no investor's forecast of her own revisions entering* *the price system.* (iii) (Levels.) *Exact implementation of consumption* levels *by self-financing naive re-solving—each plan chosen from the market value of the previous plan's continuation—holds if and only if at every τ_k the market value of the frozen-model continuation equals that of the equilibrium continuation; at logarithmic utility ($\alpha = 1$) the equality holds identically, and for $\alpha \neq 1$ it is an additional restriction that equilibrium does not impose. In equilibrium, wealth at each revision equals the equilibrium continuation value—exactly the level at which the naive re-solve reproduces the path.*

The lemma locates exactly what naivety delivers and what it does not. Between revisions the investor is behaviorally naive: her realized path satisfies the Euler equation of the model she currently holds treated as permanent, and no deviation she can finance and evaluate improves it—so the state-price density, and with it every option and every calendar spread, is supported by naive first-order conditions concatenated along the realized re-adoptions. What naivety alone does not pin is the consumption *level* at each re-solve: a fully self-financing naive sequence prices its own frozen tail, not the equilibrium one, and the two values coincide only at the logarithmic knife-edge. In the equilibrium the gap does not open: the sophisticated investor's wealth at each revision *is* the equilibrium continuation value, exactly the level the naive re-solve needs—the naive reading is exact in valuation and local optimality, the sophisticated one in levels—and she carries no hedge against her own revision; none enters the price system.

Proof of Lemma A.1. (i) With $c_{i,s} = \zeta_i^{-1/\alpha}(M_s^i)^{1/\alpha} C_s/\Xi_s$ (Proposition 1) and $\xi_s = e^{-\delta s} C_s^{-\alpha}\Xi_s^{\alpha}$, the held model is constant on $[\tau_k, \tau_{k+1})$, so $M_s^i/M_t^i = M_{t\to s}^{(k)}$ there and

$$e^{-\delta(s-t)}\left(\frac{c_{i,s}}{c_{i,t}}\right)^{-\alpha} M_{t\to s}^{(k)} = e^{-\delta(s-t)}\,\frac{M_t^i}{M_s^i}\left(\frac{C_s}{C_t}\right)^{-\alpha}\left(\frac{\Xi_s}{\Xi_t}\right)^{\alpha} M_{t\to s}^{(k)} = \frac{\xi_s}{\xi_t},$$

which is (A7); equivalently the path satisfies the frozen problem's first-order condition $e^{-\delta(s-\tau_k)} c_{i,s}^{-\alpha} M_{\tau_k\to s}^{(k)} = \eta_k\,\xi_s/\xi_{\tau_k}$ with multiplier $\eta_k = \zeta_i\,\xi_{\tau_k}\,e^{\delta\tau_k}/M_{\tau_k}^i$. For local optimality, let Δc be supported in $[\tau_k, \tau_{k+1})$ and self-financing at market prices, $\mathbb{E}_{\tau_k}\int \xi_s\,\Delta c_s\,ds = 0$ (state prices are objective;

beliefs enter the objective, not the budget). Its first-order effect on the frozen objective is $\mathbb{E}^{(k)}_{\tau_k}\int e^{-\delta(s-\tau_k)}u'(c_{i,s})\,\Delta c_s\,ds = \eta_k\,\mathbb{E}_{\tau_k}\int(\xi_s/\xi_{\tau_k})\,\Delta c_s\,ds = 0$, since on the support of Δc the frozen density coincides with the realized one; the second-order effect is strictly negative by strict concavity. *(ii)* Rearranging (A7) recovers ξ_s/ξ_{τ_k} on each interval from the naive first-order condition alone; multiplying the interval densities $M^{(k)}$ along the realized re-adoptions is precisely the realized-density construction of Appendix A, and dynamic completeness (Assumption 4) prices every claim from this ξ by absence of arbitrage. *(iii)* By strict concavity, the naive optimum from wealth W coincides with the equilibrium path on $[\tau_k, \tau_{k+1})$ iff its budget-pinned multiplier equals η_k, i.e. iff W equals the market value of the plan with multiplier η_k—the frozen continuation $c^{(k)}_s = \zeta_i^{-1/\alpha}\big(M^i_{\tau_k}M^{(k)}_{\tau_k\to s}\big)^{1/\alpha}C_s/\Xi_s$. Self-financing hands the τ_{k+1} problem the value of plan (k)'s tail, while coincidence there requires the value of the *equilibrium* tail; the two are

$$\zeta_i^{-1/\alpha}\,\xi^{-1}_{\tau_{k+1}}\,\mathbb{E}_{\tau_{k+1}}\int_{\tau_{k+1}}^{\infty} e^{-\delta s}\,C_s^{1-\alpha}\,\Xi_s^{\alpha-1}\,D_s^{1/\alpha}\,ds, \qquad D_s \in \big\{\,M^i_{\tau_k}M^{(k)}_{\tau_k\to s},\ M^i_s\,\big\},$$

frozen versus realized density. At $\alpha = 1$ the integrand is $e^{-\delta s}D_s$, and both densities are $\mathbb{P}$-martingales with the common value $M^i_{\tau_{k+1}}$ at τ_{k+1} (M^i is continuous at a calm re-adoption, its new segment starting at ratio one), so the two integrals coincide identically. For $\alpha \neq 1$ the factor $C_s^{1-\alpha}\Xi_s^{\alpha-1}$ weights the two densities by state, and their equality is an additional restriction that equilibrium does not impose; in equilibrium the investor's wealth at each revision is the equilibrium continuation value, which supplies exactly the required level. At a disaster-triggered revision the same argument runs from the post-jump state; transversality (T) makes every displayed integral finite, and markets clear at ξ because the allocation is the equilibrium one throughout. □

Proof of Proposition 2. Step 1 (the kernel's jump ratio). At a disaster at t of size $x = J_k$, the endowment falls by $C_t = C_{t-}e^{-x}$ (Assumption 1), so

$$\frac{C_t^{-\alpha}}{C_{t-}^{-\alpha}} = \big(e^{-x}\big)^{-\alpha} = e^{\alpha x}.$$

Each density jumps by its likelihood ratio, $M_t^i = M_{t-}^i\, r_i(x;t)$ (A1), so $(M_t^i)^{1/\alpha} = (M_{t-}^i)^{1/\alpha} r_i(x;t)^{1/\alpha}$, and by (6) and (A6),

$$\begin{aligned}\frac{\Xi_t}{\Xi_{t-}} &= \frac{1}{\Xi_{t-}}\int_0^1 \zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}\,di \\ &= \frac{1}{\Xi_{t-}}\int_0^1 \zeta_i^{-1/\alpha}(M_{t-}^i)^{1/\alpha} r_i(x;t)^{1/\alpha}\,di \\ &= \int_0^1 \frac{\zeta_i^{-1/\alpha}(M_{t-}^i)^{1/\alpha}}{\Xi_{t-}}\, r_i(x;t)^{1/\alpha}\,di \\ &= \int_0^1 s_{t-}^i\, r_i(x;t)^{1/\alpha}\,di \;=\; g_t(x),\end{aligned}$$

the third line moving the constant Ξ_{t-}^{-1} under the integral and the fourth using $s_{t-}^i = \zeta_i^{-1/\alpha}(M_{t-}^i)^{1/\alpha}/\Xi_{t-}$ (A6); this is (8). Multiplying, the state-price density jumps by

$$\frac{\xi_t}{\xi_{t-}} = \frac{C_t^{-\alpha}}{C_{t-}^{-\alpha}}\left(\frac{\Xi_t}{\Xi_{t-}}\right)^{\alpha} = e^{\alpha x}\, g_t(x)^{\alpha}, \tag{A8}$$

the continuous parts of $C^{-\alpha}$ and Ξ^{α} contributing no jump.

Step 2 (risk-neutral compensator). Markets are complete (Assumption 4), so there is a unique equivalent martingale measure $\mathbb{Q}$ with density process $Z_t = \xi_t \exp(\int_0^t r_s\,ds)/\xi_0$. We verify the hypotheses of the change-of-measure theorem for marked point processes (Brémaud, 1981, VI, T2–T4). *(a) Z is a local martingale.* The equilibrium short rate is the process that deflates the state-price density to a local martingale, and this fixes the drift cancellation. Write the stochastic logarithm of the SDF (6) as

$$\frac{d\xi_t}{\xi_{t-}} = b_t\,dt + dm_t,$$

with b_t the predictable finite-variation drift coefficient—the Itô drift of $\xi_t = e^{-\delta t} C_t^{-\alpha}\Xi_t^{\alpha}$, a functional of the cross-section through Ξ—and m a local martingale collecting the Brownian and compensated-jump parts. The equilibrium short rate is, by its defining no-arbitrage property, $r_t := -b_t$; this property is logically prior to its closed form (which we do not need here). Since

$e^{\int_0^t r_s ds}$ is continuous and of finite variation it has zero covariation with ξ, so integration by parts gives

$$d\big(\xi_t e^{\int_0^t r_s ds}\big) = e^{\int_0^t r_s ds}\big(d\xi_t + \xi_{t-} r_t\, dt\big) = e^{\int_0^t r_s ds}\, \xi_{t-}\big((b_t + r_t)\, dt + dm_t\big) = e^{\int_0^t r_s ds}\, \xi_{t-}\, dm_t,$$

the drift vanishing by $r_t = -b_t$. Hence $\xi_t e^{\int_0^t r_s ds}$ is a positive $\mathbb{P}$-local-martingale, and $Z_t = \xi_t e^{\int_0^t r_s ds}/\xi_0$ is its unit normalization, with $Z_0 = 1$ and $\xi > 0$. The remaining and substantive task is to upgrade this to a *true* $(\mathcal{F}_t, \mathbb{P})$-martingale with $\mathbb{E}[Z_t] = 1$ (hence the density of $\mathbb{Q}$), which requires controlling *both* factors of Z, since ξ carries a continuous (Brownian) part as well as the disaster jumps. *Continuous part:* the diffusive loading of $\log \xi$ is $-\alpha\sigma\, dW_t$, constant because σ is constant (Assumption 1); the continuous factor of Z is then the Doléans exponential $\mathcal{E}(-\alpha\sigma W)_t$, a true martingale by Novikov's condition $\mathbb{E}[\exp(\frac{1}{2}\int_0^t (\alpha\sigma)^2 ds)] = \exp(\frac{1}{2}\alpha^2\sigma^2 t) < \infty$. *Jump part:* the bounded predictable jump ratio of (b)–(c) and the integrability bound $\int_0^t \int_{\mathcal{J}} |Y_s(x) - 1|\, \nu(ds, dx) \le (e^{\alpha\overline{J}}\overline{r} + 1)\, \lambda^\star t < \infty$ place the jump factor within Brémaud (1981, VI, T4). The two factors are driven by independent sources (W and μ), so their product Z is a true martingale; completeness independently yields a unique equivalent martingale measure, of which Z is the density. *(b)* The bank account $\exp(\int_0^t r_s\, ds)$ is continuous and of finite variation and ξ_0 is constant, so Z and ξ have the same jump ratio: at a disaster of size x, $Z_t/Z_{t-} = \xi_t/\xi_{t-} = e^{\alpha x} g_t(x)^\alpha =: Y_t(x)$ by (A8). *(c)* Y_t is strictly positive and predictable (it is built from the pre-jump shares s^i_{t-}, of which $g_t(\cdot)$ is a functional) and bounded on the compact $\mathcal{J}$, since $e^{\alpha x} \le e^{\alpha\overline{J}}$ and $0 < g_t(x) \le \overline{r}^{1/\alpha}$ by (8). Under (a)–(c) the theorem gives the $\mathbb{Q}$-compensator of μ as $\nu^{\mathbb{Q}}(dt, dx) = Y_t(x)\, \nu(dt, dx)$; with $\nu(dt, dx) = \lambda^\star\, dt\, F_J(dx)$ (2),

$$\nu^{\mathbb{Q}}(dt, dx) = \lambda^\star\, e^{\alpha x}\, g_t(x)^\alpha\, dt\, F_J(dx), \tag{A9}$$

which is (9); Y_t bounded makes $\nu^{\mathbb{Q}}$ a finite measure, so the change of measure is admissible.

Step 3 (size law). Using $F_J(dx) = f(x)\, dx$ (Assumption 1), the finite measure (A9) is $\nu^{\mathbb{Q}}(dt, dx) = \lambda^\star\, e^{\alpha x} g_t(x)^\alpha f(x)\, dt\, dx$. Factor it as intensity times a normalized size density,

$\nu^{\mathbb{Q}}(dt, dx) = \lambda_t^{\mathbb{Q}}\, dt\, f_t^{\mathbb{Q}}(x)\, dx$, by integrating out x for the intensity and dividing for the density:

$$\begin{aligned} \lambda_t^{\mathbb{Q}} &= \int_{\mathcal{J}} \frac{\nu^{\mathbb{Q}}(dt, dx)}{dt} = \lambda^\star \int_{\mathcal{J}} e^{\alpha x} g_t(x)^\alpha f(x)\, dx = \lambda^\star\, \mathbb{E}_f\big[e^{\alpha J} g_t(J)^\alpha\big], \\ f_t^{\mathbb{Q}}(x) &= \frac{\lambda^\star\, e^{\alpha x} g_t(x)^\alpha f(x)}{\lambda_t^{\mathbb{Q}}} = \frac{e^{\alpha x} g_t(x)^\alpha f(x)}{\mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]}, \end{aligned} \tag{A10}$$

where $\mathbb{E}_f[h] = \int_{\mathcal{J}} h\, f$; the denominator is finite and positive by the bounds in Step 2, so $f_t^{\mathbb{Q}}$ is a proper density on $\mathcal{J}$. This is (10) and proves (i).

Step 4 (the bridge). Multiply and divide $U_t = \lambda_t^{\mathbb{Q}}$ by $\mathbb{E}_f[e^{\alpha J}]$:

$$U_t = \lambda^\star\, \mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha] = \underbrace{\lambda^\star\, \mathbb{E}_f[e^{\alpha J}]}_{U^{\mathrm{RA}}} \cdot \underbrace{\frac{\mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]}{\mathbb{E}_f[e^{\alpha J}]}}_{\kappa_t}, \tag{A11}$$

which is (11). The factor $U^{\mathrm{RA}} = \lambda^\star \mathbb{E}_f[e^{\alpha J}]$ contains only the physical $(\lambda^\star, F_J)$ and α; the factor κ_t contains the cross-section only, through g_t. This proves (ii).

Step 5 (belief-driven variation). Holding $(\alpha, \lambda^\star, F_J)$ fixed, U^{RA} is constant in t, so taking logs of (A11) and differentiating,

$$d \log U_t = d \log U^{\mathrm{RA}} + d \log \kappa_t = d \log \kappa_t,$$

which is (12). In (A10) the only cross-section-dependent object is g_t, so $f_t^{\mathbb{Q}}$, $\bar{U}_t$, and U_t are functionals of g_t alone—of the held models ϕ_i and shares s_{t-}^i of the cross-section (Internet Appendix Proposition **??**)—while $(\lambda^\star, F_J)$ stay fixed. This proves (iii). □

Appendix B. Proofs for Section 5

A. Where each instantaneous price sits in the decomposition

The risk-neutral size law is the BT object (Bollerslev and Todorov, 2011; Breeden and Litzenberger, 1978): the physical F_J bent two ways—by risk aversion $e^{\alpha x}$, which fattens the tail because a representative agent dreads large disasters, and by the belief factor $g_t(x)^{\alpha}$, the only channel through which beliefs enter the price. That factor is a wealth-weighted *power* mean of the *joint* (intensity-and-size) disaster likelihood ratios—not their head-count or arithmetic mean, and not a mean of size beliefs alone, since intensity beliefs enter even when the size law is common. With $(\lambda^{\star}, F_J)$ fixed, $d \log U_t = d \log \kappa_t$: any movement in the multiplier maps one-for-one into the fear factor, while a self-re-adoption or offsetting cross-sectional change leaves it unmoved—price moves, physical quantity fixed.

The three option-tail objects the literature measures are each a functional of the *current* cross-section S_t (Proposition B.4): the far-tail level wedge is the level of κ_t (Proposition 2), the investor-fears index is the front wedge $U_t - \lambda^{\star}$ (Bollerslev and Todorov, 2011), and the unspanned tail factor is U_t itself (Andersen et al., 2015b). Any belief-sufficient benchmark reproduces all three, so matching them certifies a sensible time-varying intensity but does not separate lazy revision from the Markov class; the object that does is the forward fear factor's *slope* (Theorem 1), S_t-measurable in level but carrying a Ψ_t-part in slope.

The mapping also places one of their findings structurally. Andersen et al. (2015b) find the compensation for tail risk reacts far more persistently after a crisis than the underlying risk, and fit it with a persistent latent factor; here that persistence is not a free parameter but the survival history, the fear decaying on the staggered revision schedule—the record channel, the S_t-measurable selection component being shared by any Markov benchmark—rather than at the speed of the underlying. The discriminating content is once more the slope—a function of the current factor under any Markov persistence, a function of how long the calm has run under lazy revision.

Frequency. The record clock here is macroeconomic: at $\lambda^\star \approx 3\%$ a disaster arrives once a generation, while option samples span decades. The bridge is that the rule is not tied to one clock. An investor testing a crash model of the equity tail runs the same deadline construction with the crash arrival rate in place of $\lambda^\star$, and records, deadlines, triggers, and destinations carry over to the higher-frequency event set; at those frequencies the proxies turn over within available option histories. The macro calibration disciplines the mechanism's magnitudes; an international disaster panel supplies the cross-sectional power the single-country clock lacks.

The short rate, the price–dividend ratio, and the disaster premium are derived below; here we record where each sits in the decomposition of Proposition 5.

The short rate is pure $B_t(S_t)$. The instantaneous safe rate is $r_t = \delta + \alpha\mu_0 - \frac{1}{2}\alpha(1+\alpha)\sigma^2 + \bar{\lambda}_t^S - U_t$ (Proposition B.1): a constant power-utility diffusive part plus the belief-driven disaster part $\bar{\lambda}_t^S - U_t$, both functionals of S_t, so the rate carries no survival-history residual and a Markov-in-cross-section benchmark matches it exactly. Beliefs enter only through the difference, so a fear wave lowers the rate precisely when the priced-tilt response dominates the mean-intensity response, $dU_t > d\bar{\lambda}_t^S$—the flight-to-safety of the lazy economy, purely belief-driven since $(\mu_0, \sigma, \lambda^\star, F_J)$ never move. The same fear factor that lowers the rate raises the disaster premium, so one channel delivers a low safe rate and a high premium together, the pairing the rare-disaster literature uses to address both puzzles (Barro, 2006). Under power utility this passes a volatile fear wave straight into r_t, against smooth real rates; recursive-preference disaster models keep r_t smooth (Wachter, 2013), a device we forgo, so the belief block must be disciplined (Section 6).

The price–dividend ratio carries the residual. The ratio is $v(\Psi_t)$ (Proposition B.2), a Feynman–Kac functional of the full record-resolved state—not autonomous in S_t in any channel and, under size disagreement, with no finite-dimensional instantaneous reduction (Proposition 3); in the decomposition it splits as $v_B(S_t) + \mathcal{H}_t^v(\Psi_t)$, the residual driven by the same staggered-revision channels as the forward-fear residual below and zero for any Markov benchmark. Its comparative static is *two-signed*: the valuation-growth rate loads on the cross-section through two non-interchangeable aggregates—the *arithmetic* mean $\bar{\lambda}_t^S$ in the reallocation drift and the *power* mean $G_t^\alpha \leq \bar{\lambda}_t^S/\lambda^\star$ in the priced-jump term—so whether a more fearful cross-section raises valuation growth and *expands* the price–dividend ratio, against the naive "higher discount, lower price" reading, is a quantitative condition met at our calibration, not a consequence of Jensen. A finite-menu Feynman–Kac equation and a first-order belief-dispersion expansion follow (Corollaries IA.1 and IA.2).

The premium loads on the fear factor, and predicts returns. The instantaneous excess return is $\alpha\sigma^2 + \mathrm{DRP}_t$ (Proposition B.3): the diffusive premium $\alpha\sigma^2$ is constant, so all variation is the disaster premium, which loads on the fear factor $U_t = U^{\mathrm{RA}}\kappa_t$. Priced disaster fear therefore forecasts excess returns while carrying no information about diffusive volatility, pinned at the constant σ—the return-premium leg of the Andersen et al. (2015b) unspanned-tail finding, with the caveat that, σ constant, the non-prediction of volatility is an identity rather than a match. The premium's *level* runs through the price–dividend jump ratio ψ_t, hence $v(\Psi_t)$, so it too carries the survival-history residual.

Lemma B.1 (Predictable revisions carry no price jump): *Let τ be a calm-deadline revision time—a predictable time with $\Delta N_\tau = 0$ at which the revising cohort's held model jumps from ϕ to the common destination $\phi_\tau^* = \mathfrak{m}(\rho_\tau)$ (Internet Appendix Lemma* **??***), possibly for a positive G-mass synchronizing on a flat level set of $\Delta^\star(\phi,\cdot)$. Then* (i) *every belief density is continuous at τ, $M_\tau^i = M_{\tau-}^i$, so Ξ_t, the kernel ξ_t, and every consumption share s_t^i are continuous at τ and each reviser carries its pre-jump share with no instantaneous wealth transfer;* (ii) *the share-weighted held-model distribution S_t jumps in total variation, so the instantaneous functionals $g_t(\cdot)$, U_t, the short rate, and the* level *of the disaster premium carry a* predictable *jump at τ;* (iii) *the traded price is continuous, $P_{\tau-} = P_\tau$, equivalently $v(\Psi_{\tau-}, \rho_{\tau-}) = v(\Psi_\tau, \rho_\tau)$ on the record-resolved state of Definition 4, and P jumps only at the totally inaccessible disaster epochs. The conclusion uses no atomlessness of the deadline law.*

Proof. (i). In (A1) the factor $\exp(-\int_0^t(\lambda_i(s-)-\lambda^\star)\,ds)$ is absolutely continuous in t, and the product $\prod_{T_k\le t} r_i(J_k;T_k)$ acquires a factor only when t crosses a disaster epoch T_k; since $\Delta N_\tau=0$, none is acquired at τ, and the revision changes $\lambda_i(\cdot-)$ and the forward likelihood ratios only on $\{s\ge\tau\}$—a Lebesgue-null contribution to M^i_τ. Hence $M^i_\tau=M^i_{\tau-}$: the held-model switch moves the predictable forward drift of M^i, not its level. (Equivalently, M^i is a true $(\mathcal{F}_t,\mathbb{P})$-martingale built from the predictable, left-continuous held-model compensator (Brémaud, 1981, VI, T4), and the driving Brownian-plus-Poisson filtration is quasi-left-continuous, so $\mathcal{F}_\tau=\mathcal{F}_{\tau-}$ and $M^i_\tau=\mathbb{E}[M^i_\tau\mid\mathcal{F}_{\tau-}]=M^i_{\tau-}$.) On a disaster-free neighbourhood of τ the ratio M^i_t/M^i_τ is bounded above and below uniformly in i (Λ compact), so dominated convergence gives $\Xi_t\to\Xi_\tau$ in (6); with C_t continuous in calm, $\xi_t=e^{-\delta t}C_t^{-\alpha}\Xi_t^\alpha$ and the shares (7) are continuous, so $s^i_\tau=s^i_{\tau-}$.

(ii). A positive mass relocates $\phi\to\phi^*_\tau$ while s^i is continuous, so $S_t=\int_0^1 s^i_t\,\delta_{\phi_i(t)}\,di$ jumps in total variation; the belief aggregator $g_t(x)=\int_\Phi r(\phi,x)^{1/\alpha}S_t(d\phi)$ and the functionals it determines inherit a predictable jump.

(iii). The equilibrium gain process is a $\mathbb{P}$-martingale (Proposition 1), so $\xi_tP_t=\mathbb{E}_t[\int_t^\infty\xi_sC_s\,ds]$. The dividend carries no mass at the single time τ, and by quasi-left-continuity $\mathcal{F}_\tau=\mathcal{F}_{\tau-}$ with $\rho_\tau=\rho_{\tau-}$ (the public posterior is continuous in calm), so the two conditional expectations coincide and $\xi_{\tau-}P_{\tau-}=\xi_\tau P_\tau$; dividing by the continuous, strictly positive ξ gives $P_{\tau-}=P_\tau$. Writing $P_t=C_t\,v(\Psi_t,\rho_t)$ with $C_\tau=C_{\tau-}$ yields the value-matching form. The destination $\phi^*_\tau=\mathfrak{m}(\rho_\tau)$ is read from ρ_τ, so the relocation map is $\mathcal{F}_{\tau-}$-measurable precisely because the forward-looking state carries ρ_t; and $\Psi_{\tau-}$ already encodes the imminent deadline (the ages and records that pin $\Delta^\star$), so the forward value already incorporates the scheduled relocation. The disaster epochs are totally inaccessible, so P jumps only there. $\square$

Proof of Proposition 5. The decomposition (18) itself is self-contained: $B_t=\mathbb{E}[F_t\mid S_t]$ is the conditional expectation of the integrable functional F_t given $\sigma(S_t)$—the L^2 projection when F_t is square-integrable, so it is S_t-measurable, and the residual $\mathcal{H}_t=F_t-B_t$ has $\mathbb{E}[\mathcal{H}_t\mid S_t]=0$ by the tower property and is measurable in the finer Ψ_t, since $S_t = \pi_\phi \Psi_t$. That each price object F_t is itself measurable in the pricing state—in (Ψ_t, ρ_t), with ρ_t entering through re-selection—is the content of the state-sufficiency results (Propositions B.4 and B.2); the projection argument takes it as input. This holds for any equilibrium price object F_t and uses nothing beyond the definitions of the two states. The classification (i)–(iii) of which objects carry a zero versus a nonzero residual is read off the pricing results established later in this appendix; we record it here only to fix the organizing split, and it carries no logical weight in the present argument. Throughout, the residual's leading term is the *centered* channel—the staggered-revision channels minus their S_t-conditional mean; the raw channels coincide with the residual only where their S_t-projection vanishes. *(i)* The fear factor and multiplier, the risk-neutral disaster law, and the short rate are S_t-measurable (Proposition B.4), so each equals its own projection and $\mathcal{H}_t \equiv 0$. *(ii)* The price–dividend ratio and the premium level are functionals of the full Ψ_t, not autonomous in S_t (Proposition B.2); and the forward fear factor has $F_t(0) = U_t$ measurable in S_t with right-hand slope $\partial_s^+ F_t(0) = \mathrm{b}_S(S_t) + U_t\, c_{\mathrm{sw}}(\Phi_t^{\mathrm{sw}}) + \lambda^\star \Theta_{\mathrm{sw}}(\Psi_t)$, whose last two terms are measurable in Ψ_t and not in S_t (Theorem 1), so $\mathcal{H}_t = s\big(U_t\, c_{\mathrm{sw}}(\Phi_t^{\mathrm{sw}}) + \lambda^\star \Theta_{\mathrm{sw}}(\Psi_t) - \mathbb{E}[U_t\, c_{\mathrm{sw}} + \lambda^\star \Theta_{\mathrm{sw}} \mid S_t]\big) + o(s)$—the *centered* staggered-revision channels, the raw channels net of their S_t-conditional mean—zero at horizon zero and first order in s. *(iii)* For any benchmark whose pricing state is Markov in the current cross-section S_t, the conditional law of the future is by construction a functional of S_t, so every forward object $\mathbb{E}[\,\cdot \mid \mathcal{F}_t]$ equals its own S_t-projection and $\mathcal{H}_t \equiv 0$. The worked instance is Theorem 1(iv) with Internet Appendix IA.9: under Bayesian updating the per-investor state is the posterior, its own sufficient statistic, so the cross-section is Markov, the deadline flow Φ_t^{sw} and trigger propensity π are absent, and every forward object closes in S_t; an affine latent-state specification and a representative filtering economy are Markov in their own states by construction, with the same conclusion. $\square$

Proposition B.1 (The equilibrium short rate): *In the equilibrium of Proposition 1, the instan-*

taneous risk-free rate is

$$r_t = \delta + \alpha\mu_0 - \tfrac{1}{2}\alpha(1+\alpha)\sigma^2 + \bar{\lambda}^S_t - U_t, \tag{B1}$$

with $\bar{\lambda}^S_t = \int_\Phi \lambda\, dS_t$ the wealth-weighted held intensity (Definition 3) and U_t the disaster-fear factor (10). The diffusive part $\delta + \alpha\mu_0 - \frac{1}{2}\alpha(1+\alpha)\sigma^2$ is the constant power-utility rate; the disaster part $\bar{\lambda}^S_t - U_t$ is belief-driven, and at correct beliefs ($\bar{\lambda}^S_t = \lambda^\star$, $U_t = U^{\mathrm{RA}}$) collapses to the rare-disaster precautionary term $-\lambda^\star(\mathbb{E}_f[e^{\alpha J}] - 1) < 0$.

Proof of Proposition B.1. Take logs of $\xi_t = e^{-\delta t} C_t^{-\alpha} \Xi_t^{\alpha}$ (6):

$$\log \xi_t = -\delta t - \alpha \log C_t + \alpha \log \Xi_t. \tag{B2}$$

By Assumption 1 (the jump multiplies C by e^{-x}, so $\Delta \log C = -x$),

$$d\log C_t = \big(\mu_0 - \tfrac{1}{2}\sigma^2\big)dt + \sigma\, dW_t - \int_{\mathcal{J}} x\, \mu(dt, dx), \tag{B3}$$

and by Internet Appendix Proposition **??**—the calm drift $d\log \Xi_t/dt = -\frac{1}{\alpha}(\bar{\lambda}^S_t - \lambda^\star)$ (from Internet Appendix (**??**) divided by Ξ_t) and the disaster jump $\Xi_t/\Xi_{t-} = g_t(x)$ (Internet Appendix (**??**)), so $\Delta \log \Xi = \log g_t(x)$, with no diffusion since M^i carries none—

$$d\log \Xi_t = -\tfrac{1}{\alpha}(\bar{\lambda}^S_t - \lambda^\star)\, dt + \int_{\mathcal{J}} \log g_t(x)\, \mu(dt, dx). \tag{B4}$$

Differentiating (B2) and substituting (B3)–(B4),

$$\begin{aligned} d\log \xi_t &= -\delta\, dt - \alpha\, d\log C_t + \alpha\, d\log \Xi_t \\ &= \big[-\delta - \alpha(\mu_0 - \tfrac{1}{2}\sigma^2) - (\bar{\lambda}^S_t - \lambda^\star)\big]dt - \alpha\sigma\, dW_t + \int_{\mathcal{J}} \big[\alpha x + \alpha \log g_t(x)\big]\mu(dt, dx), \end{aligned} \tag{B5}$$

the dt, dW, and μ terms collected from (B3)–(B4).

Now pass from $\log \xi$ to $\xi = e^{\log \xi}$ by Itô's formula for the exponential of a semimartingale. Write $L = \log \xi$; by (B5) its continuous martingale part L^c satisfies $dL^c_t = -\alpha\sigma\, dW_t$, so, with $d\langle W\rangle_t = dt$, its quadratic variation and the resulting Itô correction are

$$d\langle L^c\rangle_t = (-\alpha\sigma)^2\, d\langle W\rangle_t = \alpha^2\sigma^2\, dt, \qquad \tfrac{1}{2}\, d\langle L^c\rangle_t = \tfrac{1}{2}\alpha^2\sigma^2\, dt, \tag{B6}$$

and its jump at a disaster of size x is $\Delta L_t = \alpha x + \alpha \log g_t(x)$ (the μ-integrand of (B5)). Itô's exponential formula then gives

$$\frac{d\xi_t}{\xi_{t-}} = dL^c_t + \big(\text{drift of } L\big)dt + \tfrac{1}{2}\, d\langle L^c\rangle_t + \int_{\mathcal{J}} \big(e^{\Delta L_t(x)} - 1\big)\, \mu(dt, dx), \tag{B7}$$

in which the jump factor evaluates as

$$e^{\Delta L_t} - 1 = e^{\alpha x + \alpha \log g_t(x)} - 1 = e^{\alpha x}\, e^{\alpha \log g_t(x)} - 1 = e^{\alpha x} g_t(x)^{\alpha} - 1$$

(consistent with (A8)). Collecting the dt terms of (B7)—the drift of L read from (B5) plus the correction (B6)—the continuous drift of $d\xi_t/\xi_t$ is therefore

$$\Big(\frac{d\xi_t}{\xi_t}\Big)^{\text{cont}}_{\text{drift}} = -\delta - \alpha(\mu_0 - \tfrac{1}{2}\sigma^2) - (\bar{\lambda}^S_t - \lambda^\star) + \tfrac{1}{2}\alpha^2\sigma^2. \tag{B8}$$

The mean jump rate is the jump factor integrated against the disaster compensator $\nu(dt, dx) = \lambda^\star\, dt\, f(x)\, dx$ (2):

$$\begin{aligned} \frac{1}{dt}\mathbb{E}_t\Big[\big(e^{\alpha x} g_t(x)^{\alpha} - 1\big)\, \mu\Big] &= \lambda^\star \int_{\mathcal{J}} \big(e^{\alpha x} g_t(x)^{\alpha} - 1\big) f(x)\, dx \\ &= \lambda^\star\, \mathbb{E}_f\big[e^{\alpha J} g_t(J)^{\alpha}\big] - \lambda^\star \qquad \text{(B9)} \\ &= U_t - \lambda^\star, \qquad \text{(B10)} \end{aligned}$$

line (B9) splitting the integral and (B10) using $U_t = \lambda^\star \mathbb{E}_f[e^{\alpha J} g_t(J)^{\alpha}]$ (10). The safe rate is minus the total expected return on ξ, $r_t = -\frac{1}{dt}\mathbb{E}_t[d\xi_t/\xi_t]$, i.e. minus the sum of (B8) and (B10):

$$\begin{aligned} r_t &= \delta + \alpha(\mu_0 - \tfrac{1}{2}\sigma^2) - \tfrac{1}{2}\alpha^2\sigma^2 + (\bar\lambda_t^S - \lambda^\star) - (U_t - \lambda^\star) \\ &= \delta + \alpha\mu_0 - \tfrac{1}{2}\alpha\sigma^2 - \tfrac{1}{2}\alpha^2\sigma^2 + \bar\lambda_t^S - U_t \qquad \text{(B11)} \\ &= \delta + \alpha\mu_0 - \tfrac{1}{2}\alpha(1+\alpha)\sigma^2 + \bar\lambda_t^S - U_t, \qquad \text{(B12)} \end{aligned}$$

line (B11) expanding $\alpha(\mu_0 - \frac{1}{2}\sigma^2)$ and cancelling $\mp\lambda^\star$, and (B12) factoring $-\frac{1}{2}\alpha\sigma^2 - \frac{1}{2}\alpha^2\sigma^2 = -\frac{1}{2}\alpha(1+\alpha)\sigma^2$; this is (B1). At correct beliefs $\bar\lambda_t^S = \lambda^\star$ and $U_t = U^{\mathrm{RA}} = \lambda^\star\mathbb{E}_f[e^{\alpha J}]$, so $\bar\lambda_t^S - U_t = -\lambda^\star(\mathbb{E}_f[e^{\alpha J}] - 1)$. □

Lemma B.2 (Payoff-growth domination): *Under Assumption 1, for every horizon $h > 0$ and every $K \geq 0$, $\mathbb{E}\big[\sup_{t\leq s\leq t+h}(C_s/C_t)^{1-\alpha}\, e^{KN_{(t,s]}}\big] < \infty$.*

Proof. Write $(C_s/C_t)^{1-\alpha} = e^{(1-\alpha)(\mu_0-\sigma^2/2)(s-t)}\, e^{(1-\alpha)\sigma(W_s-W_t)}\, e^{-(1-\alpha)\sum_{t<u\leq s} x_u}$. The drift factor is bounded on $[0,h]$; $\mathbb{E}[\sup_{s\leq t+h} e^{(1-\alpha)\sigma(W_s-W_t)}] < \infty$ by the reflection principle, the running maximum of Brownian motion having Gaussian tails; and the jump factor is bounded by $e^{(|1-\alpha|\bar J+K)N_{(t,t+h]}}$, $\bar J$ the compact support bound, with finite compound-Poisson exponential moment $\exp\{\lambda^\star h(e^{|1-\alpha|\bar J+K}-1)\}$. The factors are independent, so the product is integrable. □

Lemma B.3 (Pathwise total-variation stability of the flow): *Throughout, the economy is the finite-menu economy with turnover parameter $\Omega \geq 0$ ($\Omega = 0$ the baseline of Section 5, where the death-clock clause below is void) and $\lambda_{hi} := \max_\Phi \lambda$. Let Ψ, Ψ' be record-resolved states with $\|\Psi - \Psi'\|_{\mathrm{TV}} = \varepsilon$, and couple the two economies on a common driving path—the same aggregate marked Poisson measure and, under a TV-optimal coupling of the initial states, the same idiosyncratic death clocks and type draws, and sharing the public record to date, hence the posterior ρ (the comparison varies the cross-section on a common public history)—so that investors of joint mass $1-\varepsilon$ are* shared. *Then for every horizon $h > 0$: (i) each shared investor's state coincides in the two economies on $[t, t+h]$; (ii) pathwise, $\|S_s - S'_s\|_{\mathrm{TV}} \leq \varepsilon\, q^{-N_{(t,s]}} e^{c(s-t)}$ for $s \leq t+h$, with $q = (\underline{r}/\overline{r})^{1/\alpha}$ and $c = 2\lambda_{hi}/\alpha + \Omega$; and (iii) under the margin $\delta > \kappa_L :=$*

$\bar{m}+c+\bar{\Lambda}(1/q-1)$—with $\bar{m} := \sup_S m(S)$ the transversality bound and $\bar{\Lambda} := \lambda^\star e^{(\alpha-1)\overline{J}}\,\bar{r}$ a crude bound on the payoff-tilted disaster intensity—$|v(\Psi)-v(\Psi')| \le L\,\varepsilon$ for an explicit finite L; κ_L is a crude sufficient constant, and parts (i)–(ii) need no condition.

Proof. (i) A shared investor carries the same held model, record, age, and tolerance in both economies at t, faces the same disasters and death clock, and—the one aggregate input to her dynamics—relocates, when she revises, to the *public* destination $\mathfrak{m}(\rho_s)$, a functional of the common driving path (Internet Appendix Lemma **??**). Her deadline, trigger evaluations, and destinations therefore coincide, and her state does, by induction over her revision times. *(ii)* Write the discrepancy as the unshared weight. In calm, each side's shares move by the replicator at log-rates bounded by λ_{hi}/α, and turnover retires unshared mass at rate Ω while entrants are shared; so the unshared TV-mass grows at most at rate $2\lambda_{hi}/\alpha+\Omega$ through the relative renormalisation. At a disaster the reweight multiplies every relative weight by a factor in $[q,1/q]$ (8), inflating the discrepancy by at most q^{-1}; relocations move mass to the common public destination and do not increase total variation. Compounding gives the display. *(iii)* Write $Y_\tau := e^{-\delta\tau}(C_{t+\tau}/C_t)^{1-\alpha}(\Xi_{t+\tau}/\Xi_t)^{\alpha}$ for the discounted payoff; on the common driving path its consumption factor is identical across the pair, so only the aggregator ratio differs. Pathwise,

$$A_\tau := \left|\log\frac{\Xi_{t+\tau}}{\Xi_t} - \log\frac{\Xi'_{t+\tau}}{\Xi'_t}\right| \le \int_0^\tau \frac{\Delta_\lambda}{\alpha}\,\|S_u - S'_u\|_{\mathrm{TV}}\,du + \sum_{\text{disasters}\le\tau} L_g\,\|S_{u-}-S'_{u-}\|_{\mathrm{TV}} \le \varepsilon\, q^{-N_\tau} e^{c\tau}\Big(\frac{\Delta_\lambda \tau}{\alpha} + L_g N_\tau\Big) =$$

using (ii) and the bounded-Lipschitz aggregate inputs on the compact $\Phi\times\mathcal{J}$ ($\Delta_\lambda := \overline{\lambda}-\underline{\lambda}$; $L_g := (\bar{r}^{1/\alpha}-\underline{r}^{1/\alpha})/\underline{r}^{1/\alpha}$). Since $|Y_\tau - Y'_\tau| \le \min(Y_\tau, Y'_\tau)(e^{\alpha A_\tau}-1)$ and also $\le Y_\tau + Y'_\tau$, the elementary bound $1\wedge(e^x-1) \le e\,x$ for $x\ge 0$ gives $|Y_\tau - Y'_\tau| \le e\alpha\,(Y_\tau+Y'_\tau)\,A_\tau \le e\alpha\varepsilon\,(Y_\tau+Y'_\tau)H_\tau$. Change measure to the payoff-tilted law at horizon τ, $d\mathbb{Q}^\tau \propto Y_\tau\,d\mathbb{P}$: under it the disaster count is dominated by a Poisson process of intensity $\bar{\Lambda}$ (the tilt weights each mark by $e^{-(1-\alpha)x}g(x)^\alpha \le e^{(\alpha-1)\overline{J}}\bar{r}$), so $\mathbb{E}^{\mathbb{Q}^\tau}[q^{-N_\tau}(a+bN_\tau)] \le (a+b\bar{\Lambda}\tau/q)\,e^{\bar{\Lambda}\tau(1/q-1)}$, while $\mathbb{E}[Y_\tau] \le e^{-(\delta-\bar{m})\tau}$; the same bounds hold for Y'. Hence

$$\mathbb{E}\big[|Y_\tau - Y'_\tau|\big] \le C_0\,\varepsilon\,(1+\tau)\,e^{-(\delta-\bar{m}-c-\bar{\Lambda}(1/q-1))\tau}, \qquad C_0 := 2e\alpha\Big(\frac{\Delta_\lambda}{\alpha} + L_g\Big(1+\frac{\bar{\Lambda}}{q}\Big)\Big),$$

and integrating over τ gives $|v(\Psi) - v(\Psi')| \le L\varepsilon$ with $L = C_0 \int_0^\infty (1+\tau)e^{-(\delta-\kappa_L)\tau} d\tau < \infty$ under $\delta > \kappa_L$. □

Proposition B.2 (The price–dividend ratio): *Let $v_t = P_t/C_t$ be the tree's price–dividend ratio, and let Ψ_t denote the* record-resolved pricing state*—the wealth-weighted distribution over per-investor states $(\phi, \mathcal{R}, \tau, \varepsilon)$ (held model, since-adoption record, regime age, tolerance), whose ϕ-marginal is the pricing state S_t (Definition 3). Then $v_t = v(\Psi_t)$ with*

$$v(\Psi_t) = \mathbb{E}_t\Big[\int_t^\infty e^{-\delta(s-t)}\Big(\tfrac{C_s}{C_t}\Big)^{1-\alpha}\Big(\tfrac{\Xi_s}{\Xi_t}\Big)^{\alpha} ds\Big]. \tag{B13}$$

The local expected growth rate m_t of $Y_t := C_t^{1-\alpha}\Xi_t^\alpha$ is a function of the marginal S_t alone,

$$m(S_t) = (1-\alpha)\mu_0 - \tfrac{1}{2}\alpha(1-\alpha)\sigma^2 - \big(\bar{\lambda}_t^S - \lambda^\star\big) + \lambda^\star\big(\mathbb{E}_f[e^{-(1-\alpha)J}g_t(J)^\alpha] - 1\big), \tag{B14}$$

but the valuation v depends on the full Ψ_t: scale-freeness removes only the level C_t, and the conditional law of the future path $\{S_s\}_{s\ge t}$—hence of $\{m(S_s)\}$ and of Ξ_s/Ξ_t—is governed by the record-resolved dynamics of Proposition 3, which are Markov in the pair (Ψ_t, ρ_t)—the calm re-selection destination $\phi_s^* = \mathfrak{m}(\rho_s)$ being read from ρ_s—and not *autonomous in S_t. Throughout, $v(\Psi_t)$ abbreviates $v(\Psi_t, \rho_t)$: re-selection reads ρ_t, which since-adoption records do not recover. Under the transversality condition*

$$\delta > \sup m := \sup\nolimits_S m(S) \tag{T}$$

the integral (B13) converges, $0 < v(\Psi_t) \le (\delta - \sup m)^{-1} < \infty$, and v is the unique bounded valuation. It is not *semi-closed: v is a Feynman–Kac functional of the full record-resolved* *state Ψ_t, whose conditional law is Markov in (Ψ_t, ρ_t) and not autonomous in its marginal S_t (above), and even the instantaneous disaster tilt has no finite-dimensional sufficient statistic (Proposition 3); so v is not summarized by finitely many moments of S_t. In the decomposition of Proposition 5 the valuation splits as $v(\Psi_t) = v_B(S_t) + \mathcal{H}^v_t(\Psi_t)$: the cross-section part $v_B(S_t) = \mathbb{E}[v \mid S_t]$ satisfies $\mathbb{E}[\mathcal{H}^v_t \mid S_t] = 0$, and the survival-history residual $\mathcal{H}^v_t(\Psi_t)$ is generated by the record-resolved conditional law of the payoff path: its leading local component is the forward valuation-growth residual $\mathbb{E}_t[m(S_{t+\tau}) \mid \Psi_t] - \mathbb{E}[m(S_{t+\tau}) \mid S_t]$, entering through the discounted Feynman–Kac accumulation (B13), and the higher-order components enter through the same staggered-revision channels' $(\Phi^{\mathrm{sw}}_t, \pi)$ effect on the path law; it is zero for any Markov benchmark. In the intensity channel ($f_\theta \equiv f$, $g_t \equiv G_t = \mathbb{E}_{S_t}[(\lambda/\lambda^\star)^{1/\alpha}]$), (B14) reads*

$$m(S_t) = const - (\bar{\lambda}^S_t - \lambda^\star) + \lambda^\star\big(\mathbb{E}_f[e^{-(1-\alpha)J}]\, G^\alpha_t - 1\big), \tag{B15}$$

which carries the cross-section through two distinct *aggregates: the wealth-weighted* arithmetic *mean $\bar{\lambda}^S_t$ in the reallocation drift $-(\bar{\lambda}^S_t - \lambda^\star)$, and the* power *mean G^α_t in the priced-jump term, with $G^\alpha_t \le \bar{\lambda}^S_t/\lambda^\star$, strict by Jensen for $\alpha > 1$ unless λ is S_t-a.s. constant (21).*

Reading the proposition. The transversality margin is non-vacuous because each term of (B14) is a bounded functional of S_t over the compact menu (g_t bounded, $\mathcal{J}$ compact), so $-\infty < \inf_S m(S) \le \sup_S m(S) < \infty$. The projection v_B is the best S_t-measurable summary of the valuation—the component any Markov-in-cross-section benchmark can at most reproduce, its own valuation being S_t-measurable, so the defining property $\mathbb{E}[\mathcal{H}^v_t \mid S_t] = 0$ makes the residual orthogonal to every such benchmark. A Markov-in-S_t benchmark prices forward objects by its *own* functional of S_t, which need not equal the projection B_t; the decomposition says only that whatever that functional is, it leaves the same residual. Capitalization is at the discount net of growth $\delta - m(S_t)$, not at a gross return. The comparative static is two-signed. The two aggregates are not interchangeable, and a more fearful cross-section raises both; whether the priced-jump (marginal-utility) term dominates the frequency drag—so that m *rises* and the constant-m ratio *expands*—is the quantitative condition above, met at our calibration; compression of *levered* equity would in any case require a leverage channel, absent here. That S_t is genuinely insufficient—not merely non-autonomous—has a witness. Take the aged/fresh pair of Theorem 1(iii): a common cross-section S_t and a common public posterior ρ_t, calm deadline flows $\Phi_t'^{\mathrm{sw}} \geq \Phi_t^{\mathrm{sw}}$ setwise with the difference charging the switch set. On the calm-conditional paths the dynamics are deterministic, and at the origin the selection terms cancel (both read the shared S_t), leaving only the flow difference: with $Z := (\lambda/\lambda^\star)^{1/\alpha}$,

$$\frac{d}{du}\big[m(S_u) - m(S_u')\big]\Big|_{u=t} = c_1 := -\Delta_1 + \alpha\lambda^\star \mathbb{E}_f[e^{-(1-\alpha)J}]\, G_t^{\alpha-1}\Delta_G, \tag{B16}$$

where $\Delta_1 := \int_\Phi (\lambda - \lambda(\phi_t^*))\, d(\Phi' - \Phi)_t^{\mathrm{sw}} \geq 0$ and $\Delta_G := \int_\Phi (Z - Z(\phi_t^*))\, d(\Phi' - \Phi)_t^{\mathrm{sw}} \geq 0$, both non-negative because monotone re-selection gives $\lambda(\phi_t^*) \leq \lambda$ pointwise on the switch set and Z is increasing in λ. The two terms of c_1 carry opposite signs—the frequency-drag relief against the priced-jump decline, the two-signed comparative static of (B15)—so $m(S_u) - m(S_u') = c_1(u-t) + o(u-t)$ and the calm-conditional growth paths separate at first order whenever $c_1 \neq 0$—the conditional law of $\{m(S_s)\}_{s \geq t}$ given Ψ_t is not a function of S_t, which is the insufficiency claimed; the knife-edge is the exact offset $\Delta_1 = \alpha\lambda^\star \mathbb{E}_f[e^{-(1-\alpha)J}] G_t^{\alpha-1}\Delta_G$, and the sign of c_1 is the same dominance comparison as (B15). The valuation step below turns the path separation into a strict value separation—$v(\Psi_t) \neq v(\Psi_t')$ whenever the dominance holds with one sign on the swept range, met with slack at our calibration: two states, one marginal, two values.

The valuation step. The separation (B16) concerns the conditional law of $\{m(S_s)\}$; the claim of *values* needs the Feynman–Kac integral (B13) to inherit it, and later portions of the path could in principle offset the local wedge. They cannot, in two moves. *First, a pathwise order.* Embed the pair in the one-parameter family that advances the regime ages of the differing switch-set mass by $a \geq 0$—the aged/fresh pair is an instance—and couple all members on a common driving path. Each investor's unnormalised weight is her entry weight times the product of her held models' factors over her holding history (the $(M^i)^{1/\alpha}$ representation of Proposition 1); advancing a calm re-selection by da exchanges, over the advanced window, the incumbent's calm factor for the destination's, so

$$\frac{d}{da}\big(\text{weight at } u\big) \;=\; \big(\text{weight}\big) \times \sum_k \frac{\lambda(\phi_k) - \lambda(\phi_{k+1})}{\alpha} \;\geq\; 0,$$

the sum over her advanced calm switches, each term non-negative under the maintained monotone re-selection (disaster-triggered relocations occur at common epochs and contribute no exchange; disasters landing in the shrinking window contribute zero to the derivative almost surely). Weight therefore migrates weakly from higher to lower intensities at every u on every branch: $\delta\bar{\lambda}_u := -\,\partial_a \bar{\lambda}^S_u\big|_{a=0} \geq 0$ and $\delta G_u := -\,\partial_a G_u\big|_{a=0} \geq 0$ pathwise, strictly positive from the first advanced deadline onward on the positive-probability event that no disaster precedes it, and strictness propagates through disasters—the reweight acts investor-by-investor and a triggered relocation is common to the coupled economies. *Second, the value derivative.* Pathwise Y-order is unavailable—the calm drift and the jump factor load on the cross-section with opposite signs—so condition on the first disaster: on the disaster-free event the dynamics are deterministic and (B13) takes the renewal form

$$v(\Psi_t) = \int_t^\infty e^{-(\delta+\lambda^\star)(u-t)}\, e^{\int_t^u m_c(S_r)\,dr} \Big[1 + \lambda^\star \int_{\mathcal{J}} e^{-(1-\alpha)x}\, G_u^\alpha\, v(\Psi_u^x)\, f(x)\, dx\Big] du,$$

with m_c the jump-stripped calm growth. Differentiating in a at 0 yields a linear renewal equation $\delta v = \mathfrak{F} + \mathcal{K}[\delta v]$: $\mathcal{K}$ carries the derivative to the post-jump states, positive, with norm below one under the margin $\delta > \kappa_L$ already imposed for the valuation step (Lemma B.3); and the forcing collects the two channels against strictly positive valuation weights,

$$\mathfrak{F} = \int_t^\infty e^{-(\delta+\lambda^\star)(u-t)}\, e^{\int_t^u m_c} \Big[-W_1(u)\,\delta\bar{\lambda}_u + \lambda^\star \mathbb{E}_f[e^{-(1-\alpha)J}]\,\alpha G_u^{\alpha-1}\, W_2(u)\,\delta G_u\Big] du,$$

W_1 the (Fubini-transferred) discounted continuation weight and W_2 the jump-value weight.

Hence $\delta v = \sum_{n\geq 0} \mathcal{K}^n \mathfrak{F}$ carries $\mathfrak{F}$'s sign: whenever the dominance between the two non-negative processes $(\delta\bar{\lambda}, \delta G)$ holds with one sign on the swept range—the comparison of c_1 and (B15), weighted by (W_1, W_2)—the derivative is strictly signed. The exchange-sum computation is uniform in the base a, so v is strictly monotone along the family, and the finite pair's values differ: the witness prices.

Proof of Proposition B.2. The valuation and its state. By (6), $\xi_u = e^{-\delta u} C_u^{-\alpha} \Xi_u^{\alpha}$, so

$$\begin{aligned}\frac{\xi_s C_s}{\xi_t C_t} &= \frac{e^{-\delta s} C_s^{-\alpha} \Xi_s^{\alpha}\, C_s}{e^{-\delta t} C_t^{-\alpha} \Xi_t^{\alpha}\, C_t} \\ &= e^{-\delta(s-t)} \Big(\frac{C_s}{C_t}\Big)^{1-\alpha} \Big(\frac{\Xi_s}{\Xi_t}\Big)^{\alpha}, \end{aligned} \tag{B17}$$

line (B17) using $C_s^{-\alpha} C_s = C_s^{1-\alpha}$ (likewise at t) and collecting the $e^{-\delta}$ and Ξ factors. The tree pays the dividend C_s, so by the pricing identity $P_t = \mathbb{E}_t[\int_t^{\infty} (\xi_s/\xi_t) C_s\, ds]$; inserting C_t/C_t and using (B17),

$$\begin{aligned} P_t &= C_t\, \mathbb{E}_t\Big[\int_t^{\infty} \frac{\xi_s C_s}{\xi_t C_t}\, ds\Big] \\ &= C_t\, \mathbb{E}_t\Big[\int_t^{\infty} e^{-\delta(s-t)} \Big(\tfrac{C_s}{C_t}\Big)^{1-\alpha} \Big(\tfrac{\Xi_s}{\Xi_t}\Big)^{\alpha} ds\Big]. \end{aligned}$$

Dividing by C_t gives $v_t = P_t/C_t =$ (B13). The integrand is homogeneous of degree zero in the level C_t, so v_t does not depend on C_t; it depends on the conditional law given $\mathcal{F}_t$ of the path $(C_s/C_t,\, \Xi_s/\Xi_t)_{s\geq t}$. The ratios C_s/C_t are functions of the primitive increments only (Assumption 1), independent of the cross-section. The ratios Ξ_s/Ξ_t are driven by the future cross-section path, whose conditional law is Markov in the record-resolved state Ψ_t and *not* autonomous in its ϕ-marginal S_t (Proposition 3: the switch flow depends on the within-class record distribution). Hence $v_t = v(\Psi_t)$.

The growth rate $m(S_t)$. Take logs of $Y_t = C_t^{1-\alpha} \Xi_t^{\alpha}$, $\log Y_t = (1-\alpha)\log C_t + \alpha \log \Xi_t$, and substitute (B3)–(B4):

$$d\log Y_t = (1-\alpha)\,d\log C_t + \alpha\,d\log \Xi_t$$
$$= \big[(1-\alpha)(\mu_0 - \tfrac{1}{2}\sigma^2) - (\bar{\lambda}_t^S - \lambda^\star)\big]dt + (1-\alpha)\sigma\,dW_t + \int_{\mathcal{J}} \big[-(1-\alpha)x + \alpha\log g_t(x)\big]\mu(dt,dx). \tag{B18}$$

By Itô's formula for $Y = e^{\log Y}$, the continuous martingale part of (B18) is $(1-\alpha)\sigma\,dW_t$, with quadratic variation and Itô correction

$$\big((1-\alpha)\sigma\big)^2\,d\langle W\rangle_t = (1-\alpha)^2\sigma^2\,dt, \qquad \tfrac{1}{2}(1-\alpha)^2\sigma^2\,dt,$$

and the jump factor at a disaster of size x is $e^{-(1-\alpha)x+\alpha\log g_t(x)} - 1 = e^{-(1-\alpha)x}g_t(x)^\alpha - 1$; hence

$$\frac{dY_t}{Y_{t-}} = \Big[(1-\alpha)(\mu_0-\tfrac{1}{2}\sigma^2)-(\bar{\lambda}_t^S-\lambda^\star)+\tfrac{1}{2}(1-\alpha)^2\sigma^2\Big]dt+(1-\alpha)\sigma\,dW_t+\int_{\mathcal{J}}\big(e^{-(1-\alpha)x}g_t(x)^\alpha-1\big)\mu(dt,dx).$$

Take $\mathbb{E}_t[\cdot]/dt$ (conditioning on $\mathcal{F}_t$): the dW_t term has mean zero, and the random measure μ has compensator $\nu = \lambda^\star\,dt\,f\,dx$ (2), so

$$m(S_t) = \tfrac{1}{dt}\mathbb{E}_t\Big[\tfrac{dY_t}{Y_{t-}}\Big] = (1-\alpha)(\mu_0-\tfrac{1}{2}\sigma^2)-(\bar{\lambda}_t^S-\lambda^\star)+\tfrac{1}{2}(1-\alpha)^2\sigma^2+\lambda^\star\int_{\mathcal{J}}\big(e^{-(1-\alpha)x}g_t(x)^\alpha-1\big)f(x)\,dx. \tag{B19}$$

The diffusion terms in (B19) collect as

$$(1-\alpha)(\mu_0-\tfrac{1}{2}\sigma^2)+\tfrac{1}{2}(1-\alpha)^2\sigma^2 = (1-\alpha)\mu_0 + \tfrac{1}{2}\sigma^2\big[-(1-\alpha)+(1-\alpha)^2\big] \tag{B20}$$
$$= (1-\alpha)\mu_0 + \tfrac{1}{2}\sigma^2(1-\alpha)\big[(1-\alpha)-1\big] \tag{B21}$$
$$= (1-\alpha)\mu_0 + \tfrac{1}{2}\sigma^2(1-\alpha)(-\alpha) \tag{B22}$$
$$= (1-\alpha)\mu_0 - \tfrac{1}{2}\alpha(1-\alpha)\sigma^2, \tag{B23}$$

(B20) factoring $\frac{1}{2}\sigma^2$ from the two variance terms, (B21) factoring $(1-\alpha)$ from the bracket, (B22) using $(1-\alpha)-1=-\alpha$, and (B23) multiplying out. The jump integral splits, using $\int_{\mathcal{J}} f\,dx = 1$, as

$$\begin{aligned}\lambda^\star \int_{\mathcal{J}} \big(e^{-(1-\alpha)x} g_t(x)^\alpha - 1\big) f(x)\,dx &= \lambda^\star \int_{\mathcal{J}} e^{-(1-\alpha)x} g_t(x)^\alpha f(x)\,dx - \lambda^\star \int_{\mathcal{J}} f(x)\,dx \\ &= \lambda^\star\, \mathbb{E}_f\big[e^{-(1-\alpha)J} g_t(J)^\alpha\big] - \lambda^\star \\ &= \lambda^\star\big(\mathbb{E}_f[e^{-(1-\alpha)J} g_t(J)^\alpha] - 1\big). \end{aligned} \tag{B24}$$

Substituting (B23) and (B24) into (B19),

$$m(S_t) = (1-\alpha)\mu_0 - \tfrac{1}{2}\alpha(1-\alpha)\sigma^2 - (\bar{\lambda}_t^S - \lambda^\star) + \lambda^\star\big(\mathbb{E}_f[e^{-(1-\alpha)J} g_t(J)^\alpha] - 1\big), \tag{B25}$$

which depends on the cross-section only through $\bar{\lambda}_t^S$ and g_t, i.e. through S_t; this is (B14).

Convergence, bounds, uniqueness. Set $Z_s = e^{-\sup m\,(s-t)} Y_s/Y_t$. Its local drift rate is $(m(S_s) - \sup m) Z_s \le 0$ (by the definition of $\sup m$), and its jumps have finite mean (g_t is bounded by $\bar{r}^{1/\alpha}$ (8) and $\mathcal{J}$ is compact), so Z is a supermartingale and $\mathbb{E}_t[Y_s/Y_t] \le e^{\sup m\,(s-t)}$. Therefore, under (T),

$$v(\Psi_t) \le \int_t^\infty e^{-\delta(s-t)} e^{\sup m\,(s-t)}\,ds = \int_0^\infty e^{-(\delta - \sup m)u}\,du = \frac{1}{\delta - \sup m} < \infty, \tag{B26}$$

and $v > 0$ since the integrand is positive. *The lower bound is uniform.* Set $\inf m := \inf_S m(S) > -\infty$, finite because the jump integrand in (B25) is bounded ($g_t \in [\underline{r}^{1/\alpha}, \bar{r}^{1/\alpha}]$ by (8) and $\mathcal{J}$ is compact). Then $\widetilde{Z}_s = e^{-\inf m\,(s-t)} Y_s/Y_t$ has local drift $(m(S_s) - \inf m)\widetilde{Z}_s \ge 0$, a nonnegative-drift local submartingale that is uniformly integrable on compacts by Lemma B.2, hence a true submartingale, so $\mathbb{E}_t[Y_s/Y_t] \ge e^{\inf m\,(s-t)}$ and

$$v(\Psi_t) \ge \int_t^\infty e^{-\delta(s-t)} e^{\inf m\,(s-t)}\,ds = \frac{1}{\delta - \inf m} > 0,$$

a positive lower bound independent of Ψ_t (it uses only $\inf m$ and δ). With (B26) this makes

$\psi_t(x) = v(\Psi_t^x)/v(\Psi_{t-})$ bounded above and bounded below away from zero—the bound used in Proposition B.3 and the excess-volatility result of the Internet Appendix (whose Dynkin step is licensed by the PDMP extended-generator characterization, v absolutely continuous along the calm flow by its Feynman–Kac form; Davis 1993). *Uniqueness among bounded valuations.* If v_1, v_2 are bounded solutions of the pricing identity, then for every horizon $T > t$ their difference $D := v_1 - v_2$ satisfies $D(\Psi_t) = e^{-\delta(T-t)}\mathbb{E}_t[(Y_T/Y_t)\, D(\Psi_T)]$; bounding $|D| \leq \|D\|_\infty$ and using $\mathbb{E}_t[Y_T/Y_t] \leq e^{\sup m\,(T-t)}$ (the supermartingale bound above),

$$|D(\Psi_t)| \leq \|D\|_\infty\, e^{-(\delta - \sup m)(T-t)} \xrightarrow[T\to\infty]{} 0 \quad \text{under (T)},$$

so $D \equiv 0$ and the bounded valuation is unique. That the bounded valuation depends on the full record-resolved state Ψ_t, not on finitely many moments of S_t, has two sources, in order of generality: the law of motion is non-autonomous—the switch flow and revision deadlines are functionals of Ψ_t (Internet Appendix Propositions **??**, 3)—so the forward expectation cannot collapse to a function of S_t alone in *any* channel; and in the size channel the instantaneous disaster tilt is already infinite-dimensional in S_t (Proposition 3). □

Proposition B.3 (Equity premium and disaster premium): *In the equilibrium of Proposition 1, the tree's instantaneous expected excess return decomposes into a diffusive and a disaster part,*

$$\mathbb{E}_t[dR_t^e] - r_t\, dt = \underbrace{\alpha\sigma^2}_{\textit{diffusive}}\; dt + \mathrm{DRP}_t\, dt, \tag{B27}$$

where, with the tree's disaster return $\Delta R_t^e(x) = e^{-x}\psi_t(x) - 1$ and $\psi_t(x) = v(\Psi_t^x)/v(\Psi_{t-})$ the price–dividend jump ratio (the record-resolved state moving from Ψ_{t-} to its post-disaster value Ψ_t^x), the disaster risk premium is the gap between the physical and risk-neutral expected disaster returns,

$$\mathrm{DRP}_t = \lambda^\star\, \mathbb{E}_f\big[\Delta R_t^e(J)\big] - U_t\, \mathbb{E}_{f_t^{\mathbb{Q}}}\big[\Delta R_t^e(J)\big]. \tag{B28}$$

The diffusive premium $\alpha\sigma^2$ is constant. The disaster premium is priced by the risk-neutral *disaster law* $(U_t, f_t^{\mathbb{Q}})$*—the fear factor of Definition 2—so, holding* $(\alpha, \lambda^\star, F_J)$ *fixed, its belief-driven variation runs through* $(U_t, f_t^{\mathbb{Q}})$*. In the intensity channel (shape fixed) the* U_t*-loading component of this variation runs through* U_t *alone (Proposition 2); in the general joint model the premium also moves with the shape* $f_t^{\mathbb{Q}}$ *and the state-dependent return jump. Even in the intensity channel, though, the premium's* level *and part of its variation run through the jump ratio* ψ_t*, hence through* $v(\Psi_t)$ *(Proposition B.2), which varies with the full cross-section and is not summarized by* U_t*. In the decomposition of Proposition 5,* $\mathrm{DRP}_t = B_t^{\mathrm{DRP}}(S_t) + \mathcal{H}_t^{\mathrm{DRP}}(\Psi_t)$*: the cross-section part carries the* S_t*-measurable fear-factor intensity, and the survival-history residual enters through the price–dividend jump ratio* ψ_t *and* $v(\Psi_t)$ *(Proposition B.2), zero for any Markov benchmark.*

Proof of Proposition B.3. The no-arbitrage relation. The consumption claim priced by ξ has gain process $\xi_t P_t + \int_0^t \xi_s C_s\, ds$, a $\mathbb{P}$-martingale (Proposition 1). Conditioning on $\mathcal{F}_t$, its drift is zero:

$$\mathbb{E}_t[d(\xi_t P_t)] + \xi_t C_t\, dt = 0. \tag{B29}$$

By Itô's product rule for semimartingales,

$$d(\xi_t P_t) = \xi_{t-}\, dP_t + P_{t-}\, d\xi_t + d[\xi, P]_t. \tag{B30}$$

Divide (B30) by $\xi_{t-}P_{t-}$:

$$\frac{d(\xi_t P_t)}{\xi_{t-}P_{t-}} = \frac{dP_t}{P_{t-}} + \frac{d\xi_t}{\xi_{t-}} + \frac{d[\xi, P]_t}{\xi_{t-}P_{t-}}. \tag{B31}$$

The cum-dividend return is $dR_t^e = dP_t/P_{t-} + (C_t/P_t)\, dt$, so $dP_t/P_{t-} = dR_t^e - (C_t/P_t)\, dt$; substitute this into (B31), take $\mathbb{E}_t[\cdot]$, and use (B29) divided by the left limit $\xi_{t-}P_{t-}$ (i.e. $\mathbb{E}_t[d(\xi_t P_t)/(\xi_{t-}P_{t-})] = -(C_t/P_{t-})\, dt$ to first order, the price jump entering the martingale part, not the drift):

$$\mathbb{E}_t[dR_t^e] - \frac{C_t}{P_{t-}}dt + \mathbb{E}_t\Big[\frac{d\xi_t}{\xi_{t-}}\Big] + \mathbb{E}_t\Big[\frac{d[\xi, P]_t}{\xi_{t-}P_{t-}}\Big] = -\frac{C_t}{P_{t-}}dt.$$

Cancel $-(C_t/P_t)\, dt$ from both sides, then use $\mathbb{E}_t[d\xi_t/\xi_{t-}] = -r_t\, dt$ (Proposition B.1) and

$\mathbb{E}_t[d[\xi, P]_t/(\xi_{t-}P_{t-})] = \mathrm{Cov}_t(d\xi_t/\xi_{t-}, dR^e_t)$ (the predictable covariation, the $(C_t/P_t)dt$ yield term being finite-variation and contributing none):

$$\mathbb{E}_t[dR^e_t] - r_t\,dt = -\,\mathrm{Cov}_t\Big(\frac{d\xi_t}{\xi_{t-}},\,dR^e_t\Big).$$

Predictability and the jump structure. Placing the price jump in the martingale part of (B29) requires the jumps of P to be *inaccessible*, and they are. Between disasters the belief densities M^i are continuous—they jump only at the points of the disaster measure N, by (A1)—so the kernel ξ and, by Lemma B.1, the price P are continuous across every predictable revision time, including a synchronized calm deadline at which a positive G-mass re-selects on a flat level set. There the share-weighted distribution S_t carries a predictable finite-variation jump, but the *traded* price does not: the re-selecting cohort relocates to the public and *predictable* destination $\phi^*_t = \mathfrak{m}(\rho_t)$, an anticipated move that leaves the forward-looking value unchanged (Lemma B.1(iii)). The only jumps in P are at disasters, totally inaccessible stopping times of the Poisson filtration, so (B29) holds without any atomlessness restriction on the deadline law.

The two covariation channels. The martingale part of $d\xi_t/\xi_t$ is $-\alpha\sigma\,dW_t$ plus the compensated jump $\int_{\mathcal{J}}(e^{\alpha x}g_t(x)^\alpha - 1)(\mu - \nu)(dt, dx)$ ((B5), (A8)); the martingale part of dR^e_t is $\sigma\,dW_t$ (the record-resolved state Ψ_t has no Brownian part by Proposition 3, so the deterministic $v(\Psi_t)$ is diffusion-free and the sole Brownian source is $\sigma\,dW_t$ from C_t) plus the compensated jump $\int_{\mathcal{J}} \Delta R^e_t(x)(\mu - \nu)(dt, dx)$ with $\Delta R^e_t(x) = e^{-x}\psi_t(x) - 1$. The Brownian and Poisson sources are independent, so the covariation splits into a diffusion and a jump part. The diffusion part, using $d\langle W\rangle_t = dt$,

$$\mathrm{Cov}_t(-\alpha\sigma\,dW_t,\ \sigma\,dW_t) = -\alpha\sigma \cdot \sigma\,d\langle W\rangle_t = -\alpha\sigma^2\,dt.$$

For the jump part, the predictable covariation of two compensated integrals $\int H(\mu - \nu)$ and $\int K(\mu - \nu)$ against the same random measure is $\int HK\,\nu$, valid for predictable H, K with

$\int H^2\nu, \int K^2\nu < \infty$ (Brémaud, 1981); here $H = e^{\alpha x} g_t^\alpha - 1$ and $K = \Delta R_t^e$ are bounded (g_t bounded by (8), ψ_t bounded since v is bounded above and below away from zero under (T), $\mathcal{J}$ compact), so with $\nu = \lambda^\star\, dt\, f\, dx$,

$$\mathrm{Cov}_t^{\mathrm{jump}} = \lambda^\star \int_{\mathcal{J}} (e^{\alpha x} g_t(x)^\alpha - 1)\, \Delta R_t^e(x)\, f(x)\, dx\, dt.$$

Hence the premium is $-\mathrm{Cov}_t = -(\text{diffusion}) - (\text{jump})$:

$$\frac{\mathbb{E}_t[dR_t^e]}{dt} - r_t = \alpha\sigma^2 - \lambda^\star \int_{\mathcal{J}} (e^{\alpha x} g_t(x)^\alpha - 1)\, \Delta R_t^e(x)\, f(x)\, dx. \tag{B32}$$

Split the integral over the factor $(e^{\alpha x} g_t^\alpha - 1)$:

$$\begin{aligned} -\lambda^\star \int_{\mathcal{J}} (e^{\alpha x} g_t^\alpha - 1) \Delta R_t^e f\, dx &= \lambda^\star \int_{\mathcal{J}} \Delta R_t^e f\, dx - \lambda^\star \int_{\mathcal{J}} e^{\alpha x} g_t^\alpha \Delta R_t^e f\, dx \\ &= \lambda^\star\, \mathbb{E}_f[\Delta R_t^e] - U_t\, \mathbb{E}_{f_t^{\mathbb{Q}}}[\Delta R_t^e], \end{aligned} \tag{B33}$$

line (B33) using $\int \Delta R_t^e f\, dx = \mathbb{E}_f[\Delta R_t^e]$ and the risk-neutral disaster compensator $\lambda^\star e^{\alpha x} g_t(x)^\alpha f(x) = \nu_t^{\mathbb{Q}}(dx)/dt = U_t f_t^{\mathbb{Q}}(x)$ (10), so that $\lambda^\star \int e^{\alpha x} g_t^\alpha \Delta R_t^e f\, dx = U_t \int \Delta R_t^e f_t^{\mathbb{Q}}\, dx = U_t \mathbb{E}_{f_t^{\mathbb{Q}}}[\Delta R_t^e]$. Substituting (B33) into (B32) gives the diffusive term $\alpha\sigma^2$ and $\mathrm{DRP}_t = \lambda^\star \mathbb{E}_f[\Delta R_t^e] - U_t \mathbb{E}_{f_t^{\mathbb{Q}}}[\Delta R_t^e]$, which are (B27)–(B28). $\square$

Proposition B.4 (State reduction to the wealth-weighted distribution): *In the equilibrium of Proposition 1, each of the following instantaneous equilibrium price quantities is a measurable functional of the pricing state S_t (Definition 3) alone: the belief aggregator $g_t(\cdot)$ (17), the kernel calm drift and disaster jump ratio (Internet Appendix* (??)–(??)*) (jump quantities at a disaster epoch evaluated at the predictable left limit S_{t-}), the fear factor U_t and multiplier κ_t and the risk-neutral disaster law $f_t^{\mathbb{Q}}$* (10)–(11)*, the short rate r_t* (B1)*, the local valuation growth $m(S_t)$* (B14)*, and the diffusive premium $\alpha\sigma^2$* (B27)*. The disaster-premium* level *is excluded: it carries ψ_t (Proposition B.3). Equivalently, any two record-resolved states Ψ_t, Ψ_t' (Definition 4) with a* *common held-model marginal $\pi_\phi \Psi_t = \pi_\phi \Psi'_t = S_t$ yield identical values of all these quantities at t: the since-adoption records, regime ages, surprise thresholds, and second-order posteriors enter instantaneous prices only through S_t. These prices factor through the aggregator $g_t(\cdot)$, itself a functional of S_t coarser than the full measure; what fails under size disagreement is any* finite-dimensional *reduction—by Proposition 3 no finite collection of linear moments of S_t determines $g_t(\cdot)$; in the intensity channel $g_t \equiv G_t$ is a scalar, and on a finite menu S_t is itself a point of the simplex, so the non-reduction is a size-disagreement statement. The forward-looking objects do* not *reduce to S_t: the price–dividend ratio $v(\Psi_t)$ (B13), the price–dividend jump ratio ψ_t, and hence the* level *of the disaster premium* DRP_t (B28) *are functionals of the full record-resolved state Ψ_t (Proposition B.2).*

Proof of Proposition B.4. Each instantaneous quantity is, by its defining display, a functional of S_t.

(a) By (17), $g_t(x) = \int_\Phi (\lambda/\lambda^\star)^{1/\alpha} \big(f_\theta(x)/f(x)\big)^{1/\alpha} S_t(d\phi)$ is the integral of a fixed (t, Ψ)-independent kernel against S_t; so the whole function $g_t(\cdot)$ depends on Ψ_t only through $S_t = \pi_\phi \Psi_t$.

(b) By Definition 3, $\bar{\lambda}^S_t = \int_\Phi \lambda \, S_t(d\phi)$. Hence the calm drift (Internet Appendix (**??**)), $-\frac{1}{\alpha}(\bar{\lambda}^S_t - \lambda^\star)$, and the disaster jump ratio (Internet Appendix (**??**)), $g_t(x)$, are functionals of S_t by (a).

(c) By (10)–(11), $U_t = \lambda^\star \, \mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]$, $\kappa_t = U_t/U^{\mathrm{RA}}$, and $f^{\mathbb{Q}}_t(x) = e^{\alpha x} g_t(x)^\alpha f(x)/\mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]$ are fixed functionals of $g_t(\cdot)$, hence of S_t by (a).

(d) By (B1), $r_t = \delta + \alpha\mu_0 - \frac{1}{2}\alpha(1+\alpha)\sigma^2 + \bar{\lambda}^S_t - U_t$, and by (B14) $m(S_t)$, are functionals of $\bar{\lambda}^S_t$ and $g_t(\cdot)$ by (b)–(c); the diffusive premium $\alpha\sigma^2$ (B27) is constant.

The displays in (a)–(d) involve Ψ_t only through S_t; the coordinates $(\mathcal{R}, \tau, \varepsilon)$ of Definition 4 appear in none of them. Therefore $\pi_\phi \Psi_t = \pi_\phi \Psi'_t$ forces equality of every quantity in (a)–(d), the stated equivalence. That the instantaneous prices admit no *finite-dimensional* reduction—though they factor through the coarser aggregator $g_t(\cdot)$—is Proposition 3: the map $S \mapsto g_t(\cdot)$

in (a) is into an infinite-dimensional image (by the spanning condition, the family $\{(f_\theta/f)^{1/\alpha}\}$ is linearly independent), so no finite collection of *linear* moments of S_t is sufficient (nonlinear summaries are not excluded by this argument, and the intensity channel and the finite menu do reduce); the claim is the size-disagreement one. For the forward-looking objects, $v(\Psi_t) = \mathbb{E}_t[\int_t^\infty e^{-\delta(s-t)}(C_s/C_t)^{1-\alpha}(\Xi_s/\Xi_t)^\alpha ds]$ (B13) is a conditional expectation of a functional of the future path $\{S_s\}_{s\geq t}$, whose conditional law is governed by the record-resolved generator of Proposition 3, Markov in Ψ_t and not autonomous in S_t (Proposition B.2); so v is a conditional expectation of a functional of the future path whose law is fixed by Ψ_t and not by S_t alone; v is therefore a functional of the full record-resolved state, as claimed, and where two such states with a common marginal produce different forward laws the values can differ—the constructive instance is the deadline channel of Theorem 1(iii), established independently of this proposition's forward clause. The jump ratio $\psi_t = v(\Psi_t^J)/v(\Psi_{t-})$ and the premium level (B28) inherit this dependence. □

REFERENCES

Abel, Andrew B., Janice C. Eberly, and Stavros Panageas, 2013, Optimal inattention to the stock market with information costs and transactions costs, *Econometrica* 81, 1455–1481.

Andersen, Torben G., Nicola Fusari, and Viktor Todorov, 2015a, Pricing of tail risk and the equity premium, *Working Paper* .

Andersen, Torben G., Nicola Fusari, and Viktor Todorov, 2015b, The risk premia embedded in index options, *Journal of Financial Economics* 117, 558–584.

Andrès, Hervé, Alexandre Boumezoued, and Benjamin Jourdain, 2024, The implied volatility surface (also) is path-dependent, Working paper, arXiv:2312.15950.

Bacchetta, Philippe, and Eric van Wincoop, 2010, Infrequent portfolio decisions: A solution to the forward discount puzzle, *American Economic Review* 100, 870–904.

Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters implied by equity index options, *Journal of Finance* 66, 1969–2012.

Baron, Matthew, Wei Xiong, and Zhijiang Ye, 2024, Measuring time-varying disaster risk: An empirical analysis of "dark matter" in asset prices, Working paper, Cornell University and Princeton University.

Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.

Barro, Robert J., and Tao Jin, 2011, On the size distribution of macroeconomic disasters, *Econometrica* 79, 1567–1589.

Barro, Robert J., and José F. Ursúa, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity* 2008, 255–335.

Basak, Suleyman, 2005, Asset pricing with heterogeneous beliefs, *Journal of Banking and Finance* 29, 2849–2881.

Bastianello, Francesca, and Paul Fontanier, 2026, Partial equilibrium thinking, extrapolation, and bubbles, *Review of Financial Studies* forthcoming.

Berk, Robert H., 1966, Limiting behavior of posterior distributions when the model is incorrect, *Annals of Mathematical Statistics* 37, 51–58.

Blanchard, Olivier J., 1985, Debt, deficits, and finite horizons, *Journal of Political Economy* 93, 223–247.

Blume, Lawrence, and David Easley, 2006, If you're so smart, why aren't you rich? belief selection in complete and incomplete markets, *Econometrica* 74, 929–966.

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears, and risk premia, *Journal of Finance* 66, 2165–2211.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2020, Memory, attention, and choice, *Quarterly Journal of Economics* 135, 1399–1442.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621–651.

Brémaud, Pierre, 1981, *Point Processes and Queues: Martingale Dynamics* (Springer).

Campbell, John Y., 2018, *Financial Decisions and Markets: A Course in Asset Pricing* (Princeton University Press, Princeton, NJ).

Chen, Hui, Winston Wei Dou, and Leonid Kogan, 2024, Measuring the "dark matter" in asset pricing models, *Journal of Finance* 79, 843–902.

Chen, Hui, Scott Joslin, and Ngoc-Khanh Tran, 2012, Rare disasters and risk sharing with heterogeneous beliefs, *Review of Financial Studies* 25, 2189–2224.

Cho, In-Koo, and Kenneth Kasa, 2015, Learning and model validation, *Review of Economic Studies* 82, 45–82.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer, 2016, Parameter learning in general equilibrium: The asset pricing implications, *American Economic Review* 106, 664–698.

Cox, John C., and Chi-fu Huang, 1989, Optimal consumption and portfolio policies when asset prices follow a diffusion process, *Journal of Economic Theory* 49, 33–83.

Cvitanić, Jakša, and Semyon Malamud, 2011, Price impact and portfolio impact, *Journal of Financial Economics* 100, 201–225.

d'Arienzo, Daniele, 2020, Maturity increasing overreaction and bond market puzzles, Working paper, Bocconi University.

Davis, Mark H. A., 1993, *Markov Models and Optimization* (Chapman & Hall, London).

Dew-Becker, Ian, Stefano Giglio, and Pooya Molavi, 2025, Learning and the emergence of nonlinearity in financial markets, NBER Working Paper 34584, National Bureau of Economic Research.

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.

Gârleanu, Nicolae, and Stavros Panageas, 2015, Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing, *Journal of Political Economy* 123, 670–685.

Ghaderi, Mohammad, Mete Kilic, and Sang Byung Seo, 2022, Learning, slowly unfolding disasters, and asset prices, *Journal of Financial Economics* 143, 527–549.

Giglio, Stefano, and Bryan Kelly, 2018, Excess volatility: Beyond discount rates, *Quarterly Journal of Economics* 133, 71–127.

Gormsen, Niels Joachim, and Christian Skov Jensen, 2025, Higher-moment risk, *Journal of Finance* forthcoming.

Guyon, Julien, and Jordan Lekeufack, 2023, Volatility is (mostly) path-dependent, *Quantitative Finance* 23, 1221–1258.

Hansen, Lars Peter, and Thomas J. Sargent, 2010, Fragile beliefs and the price of uncertainty, *Quantitative Economics* 1, 129–162.

Jacod, Jean, and Albert N. Shiryaev, 2003, *Limit Theorems for Stochastic Processes*, second edition (Springer).

Jouini, Elyes, and Clotilde Napp, 2007, Consensus consumer and intertemporal asset pricing with heterogeneous beliefs, *Review of Economic Studies* 74, 1149–1174.

Karatzas, Ioannis, John P. Lehoczky, and Steven E. Shreve, 1990, Existence and uniqueness of multi-agent equilibrium in a stochastic, dynamic consumption/investment model, *Mathematics of Operations Research* 15, 80–128.

Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield, 2006, The price impact and survival of irrational traders, *Journal of Finance* 61, 195–229.

Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield, 2017, Market selection, *Journal of Economic Theory* 168, 209–236.

Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran, 2020, The tail that wags the economy: Belief-driven business cycles and persistent stagnation, *Journal of Political Economy* 128, 2839–2879.

Lanzani, Giacomo, 2025, Dynamic concern for misspecification, *Econometrica* 93, 1333–1370.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: Do macroeconomic experiences affect risk taking?, *Quarterly Journal of Economics* 126, 373–416.

Menkhoff, Manuel, 2025, The devil is in the tail: Macroeconomic tail risk expectations of firms, Working paper, CESifo.

Menzies, Gordon D., and Daniel John Zizzo, 2009, Inferential expectations, *The B.E. Journal of Macroeconomics* 9.

Nagel, Stefan, and Zhengyang Xu, 2022, Asset pricing with fading memory, *Review of Financial Studies* 35, 2190–2245.

Ortoleva, Pietro, 2012, Modeling the change of paradigm: Non-bayesian reactions to unexpected news, *American Economic Review* 102, 2410–2436.

Sandroni, Alvaro, 2000, Do markets favor agents able to make accurate predictions?, *Econometrica* 68, 1303–1341.

Seo, Sang Byung, and Jessica A. Wachter, 2019, Option prices in a model with stochastic disaster risk, *Management Science* 65, 3449–3469.

Teguia, Alberto Mokak, 2026, Endogenous belief hazards, Working paper, University of British Columbia.

Wachter, Jessica A., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *Journal of Finance* 68, 987–1035.

Wachter, Jessica A., and Michael J. Kahana, 2024, A retrieved-context theory of financial decisions, *Quarterly Journal of Economics* 139, 1095–1147.

Wachter, Jessica A., and Yicheng Zhu, 2025, Learning with rare disasters, *Quantitative Economics* 16, 1189–1221.

Wang, Chen, 2021, Under- and overreaction in yield curve expectations, Working paper, University of Notre Dame.

Weinstein, Jonathan, 2017, Bayesian inference tempered by classical hypothesis testing, Working paper, Washington University in St. Louis.

Yaari, Menahem E., 1965, Uncertain lifetime, life insurance, and the theory of the consumer, *Review of Economic Studies* 32, 137–150.

Yan, Hongjun, 2008, Natural selection in financial markets: Does it work?, *Management Science* 54, 1935–1950.

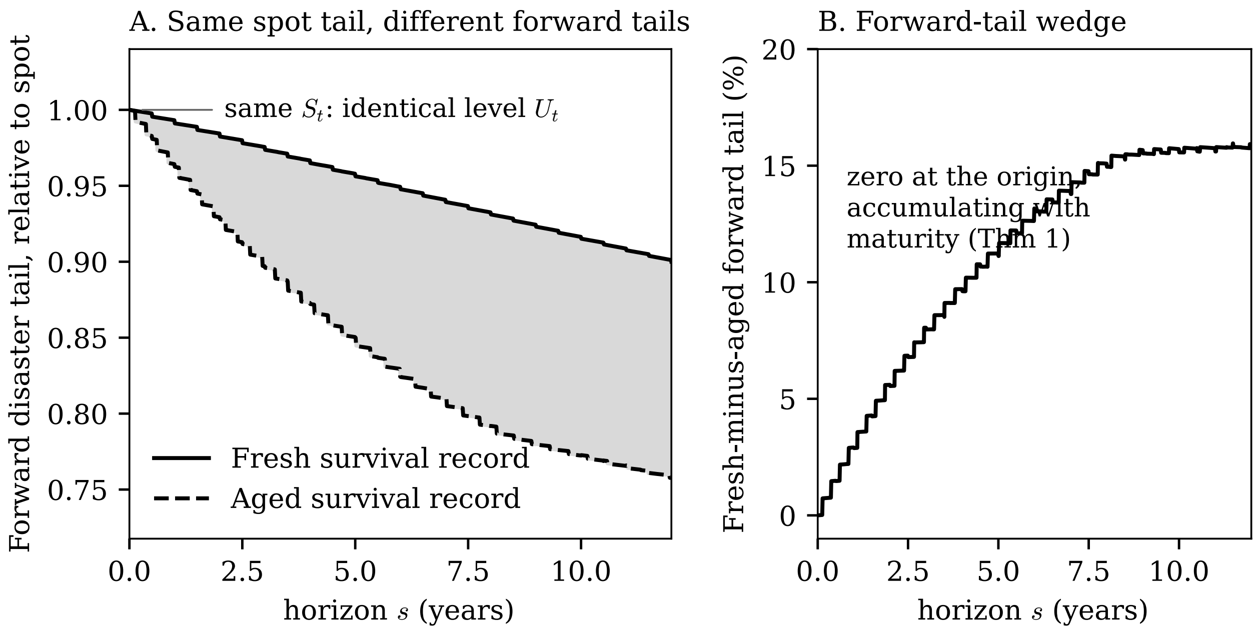

Figure 2: **Two histories, one cross-section (Theorem 1(iii)).** At $t^* = 44$ along a realized disaster record—fourteen years into the calm that follows a disaster at $t = 30$—two record-resolved states share the same held-model cross-section S_{t^*}, hence the same spot tail U_{t^*}, and differ only in their records: *aged* is the state the realized history delivers, its records as the path left them—the disaster's re-anchoring wave, later calm re-adoptions, and records carried through the disaster untriggered; *fresh* is the same cross-section with every since-adoption record reset. Left: the calm-conditional forward tail $F_t^0(s) := \mathbb{E}[U_{t+s} \mid \mathcal{F}_t$, no disaster in $(t, t+s]]$, normalized by the common level. Conditional on no disaster the path is a deterministic functional of the record-resolved state, so both curves are deterministic numerical transports of the model laws, not simulations. Right: the wedge between the curves—zero at the origin and accumulating with maturity as the aged state's calm deadlines arrive, reaching 15.8–16.2% of the level at a twelve-year horizon across converged configurations, under the calibrated belief block of Section 6.B. This is the survival-history residual formalized below (Proposition 5), the object any belief-sufficient benchmark sets to zero.

Table II: Confronting the option-tail evidence. The physical block is fixed at the macro-disaster calibration (Table I), giving $U^{\mathrm{RA}} = 0.085/\mathrm{yr}$; the belief block is pinned to the crisis jump and the calm moderation; the persistence is achieved, not targeted. The model entries are closed-form in the fear factor $U_t = U^{\mathrm{RA}}\kappa_t$, the premium rows also carrying the jump ratio ψ_t. The Status column separates what is calibrated, what is directional, what is structural, and what is achieved; the half-life is generated by the deadline and turnover blocks rather than targeted. The computed record is the single event-study path of this subsection: path statistics, not stationary averages.

Evidence	Documented finding	Model (closed form)	Status
Backus et al. (2011)	options imply a more modest disaster law than the macro record	deep-calm $\kappa_t \approx 0.53$–$0.63 < 1$: U_t falls below the benchmark U^{RA} late in calm; the point level $U_t \approx \lambda^\star$ ($\bar{\kappa} = 0.35$) lies below the two-point menu's re-anchoring floor and is not attained	direction matched; level not attained
Bollerslev and Todorov (2011)	the risk-neutral jump tail exceeds the statistical one (a fear premium), time-varying	mean $U_t/\lambda^\star \approx 3$ over the computed record—a single-path statistic, not a stationary mean (1.5–1.6, $\kappa_t \approx 0.53$–0.57, entering disasters); spikes 4.5–4.8× ($\kappa_t \approx 1.60$–1.70), inside the documented 3.2–4.8 crisis band	calibrated target
Bollerslev and Todorov (2011)	rare-event compensation is a large fraction of the equity premium	diffusive compensation is fixed at $\alpha\sigma^2 = 0.14\%/\mathrm{yr}$, so any sizable equity premium here is disaster compensation; its level and share carry the valuation-jump ratio ψ_t (B28), not solved quantitatively	structural implication; premium level unsolved
Andersen et al. (2015b)	the tail factor predicts the equity premium; volatility does not	DRP_t loads on U_t while σ is held constant, so the volatility loading is zero by construction	qualitative implication of the restricted calibration
Andersen et al. (2015a)	the tail factor is far more persistent than volatility	post-crisis wave half-life of $U_t \approx 12$ years on the computed record $\gg$ the far shorter-lived volatility factors	achieved moment

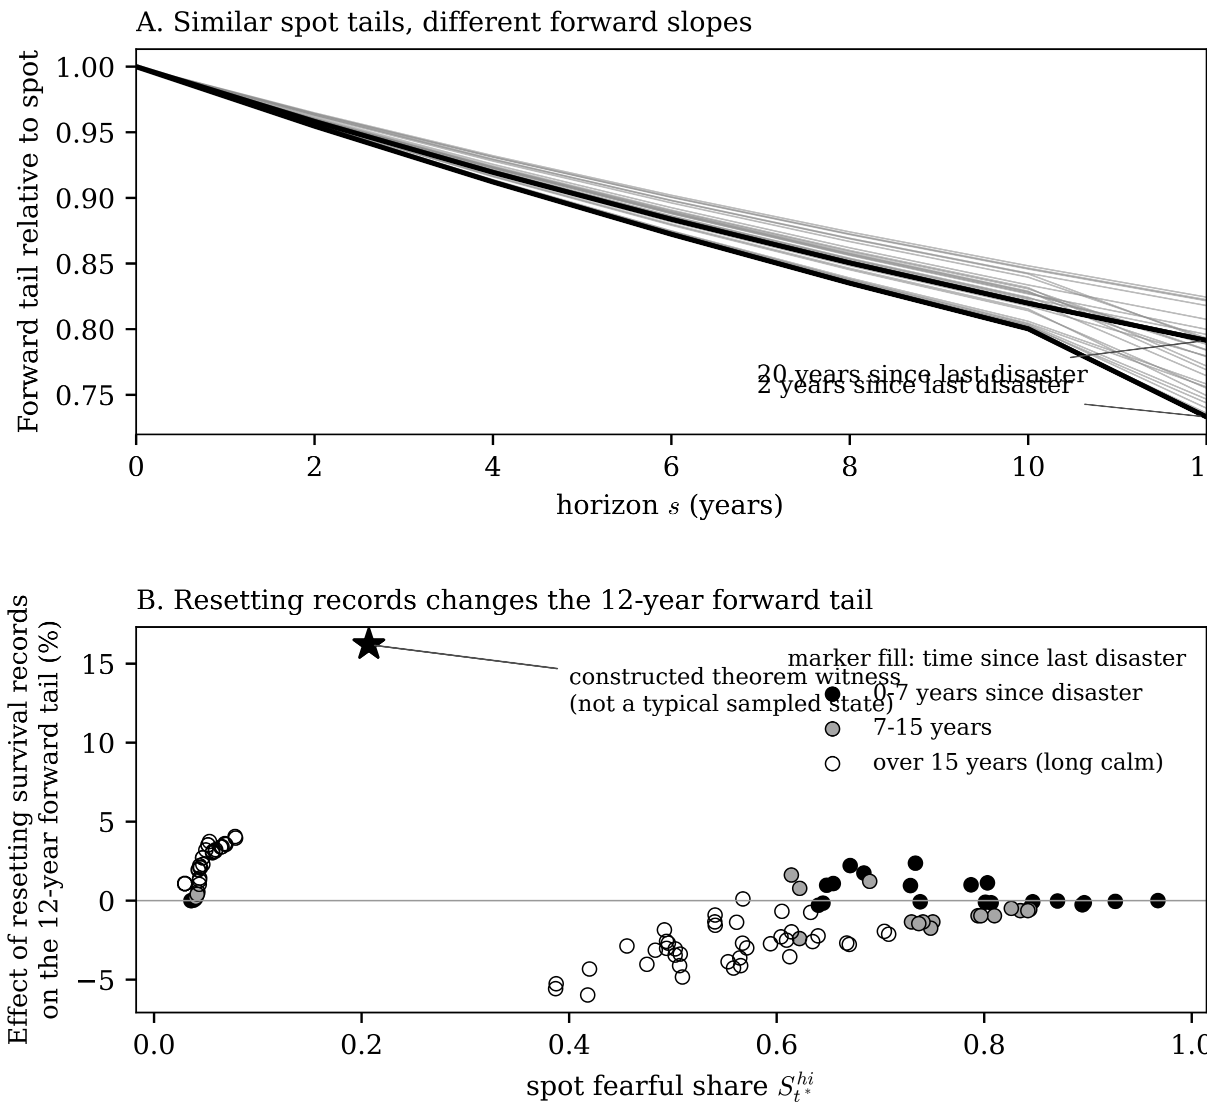

Figure 3: **The forward slope across records.** Top: the calm-conditional forward tail $F^0_{t^*}(s)/U_{t^*}$ for the twenty-nine sampled sixty-year histories in a narrow comparison window inside the post-disaster band, $S^{hi}_{t^*} \in [0.55, 0.70)$; the sharpest near-match (bold: $\Delta S^{hi} = 0.001$, two and twenty years since the last disaster) differs by 5.8 percentage points in the twelve-year slope. Bottom: each history's twelve-year record wedge $[F^{\text{reset}}_{t^*} - F^{\text{actual}}_{t^*}]/F^{\text{reset}}_{t^*}$ against its spot state, marker fill coding time since the last disaster (filled: within seven years; gray: seven to fifteen; open: longer); positive wedges mean resetting raises the forward tail, negative that it lowers it; the star is the constructed theorem witness of Figure 2, not a typical sampled state. One hundred fifty independent sixty-year histories from the calibrated initial condition (a finite-history ensemble, not draws from the invariant law) at the calibration of Table I; construction, regressions, and the locked-destination regime in Internet Appendix ??.

Table III: **The forward slope across sampled records.** One hundred fifty independent sixty-year histories at the calibration of Table I. Slope is $F^0_{t^*}(12)/U_{t^*} - 1$; cross-record dispersion is its standard deviation within band after residualizing on $S^{hi}_{t^*}$; the record wedge is $[F^{\text{reset}} - F^{\text{actual}}]/F^{\text{reset}}$ at twelve years. Panel A reports moments of the simulated histories; the constructed witness appears only in the note.

Band ($S^{hi}_{t^*}$)	Histories	Median slope	Cross-record dispersion	Wedge, % (q10/5
Deep calm, $[0, 0.10)$	73	-0.001	0.007	0.4 / 1.1 / 3
Transient middle, $[0.10, 0.40)$	2	-0.133	—	-5.5, -5.4 (two h
Post-disaster, $[0.40, 1]$	75	-0.228	0.019	-3.8 / -1.3 /
Panel B: within-band regressions of the twelve-year slope (model-implied R^2, sampled histories)				
Region	Baseline	Added record proxy	R^2: base $\rightarrow$ augmented	IQR effect
Post-disaster	S^{hi}	time since disaster	$0.64 \rightarrow 0.92$	+5.3 pp
Deep calm	S^{hi}	deadline-proximity mass	$0.87 \rightarrow 0.99$	-0.10 pp

The constructed witness of Figure 2 lies in the transient band: $S^{hi} = 0.21$, wedge 16.2%—a constructed benchmark, not a sampled history.