Internet Appendix for
"Belief Records and the Price of Disaster Fear"

July 2026

This Internet Appendix collects Appendices C–K of the paper and, in a supplementary part, the demographic and stationarity proofs, secondary results, the benchmark comparisons, and the numerical stability checks: the proofs for the setup, the belief dynamics, and the revision hazard and law of motion of the pricing state; the dynamics around a disaster; the record-independent-clock benchmark; the discussion of what history prices; the quantitative development; and the demographic turnover and stationary constructions. Appendix lettering and result numbering are continuous with Appendices A–B of the paper, and cross-references resolve across the two documents.

Appendix C. Proofs for Section ??

Two facts about the disputed object are used throughout. First, the held and physical laws are mutually absolutely continuous on every finite horizon (Lemma C.1), so the likelihood ratios that drive revision are well posed. Second, the truth diverges from a held model ϕ at a per-period Kullback–Leibler rate $\mathcal{K}(\phi)$ (Definition C.1) that splits additively into a timing part and a size part and is minimized at a unique pseudo-true model $\phi^\star = (\lambda^\star, \theta^\star)$ (Lemma C.2): the held intensity is asymptotically pinned to the truth $\lambda^\star$, while whether the size family is correctly specified is left open and matters only in Section **??**.

Proof of the example proposition (Section **??***).* (i) κ_0 is (**??**) at the common s_0. (ii) Three timing checks make the window clean. The complacent model's exact quiet-spell deadline, ≈ 72 years, lies beyond every horizon here, so only the fearful cohorts can revise; before either does, both economies run the replicator (**??**), a function of the cross-section alone, from the common s_0. In Aged the fearful deadline $\Delta^\star - a \approx 3.9$ years precedes the window (any adoption age in $(\Delta^\star - 5, \Delta^\star) \approx (6.9, 11.9)$ years gives the same split): its wealth relocates to the posterior mode—complacent under the even public prior after any quiet spell—so the economy is degenerate at λ_{lo} from year four on and $\kappa_5 = \lambda_{lo}/\lambda^\star = 0.333$. In Fresh the fearful deadline $\Delta^\star \approx 11.9$ years lies beyond the window, only the replicator acts, and $s_5 = 0.237$, $\kappa_5' = 0.558$. (iii) A functional of the date-0 state takes one value on the two economies, which share that state. $\square$

This appendix states and proves Lemma C.1, Lemma C.2, and Lemma IA.1. The first is a direct computation from the change-of-measure density; the second is the Kullback–Leibler rate calculation; and the third invokes the high-frequency threshold theory of Mancini (2009), whose hypotheses we verify by name before assembling the lemma.

The identification and regularity conditions on the belief family—Assumption **??** of the main text, Section **??**'s deferred conditions—are maintained throughout and verified below for the families used in the paper.

The conditions invoked in Assumption **??** play the following roles.

Condition (iv) is a separate identification requirement from (iii), neither implying the other: where (iii) asks injectivity of $\theta \mapsto F_\theta$, (iv) asks uniqueness of the Kullback–Leibler projection of the truth onto the family. It holds, for instance, whenever $\{F_\theta\}$ is a regular exponential family in θ (where $\theta \mapsto \mathrm{KL}(F_J \| F_\theta)$ is strictly convex), and it is what makes the pseudo-true model a point rather than a set. It is invoked wherever a single pseudo-true point is needed—the pseudo-true model $\phi^\star$ of (C3), the smooth pseudo-true/asymptotic-normality setup of Assumption **??**(vi), and the concentration limits of Remark E.1 and Proposition K.1; without it each such statement comes with $\arg\min_\theta \mathrm{KL}(F_J \| F_\theta)$ in place of a point. Condition (v) is the priced-aggregator analogue of identifiability: it ensures distinct *finite* size-belief mixtures leave distinct imprints on the risk-neutral tail (finite-mixture separation; it does not by itself give injectivity of the full mixture map over arbitrary continuum-supported states, which would need a completeness or transform-uniqueness condition). It holds for the families used here by distinct-exponent independence, not by analyticity alone. For the power-law size family $f_\theta(x) \propto x^{-(\theta+1)}$ on $\mathcal{J}$, $(f_\theta/f)^{1/\alpha} = c_\theta\, x^{-(\theta+1)/\alpha} f(x)^{-1/\alpha}$, which reduces to the monomial $c_\theta\, x^{(\theta_0-\theta)/\alpha}$ when $f = f_{\theta_0}$ lies in the family; in the misspecified case the common factor $f^{-1/\alpha}$ is shared across θ and so leaves independence intact—distinct indices give distinct real exponents, linearly independent on any non-degenerate interval (their Wronskian is a nonzero Vandermonde determinant in the exponents); the same holds for any family reducible to distinct-exponent or regular-exponential form. It is used only in Proposition **??** (the infinite-dimensionality of the pricing state). Condition (vi) is the standard regularity under which a possibly-misspecified maximum-likelihood estimator is $\sqrt{n}$-asymptotically normal (White, 1982; van der Vaart, 1998); it strengthens (ii) and, under correct specification, has its sandwich collapse ($H = B = I(\theta^\star)$), recovering the Fisher-information form of (ii). It is used only in Lemma IA.1(iii).

Conventions carried from Section **??**. The rejection criterion is the surprise $\bar{\varepsilon}$ (**??**): the held model is rejected iff $\bar{\varepsilon} \leq \varepsilon^i$, the boundary counted as a rejection. The closed-threshold form $\{D_{\phi,\Delta} \geq r_{\phi,\Delta}(\varepsilon)\}$ used in later expressions is shorthand for this p-value criterion; the operative criterion is always the p-value rule $\bar{\varepsilon} \leq \varepsilon^i$. Along a disaster-free stretch $\bar{\varepsilon}$ is piecewise continuous in the age, stepping downward where an integer count leaves the tail set, and need not attain the tolerance; the calm deadline $\Delta^\star$ is the first-exit time $\inf\{\Delta > 0 : \bar{\varepsilon} \leq \varepsilon^i\}$, finite (Proposition D.1), and rejection is effective at $\Delta^\star$ whether the tolerance is attained or passed by a step. The predictive law of D carries an atom at the zero-count value ($e^{-\lambda\Delta} > 0$), so the shorthand's agreement with a strict quantile cutoff is a convention, not a continuity fact, and at any atom the p-value rule governs. The test is two-sided: the deviance is large whether disasters run too frequent or too rare, and whether the realized sizes pull $\hat{\theta}$ above or below θ. Two divergence rates also stay distinct: the full KL (C2) fixes second-order posterior concentration, the pseudo-true model (C3), and survival (Lemma C.2), while the since-adoption rejection deviance that triggers revision drifts at the Poisson KL in timing and at the *excess* in-family size divergence relative to the pseudo-true fit, which can vanish at $\theta^\star$ even when the full size KL is positive—so the two coincide only when the size family is correctly specified.

Why parameters, not outcomes. The disputed primitives are parameters of laws, not realized outcomes: a Dirac mass on a fixed disaster size is falsified by the first continuous draw, whereas disputing the size *law* and the arrival *rate* through their parameters is the faithful object, the one the data reveal only slowly (full support, Assumption **??**(i)).

The rule in the literature. The rule is Ortoleva's, applied in real time to the accumulating record as in the model-validation learning of Cho and Kasa (2015), so the rejection time is a first exit; by Sanov's theorem (Dembo and Zeitouni, 1998; Cover and Thomas, 2006) the acceptance region is the empirical law lying within a Kullback–Leibler ball of the held model. Fixed-tolerance testing makes recurrence the *foregone conclusion* of repeated significance testing under continuous monitoring (Anscombe, 1954; Armitage et al., 1969; Robbins, 1970; Robbins and Siegmund, 1970): the standardized record crosses a fixed band infinitely often by the law of the iterated logarithm. Cho and Kasa (2015), who share the test-and-reselect primitive, instead engineer this away with a time-varying threshold and vanishing gain; recurrent regime change also arises as constant-gain escape dynamics (Kasa, 2004; Cho and Kasa, 2008). The consolidating side is classical: under countable additivity a Bayesian cannot be driven to a foregone conclusion (Kadane et al., 1996) and the posterior merges with the truth (Blackwell and Dubins, 1962; Kalai and Lehrer, 1993). That the since-adoption sufficient statistic—the count in the intensity channel; with size disagreement the (count, marks) pair, low-dimensional only for particular exponential-family size laws—stays sufficient for the held model at the *endogenous* revision time follows from the exponential-family (Wald) factorization of the disaster-count likelihood, the disaster channel being the last-observation-sufficient member of the families of Küchler and Lauritzen (1989) and Küchler and Sørensen (1989), for which sufficiency survives optional stopping (Küchler and Sørensen, 1998). The architecture itself—testing the held model on its since-adoption record while re-selecting on the full-history posterior—is the axiomatized primitive of the companion paper (Teguia, 2026): the acceptance region's interval form is forced (an essentially complete class), the test size is the single behavioral scalar, the per-test reading carries a per-shift consistency control (Weinstein, 2017), and stakes-invariance separates the surprise-governed hazard from cost-based rivals. We import the rule; the equilibrium is this paper's object.

We first record the change-of-measure and divergence facts.

Lemma C.1 (Equivalence of held and physical laws on finite horizons): *Under Assumptions* **??** *and* **??**, *for every* $\phi = (\lambda, \theta) \in \Phi$ *and every* $t > 0$, $\mathbb{P}|_{\mathcal{F}_t} \sim \mathbb{P}^\phi|_{\mathcal{F}_t}$, *with*

$$\frac{d\mathbb{P}^\phi}{d\mathbb{P}}\Big|_{\mathcal{F}_t} = \exp\big(-(\lambda - \lambda^\star)t\big) \prod_{k:\, T_k \le t} \frac{\lambda\, f_\theta(J_k)}{\lambda^\star\, f(J_k)} \;\in\; (0, \infty) \quad \mathbb{P}\text{-}a.s. \tag{C1}$$

Proof of Lemma C.1. Imported result (marked-Poisson change of measure; Brémaud, 1981, Ch. VI; Jacod and Shiryaev, 2003, III.3). If $\mathbb{P}, \mathbb{P}^\phi$ are the laws on $(\Omega, \mathcal{F}_t)$ of marked Poisson processes with (intensity, mark law) $(\lambda^\star, F_J)$ and (λ, F_θ), $\lambda^\star, \lambda \in (0, \infty)$, and if $F_J \sim F_\theta$, then $\mathbb{P}|_{\mathcal{F}_t} \sim \mathbb{P}^\phi|_{\mathcal{F}_t}$ with Radon–Nikodym density (C1). *Hypotheses in our model.* (1) $\lambda^\star > 0$ (Assumption **??**) and $\lambda \in [\underline{\lambda}, \overline{\lambda}] \subset (0, \infty)$ (Assumption **??**); (2) f and f_θ are continuous and strictly positive on the common compact support $\mathcal{J}$ (Assumption **??**; Assumption **??**(i)), so F_J and F_θ have identical null sets, i.e. $F_J \sim F_\theta$. Both hypotheses hold, so the import applies and yields (C1). Finally N has finitely many jumps on $[0, t]$ $\mathbb{P}$-a.s. (Assumption **??**), so the product is a finite product of strictly positive factors and the exponential is strictly positive; hence the density lies in $(0, \infty)$ $\mathbb{P}$-a.s., and $\mathbb{P}, \mathbb{P}^\phi$ share their null sets.

On the full filtration. Lemma C.1 is stated on the consumption-path filtration $\mathcal{F}_t$, which carries both the Brownian component σW and the marked-Poisson disaster component. The diffusion (μ_0, σ) is undisputed (common knowledge, Assumption **??**): every held law $\mathbb{P}^\phi$ shares the law of the continuous part with $\mathbb{P}$ and differs only in the disaster law (λ, F_θ). Because the Brownian and marked-Poisson components are independent under each law, the Radon–

Nikodym density on $\mathcal{F}_t$ factors as the (identity) density of the unchanged Brownian law times the marked-Poisson density (C1); the change of measure acts only on the jump factor, and the equivalence $\mathbb{P}|_{\mathcal{F}_t} \sim \mathbb{P}^{\phi}|_{\mathcal{F}_t}$ holds on the full filtration. (The factorization is legitimate there because disaster times and sizes are recovered pathwise from the consumption path—the jumps of $\log C$—with the Brownian component the remainder, so the marked-Poisson record is measurable on the consumption-path filtration; the usual $\mathbb{P}$-augmentation adds only null sets, under which the densities and the factorization persist.) $\square$

Definition C.1 (Per-period Kullback–Leibler divergence rate): *For a held model $\phi = (\lambda, \theta) \in \Phi$, let $\mathbb{P}^{\phi}$ be its law—the marked Poisson process of intensity λ with i.i.d. marks F_θ (Assumption **??**)—and $\mathbb{P}$ the physical law (Assumption **??**). The* per-period Kullback–Leibler rate *is the divergence of the truth from the held model, measured on the unit interval $\mathcal{F}_1$,*

$$\mathcal{K}(\phi) \;:=\; \mathrm{KL}\big(\mathbb{P} \,\big\|\, \mathbb{P}^{\phi}\big)\big|_{\mathcal{F}_1} \;=\; \mathbb{E}_{\mathbb{P}}\Big[\log \tfrac{d\mathbb{P}}{d\mathbb{P}^{\phi}}\Big|_{\mathcal{F}_1}\Big],$$

the $\mathbb{P}$-expectation of the log-likelihood ratio of the two laws restricted to $\mathcal{F}_1$.

Lemma C.2 (Divergence rate: decomposition and pseudo-true model): *Under Assumptions **??** and **??**, $\mathcal{K}(\phi) \in [0, \infty)$ for every $\phi \in \Phi$ and decomposes additively into a timing part and a size part,*

$$\mathcal{K}(\phi) \;=\; \underbrace{\Big[\lambda^\star \log \tfrac{\lambda^\star}{\lambda} - \lambda^\star + \lambda\Big]}_{\textit{timing}} \;+\; \underbrace{\lambda^\star\,\mathrm{KL}\big(F_J \,\big\|\, F_\theta\big)}_{\textit{size}}, \qquad \mathrm{KL}\big(F_J \,\big\|\, F_\theta\big) = \int_{\mathcal{J}} \log \frac{f(x)}{f_\theta(x)}\, F_J(dx). \tag{C2}$$

Each term is nonnegative; the timing term is zero only at $\lambda = \lambda^\star$, and the size term only when $f_\theta = f$, F_J-a.e. $\mathcal{K}$ is continuous on the compact Φ and attains its minimum, and because the two coordinates separate the minimizer factorizes,

$$\phi^\star = (\lambda^\star, \theta^\star), \qquad \theta^\star = \arg\min_{\theta \in \Theta} \mathrm{KL}\big(F_J \,\big\|\, F_\theta\big), \tag{C3}$$

with $\theta^\star$ unique by Assumption ??(iv). Finally $\mathcal{K}$ is a genuine per-unit-time rate: on any horizon $\mathrm{KL}(\mathbb{P}\|\mathbb{P}^\phi)\big|_{\mathcal{F}_t} = t\,\mathcal{K}(\phi)$.

Proof of Lemma C.2. Step 0 (the log-density is well defined, finite, and integrable). By Lemma C.1, $\mathbb{P}|_{\mathcal{F}_1} \sim \mathbb{P}^\phi|_{\mathcal{F}_1}$, and evaluating (C1) at $t=1$ and inverting,

$$\frac{d\mathbb{P}}{d\mathbb{P}^\phi}\Big|_{\mathcal{F}_1} = \exp\!\big((\lambda-\lambda^\star)\big) \prod_{k:\,T_k\le 1} \frac{\lambda^\star f(J_k)}{\lambda\, f_\theta(J_k)}.$$

Each factor is strictly positive: $\lambda, \lambda^\star \in (0,\infty)$ gives $e^{(\lambda-\lambda^\star)} > 0$; $f(J_k) > 0$ and $f_\theta(J_k) > 0$ (Assumption ??(i)) give $\lambda^\star f(J_k)/(\lambda f_\theta(J_k)) > 0$; and $N_1 < \infty$ $\mathbb{P}$-a.s. (Assumption ??) makes the product a finite product of strictly positive factors. Hence $d\mathbb{P}/d\mathbb{P}^\phi|_{\mathcal{F}_1} \in (0,\infty)$ $\mathbb{P}$-a.s. For integrability, the integrand of $\mathcal{K}(\phi)$ is, by (C4) below, $\log \frac{d\mathbb{P}}{d\mathbb{P}^\phi}|_{\mathcal{F}_1} = (\lambda-\lambda^\star) + \sum_{k:T_k\le 1} \log\frac{\lambda^\star f(J_k)}{\lambda f_\theta(J_k)}$, so, with $B := \sup_{(x,\theta)\in\mathcal{J}\times\Theta} \big|\log\frac{f(x)}{f_\theta(x)}\big| < \infty$ (finite because $(x,\theta)\mapsto \log\frac{f(x)}{f_\theta(x)}$ is continuous on the compact $\mathcal{J}\times\Theta$, Assumption ??(i), by Weierstrass) and $C := \big|\log\frac{\lambda^\star}{\lambda}\big| + B < \infty$, the triangle inequality gives

$$\Big|\log \tfrac{d\mathbb{P}}{d\mathbb{P}^\phi}\Big|_{\mathcal{F}_1}\Big| \;\le\; |\lambda-\lambda^\star| + \sum_{k:\,T_k\le 1}\Big(\big|\log\tfrac{\lambda^\star}{\lambda}\big| + \big|\log\tfrac{f(J_k)}{f_\theta(J_k)}\big|\Big) \;\le\; |\lambda-\lambda^\star| + C\,N_1.$$

Taking $\mathbb{P}$-expectations and using $\mathbb{E}[N_1] = \lambda^\star < \infty$,

$$\mathbb{E}_{\mathbb{P}}\Big|\log\tfrac{d\mathbb{P}}{d\mathbb{P}^\phi}\Big|_{\mathcal{F}_1}\Big| \;\le\; |\lambda-\lambda^\star| + C\,\mathbb{E}[N_1] = |\lambda-\lambda^\star| + C\lambda^\star < \infty,$$

so the integrand is $\mathbb{P}$-integrable and $\mathcal{K}(\phi) = \mathbb{E}_{\mathbb{P}}\big[\log\frac{d\mathbb{P}}{d\mathbb{P}^\phi}\big|_{\mathcal{F}_1}\big] = \mathrm{KL}\big(\mathbb{P}\,\|\,\mathbb{P}^\phi|_{\mathcal{F}_1}\big)$ is finite and well defined—the Kullback–Leibler divergence of the physical law from the held-model law, written throughout with the single ratio $d\mathbb{P}/d\mathbb{P}^\phi$.

Step 1 (the integrand). By Lemma C.1, inverting (C1) at $t=1$ and taking logs, the integrand of $\mathcal{K}$ is, with each jump factor split into a timing and a size piece,

$$\log \frac{d\mathbb{P}}{d\mathbb{P}^{\phi}}\Big|_{\mathcal{F}_1} = (\lambda - \lambda^\star) + \sum_{k:\, T_k \le 1} \Big(\log \frac{\lambda^\star}{\lambda} + \log \frac{f(J_k)}{f_\theta(J_k)} \Big). \tag{C4}$$

Step 2 (expectation of a jump sum). For any g with $\mathbb{E}_{F_J}|g| < \infty$, condition on $N_1 = n$; under $\mathbb{P}$ the marks are then i.i.d. F_J and independent of the count (Assumption **??**), so

$$\mathbb{E}_{\mathbb{P}}\Big[\sum_{k:\, T_k \le 1} g(J_k) \,\Big|\, N_1 = n \Big] = n\, \mathbb{E}_{F_J}[g(J)],$$
$$\mathbb{E}_{\mathbb{P}}\Big[\sum_{k:\, T_k \le 1} g(J_k) \Big] = \sum_{n \ge 0} \mathbb{P}(N_1 = n)\, n\, \mathbb{E}_{F_J}[g] = \mathbb{E}[N_1]\, \mathbb{E}_{F_J}[g] = \lambda^\star\, \mathbb{E}_{F_J}[g],$$

using $\mathbb{E}[N_1] = \lambda^\star$ on the unit interval.

Step 3 (apply Step 2 to the split term). With $g(x) = \log \frac{\lambda^\star}{\lambda} + \log \frac{f(x)}{f_\theta(x)}$ (integrable by Step 0),

$$\mathbb{E}_{\mathbb{P}}\Big[\sum_{k:\, T_k \le 1} \log \tfrac{\lambda^\star f(J_k)}{\lambda f_\theta(J_k)} \Big] = \lambda^\star \log \frac{\lambda^\star}{\lambda} + \lambda^\star \int_{\mathcal{J}} \log \frac{f(x)}{f_\theta(x)}\, F_J(dx) = \lambda^\star \log \frac{\lambda^\star}{\lambda} + \lambda^\star\, \mathrm{KL}(F_J \| F_\theta). \tag{C5}$$

Step 4 (assemble). Taking $\mathbb{E}_{\mathbb{P}}$ of (C4) ($\lambda - \lambda^\star$ is its own expectation) and substituting (C5), then regrouping the timing terms,

$$\mathcal{K}(\phi) = (\lambda - \lambda^\star) + \lambda^\star \log \tfrac{\lambda^\star}{\lambda} + \lambda^\star\, \mathrm{KL}(F_J \| F_\theta) = \Big[\lambda^\star \log \tfrac{\lambda^\star}{\lambda} - \lambda^\star + \lambda \Big] + \lambda^\star\, \mathrm{KL}(F_J \| F_\theta),$$

which is (C2).

Signs. Write $T(\lambda) = \lambda^\star \log \frac{\lambda^\star}{\lambda} - \lambda^\star + \lambda = \lambda^\star \log \lambda^\star - \lambda^\star \log \lambda - \lambda^\star + \lambda$. Differentiating term by term in λ ($\lambda^\star \log \lambda^\star$ and $-\lambda^\star$ constant),

$$T'(\lambda) = -\lambda^\star \cdot \tfrac{1}{\lambda} + 1 = 1 - \frac{\lambda^\star}{\lambda}, \qquad T''(\lambda) = \frac{d}{d\lambda}\Big(1 - \frac{\lambda^\star}{\lambda} \Big) = \frac{\lambda^\star}{\lambda^2}. \tag{C6}$$

By (C6) T' vanishes only at $\lambda = \lambda^\star$, and $T''(\lambda) = \lambda^\star/\lambda^2 > 0$ on $\Lambda \subset (0, \infty)$, so T is strictly convex with unique minimiser $\lambda = \lambda^\star$ and $T(\lambda^\star) = 0$; hence the timing term is ≥ 0, zero only at $\lambda = \lambda^\star$. By Gibbs' inequality $\mathrm{KL}(F_J \| F_\theta) \geq 0$, zero iff $f_\theta = f$, F_J-a.e.; hence the size term is ≥ 0 with the stated equality. Both terms being ≥ 0, $\mathcal{K} \geq 0$.

Continuity, minimiser, factorisation. T is continuous on Λ (bounded away from 0). For the size term, the KL integrand $x \mapsto \log\big(f(x)/f_\theta(x)\big)$ is continuous in θ (Assumption **??**(i)) and dominated by the integrable envelope B of Step 0 (which bounds the *log-ratio*, not $\log f_\theta$ alone), so by dominated convergence $\theta \mapsto \mathrm{KL}(F_J \| F_\theta)$ is continuous on the compact Θ; thus $\mathcal{K}$ is continuous on the compact Φ and attains its minimum (extreme-value theorem). Since $\mathcal{K}(\lambda, \theta) = T(\lambda) + \lambda^\star \mathrm{KL}(F_J \| F_\theta)$ separates additively, the joint minimiser factorises into $\arg\min_\lambda T = \{\lambda^\star\}$ and $\arg\min_\theta \mathrm{KL}(F_J \| F_\theta) = \{\theta^\star\}$, the latter a singleton by Assumption **??**(iv); this is (C3).

Rate. The marked Poisson record has stationary independent increments (Assumption **??**), so the densities multiply across disjoint unit blocks: on $\mathcal{F}_t$ with integer t, $\log \frac{d\mathbb{P}}{d\mathbb{P}^\phi} = \sum_{j=1}^t \log \frac{d\mathbb{P}}{d\mathbb{P}^\phi}|_{(j-1,j]}$, each block term i.i.d. with mean $\mathcal{K}(\phi)$ by Steps 0–4; hence $\mathrm{KL}(\mathbb{P} \| \mathbb{P}^\phi)|_{\mathcal{F}_t} = t\,\mathcal{K}(\phi)$, and the same expectation computation with $\mathbb{E}[N_t] = \lambda^\star t$ extends it to all $t > 0$. Thus $\mathcal{K}$ is a per-unit-time rate. $\square$

The imported results, and our hypotheses. Mancini (2009, Thm. 3.1 and Cor. 3.2) consider $dX_t = a_t\,dt + \sigma_t\,dW_t + dJ_t$ with J a finite-activity jump process $J_t = \sum_{j=1}^{N_t} \gamma_j$, observed on a grid of mesh h over a fixed horizon. Under (1) a pathwise bound on the drift integral, which holds whenever a is càdlàg; (2) $\int_0^T \sigma_s^2\,ds < \infty$ with σ pathwise bounded, which holds whenever σ is càdlàg; (3) a deterministic threshold $r(h)$ with $r(h) \to 0$ and $h \log(1/h)/r(h) \to 0$; and (4) $\mathbb{P}\{\Delta N_t \neq 0,\ \gamma_{N_t} = 0\} = 0$, the threshold a.s. isolates the jump-free intervals, $\mathbf{1}\{\Delta_i N = 0\} = \mathbf{1}\{(\Delta_i X)^2 \leq r(h)\}$ for small h, and the truncated realized variation is consistent for the integrated variance,

$$\operatorname*{plim}_{h \to 0} \sum_i (\Delta_i X)^2\, \mathbf{1}\{(\Delta_i X)^2 \leq r(h)\} = \int_0^T \sigma_s^2\,ds.$$

Mancini (2009, Thm. 3.7) adds that if J is a *compound Poisson* process, if σ is adapted, continuous, and *independent of* W *and* N (no leverage), with $\mathbb{E}[\int_0^T \sigma_s^2\, ds] < \infty$, then the truncation estimator of each realized size, $\hat{\gamma}^{(i)} = \Delta_i X\, \mathbf{1}\{(\Delta_i X)^2 > r(h)\}$, satisfies

$$\sqrt{n}\sum_i \big(\hat{\gamma}^{(i)} - \gamma^{(i)}\big)\mathbf{1}\{\Delta_i N \geq 1\} \xrightarrow{d} MN\Big(0,\ T\int_0^T \sigma_s^2\, dN_s\Big). \tag{C7}$$

Our endowment (**??**) is the special case $a_t \equiv \mu_0 - \frac{1}{2}\sigma^2$, $\sigma_t \equiv \sigma$, $J_t = -\sum_{k:T_k\leq t} J_k$ with N Poisson$(\lambda^\star)$ and marks $\{J_k\}$ i.i.d. F_J. We verify each hypothesis. The drift is constant, hence càdlàg, so (1) holds; σ is constant, hence càdlàg with $\int_0^T \sigma^2\, ds = \sigma^2 T < \infty$, so (2) holds; take $r(h) = h^{0.99}$, for which $r(h) \to 0$ and $h\log(1/h)/r(h) = h^{0.01}\log(1/h) \to 0$, so (3) holds; the marks satisfy $J_k \in [\underline{J}, \overline{J}]$ with $\underline{J} > 0$, so $\gamma_{N_t} = -J_{N_t} \neq 0$ a.s. and (4) holds. For (C7): J is compound Poisson by construction; σ is constant, hence adapted, continuous, and (being deterministic) independent of W and N, so the no-leverage hypothesis holds; and $\mathbb{E}[\int_0^T \sigma^2 ds] = \sigma^2 T < \infty$. All hypotheses of both results hold in our model.

What the path identifies, and what it does not. A model could in principle differ from the truth in the diffusion coefficient—an admissibility counterfactual outside the maintained space $\phi = (\lambda, \theta)$—in the arrival intensity, or in the size law, and these have different fates. The diffusion coefficient and the realized disaster sizes are *pathwise functionals* of the record: recovered, in the high-frequency limit, from the path on any interval, so a wrong diffusion is refuted on every interval. The intensity and the size *law* are not functionals of any finite record; they are properties of the generating law, learned only from the disasters that accrue at the rare-disaster rate. That asymmetry is what makes the joint (λ, θ) the admissible object of persistent disagreement.

In the high-frequency limit the path delivers the diffusion coefficient and every realized disaster size exactly (Mancini, 2009; Aït-Sahalia and Jacod, 2009), leaving no room to disagree past the first interval; in the long-horizon limit the intensity and the size law are still being learned, at the common rate $(\lambda^\star t)^{-1/2}$ in relative (log-intensity) error—the effective sample is the disaster count—because both are inferred from a sample that grows only as fast as disasters arrive (Lemma IA.1). A disagreement about σ cannot survive a single interval; a disagreement about (λ, θ) survives for a length of time set by $1/\lambda^{\star}$. That is the property a disputed primitive must have. The intensity coordinate is the object of the companion foundation paper (Teguia, 2026); here it is carried jointly with the size law, the two learned in parallel from the same rare disasters. The identification asymmetry between the physical law and the held cross-section is stated and proved in Section IA.3.

Appendix D. Proofs for Section ??

The rate at which an investor abandons her model is the *revision hazard*: the first passage of the deviance $D_{\phi,\tau}(\mathcal{R})$ (**??**) across its acceptance threshold (Propositions D.1 and E.2). Between disasters, with the record fixed, the deviance drifts deterministically and meets the threshold at a scheduled *calm deadline* $\Delta^{\star}(\phi, \varepsilon)$ (D4)—a regime age at which a disaster-free record is rejected for a process expected to deliver disasters; a disaster appends a mark and triggers rejection when the updated record leaves the typical set. The single-agent foundation—that the surviving record, not the current belief, is the sufficient state for this hazard—is the companion paper's result (Teguia, 2026); the record-resolved passage law, which does not close in general, is in Appendix E. In the baseline ($f_\theta \equiv f$, Corollary E.1) the deviance depends on the record through the count alone, and the hazard closes in the low-dimensional count-and-age state.

Between revisions the cross-section drifts by *selection* (Proposition E.1): in calm the share of a below-average-intensity investor grows and that of an above-average one shrinks, while a disaster rewards whoever held a model that fits it—the same selection a Barro (2006)-style truth would impose, here on *held* models. This is the wealth channel that turns the static tilt of Proposition **??** into a moving one. The revisions themselves enter through two flows the hazard defines: a staggered *calm switch flow*, as tolerance-indexed deadlines arrive and each reviser relocates to the common destination $\phi_t^* = \mathfrak{m}(\rho_t)$ (Lemma D.1), and a synchronized *disaster jump*, at which every triggered investor relocates to that same destination at once (Proposition D.2). The law of motion collects the three.

The disaster channel (a) needs only a partial rejection with a moved posterior mode; the calm channel (b) needs the regularity of Remark D.1. The point is sharp against the merging benchmark: common priors updated on common data converge—predictive laws merge (Blackwell and Dubins, 1962) and rational learners come to agree (Kalai and Lehrer, 1993)—so persistent disagreement ordinarily needs heterogeneous priors or private signals. Here it needs neither: identical priors, identical data, and a single non-Bayesian rule whose staggered rejection timing holds the cross-section apart at any finite date. Long-run persistence is a separate matter: without demographic refresh the survivors eventually agree (Proposition K.1(ii)), and the non-degenerate stationary cross-section is supplied by the turnover economy (Propositions J.1–K.1). With common data the population is the exact deterministic transport of the tolerance distribution along the path—the cross-section the image of G, not a continuum average. That is the finite-horizon statement: dispersion emerges at a date from a homogeneous start. Left to run, the survivors' convergence re-concentrates it; the *stationary* dispersion the quantitative sections use is the turnover object of Proposition K.1.

Sufficiency of the since-adoption record. For an exponential-family size law with natural sufficient statistic T, the pair $\left(n, \sum_j T(x_j)\right)$ is jointly sufficient for (λ, θ) by the Fisher–Neyman factorisation criterion, so the deviance (**??**) depends on the record only through it (the predictive density (**??**) additionally carries the marks' base-measure factors, which cancel in the deviance ratio); for a general family the test uses a goodness-of-fit statistic of (**??**) and no sufficiency is claimed. It is recorded for completeness—it is why exponential families are computationally special—but no result in the main text relies on it.

Proof of Proposition **??**. By Proposition D.3, $m_t = \mathsf{T}_{t\#}G$ with $\mathsf{T}_t(\cdot)$ deterministic, so the held-model marginal is the image of G under the threshold-to-trajectory map. Under (a): at T_k the rejecting cohort relocates to the common $\phi^{(k)}$ (Proposition D.2) while the non-rejecting cohort of positive G-mass retains ϕ_0; since the destination differs in intensity by hypothesis, $\lambda(\phi^{(k)}) \neq \lambda(\phi_0)$, the held-model marginal places positive mass on two distinct intensities, whence $\mathrm{Var}_{S_{T_k}}(\lambda) > 0$. (We invoke only distinctness: the revealed-preference monotonicity of Lemma D.1 is a fixed-count comparison and does not by itself sign $\lambda(\phi^{(k)}) - \lambda(\phi_0)$ once the elapsed calm time at T_k enters.) Under (b): along the zero-count history the reviser of threshold ε adopts the calm mode $\lambda^{\mathrm{m}}_{\Delta^\star(\phi_0,\varepsilon)}$ at her deadline; with the deadline map non-constant on a positive-mass set and $\tau \mapsto \lambda^{\mathrm{m}}_\tau$ strictly decreasing, two positive-mass groups of thresholds adopt distinct intensities before the first disaster, whence $\mathrm{Var}_{S_t}(\lambda) > 0$ for t past the earliest such deadline. *From head count to consumption shares.* Each case produces a split of positive G-mass, hence positive head-count mass m_t (Proposition D.3); the variance claim is in the consumption-share measure S_t. The two are mutually absolutely continuous: with the symmetric allocation (Assumption **??**(iv)) the share is $s^i_t = (M^i_t)^{1/\alpha} / \int (M^j_t)^{1/\alpha}\, dj$, and M^i_t is bounded and bounded away from zero on the compact menu: by (**??**), with $N_t < \infty$ (Assumption **??**) the density is a finite product of reweights $r_i = \lambda_i f_{\theta_i}/(\lambda^\star f)$—each continuous and strictly positive on the compact $\Phi \times \mathcal{J}$ (Assumption **??**), hence in a fixed positive range—times the intensity factors accrued segment-by-segment over the full history—the density is the product across held-model segments—each of the form $e^{-(\lambda_i(u)-\lambda^\star)\times \mathrm{length}}$ and bounded for $\lambda_i \in [\underline{\lambda}, \overline{\lambda}]$ (a revision only relabels ϕ_i and restarts the segment, preserving the two-sided bound); so there are $0 < \underline{c} \le \overline{c} < \infty$ with $\underline{c}\, m_t(A) \le S_t(A) \le \overline{c}\, m_t(A)$ for every Borel A. A positive-m_t-mass split into two intensity classes is therefore a positive-S_t-mass split. In either case the dispersion is strictly positive at a finite t, from the homogeneous start. $\square$

Lemma D.1 (The second-order posterior is common): *For every pair of investors i, j and every t, $\rho^i_t = \rho^j_t =: \rho_t$. The common posterior is continuous in t between disasters and jumps at each disaster epoch; during a disaster-free spell its maximum-a-posteriori intensity* $\lambda^{\mathrm{m}}_t := \min \arg\max_\lambda \max_\theta \rho_t(\lambda, \theta)$ *(the intensity coordinate of the tie-broken mode $\mathfrak{m}(\rho_t)$ of Definition* **??***) is non-increasing in t.*

Proof of Lemma D.1. By (**??**), $\rho_t^i(\phi) \propto \rho_0(\phi)\, L_t(\phi)$ with $L_t(\phi) = e^{-\lambda t}\lambda^{N_t}\prod_{k:T_k\le t} f_\theta(J_k)$. The prior ρ_0 is common across investors (Assumption **??**) and L_t is a function of the common path $(t, N_t, \{J_k\})$ alone, with no dependence on i or ε^i; hence $\rho_t^i = \rho_t^j =: \rho_t$. Continuity in t follows from L_t and its normalizer: $L_t(\phi)$ is continuous in t off the disaster epochs and jumps at them, and the normalizing constant $Z_t = \int_\Phi \rho_0(\phi) L_t(\phi)\, d\phi$ is continuous and strictly positive *on disaster-free intervals* (a continuous, strictly positive integrand on the compact Φ) and jumps at each disaster epoch. At a disaster of size J_k, $L_t(\lambda, \theta)$ is multiplied pointwise by the model-specific factor $\lambda f_\theta(J_k)$, while the normalizer Z_t is multiplied by its scalar posterior-predictive average $\int_\Phi \lambda f_\theta(J_k)\, \rho_{t-}(d\phi)$, *not* by any model-dependent factor; the posterior $\rho_t = \rho_0 L_t/Z_t$ therefore jumps precisely because the numerator updates model by model and the denominator by this common scalar average, and is continuous between disasters. This gives the continuity-between-disasters and disaster-jump claims of the statement.

For the monotonicity, fix a disaster-free interval, on which N_t and the sizes are constant. By the product prior $\rho_0 = \rho_0^\Lambda \otimes \rho_0^\Theta$ (Assumption **??**), the log-posterior separates,

$$\log \rho_t(\lambda, \theta) = \underbrace{\Big[\log \rho_0^\Lambda(\lambda) - \lambda t + N_t \log \lambda\Big]}_{=:\ell_t^\Lambda(\lambda)} + \underbrace{\Big[\log \rho_0^\Theta(\theta) + \textstyle\sum_k \log f_\theta(J_k)\Big]}_{\theta\text{-only}} + \text{const.}$$

The θ-only term does not depend on λ, so $\max_\theta \log \rho_t(\lambda, \theta) = \ell_t^\Lambda(\lambda) + \text{const}$ and the maximum-a-posteriori intensity is $\lambda_t^{\mathrm{m}} = \min \arg\max_{\lambda\in\Lambda} \ell_t^\Lambda(\lambda)$. The maximizer set is non-empty because ℓ_t^Λ is continuous (hence upper semicontinuous) on the compact interval Λ, so $\lambda_t^{\mathrm{m}} = \min \arg\max_{\lambda\in\Lambda} \ell_t^\Lambda(\lambda)$ is well defined. We prove monotonicity directly, by revealed preference. Fix $t_2 > t_1$ in the disaster-free interval and write $\lambda_1 := \lambda_{t_1}^{\mathrm{m}}$, $\lambda_2 := \lambda_{t_2}^{\mathrm{m}}$. Each is optimal at its own date:

$$\ell_{t_1}^\Lambda(\lambda_1) \ge \ell_{t_1}^\Lambda(\lambda_2), \qquad \ell_{t_2}^\Lambda(\lambda_2) \ge \ell_{t_2}^\Lambda(\lambda_1).$$

Adding the two inequalities and regrouping,

$$\left[\ell^{\Lambda}_{t_2}(\lambda_2) - \ell^{\Lambda}_{t_1}(\lambda_2)\right] - \left[\ell^{\Lambda}_{t_2}(\lambda_1) - \ell^{\Lambda}_{t_1}(\lambda_1)\right] \;\ge\; 0. \tag{D1}$$

Since N_t and the sizes are constant on the interval (write N for the common count), the prior and count terms of ℓ^{Λ} cancel in the difference at any fixed λ:

$$\ell^{\Lambda}_{t_2}(\lambda) - \ell^{\Lambda}_{t_1}(\lambda) = \left[\log \rho^{\Lambda}_0(\lambda) - \lambda t_2 + N \log \lambda\right] - \left[\log \rho^{\Lambda}_0(\lambda) - \lambda t_1 + N \log \lambda\right] = -\lambda(t_2 - t_1). \tag{D2}$$

Evaluating (D2) at $\lambda = \lambda_2$ and $\lambda = \lambda_1$, the left-hand side of (D1) equals

$$\left[-\lambda_2(t_2 - t_1)\right] - \left[-\lambda_1(t_2 - t_1)\right] = -(t_2 - t_1)(\lambda_2 - \lambda_1).$$

Thus

$$-(t_2 - t_1)(\lambda_2 - \lambda_1) \;\ge\; 0,$$

and since $t_2 > t_1$ this forces $\lambda_2 \le \lambda_1$, i.e. $\lambda^{\mathrm{m}}_{t_2} \le \lambda^{\mathrm{m}}_{t_1}$. Hence λ^{m}_t is non-increasing in t through the disaster-free spell. □

Proposition D.1 (Law of motion; scheduled intensity relief): *Fix investor i. Between adoptions the held model ϕ_i is constant and the regime age advances at unit rate, $\dot{\tau}^i_t = 1$; the revision time is the first exit of the record from the typical set,*

$$\tau^{\mathrm{rev}}_i = \inf\left\{t > a^i :\ \bar{\varepsilon}(\mathcal{R}^i_t; \phi_i, \tau^i_t) \le \varepsilon^i\right\}. \tag{D3}$$

Two exit channels generate it. (Calm.) *If no disaster occurs after adoption, the record is the zero-count history, whose surprise $\bar{\varepsilon}$ tends to 0 as the age grows—the zero-count record becomes atypical of a process expected to deliver disasters; hence there is a finite, deterministic calm deadline*

$$\Delta^{\star}(\phi_i, \varepsilon^i) = \inf\left\{\Delta > 0 :\ \bar{\varepsilon}\big((0, \varnothing); \phi_i, \Delta\big) \le \varepsilon^i\right\} < \infty, \tag{D4}$$

at which the held model is rejected absent any disaster; the two-argument $\Delta^\star(\phi,\varepsilon)$ denotes this zero-count case. (Disaster.) *At a disaster the record jumps—count +1 and a new size—and $\bar\varepsilon$ is re-evaluated, so a disaster is an exit if it lands the record outside the typical set. The revision time is the first exit of the* updated *record—the earlier of the current calm deadline and the next disaster that lands outside the typical set; after a non-rejecting disaster the count is $n \geq 1$, so the relevant calm deadline is the first passage of the deviance for that updated record, not the original zero-count deadline. At a calm rejection reached along a spell that is* disaster-free since adoption, *the re-selected intensity satisfies $\lambda^{\mathrm{m}}_{\tau_i^{\mathrm{rev}}} \leq \lambda_i$; an intervening non-rejecting disaster updates the record to count $n \geq 1$ and can raise the mode, so this scheduled relief is the disaster-free statement, relative to the start of the current disaster-free spell.*

Proof of Proposition D.1. By Definition **??** the held model equals the second-order mode fixed at the last adoption and is unchanged until the next rejection, while the regime age advances at unit rate; this gives constancy and $\dot\tau^i_t = 1$. By Definition **??** a rejection is the first time the deviance exceeds its threshold, equivalently the first time $\bar\varepsilon \leq \varepsilon^i$, which is (D3).

Calm channel. Suppose no disaster occurs after adoption, so the record at age $\Delta = \tau^i_t$ is the zero-count history $(0,\varnothing)$. Its deviance is the pure count term, by (**??**) and $d_P(0\|\lambda) = \lambda$,

$$D_{\phi_i,\Delta}\big((0,\varnothing)\big) = \Delta\, d_P(0\|\lambda_i) = \lambda_i\Delta,$$

and by (**??**) its surprise is the predictive mass of records at least as deviant, which Markov's inequality (valid as $D \geq 0$) bounds by the mean deviance under the held model,

$$\bar\varepsilon\big((0,\varnothing);\phi_i,\Delta\big) = Q_{\phi_i,\Delta}\big(\{\mathcal{R}' : D_{\phi_i,\Delta}(\mathcal{R}') \geq \lambda_i\Delta\}\big) \leq \frac{\mathbb{E}_{Q_{\phi_i,\Delta}}\big[D_{\phi_i,\Delta}\big]}{\lambda_i\Delta}. \tag{D5}$$

We show the bound vanishes as $\Delta \to \infty$. Under $Q_{\phi_i,\Delta}$ the count is $n \sim \mathrm{Poisson}(m)$ with $m := \lambda_i\Delta$, and given n the sizes are i.i.d. F_{θ_i}; by (**??**) the deviance splits as $D = c(n) + \Lambda_n$ into the count deviance $c(n) = n\log(n/m) - n + m \geq 0$ (with $c(0) = m$) and the size deviance

$\Lambda_n = \sum_{k=1}^n \log \frac{f_{\hat\theta}(J_k)}{f_{\theta_i}(J_k)} = \sup_{\theta\in\Theta} \sum_{k=1}^n \log \frac{f_\theta(J_k)}{f_{\theta_i}(J_k)} \ge 0$, and we bound each mean.

Count mean: bounded. Write the count deviance as $c(n) = n\log(n/m) - n + m = n\log n - n\log m - n + m$. Taking the Q-expectation with $n \sim \text{Poisson}(m)$, so $\mathbb{E}[n] = m$ and hence $\mathbb{E}[-n+m] = 0$,

$$\mathbb{E}[c(n)] = \mathbb{E}[n\log n] - (\log m)\,\mathbb{E}[n] + \mathbb{E}[-n+m] = \mathbb{E}[n\log n] - m\log m. \tag{D6}$$

For the first term, the $n = 0$ summand is zero and $n/n! = 1/(n-1)!$, so shifting the index by $j = n-1$ (with $M \sim \text{Poisson}(m)$),

$$\begin{aligned}\mathbb{E}[n\log n] &= \sum_{n\ge1} e^{-m}\frac{m^n}{n!}\, n\log n = \sum_{n\ge1} e^{-m}\frac{m^n}{(n-1)!}\log n \\ &= m\sum_{j\ge0} e^{-m}\frac{m^j}{j!}\log(j+1) = m\,\mathbb{E}[\log(M+1)].\end{aligned} \tag{D7}$$

Substituting (D7) into (D6) and combining the logarithms,

$$\mathbb{E}[c(n)] = m\,\mathbb{E}[\log(M+1)] - m\log m = m\,\mathbb{E}\Big[\log \tfrac{M+1}{m}\Big] \le m\log\tfrac{m+1}{m} \le m\cdot\tfrac{1}{m} = 1,$$

the first inequality by Jensen (log concave, $\mathbb{E}[M+1] = m+1$) and the second by $\log(1+x) \le x$. Thus $\mathbb{E}[c(n)] \le 1$ for every $\Delta > 0$.

Size mean: $o(m)$. The map $(\theta, x) \mapsto \log f_\theta(x)$ is continuous on the compact $\Theta \times \mathcal{J}$ (Assumption **??**(i)), hence bounded, so $B := \sup_{\theta,\theta'\in\Theta,\, x\in\mathcal{J}} \big|\log \frac{f_\theta(x)}{f_{\theta'}(x)}\big| < \infty$ and $0 \le \Lambda_n \le Bn$. We invoke the uniform strong law of large numbers in the form underlying maximum-likelihood consistency (Wald, 1949): if Θ is compact, $\theta \mapsto \log f_\theta(x)$ is continuous for each x, and the envelope is F_{θ_i}-integrable, then

$$\sup_{\theta\in\Theta}\Big|\tfrac{1}{n}\textstyle\sum_{k\le n}\log\frac{f_\theta(J_k)}{f_{\theta_i}(J_k)} - \big(-\mathrm{KL}(F_{\theta_i}\|F_\theta)\big)\Big| \xrightarrow{\text{a.s.}} 0 \qquad (n\to\infty). \tag{D8}$$

Each hypothesis holds in our model: Θ is compact (Assumption **??**), $\log f_\theta(x)$ is continuous in θ (Assumption **??**(i)–(ii)), and the constant envelope B is F_{θ_i}-integrable, $\int B\,dF_{\theta_i} = B < \infty$. Since $\sup_\theta\big(-\mathrm{KL}(F_{\theta_i}\|F_\theta)\big) = 0$, attained at $\theta = \theta_i$, (D8) gives $\Lambda_n/n = \sup_\theta \frac{1}{n}\sum_{k\le n}\log\frac{f_\theta}{f_{\theta_i}}(J_k) \to 0$ a.s., with $0 \le \Lambda_n/n \le B$. As $\Delta \to \infty$, $m \to \infty$ and $n \to \infty$ a.s., and $n/m \to 1$ a.s. with $\{n/m\}$ uniformly integrable ($\mathbb{E}[(n/m)^2] = 1 + 1/m$ is bounded). Hence $\Lambda_n/m = (\Lambda_n/n)(n/m) \to 0$ a.s. and is dominated by the uniformly integrable $B\,(n/m)$, so by Vitali's convergence theorem $\mathbb{E}[\Lambda_n]/m \to 0$.

Conclusion. Combining the two means, $\mathbb{E}_{Q_{\phi_i,\Delta}}[D]/(\lambda_i\Delta) \le 1/m + \mathbb{E}[\Lambda_n]/m \to 0$, so by (D5) $\bar\varepsilon\big((0,\varnothing);\phi_i,\Delta\big) \to 0$ as $\Delta \to \infty$. As $\varepsilon^i > 0$, the set $\{\Delta > 0 : \bar\varepsilon((0,\varnothing);\phi_i,\Delta) \le \varepsilon^i\}$ is non-empty (it contains the non-empty strict sublevel set), so its infimum—the calm deadline (D4)—is finite, $\Delta^\star(\phi_i,\varepsilon^i) < \infty$. By definition of the infimum, $\bar\varepsilon \ge \varepsilon^i$ for every $\Delta < \Delta^\star$, so the held model survives up to $\Delta^\star$ and is rejected there: $\Delta^\star$ is the first crossing, and no monotonicity of $\bar\varepsilon$ in the age is needed. In the intensity channel the tail set has an exact form: the zero-count deviance is $\Delta\, d_{\mathrm{P}}(0\|\lambda) = \lambda\Delta$, and $d_{\mathrm{P}}(x\|\lambda) \ge \lambda$ for $x > 0$ iff $x \ge e\lambda$, so

$$\bar\varepsilon\big((0,\varnothing);\phi,\Delta\big) \;=\; e^{-\lambda\Delta} + \mathbb{P}\big(\mathrm{Pois}(\lambda\Delta) \ge e\lambda\Delta\big) \;\le\; 2e^{-\lambda\Delta},$$

the bound by the Poisson–Chernoff inequality at the threshold $e\lambda\Delta$; hence $\Delta^\star$ is finite, with $\lambda^{-1}\log(1/\varepsilon) \le \Delta^\star \le \lambda^{-1}\log(2/\varepsilon)$. The map $\Delta \mapsto \bar\varepsilon$ is continuous between the count thresholds $\Delta_n = n/(e\lambda)$ and steps *downward* as Δ passes each of them—the count n leaves the tail set—so it is neither continuous nor monotone globally, the sublevel set $\{\Delta : \bar\varepsilon \le \varepsilon^i\}$ may be left-open, and the tolerance need not be attained. The closed-threshold convention of Definition **??**—reject at the first age with $\bar\varepsilon \le \varepsilon^i$—is thus a first-exit convention: rejection is effective at the infimum $\Delta^\star$ whether the tolerance is attained or passed by a step; a disaster before $\Delta^\star$ updates the record and the deadline is re-evaluated on it. With a nondegenerate size family the size deviances are continuously distributed, $\bar\varepsilon$ is continuous in the age, and the crossing is attained; the steps are the intensity channel's. It is deterministic, since $\bar\varepsilon((0,\varnothing);\phi_i,\cdot)$ is a function of

(ϕ_i, Δ) alone. Finally the deadline is non-increasing in ε^i: a larger threshold enlarges the rejection set $\{\Delta : \bar{\varepsilon} \leq \varepsilon^i\}$, lowering its infimum—which, as all investors share ϕ_0 at the start, is what staggers calm revision across the population.

Disaster channel. At a disaster the record gains a count and a size, so the deviance jumps and is re-evaluated; the record exits the typical set there iff the post-disaster deviance exceeds the threshold. The revision time is the earlier of the calm deadline and the first such disaster.

Relief (disaster-free spell). Consider a calm rejection reached along a spell disaster-free since adoption a^i, with re-selected intensity the mode $\lambda^{\mathrm{m}}_{\tau_i^{\mathrm{rev}}}$ and incumbent the mode at adoption $\lambda_i = \lambda^{\mathrm{m}}_{a^i}$. By Lemma D.1 the mode is non-increasing *through a disaster-free spell*, so $\lambda^{\mathrm{m}}_{\tau_i^{\mathrm{rev}}} \leq \lambda_i$. If instead a non-rejecting disaster intervenes it appends a mark and can raise the mode—Lemma D.1 bounds the mode only along disaster-free motion—so the bound then holds relative to the mode at the start of the *current* disaster-free spell, not necessarily to λ_i. This is the scheduled relief that erodes the held intensity *between* disasters; disasters themselves enter through the disaster reweight and relocation, not the calm relief. □

Proposition D.2 (Synchronization at disasters; staggered calm): *At a disaster epoch T_k every investor's record receives the same increment at the same instant. Within the cohort of investors sharing a held model ϕ and adoption date a at T_k^- (hence the same record and predictive law), the post-disaster surprise depth is a common value $\bar{\varepsilon}_k(\phi, a)$, so the investors of that cohort who revise are exactly those with $\varepsilon^i \geq \bar{\varepsilon}_k(\phi, a)$; and* every *investor who revises at T_k, across all cohorts, adopts the common mode $\phi^{(k)} = \mathfrak{m}(\rho_{T_k})$ (Definition* **??***). Revision at a disaster is thus synchronized in timing and destination, with a cohort-specific trigger. Away from disasters, investor i revises at her calm deadline $a^i + \Delta^\star(\phi_i, \varepsilon^i)$, evaluated on the record accrued since her current adoption (a post-adoption disaster that does not itself trigger rejection resets the deadline on the updated, positive-count record; the zero-count expression is the disaster-free special case); we write $\Delta^\star(\phi, \varepsilon)$ for the zero-count deadline and leave the record argument implicit on positive-*

count records—a time non-increasing *in ε^i (Proposition D.1), so calm revisions are staggered across thresholds wherever that deadline map is non-constant. Starting from the homogeneous configuration of Assumption* **??***, both channels generate dispersion rather than assume it—the synchronized disaster split and the staggered calm relief each disperse the held intensity, under conditions made precise in Proposition* **??** *(see also Remark D.1).*

Proof of Proposition D.2. At a disaster T_k the path increments by a single count and the common size J_k, which enters every investor's record simultaneously. Two investors of the same cohort—same held ϕ and same adoption date a at T_k^-—have, by construction, the identical record on $(a, T_k]$ and hence the identical post-disaster record; since the deviance (**??**) and the surprise (**??**) are deterministic functions of $(\mathcal{R}, \phi, \Delta)$, both share a common post-disaster surprise $\bar{\varepsilon}_k(\phi, a)$. By Definition **??** cohort member i rejects iff $\varepsilon^i \geq \bar{\varepsilon}_k(\phi, a)$. Every investor who rejects at T_k, of any cohort, re-selects $\arg\max_\phi \rho^i_{T_k}(\phi)$, which by Lemma D.1 is the common posterior ρ_{T_k}; the deterministic tie-break $\mathfrak{m}$ of Definition **??** then selects the single common point $\phi^{(k)} = \mathfrak{m}(\rho_{T_k}) \in \arg\max_\phi \rho_{T_k}(\phi)$. The destination depends on the common full-history posterior alone—not on the reviser's held model or adoption date—so it is common to all revisers, while the trigger $\bar{\varepsilon}_k(\phi, a)$ is constant only within a cohort.

In calm, by Proposition D.1 investor i revises at $a^i + \Delta^\star(\phi_i, \varepsilon^i)$, a deadline *non-increasing* in ε^i; investors of a common cohort therefore revise at times ordered (weakly) by ε^i—staggered across thresholds with *distinct* deadlines, and simultaneous on any level set where $\Delta^\star(\phi, \cdot)$ is flat. From the homogeneous start (Assumption **??**) all hold ϕ_0 at $t = 0$ with common second-order prior; wherever the deadline map $\Delta^\star(\phi_0, \cdot)$ is non-constant across thresholds (Remark D.1) the population begins to disperse in calm before any disaster, and each disaster induces a synchronized split—dispersion is generated, not assumed. □

Remark D.1 (When the held-model marginal has continuum support)*.* The size-only design produced a finitely supported held-model marginal because all revision was synchronized at disasters. The calm channel can enlarge the support: if, on a sub-interval of $(0, 1)$, the calm deadline $\varepsilon \mapsto \Delta^\star(\phi_0, \varepsilon)$ is continuous and strictly monotone and the calm intensity mode $\tau \mapsto \lambda^{\mathrm{m}}_\tau$ is strictly decreasing, then the first-rejection cohort adopts a continuum of distinct intensities and the held-model marginal is not finitely supported. In pure calm $\ell^\Lambda_\tau(\lambda) = \log \rho^\Lambda_0(\lambda) - \lambda\tau$ has strictly decreasing differences in (λ, τ)—for $\lambda' > \lambda$, $\ell^\Lambda_\tau(\lambda') - \ell^\Lambda_\tau(\lambda) = \log \frac{\rho^\Lambda_0(\lambda')}{\rho^\Lambda_0(\lambda)} - (\lambda' - \lambda)\tau$ is strictly decreasing in τ from the linear $-\lambda\tau$ term alone—so by Topkis's theorem the maximiser $\lambda^{\mathrm{m}}_\tau$ is non-increasing in τ, with no condition on ρ^Λ_0. The *strict* interior decline is the separate, stronger property: it holds when ρ^Λ_0 is strictly log-concave (the first-order condition then has a strictly decreasing interior root), while at the boundary $\lambda^{\mathrm{m}}_\tau$ can plateau (no interior maximizer is assumed). The strict monotonicity of $\Delta^\star(\phi_0, \cdot)$ requires, in *addition*, that the zero-count surprise $\bar{\varepsilon}$ be strictly decreasing in the age—a deadline-regularity condition not implied by log-concavity alone, which we impose where this remark is invoked. Absent either regularity the support may remain finite; the claim is scoped to these conditions and is not used downstream.

Proposition D.3 (The cross-section is a pushforward of G): *Fix the common path. Let $\mathsf{T}_t(\varepsilon)$ be the held model and regime age that Definition* **??** *produces at t for an investor of threshold ε, started from ϕ_0 at age 0. Then T_t is a deterministic measurable map and the cross-sectional distribution of (held model, regime age) is the pushforward*

$$m_t \;=\; \mathsf{T}_{t\#}\, G, \tag{D9}$$

with no law of large numbers invoked: G is the primitive continuum type law of Assumption **??***, not a random sample, so the transport is exact. The held-model marginal jumps at disasters, where all revisers move to the common $\phi^{(k)}$, and changes at the staggered calm deadlines $a^i + \Delta^\star(\phi_i, \varepsilon^i)$, where the reviser of threshold ε moves to the common mode prevailing at her deadline.*

Proof of Proposition D.3. Fix the common path. By Definition **??** and Proposition D.1, given the path an investor's revision times are determined by her threshold ε alone—the calm deadline (D4) and the disaster rejections $\{T_k : \varepsilon > \bar{\varepsilon}_k\}$ are deterministic functions of ε—and the re-selected models are the common modes prevailing at those instants (Lemma D.1), also path-determined. Hence $\varepsilon \mapsto \mathsf{T}_t(\varepsilon) = (\phi(t;\varepsilon), \tau(t;\varepsilon))$ is a deterministic map of the threshold. It is Borel measurable. On $[0,t]$ there are a.s. finitely many disasters. Let $K_t(\varepsilon)$ denote the number of rejections the type-ε investor makes on $[0,t]$—the object the partition below runs on. It is finite for every ε: a rejection on a disaster-free-since-adoption holding cannot occur before the fresh-record deadline, bounded below by $\delta(\varepsilon) > 0$ uniformly over Φ (the fresh record itself is the deepest point of its own deviance tail, so its surprise is at least $e^{-\overline{\lambda}\Delta}$, exceeding ε for $\Delta < \overline{\lambda}^{-1}\log(1/\varepsilon)$—a short no-rejection window after every reset), while a rejection on a holding whose record contains a disaster is charged to that holding's most recent disaster, distinct across holdings; hence $K_t(\varepsilon) \leq \lfloor t/\delta(\varepsilon)\rfloor + 2N_t + 1$. Thus $\varepsilon \mapsto \mathsf{T}_t(\varepsilon)$ is measurable by countable partition on the reset count: on $\{K_t = k\}$ it is the k-fold composition of the calm flow and the jump update, alternating calm and disaster steps—each measurable: the calm deadline $\varepsilon \mapsto \Delta^\star(\phi,\varepsilon)$ is monotone hence Borel (Proposition D.1), and the re-selected mode is a measurable function of the record—while $\{K_t \leq k\}$ is a finite intersection and union of measurable deadline and trigger comparisons along the first k steps, by induction on k; the finiteness of $K_t(\varepsilon)$ makes the cells $\{K_t = k\}$, $k \geq 0$, exhaust $(0,1)$, and T_t is Borel on their union (the piecewise-deterministic construction of Davis, 1993; Jacobsen, 2006): the calm deadline $\varepsilon \mapsto \Delta^\star(\phi,\varepsilon)$ is monotone in ε (Proposition D.1), hence Borel, and the calm re-selected mode is a measurable function of the path-determined record at the deadline; the disaster-rejection threshold at T_k is not a population constant but the state-dependent surprise depth $\bar{\varepsilon}_k\big(\phi(T_k^-;\varepsilon), a(\varepsilon)\big)$, a measurable function of the (already measurable) pre-disaster held model and adoption date, so $\{\varepsilon : \varepsilon > \bar{\varepsilon}_k(\cdot)\}$ is Borel as the preimage of a measurable map even though it need not be a fixed sub-interval of $(0,1)$; and the post-disaster mode $\phi^{(k)} = \mathfrak{m}(\rho_{T_k})$ is a measurable function of the common record. A finite alternating composition of these measurable maps is Borel—without asserting that $\phi(t;\cdot)$ is constant on threshold intervals, which the staggered calm deadlines and the state-dependent $\bar{\varepsilon}_k$ may violate. Since investors differ only in the permanent threshold, whose population distribution is G (Assumption **??**), and all share the one path, each investor's state is the deterministic image $\mathsf{T}_t(\varepsilon^i)$ of her type; the cross-sectional distribution of states is therefore the image measure $m_t = \mathsf{T}_{t\#}G$ by construction—there is no idiosyncratic randomness to average, so no law of large numbers is invoked. This is (D9). The held-model marginal jumps at disasters, where all revisers move to the common $\phi^{(k)}$ (Proposition D.2), and changes at the calm deadlines $a^i + \Delta^\star(\phi_i, \varepsilon^i)$, where the reviser of threshold ε moves to the common mode prevailing at her deadline (Lemma D.1). $\square$

Appendix E. Proofs for the revision hazard and the law of motion of the pricing state

Proposition E.1 (Selection dynamics of the kernel and shares): *In the equilibrium of Proposition* **??**, *between revisions (held models fixed):*

(i) in calm (no disaster) the aggregator and the shares drift by selection,

$$\frac{d\log \Xi_t}{dt} = -\tfrac{1}{\alpha}\big(\bar\lambda_t^S - \lambda^\star\big), \qquad \frac{d\log s_t^i}{dt} = \tfrac{1}{\alpha}\big(\bar\lambda_t^S - \lambda_i(t)\big), \tag{E1}$$

so the share of a below-average-intensity (complacent) investor grows and that of an above-average-intensity (fearful) investor shrinks;

(ii) at a disaster of size x the aggregator and the shares jump multiplicatively,

$$\frac{\Xi_t}{\Xi_{t-}} = g_t(x), \qquad \frac{s_t^i}{s_{t-}^i} = \frac{r_i(x;t)^{1/\alpha}}{g_t(x)}, \tag{E2}$$

so an investor whose held model fits the realized disaster better than the aggregate, $r_i(x;t)^{1/\alpha} > g_t(x)$, gains share.

Proof of Proposition E.1. (i) Calm. Between disasters the product in (**??**) is constant, so $M_t^i =$

$\exp(-\int_0^t(\lambda_i(s-)-\lambda^\star)ds)\times$ (const) and, with the held model fixed,

$$\frac{d\log M_t^i}{dt}=-\big(\lambda_i(t)-\lambda^\star\big),\qquad \frac{d\log(M_t^i)^{1/\alpha}}{dt}=-\tfrac{1}{\alpha}\big(\lambda_i(t)-\lambda^\star\big). \tag{E3}$$

Differentiating $\Xi_t=\int_0^1\zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}\,di$ in t: the interchange of $\frac{d}{dt}$ and $\int_0^1\cdot\,di$ is justified by dominated convergence, since on any compact t-interval both the integrand $\zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}$ and its t-derivative $\zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}\big(-\frac{1}{\alpha}(\lambda_i-\lambda^\star)\big)$ are bounded uniformly in i by $\zeta_i^{-1/\alpha}$ times path-constants (by (**??**) and $\lambda_i\in[\underline{\lambda},\overline{\lambda}]$), with $\int_0^1\zeta_i^{-1/\alpha}di<\infty$ (Assumption **??**). Using (E3) for $\frac{d}{dt}(M_t^i)^{1/\alpha}=(M_t^i)^{1/\alpha}(-\frac{1}{\alpha}(\lambda_i-\lambda^\star))$ and then $\zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}=s_t^i\Xi_t$ from (**??**),

$$\begin{aligned}\frac{d\Xi_t}{dt}&=\int_0^1\zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}\Big(-\tfrac{1}{\alpha}(\lambda_i-\lambda^\star)\Big)di\\&=-\tfrac{1}{\alpha}\,\Xi_t\int_0^1 s_t^i(\lambda_i-\lambda^\star)\,di=-\tfrac{1}{\alpha}\,\Xi_t\big(\bar{\lambda}_t^S-\lambda^\star\big),\end{aligned}\tag{E4}$$

the last equality using $\int_0^1 s_t^i\lambda_i\,di=\bar{\lambda}_t^S$ and $\int_0^1 s_t^i\,di=1$ (Definition **??**); this gives the first equation of (E1). For the share, take logs of (**??**): $\log s_t^i=-\frac{1}{\alpha}\log\zeta_i+\frac{1}{\alpha}\log M_t^i-\log\Xi_t$; the first term is constant in t, so

$$\frac{d\log s_t^i}{dt}=\tfrac{1}{\alpha}\frac{d\log M_t^i}{dt}-\frac{d\log\Xi_t}{dt}. \tag{E5}$$

By (E3), $\frac{1}{\alpha}\,d\log M_t^i/dt=-\frac{1}{\alpha}(\lambda_i-\lambda^\star)$; and dividing (E4) by Ξ_t, $d\log\Xi_t/dt=\Xi_t^{-1}\,d\Xi_t/dt=-\frac{1}{\alpha}(\bar{\lambda}_t^S-\lambda^\star)$. Substituting both into (E5),

$$\frac{d\log s_t^i}{dt}=-\tfrac{1}{\alpha}(\lambda_i-\lambda^\star)+\tfrac{1}{\alpha}(\bar{\lambda}_t^S-\lambda^\star)=\tfrac{1}{\alpha}\big(\bar{\lambda}_t^S-\lambda_i\big), \tag{E6}$$

the $\lambda^\star$ terms cancelling; this is the second equation of (E1).

(ii) Disaster. At a disaster of size x, $M_t^i=M_{t-}^i r_i(x;t)$ (**??**), so $(M_t^i)^{1/\alpha}=(M_{t-}^i)^{1/\alpha}r_i(x;t)^{1/\alpha}$ and $\Xi_t/\Xi_{t-}=g_t(x)$ exactly as in (**??**). Then from $s_t^i=\zeta_i^{-1/\alpha}(M_t^i)^{1/\alpha}/\Xi_t$, the constant $\zeta_i^{-1/\alpha}$

cancels in the ratio:

$$\frac{s_t^i}{s_{t-}^i} = \frac{(M_t^i)^{1/\alpha}/\Xi_t}{(M_{t-}^i)^{1/\alpha}/\Xi_{t-}} = \frac{(M_t^i)^{1/\alpha}}{(M_{t-}^i)^{1/\alpha}} \cdot \frac{\Xi_{t-}}{\Xi_t} = r_i(x;t)^{1/\alpha} \cdot \frac{1}{g_t(x)} = \frac{r_i(x;t)^{1/\alpha}}{g_t(x)},$$

which is (E2); the share grows iff $r_i(x;t)^{1/\alpha} > g_t(x)$. □

Proof of Proposition **??**, *the no-finite-state claim.* That the calm drift (E1) depends on S_t only through $\bar{\lambda}_t^S = \int_\Phi \lambda \, dS_t$ is read off (E4), whose right-hand side involves the cross-section only through $\bar{\lambda}_t^S$. For the disaster tilt, (**??**) written against S_t is (**??**), a linear map $S \mapsto g(\cdot) = \int_\Phi H(\cdot, \phi) \, S(d\phi)$ on signed measures, with kernel $H(x, \phi) = (\lambda/\lambda^\star)^{1/\alpha}(f_\theta(x)/f(x))^{1/\alpha}$. Since the factor $(\lambda/\lambda^\star)^{1/\alpha} > 0$ is a θ-independent scalar, $\operatorname{span}\{H(\cdot, \phi) : \phi \in \Phi\} = \operatorname{span}\{(f_\theta/f)^{1/\alpha} : \theta \in \Theta\}$, which by Assumption **??**(v) is infinite-dimensional in $L^1(\mathcal{J}, f)$; hence the linear map $T : \eta \mapsto \int_\Phi H(\cdot, \phi) \, \eta(d\phi)$ has infinite-dimensional range.

Fix any finite collection of moment functionals $L_1, \dots, L_n$ of the pricing state (each $L_j(\eta) = \int_\Phi \ell_j \, d\eta$ for a bounded ℓ_j), and adjoin the total-mass functional $L_0(\eta) = \int_\Phi d\eta$. Work in the space $\mathcal{V}$ of signed measures on Φ with bounded Lebesgue density, and let $\mathcal{N} = \{\eta \in \mathcal{V} : L_0(\eta) = L_1(\eta) = \cdots = L_n(\eta) = 0\}$, the intersection of $n+1$ hyperplanes, so $\operatorname{codim}_{\mathcal{V}} \mathcal{N} \le n+1$.

There is $\eta \in \mathcal{N}$ with $T\eta \not\equiv 0$. Suppose not, i.e. $T(\mathcal{N}) = \{0\}$, so $\mathcal{N} \subseteq \ker T$. Then $\operatorname{codim} \ker T \le \operatorname{codim} \mathcal{N} \le n+1$, and by rank–nullity $\dim \operatorname{range}(T|_{\mathcal{V}}) = \operatorname{codim} \ker T \le n+1$. But T has infinite-dimensional range already on bounded-density measures (the functions $(f_\theta/f)^{1/\alpha}$ are approximated in $L^1(\mathcal{J}, f)$ by T of bounded-density measures concentrated near each θ, and are linearly independent by (v)), a contradiction. Hence such an η exists; fix it, with bounded density and $\|\eta\|_\infty < \infty$.

Admissibility of the perturbation. Let $\mathcal{S}$ denote the *admissible instantaneous pricing states*—probability measures on Φ with a density bounded below on the compact Φ, the valid wealth-weighted held-model cross-sections. The directions η constructed above keep the perturbation inside $\mathcal{S}$ (next paragraph; the explicit m-parameter version is Proposition IA.1), so the statement is about the *instantaneous pricing map* $S \mapsto (U_t, f_t^{\mathbb{Q}})$ over $\mathcal{S}$: no finite collection of moments is a sufficient statistic for it, two states in $\mathcal{S}$ sharing $L_1, \ldots, L_n$ yet pricing disasters differently. The dynamically reachable cross-sections—wealth-weighted pushforwards of the primitive type law G (Proposition D.3)—need not lie in $\mathcal{S}$—they can be atomic, as at the homogeneous start or on a finite menu; the claim is no finite sufficient statistic for the pricing map on $\mathcal{S}$, not that the economy visits any particular separating pair.

The perturbation is a probability measure. For $0 < \varepsilon \leq c/\|\eta\|_\infty$, the density of $S_t' = S_t + \varepsilon\eta$ is $\geq c - \varepsilon\|\eta\|_\infty \geq 0$, and its total mass is $1 + \varepsilon L_0(\eta) = 1$; so S_t' is a probability measure on Φ. By construction $L_j(S_t') = L_j(S_t)$ for $j = 1, \ldots, n$ (as $L_j(\eta) = 0$), yet

$$g_t'(\cdot) - g_t(\cdot) = T(S_t' - S_t) = \varepsilon\, T\eta \not\equiv 0 \quad \text{in } L^1(\mathcal{J}, f).$$

Distinct g_t give distinct prices. The risk-neutral compensator obeys $\nu_t^{\mathbb{Q}}(dx) \propto e^{\alpha x} g_t(x)^\alpha f(x)\, dx$ (**??**). The map $y \mapsto y^\alpha$ is strictly increasing on $(0, \infty)$, so $g_t' \neq g_t$ on a set of positive f-measure forces $g_t'^{\alpha} \neq g_t^\alpha$ there, hence $e^{\alpha x} g_t'^{\alpha} f \neq e^{\alpha x} g_t^\alpha f$ on that set and $\nu_t^{\mathbb{Q}\prime} \neq \nu_t^{\mathbb{Q}}$ as measures; two finite measures that differ as measures differ in at least one of total mass or normalized shape, and U_t is the total mass while $f_t^{\mathbb{Q}}$ is the normalized shape, so $(U_t', f_t^{\mathbb{Q}\prime}) \neq (U_t, f_t^{\mathbb{Q}})$. Thus S_t and S_t' agree on $L_1, \ldots, L_n$ but price disaster risk differently: no finite collection of moments is a sufficient statistic for $(U_t, f_t^{\mathbb{Q}})$. □

Corollary E.1 (Intensity channel: a scalar fear state): *If the size law is undisputed and correct—$f_{\theta_i} = f$ for all i, the held size densities equalling the objective f—then $r_i(x;t) = \lambda_i(t-)/\lambda^\star$ is independent of x,*

$$g_t(x) \equiv G_t := \int_\Phi \left(\tfrac{\lambda}{\lambda^\star}\right)^{1/\alpha} S_t(d\phi) = \int_0^1 s_t^i \left(\tfrac{\lambda_i(t)}{\lambda^\star}\right)^{1/\alpha} di, \tag{E7}$$

displayed at the generic state S_t; as a jump compensator G_t is read at the predictable left limit S_{t-}, the two agreeing off disasters. The risk-neutral size law then collapses to the fixed risk-

aversion tilt $f_t^{\mathbb{Q}}(x) = e^{\alpha x} f(x)/\mathbb{E}_f[e^{\alpha J}]$*, and the fear factor is the single scalar*

$$U_t = G_t^{\alpha}\, U^{\mathrm{RA}}. \tag{E8}$$

The size dispute is precisely what makes $f_t^{\mathbb{Q}}$ *move; absent it, beliefs price only the disaster* frequency U_t*, through the one scalar* G_t*.*

Proof of Corollary E.1. With $f_{\theta_i} = f$ the likelihood ratio (**??**) simplifies:

$$\begin{aligned} r_i(x;t) &= \frac{\lambda_i(t-)\, f_{\theta_i}(x)}{\lambda^\star f(x)} \\ &= \frac{\lambda_i(t-)\, f(x)}{\lambda^\star f(x)} \qquad \text{(E9)} \\ &= \frac{\lambda_i(t-)}{\lambda^\star}, \qquad \text{(E10)} \end{aligned}$$

line (E9) substituting $f_{\theta_i} = f$ and (E10) cancelling $f(x)$; r_i is independent of x. Substituting $r_i(x;t) = \lambda_i(t-)/\lambda^\star$ into (**??**),

$$g_t(x) = \int_0^1 s_{t-}^i \left(\tfrac{\lambda_i(t-)}{\lambda^\star}\right)^{1/\alpha} di =: G_t,$$

independent of x; this is (E7), the predictable aggregator evaluated at the left-limit S_{t-} (equal to S_t off disasters and off the predictable deadline dates, at which the calm switch flow can move S; (E7) displays it at the generic state). For the size law, start from (**??**) and use $g_t \equiv G_t$:

$$\begin{aligned} f_t^{\mathbb{Q}}(x) &= \frac{e^{\alpha x} g_t(x)^{\alpha} f(x)}{\mathbb{E}_f[e^{\alpha J} g_t(J)^{\alpha}]} \\ &= \frac{e^{\alpha x} G_t^{\alpha} f(x)}{\mathbb{E}_f[e^{\alpha J} G_t^{\alpha}]} \qquad \text{(E11)} \\ &= \frac{e^{\alpha x} G_t^{\alpha} f(x)}{G_t^{\alpha}\mathbb{E}_f[e^{\alpha J}]} \qquad \text{(E12)} \\ &= \frac{e^{\alpha x} f(x)}{\mathbb{E}_f[e^{\alpha J}]}, \qquad \text{(E13)} \end{aligned}$$

line (E11) substituting $g_t = G_t$, (E12) pulling the constant G_t^α out of the f-expectation, and (E13) cancelling it; the law is free of G_t. For the fear factor, from $U_t = \lambda_t^{\mathbb{Q}} = \lambda^\star \mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]$ (??),

$$U_t = \lambda^\star \, \mathbb{E}_f[e^{\alpha J} G_t^\alpha] \tag{E14}$$

$$= G_t^\alpha \, \lambda^\star \, \mathbb{E}_f[e^{\alpha J}] \tag{E15}$$

$$= G_t^\alpha \, U^{\mathrm{RA}}, \tag{E16}$$

line (E14) substituting $g_t = G_t$, (E15) pulling the constant G_t^α out of the expectation, and (E16) using $U^{\mathrm{RA}} = \lambda^\star \mathbb{E}_f[e^{\alpha J}]$ (Definition **??**); this is (E8). □

Corollary E.2 (Finite menu: finite-dimensional pricing functionals): *If the menu is finite,* $\Phi = \{\phi_1, \ldots, \phi_K\}$*, then* $S_t = (S_t^1, \ldots, S_t^K)$ *lies in the simplex* Δ^{K-1} *and the* instantaneous *pricing objects* $\bar{\lambda}_t^S$*,* $g_t(\cdot)$*,* $f_t^{\mathbb{Q}}$*, and* U_t *are functions of this finite vector. (The vector is a sufficient statistic for the* instantaneous *pricing objects above; it is not a sufficient statistic for forward-looking prices such as the price–dividend ratio* $v(\Psi_t)$*, nor by itself a Markov state for the dynamics—see the proof.) Under the intensity channel of Corollary E.1 with a two-point menu* $\{\lambda_{lo}, \lambda_{hi}\}$*, the instantaneous fear factor is a function of the single share* S_t^{hi}*,*

$$U_t = \Big[S_t^{hi} \big(\tfrac{\lambda_{hi}}{\lambda^\star}\big)^{1/\alpha} + (1 - S_t^{hi}) \big(\tfrac{\lambda_{lo}}{\lambda^\star}\big)^{1/\alpha} \Big]^\alpha U^{\mathrm{RA}}, \tag{E17}$$

recovering the scalar, intensity-only fear factor as a doubly-restricted special case.

Proof of Corollary E.2. With $\Phi = \{\phi_1, \ldots, \phi_K\}$ finite the pricing state is the atomic measure $S_t = \sum_{k=1}^K S_t^k \, \delta_{\phi_k}$ with $(S_t^1, \ldots, S_t^K) \in \Delta^{K-1}$. Each aggregate is the integral of a fixed function against S_t, so it collapses to a finite sum of the shares:

$$\bar{\lambda}_t^S = \int_\Phi \lambda\, S_t(d\phi) = \sum_{k=1}^{K} S_t^k\, \lambda_k,$$
$$g_t(x) = \int_\Phi r(\phi, x)^{1/\alpha}\, S_t(d\phi) = \sum_{k=1}^{K} S_t^k\, r(\phi_k, x)^{1/\alpha}, \tag{E18}$$

with $r(\phi, x) = \lambda f_\theta(x)/(\lambda^\star f(x))$ as in (??). Through (??), $f_t^{\mathbb{Q}}(x) = e^{\alpha x} g_t(x)^\alpha f(x)/\mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]$ and $U_t = \lambda^\star \mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]$ depend on the cross-section only through $g_t(\cdot)$, which by (E18) is a function of (S_t^k); hence $\bar{\lambda}_t^S, g_t(\cdot), f_t^{\mathbb{Q}}, U_t$ are functions of $(S_t^k) \in \Delta^{K-1}$.

Under the intensity channel (Corollary E.1) with $K = 2$ and $\Phi = \{\lambda_{lo}, \lambda_{hi}\}$, $f_\theta = f$ gives $r(\phi_k, x)^{1/\alpha} = (\lambda_k/\lambda^\star)^{1/\alpha}$ (independent of x), so (E18) becomes

$$\begin{aligned} g_t(x) \equiv G_t &= S_t^{hi}\big(\tfrac{\lambda_{hi}}{\lambda^\star}\big)^{1/\alpha} + S_t^{lo}\big(\tfrac{\lambda_{lo}}{\lambda^\star}\big)^{1/\alpha} \\ &= S_t^{hi}\big(\tfrac{\lambda_{hi}}{\lambda^\star}\big)^{1/\alpha} + (1 - S_t^{hi})\big(\tfrac{\lambda_{lo}}{\lambda^\star}\big)^{1/\alpha}, \end{aligned} \tag{E19}$$

line (E19) substituting $S_t^{lo} = 1 - S_t^{hi}$ (the simplex constraint). Substituting (E19) into $U_t = G_t^\alpha U^{\mathrm{RA}}$ (E8) gives (E17).

The reduction is of the *instantaneous* pricing map only. The vector (S_t^k) summarizes current prices, but it is not by itself a Markov state: its own law of motion is not autonomous, because the calm switch flow out of each class depends on the within-class distribution of regime ages and record statistics, not on the shares alone (Proposition **??**). The full Markov state therefore carries that within-class record structure, and the valuation operator acts on it rather than on (S_t^k) alone—a point made precise where the price level is computed. □

Proposition E.2 (The revision hazard as a deviance first-passage): *Fix an investor holding $\phi = (\lambda, \theta)$ with tolerance ε and since-adoption record $\mathcal{R} = (n, (x_1, \ldots, x_n))$ at age τ; let $D_{\phi,\tau}(\mathcal{R})$ be the deviance* (??) *and $r_{\phi,\tau}(\varepsilon)$ the acceptance threshold of* (??). *The revision time $\tau^{\mathrm{rev}} = \inf\{t : D_{\phi_i,\tau_t^i}(\mathcal{R}_t^i) \geq r_{\phi_i,\tau_t^i}(\varepsilon^i)\}$ is the first passage of the deviance across r (the closed-threshold* *form is shorthand for the p-value criterion $\bar{\varepsilon} \leq \varepsilon^i$, boundary counted; the conventions preceding Definition C.1 record exactly when the two coincide), with:*

(i) (calm drift) *between disasters the record is fixed and the deviance drifts deterministically,*

$$\frac{\partial}{\partial \tau} D_{\phi,\tau}(\mathcal{R}) = \lambda - \frac{n}{\tau}, \tag{E20}$$

rising—the held model increasingly contradicted—iff the realized rate n/τ is below the held λ; with $\mathcal{R}$ fixed, D meets the moving threshold $r_{\phi,\tau}(\varepsilon)$ at a deterministic calm deadline, the record-$\mathcal{R}$ generalization of $\Delta^*(\phi, \varepsilon)$, whose $n = 0$ slice is exactly Proposition D.1;

(ii) (disaster trigger) *a disaster arrives at the true rate $\lambda^\star$ (Assumption **??**) with size $x \sim f$, appends x to the record, and triggers rejection iff the updated deviance leaves the typical set, so the disaster-triggered rejection intensity is*

$$\lambda^\star \, \pi(\phi, \tau, \mathcal{R}; \varepsilon), \qquad \pi(\phi, \tau, \mathcal{R}; \varepsilon) = \int_{\mathcal{J}} \mathbf{1}\big\{D_{\phi,\tau}(\mathcal{R} \oplus x) \geq r_{\phi,\tau}(\varepsilon)\big\} f(x) \, dx, \tag{E21}$$

a functional of the full record $\mathcal{R}$, not of the age τ alone;

(iii) (no closed form) *τ^{rev} is the first passage of the piecewise-deterministic Markov process on the record-resolved state $(\phi, \mathcal{R}, \tau)$—the scalar deviance alone is not Markov in the size-dispute model—with calm drift (E20), mark updates $\mathcal{R} \mapsto \mathcal{R} \oplus x$ at rate $\lambda^\star f(x) \, dx$, and rejection exits on the trigger set (E21); for a general regular family $\{F_\theta\}$ neither the threshold $r_{\phi,\tau}(\varepsilon)$ nor the first-passage law has an elementary form, and both are computed numerically from the process.*

Two special cases calibrate the generality. A hazard in the low-dimensional state (n, τ)—the since-adoption count $n = n_t^i = N_t - N_{a^i}$ (**??**) and the age, rather than the full mark record—and an exactly computable passage law—finite Poisson sums over the (n, τ) grid—obtain under the restriction of Corollary E.1 (intensity-only), where $f_\theta \equiv f$ makes the deviance depend on the record only through the count n (so the rejection still reads the since-adoption count n, not the age alone). The finite menu of Corollary E.2 reduces the *pricing* functionals to finite dimension but leaves the rejection hazard record-dependent—the marks x_j still enter the deviance through f_θ—so in general the rejection is governed by the full record.

Proof of Proposition E.2. (i) With the held model ϕ fixed and no disaster on the interval ending at age τ, the realized sizes and their MLE $\hat\theta$ are constant, so the size term of (**??**) is constant in τ and only the count term $C(\tau) := \tau\, d_P(n/\tau \parallel \lambda)$ varies. Substituting $a = n/\tau$ into $d_P(a\|\lambda) = a \log \frac{a}{\lambda} - a + \lambda$,

$$
\begin{aligned}
C(\tau) &= \tau\left[\tfrac{n}{\tau}\log\tfrac{n/\tau}{\lambda} - \tfrac{n}{\tau} + \lambda\right] \\
&= n\log\tfrac{n/\tau}{\lambda} - n + \lambda\tau && \text{(E22)}\\
&= n\log\tfrac{n}{\tau\lambda} - n + \lambda\tau, && \text{(E23)}
\end{aligned}
$$

line (E22) distributing the factor τ and (E23) using $\frac{n/\tau}{\lambda} = \frac{n}{\tau\lambda}$; with the convention $0\log 0 = 0$ this holds also at $n = 0$, where $C(\tau) = \lambda\tau$. To differentiate, expand the logarithm $\log\frac{n}{\tau\lambda} = \log n - \log\tau - \log\lambda$:

$$C(\tau) = n\log n - n\log\tau - n\log\lambda - n + \lambda\tau.$$

Only the $-n\log\tau$ and $\lambda\tau$ terms depend on τ, and $D_{\phi,\tau}(\mathcal{R})$ differs from $C(\tau)$ by a τ-constant, so

$$
\begin{aligned}
\frac{\partial}{\partial\tau} D_{\phi,\tau}(\mathcal{R}) = C'(\tau) &= -\frac{n}{\tau} + \lambda && \text{(E24)}\\
&= \lambda - \frac{n}{\tau}, && \text{(E25)}
\end{aligned}
$$

(for $n \geq 1$; at $n = 0$ the count deviance is $\lambda\tau$ outright and its τ-derivative is λ, no $\log n$ expansion entering) line (E24) using $\partial_\tau(-n\log\tau) = -n/\tau$ and $\partial_\tau(\lambda\tau) = \lambda$ and (E25) reordering; this is (E20). The drift is positive—the held model increasingly contradicted—iff $\lambda > n/\tau$, i.e. the realized rate n/τ lies below the held λ. With $\mathcal{R}$ fixed, $D_{\phi,\tau}(\mathcal{R})$ meets the moving threshold $r_{\phi,\tau}(\varepsilon)$ at a deterministic calm deadline (the record-$\mathcal{R}$ generalization of $\Delta^*(\phi,\varepsilon)$); at the $n=0$ slice $C(\tau)=\lambda\tau$ and $\partial_\tau D=\lambda>0$, recovering the deterministic deadline $\Delta^*(\phi,\varepsilon)$ of Proposition D.1.

(ii) Disasters arrive as a marked Poisson process with $(\mathcal{F}_t,\mathbb{P})$-compensator $\nu(dt,dx) = \lambda^\star\,dt\,f(x)\,dx$ (**??**) (Assumption **??**). A disaster of size x at age τ appends x to the record, $\mathcal{R}\mapsto\mathcal{R}\oplus x=(n+1,(x_1,\dots,x_n,x))$, and by the pass/fail rule (**??**) the test rejects at that instant iff the updated record leaves the typical set,

$$D_{\phi,\tau}(\mathcal{R}\oplus x)\geq r_{\phi,\tau}(\varepsilon).$$

The triggering marks form the set $B_{\phi,\tau,\mathcal{R}} := \{x\in\mathcal{J} : D_{\phi,\tau}(\mathcal{R}\oplus x)\geq r_{\phi,\tau}(\varepsilon)\}$, matching the closed-threshold convention of Definition **??**; its boundary $\{x : D_{\phi,\tau}(\mathcal{R}\oplus x) = r_{\phi,\tau}(\varepsilon)\}$ under the closed-threshold convention of Definition **??** the trigger set and its intensity π are unambiguous; the open convention can differ—in the intensity channel $D_{\phi,\tau}(R\oplus x)$ is mark-independent, so the equality set can carry full mark mass—and is not used. Restricting the disaster random measure to $B_{\phi,\tau,\mathcal{R}}$ predictably thins it to a point process whose (record-dependent) intensity is the ν-mass of $B_{\phi,\tau,\mathcal{R}}$:

$$\begin{aligned}
\int_{B_{\phi,\tau,\mathcal{R}}}\frac{\nu(dt,dx)}{dt} &= \lambda^\star\int_{\mathcal{J}}\mathbf{1}\{x\in B_{\phi,\tau,\mathcal{R}}\}\,f(x)\,dx && \text{(E26)}\\
&= \lambda^\star\int_{\mathcal{J}}\mathbf{1}\{D_{\phi,\tau}(\mathcal{R}\oplus x)\geq r_{\phi,\tau}(\varepsilon)\}\,f(x)\,dx && \text{(E27)}\\
&= \lambda^\star\,\pi(\phi,\tau,\mathcal{R};\varepsilon), && \text{(E28)}
\end{aligned}$$

line (E26) writing the restricted mass against $\nu=\lambda^\star\,dt\,f\,dx$, (E27) substituting the definition of $B_{\phi,\tau,\mathcal{R}}$, and (E28) the definition of π; this is the disaster-triggered rejection intensity (E21). It depends on $\mathcal{R}$ through the map $x\mapsto D_{\phi,\tau}(\mathcal{R}\oplus x)$, not on the age τ alone.

(iii) Between disasters D evolves by (E20); at disasters it jumps as in (ii); τ^{rev} is the

first time D exceeds $r_{\phi,\tau}(\varepsilon)$—the first passage of this piecewise-deterministic Markov process. The threshold is the $(1-\varepsilon)$-quantile of $D_{\phi,\tau}$ under $Q_{\phi,\tau}$ (**??**), a Poisson–mark mixture with no elementary quantile for general $\{F_\theta\}$, and the first-passage law of a PDMP across a moving boundary is likewise not elementary; both are computed numerically. The state collapses to (n,τ) only when the size law is undisputed, and to a finite vector under a finite menu; in general the rejection is governed by the full record. □

Lemma E.1 (Rejection hazard, switch hazard, and the switch flow): *Fix the public posterior ρ_t and its mode $\phi_t^* = \mathfrak{m}(\rho_t)$ (**??**). For an investor in state $(\phi,\mathcal{R},\tau,\varepsilon)$, let the* rejection hazard *$H_t^{\mathrm{rej}}(d\tau' \mid \phi,\mathcal{R},\tau,\varepsilon)$ be the hazard (intensity) measure of the first-passage time τ^{rev} of Proposition E.2—a disaster-triggered rate $\lambda^\star\pi(\phi,\tau,\mathcal{R};\varepsilon)$ (E21) together with* certain *rejection at the deterministic calm deadline $\Delta^\star$, so that τ^{rev} is the earlier of $\Delta^\star$ and the first disaster-triggered rejection, so along the fixed record the compensator of τ^{rev} is $\lambda^\star\pi\,dt$ plus a unit atom at $\Delta^\star$; the probability that the record itself survives unchanged to u is $\mathbf{1}\{u<\Delta^\star\}e^{-\lambda^\star(u-\tau)}$ and the marginal law of τ^{rev} is the recursive piecewise-deterministic object, not this local factor, dropping to zero at $\Delta^\star$ (the deterministic deadline is not represented as a unit atom of an ordinary exponential hazard). Write $\beta_t(d\phi)$ for the wealth-weighted rate at which holders of ϕ reach their scheduled calm deadlines at t. Then:*

(i) the switch hazard *is the rejection hazard restricted to the switch event, with the destination evaluated where it is resolved. Across a scheduled calm deadline the public posterior is continuous, so the calm component is restricted by the predictable indicator,*

$$H_t^{\mathrm{sw}}(d\tau' \mid \phi,\mathcal{R},\tau,\varepsilon) = \mathbf{1}\{\phi \neq \phi_{t-}^*\}\, H_t^{\mathrm{rej}}(d\tau' \mid \phi,\mathcal{R},\tau,\varepsilon) \qquad \textit{on the calm channel;} \tag{E29}$$

on the disaster channel the destination is resolved by the mark, so the switch propensity at a size-x disaster is $\mathbf{1}\{D_{\phi,\tau}(\mathcal{R}\oplus x) \geq r_{\phi,\tau}(\varepsilon)\}\,\mathbf{1}\{\phi \neq \mathfrak{m}(\rho_{t-}\oplus x)\}$—$\mathfrak{m}(\rho_{t-}\oplus x)$ the mode of the posterior on the public record augmented by x—and the predictable switch intensity integrates it against the mark law, $\lambda^\star\int_{\mathcal{J}} \mathbf{1}\{D_{\phi,\tau}(\mathcal{R}\oplus x) \geq r_{\phi,\tau}(\varepsilon)\}\,\mathbf{1}\{\phi \neq \mathfrak{m}(\rho_{t-}\oplus$

$x)\}\, f(x)\, dx$;

(ii) the switch flow has two parts on different timescales: a continuous *calm switch flow*

$$\Phi_t^{\mathrm{sw}}(A) = \int_A \mathbf{1}\{\phi \neq \phi_t^*\}\, \beta_t(d\phi), \qquad A \subseteq \Phi, \tag{E30}$$

*with β_t non-atomic where deadlines stagger (and carrying a predictable atom where a positive G-mass synchronizes, which by Lemma **??** moves no traded price), and a disaster-triggered component whose predictable rejection* intensity *(compensator, integrating over marks) is $\lambda^\star \pi(\phi, \tau, \mathcal{R}; \varepsilon)$* (E21)*, the relocations among the rejectors being those for whom the held model differs from the post-jump mode (a self-re-adoption moves nothing); the realized state jump at a disaster epoch is set by the realized mark x and its triggered set, not by a continuous rate. Every law-of-motion term runs against the switch set $\{\phi \neq \phi_t^*\}$, never the raw rejection rate, so a self-re-adoption ($\phi = \phi_t^*$) renews the test window but leaves ϕ, hence S_t, unchanged.*

On the hazard's reading: every disaster, rejecting or not, moves the state; a non-rejecting disaster appends a mark, updating both π and the next deadline $\Delta^\star$, after which the first passage continues as the piecewise-deterministic law of Proposition E.2(iii). The indicator is an instantaneous-at-t restriction defining the *current* flow; at a future rejection τ' the relevant mode is $\mathfrak{m}(\rho_{\tau'})$, and the hazard is re-read there.

Proof of Lemma E.1. (i) Switch hazard. By the re-selection map (**??**), after a rejection the held model is $\phi_t^* = \mathfrak{m}(\rho_t)$, which differs from the incumbent ϕ exactly on the switch event $\{\phi \neq \phi_t^*\}$. A rejection resets the record $\mathcal{R} \to \varnothing$ and age $\tau \to 0$ whether or not the held model moves, but the held model itself moves only on a switch; hence the held-model part of the rejection hazard is its restriction to $\{\phi \neq \phi_t^*\}$, which is (E29). The current-mode indicator is exact for the *calm* (deadline) part of H^{rej}; on the disaster-triggered part the switch event is evaluated at the post-jump mode, so the triggered *switch* hazard is $\lambda^\star \int_{\mathcal{J}} \mathbf{1}\{D_{\phi,\tau}(\mathcal{R} \oplus x) \geq r_{\phi,\tau}(\varepsilon)\}\, \mathbf{1}\{\phi \neq$

$\mathfrak{m}(\rho_{t-} \oplus x)\} \, f(x) \, dx$, by splitting the trigger indicator over the destination event—the law of motion, Proposition **??**(iii), reads the relocation exactly there, and a triggered self-re-adoption moves nothing. A self-re-adoption $\phi = \phi_t^*$ renews the test window but leaves ϕ, and therefore S_t, unchanged.

(ii) Two timescales. For a single investor the calm-deadline first passage $\Delta^\star$ is deterministic—an atom of her revision law. Across investors the induced calm-deadline law is atomless on strictly monotone ranges: the tolerance law G is continuous (Assumption **??**) and the calm deadline $\varepsilon \mapsto a + \Delta^\star(\phi, \varepsilon)$ is non-increasing in ε, and *strictly* so wherever the continuous size-deviance component of (**??**) renders the threshold $r_{\phi,\Delta}(\varepsilon)$ strictly monotone (Proposition E.2); on such ranges the deadline map has G-null level sets, its pushforward places no positive mass at any single date, the wealth-weighted deadline-arrival rate $\beta_t(d\phi)$ is non-atomic, and the calm switch flow (E30) is a *continuous* finite-variation flow. On a flat level set of $\Delta^\star(\phi, \cdot)$ of positive G-mass—possible where the discrete count deviance dominates—the deadline-coincident revisers re-select simultaneously, contributing a *predictable* finite-variation jump rather than a Brownian increment; either way the flow is finite-variation, and the *traded* price is continuous across it (Lemma **??**), so the no-arbitrage decomposition is unaffected. The disaster-triggered component, whose predictable intensity (compensator, integrating over marks) is $\lambda^\star \pi(\phi, \tau, \mathcal{R}; \varepsilon)$ (E21), can fire only at a disaster epoch; the realized jump is then set by the realized mark x, simultaneously for the whole triggered set, and is not a continuous rate. Both parts are restricted to $\{\phi \neq \phi_t^*\}$, so a self-re-adoption moves no held-model mass and leaves S_t unchanged; it does reset the reviser's record and age, a movement of the record-resolved Ψ_t that the forward-looking valuation tracks (Proposition **??**). $\square$

Remark E.1 (Why a stationary fear cross-section needs demographic refresh). The re-selection (**??**) reads the second-order posterior ρ_t^i, updated on the *full* history (**??**), whereas the test reads only the since-adoption window: the two memories differ. Under the true DGP and the identifiable family of Assumption **??**, ρ_t^i is consistent for the Kullback–Leibler model $\phi^\star$, so $\mathrm{MAP}(\rho_t^i) \to \phi^\star$; once an infinitely-lived investor holds $\phi^\star$ she re-selects near $\phi^\star$ at each rejection, and the held-model cross-section concentrates on $\phi^\star$—the fear dispersion dies out. A *nondegenerate* stationary fear cross-section is therefore not a property of the present primitives; it is sustained by demographic refresh—newborns who enter holding a *drawn* model $\phi_0 \sim \Pi_\phi$ with a reset test clock (Assumption J.1(iv)), held until their first rejection—which an overlapping-generations birth–death structure supplies, added as an explicit primitive where the stationary cross-section is required. The refresh is in the *held models*, not the posteriors: each newborn inherits the (concentrated) common posterior ρ_t, so $\mathfrak{m}(\rho_t)$ stays common across cohorts, but holds her drawn $\phi_0 \neq \mathfrak{m}(\rho_t)$ until she revises, keeping a positive mass of wealth away from $\phi^\star$. The drawn ϕ_0 is an exogenous status-quo endowment, not a birth-time MAP re-initialization: a newborn carries her inherited held model until her own test rejects it, and the MAP rule (**??**) governs only re-selection at rejection, exactly as for an incumbent. This departure from MAP-initialization at birth is the OLG primitive on which the stationary dispersion rests. We do not assert stationarity from the present primitives. *(Not proved here; the posterior-consistency step and the stationary cross-section are formalized with the demographic primitive.)*

Proof of Proposition **??**, *the law of motion.* *(i)* Between disasters and switches every held model ϕ_i is fixed. Writing S_t as the pushforward of the wealth measure $s_t^i\,di$ under $i \mapsto \phi_i$ (Definition **??**), $\langle h, S_t\rangle = \int_0^1 h(\phi_i)\, s_t^i\, di$. The interchange of $\frac{d}{dt}$ and $\int_0^1 \cdot\, di$ is justified by dominated convergence: h is bounded and, on compact t-intervals, s_t^i and its t-derivative are bounded uniformly in i (as in Proposition E.1, via (**??**) and $\lambda_i \in [\underline{\lambda}, \overline{\lambda}]$). Then

$$\begin{aligned}
\frac{d}{dt}\langle h, S_t\rangle &= \int_0^1 h(\phi_i)\,\frac{ds_t^i}{dt}\,di \\
&= \int_0^1 h(\phi_i)\, s_t^i\,\frac{d\log s_t^i}{dt}\,di && \text{(E31)}\\
&= \int_0^1 h(\phi_i)\, s_t^i\,\tfrac{1}{\alpha}\big(\bar{\lambda}_t^S - \lambda_i\big)\,di && \text{(E32)}\\
&= \tfrac{1}{\alpha}\int_\Phi \big(\bar{\lambda}_t^S - \lambda\big)\,h(\phi)\,S_t(d\phi), && \text{(E33)}
\end{aligned}$$

line (E31) using $ds_t^i/dt = s_t^i\, d\log s_t^i/dt$, (E32) substituting the calm share drift (E6), and (E33) pulling out the constant $\frac{1}{\alpha}$ and rewriting the i-integral as a ϕ-integral against S_t (Definition **??**); this is (**??**).

(ii) At a calm-deadline switch, investor i's held model jumps from ϕ_i to $\phi_i' = \mathfrak{m}(\rho_t^i)$ (**??**). By Lemma D.1 the second-order posterior is common, $\rho_t^i = \rho_t$, so $\phi_i' = \mathfrak{m}(\rho_t) = \phi_t^*$ for every investor revising at t. Her contribution to $\langle h, S_t\rangle$ changes by $s_t^i[h(\phi_t^*) - h(\phi_i)]$; the aggregate rate of switching out of ϕ is the calm switch flow $\Phi_t^{\mathrm{sw}}(d\phi)$ of Lemma E.1 (E30) (restricted to $\{\phi \neq \phi_t^*\}$). Where the deadline pushforward of the tolerance law is atomless the flow is continuous in time; on a flat deadline level set of positive G-mass it carries a predictable finite-variation atom, and (**??**) is then read in the Stieltjes sense, integrating $h(\phi_t^*) - h(\phi)$ against $d\Phi^{\mathrm{sw}}$—a predictable step in S_t that by Lemma **??** moves no traded price. Aggregating gives (**??**). Disaster-triggered switches do not enter here—they fire only at disaster epochs, in (iii).

(iii) At a disaster of size x, Proposition E.1(ii) gives $s_t^i = s_{t-}^i\, r(\phi_i, x)^{1/\alpha}/g_t(x)$ (with $r_i(x;t) = r(\phi_i, x)$). Hence for any bounded h,

$$\begin{aligned}
\langle h, S_t\rangle &= \int_0^1 h(\phi_i)\, s_t^i\, di \\
&= \int_0^1 h(\phi_i)\, s_{t-}^i\, \frac{r(\phi_i, x)^{1/\alpha}}{g_t(x)}\, di && \text{(E34)} \\
&= \frac{1}{g_t(x)} \int_0^1 h(\phi_i)\, r(\phi_i, x)^{1/\alpha}\, s_{t-}^i\, di && \text{(E35)} \\
&= \frac{1}{g_t(x)} \int_\Phi h(\phi)\, r(\phi, x)^{1/\alpha}\, S_{t-}(d\phi), && \text{(E36)}
\end{aligned}$$

line (E34) substituting the share jump, (E35) pulling out the constant $g_t(x)^{-1}$, and (E36) rewriting the i-integral against S_{t-} (Definition **??**). Throughout the display, s_t^i and $g_t(x)$ denote the *reweight-only* intermediate state—step (a) of Definition **??**, before relocation—as in Proposition E.1(ii). As this holds for every bounded h, it is the measure identity (**??**)—but only the *reweight* step (a) of Definition **??**, in which h is still read at the pre-disaster model ϕ. The full post-disaster state applies the relocation step (d) on top: the holders the disaster triggers move to ϕ_t^* at the disaster epoch (a jump, not the continuous flow of (ii)), so for them h is read at ϕ_t^*. The triggered investors form the region $\mathcal{T}_t^x := \{(\phi, \mathcal{R}, \tau, \varepsilon) : D_{\phi,\tau}(\mathcal{R} \oplus x) \geq r_{\phi,\tau}(\varepsilon)\}$ (E21) in the *record-resolved* state—*not* a subset of Φ, since the trigger depends on the full state and two investors sharing ϕ can differ in it—so the relocation is read against Ψ_{t-}, with post-disaster destination $\phi_t^*(x) := \mathfrak{m}(\rho_{t-} \oplus x)$ reading the mark-updated posterior (independent of x in the intensity channel):

$$\langle h, S_t^x \rangle = \frac{1}{g_t(x)} \int r(\phi, x)^{1/\alpha} \Big[h(\phi)\, \mathbf{1}\{(\phi, \mathcal{R}, \tau, \varepsilon) \notin \mathcal{T}_t^x\} + h(\phi_t^*(x))\, \mathbf{1}\{(\phi, \mathcal{R}, \tau, \varepsilon) \in \mathcal{T}_t^x\} \Big]\, \Psi_{t-}\big(d(\phi, \mathcal{R}, \tau, \varepsilon)\big); \tag{E37}$$

the reweight (E36) is the special case $\mathcal{T}_t^x = \varnothing$ (no triggered switchers). Since $\mathcal{T}_t^x$ and $\phi_t^*(x)$ are functionals of the record cross-section and the public posterior, the disaster map (E37) reads Ψ_{t-} and not S_{t-} alone; that no finite collection of moments of S_t suffices even for the instantaneous prices is the separate claim proved above. □

Appendix F. Dynamics around a disaster

In words, item by item. *(i)* selection erodes the mean belief at a rate equal to its cross-sectional variance, scaled by the risk tolerance $1/\alpha$, while the calm switch flow relocates the staggered revisers to the calm destination ϕ_t^*. For switchers whose since-adoption records already carry non-rejecting disasters ($n \geq 1$), a disaster can have raised the public mode, so $\lambda(\phi_t^*) \leq \lambda$ is not automatic; that is the condition's content. Its regime is the finite-history mode: along calm the public mode follows the calm evidence to the complacent point—on the computed record, about seven years after a disaster—where the condition holds; in the locked stationary limit of Proposition K.1 the destination is the menu's Kullback–Leibler projection, the fearful point at the binary menu, and calm relaxation there falls to selection and turnover. In the intensity channel, and under that condition generally, both calm channels push the mean down. *(ii)* The relocation contribution *adds* to the jump when $\lambda(\phi_t^*)$ exceeds that post-reweight mean, a cross-sectional condition we state rather than derive. The unconditional fast-in is the reweight

(**??**) alone—a covariance signed without reference to the re-selection direction, so the fear jump does not require the disaster to raise the held mode—while the relocation is a conditional amplifier. *(iii)* The calm decline of U_t carries down with it the U_t-loading component of the disaster premium, the jump ratio ψ_t held fixed—the event-study reading of Section **??**. *(iv)* The horizon is long when ε or the held intensity is small (the dependence on $\lambda^\star$ is indirect—through the held intensity coordinate and the chance a calm spell is interrupted by the next disaster—not a direct property of (D4)).

Part (i) is market selection in its sharpest local form—the replicator drift in which the mean held belief falls at a rate equal to its own cross-sectional variance scaled by risk tolerance, the transient face of the long-run selection results (Sandroni, 2000; Blume and Easley, 2006; Yan, 2008) specialized to an economy where "accurate" is path-contingent: the complacent are vindicated in calm, the fearful at disasters. The asymmetry is the content—a disaster reweights toward fear and lifts the posterior mode at one instant (fast in), while calm erodes the wave only as fast as the surviving disagreement that selection consumes, staggered across the population by the deadline spread (slow out)—the model counterpart of the persistent post-crisis tail factor of Andersen et al. (2015), an asymmetry visible in the beliefs themselves as well: surveyed firms' subjective disaster probabilities spike on a shock and decay only slowly through the calm that follows (Menkhoff, 2025). The signs are intensity-channel statements—in the joint channel they need λ and $r(\cdot, x)^{1/\alpha}$ positively associated, which we flag rather than assume—and the *magnitudes* (the jump size, the decay rate $\frac{1}{\alpha}\,\mathrm{Var}_S(\lambda)$, the horizon $\Delta^\star$) wait on the stationary cross-section of Section **??**.

The three price functionals move together at a disaster, each driven by the same jump in the fear factor: the disaster premium's U_t-component rises and drifts down through calm at the rate (**??**) (Proposition **??**); the short rate moves by $dr_t = d\bar{\lambda}_t^S - dU_t$, falling *iff* the priced-tilt response dominates the mean-intensity response (Proposition **??**); and the excess-variance term, generated by movement of the cross-section (supplementary part), bursts when the repricing ratio ψ_t amplifies the jump. Under $dU_t > d\bar{\lambda}^S_t$ and in the intensity channel one belief object thus delivers a premium jump up, a safe-rate jump down, and a variance burst, each relaxing slowly through the staggered calm recovery. The cycle's *shape* is pinned by the primitives; its *scale*, and the demographic refresh that makes it stationary, is the work of Section **??**. A forecast-data signature—lazy revision read as state-dependent information rigidity, and its place relative to the under- and over-reaction evidence—is developed in the supplementary part below.

Proof of Proposition **??**. *(i) Calm erosion.* Between disasters the calm derivative is the sum of the selection drift (**??**) and the calm switch flow (**??**) (Proposition **??**(i)–(ii)). The selection part gives, for bounded h,

$$\tfrac{d}{dt}\langle h, S_t\rangle = \tfrac{1}{\alpha}\int_\Phi (\bar{\lambda}^S_t - \lambda)\, h(\phi)\, S_t(d\phi) \tag{F1}$$

$$= \tfrac{1}{\alpha}\Big[\bar{\lambda}^S_t \int_\Phi h\, dS_t - \int_\Phi \lambda\, h\, dS_t\Big] \tag{F2}$$

$$= \tfrac{1}{\alpha}\big[\mathbb{E}_{S_t}[\lambda]\,\mathbb{E}_{S_t}[h] - \mathbb{E}_{S_t}[\lambda h]\big] \tag{F3}$$

$$= -\tfrac{1}{\alpha}\big(\mathbb{E}_{S_t}[\lambda h] - \mathbb{E}_{S_t}[\lambda]\,\mathbb{E}_{S_t}[h]\big) \tag{F4}$$

$$= -\tfrac{1}{\alpha}\,\mathrm{Cov}_{S_t}(\lambda, h), \tag{F5}$$

where (F1) is (**??**); (F2) distributes $(\bar{\lambda}^S_t - \lambda)h = \bar{\lambda}^S_t h - \lambda h$ and splits the integral; (F3) writes $\int h\, dS_t = \mathbb{E}_{S_t}[h]$, $\int \lambda h\, dS_t = \mathbb{E}_{S_t}[\lambda h]$ and substitutes $\bar{\lambda}^S_t = \mathbb{E}_{S_t}[\lambda]$; (F4) factors -1; and (F5) is the definition of covariance; the calm switch flow (**??**) adds $\int_\Phi [h(\phi^*_t) - h(\phi)]\,\Phi^{\mathrm{sw}}_t(d\phi)$. Setting $h(\phi) = \lambda$, the selection part is $-\frac{1}{\alpha}\,\mathrm{Var}_{S_t}(\lambda)$ and the switch part is $\int_\Phi (\lambda(\phi^*_t) - \lambda)\,\Phi^{\mathrm{sw}}_t(d\phi)$, which is ≤ 0 on the disaster-free-since-adoption part of the switch set—there the scheduled relief of Proposition D.1 gives $\lambda(\phi^*_t) \le \lambda$ (and $\Phi^{\mathrm{sw}}_t \ge 0$)—and non-positive on the rest under the monotone-reselection condition, so

$$\frac{d\bar{\lambda}^S_t}{dt} = -\tfrac{1}{\alpha}\,\mathrm{Var}_{S_t}(\lambda) + \int_\Phi \big(\lambda(\phi^*_t) - \lambda\big)\,\Phi^{\mathrm{sw}}_t(d\phi), \tag{F6}$$

the selection term non-positive ($\mathrm{Var}_{S_t}(\lambda) \geq 0$, $\alpha > 0$) and the switch term non-positive under that condition; this is (**??**). Equation (F6) is the absolutely continuous (dt-rate) part of the calm derivative; where a positive G-mass synchronizes at a calm deadline, Φ_t^{sw} carries a predictable atom and $\bar{\lambda}_t^S$ a predictable jump, which by Lemma **??** moves no traded price.

(ii) Disaster reweight. At a disaster of size x, before triggered switches, the cross-section reweights by $\widetilde{S}_t^x(d\phi) = r(\phi, x)^{1/\alpha} g_t(x)^{-1} S_{t-}(d\phi)$ (**??**), with $g_t(x) = \mathbb{E}_{S_{t-}}[r(\cdot, x)^{1/\alpha}]$ (**??**). Hence

$$
\begin{aligned}
\bar{\lambda}_t^{\widetilde{S}} &= \int_\Phi \lambda \, \widetilde{S}_t^x(d\phi) \\
&= \int_\Phi \lambda \, r(\phi, x)^{1/\alpha} g_t(x)^{-1} S_{t-}(d\phi) && \text{(F7)} \\
&= \frac{\mathbb{E}_{S_{t-}}[\lambda \, r(\cdot, x)^{1/\alpha}]}{\mathbb{E}_{S_{t-}}[r(\cdot, x)^{1/\alpha}]}, && \text{(F8)}
\end{aligned}
$$

(F7) substituting the reweight and (F8) pulling out the constant $g_t(x)^{-1}$ and writing numerator and denominator ($= g_t(x)$) as expectations under S_{t-}. Subtract $\bar{\lambda}_{t-}^S = \mathbb{E}_{S_{t-}}[\lambda]$:

$$
\begin{aligned}
\bar{\lambda}_t^{\widetilde{S}} - \bar{\lambda}_{t-}^S &= \frac{\mathbb{E}_{S_{t-}}[\lambda \, r(\cdot, x)^{1/\alpha}]}{\mathbb{E}_{S_{t-}}[r(\cdot, x)^{1/\alpha}]} - \mathbb{E}_{S_{t-}}[\lambda] \\
&= \frac{\mathbb{E}_{S_{t-}}[\lambda \, r(\cdot, x)^{1/\alpha}] - \mathbb{E}_{S_{t-}}[\lambda] \, \mathbb{E}_{S_{t-}}[r(\cdot, x)^{1/\alpha}]}{\mathbb{E}_{S_{t-}}[r(\cdot, x)^{1/\alpha}]} && \text{(F9)} \\
&= \frac{\mathrm{Cov}_{S_{t-}}(\lambda, \, r(\cdot, x)^{1/\alpha})}{\mathbb{E}_{S_{t-}}[r(\cdot, x)^{1/\alpha}]}, && \text{(F10)}
\end{aligned}
$$

with $\bar{\lambda}_t^{\widetilde{S}}$ the post-reweight mean; the final $\bar{\lambda}_t^S$ adds the triggered switchers' relocation term.

(F9) over the common denominator and (F10) the covariance definition; this is (**??**). To sign the numerator, expand the covariance with an independent copy. Let (λ, ϕ) and (λ', ϕ') be two independent draws from S_{t-} and write $k := r(\phi, x)^{1/\alpha}$, $k' := r(\phi', x)^{1/\alpha}$; all expectations below are under S_{t-}. Then

$$\mathbb{E}\big[(\lambda-\lambda')(k-k')\big] = \mathbb{E}[\lambda k] - \mathbb{E}[\lambda k'] - \mathbb{E}[\lambda' k] + \mathbb{E}[\lambda' k'] \tag{F11}$$

$$= \mathbb{E}[\lambda k] - \mathbb{E}[\lambda]\,\mathbb{E}[k'] - \mathbb{E}[\lambda']\,\mathbb{E}[k] + \mathbb{E}[\lambda' k'] \tag{F12}$$

$$= 2\,\mathbb{E}[\lambda k] - 2\,\mathbb{E}[\lambda]\,\mathbb{E}[k] \tag{F13}$$

$$= 2\ \mathrm{Cov}_{S_{t-}}(\lambda, k), \tag{F14}$$

(F11) expanding the product inside the expectation, (F12) factoring the cross terms $\mathbb{E}[\lambda k'] = \mathbb{E}[\lambda]\mathbb{E}[k']$ and $\mathbb{E}[\lambda' k] = \mathbb{E}[\lambda']\mathbb{E}[k]$ by independence of the two draws, (F13) using that the draws are identically distributed ($\mathbb{E}[\lambda' k'] = \mathbb{E}[\lambda k]$, $\mathbb{E}[\lambda'] = \mathbb{E}[\lambda]$, $\mathbb{E}[k'] = \mathbb{E}[k]$), and (F14) the covariance definition. Dividing (F14) by 2 and writing $k = r(\cdot, x)^{1/\alpha}$,

$$\mathrm{Cov}_{S_{t-}}\big(\lambda, r(\cdot, x)^{1/\alpha}\big) = \tfrac{1}{2}\mathbb{E}\big[(\lambda-\lambda')\big(r(\phi, x)^{1/\alpha} - r(\phi', x)^{1/\alpha}\big)\big]. \tag{F15}$$

In the intensity channel $r(\phi, x) = \lambda/\lambda^\star$, so $r(\phi, x)^{1/\alpha} = (\lambda/\lambda^\star)^{1/\alpha}$ is non-decreasing in λ; the two factors in (F15) then carry the common sign of $\lambda - \lambda'$, their product is ≥ 0 pointwise, and the expectation is ≥ 0, so $\mathrm{Cov}_{S_{t-}}(\lambda, r(\cdot, x)^{1/\alpha}) \geq 0$, strict iff $\mathrm{Var}_{S_{t-}}(\lambda) > 0$. Thus the reweight raises $\bar{\lambda}^S$. For the relocation mode: by the proof of Lemma D.1 the intensity log-posterior is $\ell_t^\Lambda(\lambda) = \log \rho_0^\Lambda(\lambda) - \lambda t + N_t \log \lambda$, and at a disaster the count N_t rises by one at fixed t. We show the intensity mode $\lambda_t^{\mathrm{m}} = \min \arg\max_\lambda \ell_t^\Lambda$ is non-decreasing in N_t directly, by revealed preference. Write $\ell^{(N)}(\lambda) := \log \rho_0^\Lambda(\lambda) - \lambda t + N \log \lambda$ and let λ_N, λ_{N+1} be the corresponding min-maximisers. Optimality at each count gives

$$\ell^{(N)}(\lambda_N) \geq \ell^{(N)}(\lambda_{N+1}), \qquad \ell^{(N+1)}(\lambda_{N+1}) \geq \ell^{(N+1)}(\lambda_N).$$

Adding the two inequalities and regrouping,

$$\big[\ell^{(N+1)}(\lambda_{N+1}) - \ell^{(N)}(\lambda_{N+1})\big] - \big[\ell^{(N+1)}(\lambda_N) - \ell^{(N)}(\lambda_N)\big]\ \geq\ 0. \tag{F16}$$

Since $\ell^{(N+1)}(\lambda) - \ell^{(N)}(\lambda) = \log\lambda$ for every λ, the left-hand side of (F16) equals $\log\lambda_{N+1} - \log\lambda_N$; hence $\log\lambda_{N+1} \geq \log\lambda_N$, and since log is strictly increasing, $\lambda_{N+1} \geq \lambda_N$. Thus the intensity mode of ϕ_t^* is non-decreasing in N_t: a disaster raises N_t by one, so the common destination ϕ_t^* weakly rises. This is a statement about the destination across counts, not about the switchers' starting point: the triggered revisers relocate *from* their held intensities *to* $\lambda(\phi_t^*)$, so the relocation adds to the reweight (F10) exactly when $\lambda(\phi_t^*)$ exceeds the switchers' mean held intensity under their post-reweight shares—a conditional mean over the *triggered set* after the reweight, not the full-cross-section covariance of (F15)—the complacent-switcher condition, a condition on that post-reweight triggered-switcher mean; a complacent newborn law Π_ϕ makes it plausible in the stationary regime but does not imply it, and we do not derive it from primitives. The reweight (F15) raises $\bar{\lambda}^S$ unconditionally, so the fast-in does not rely on this amplifier.

(iii) Fear factor. By the bridge relation (**??**), $d\log U_t = d\log\kappa_t$, and $U_t = \lambda^\star \mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]$ (**??**). In calm the held models are fixed and $g_t(x) = \langle r(\cdot,x)^{1/\alpha}, S_t\rangle$, so (F5) with $h = r(\cdot,x)^{1/\alpha}$ gives

$$\frac{dg_t(x)}{dt} = -\tfrac{1}{\alpha}\operatorname{Cov}_{S_t}\big(\lambda, r(\cdot,x)^{1/\alpha}\big). \tag{F17}$$

This is the selection component—the derivative at held models fixed; at active deadline dates the calm switch flow adds the relocation term the theorem's c_{sw} collects, and (i)'s displayed drift is the atomless-flow component.

Differentiate $U_t = \lambda^\star \int_{\mathcal{J}} e^{\alpha x} g_t(x)^\alpha f(x)\,dx$ in t. The t-derivative of the integrand is $\alpha e^{\alpha x} g_t(x)^{\alpha-1}\,dg_t(x)/dt$, bounded uniformly on the compact $\mathcal{J}$: $g_t(x) = \int_0^1 s_t^i r_i(x)^{1/\alpha} di \in [\underline{r}^{1/\alpha}, \overline{r}^{1/\alpha}]$ with $0 < \underline{r} = \min_{\mathcal{J}\times\Phi} r \leq \overline{r} = \max_{\mathcal{J}\times\Phi} r < \infty$ (continuity and positivity of the densities on the compact $\mathcal{J}\times\Phi$, Assumption **??**), and dg_t/dt is bounded by (F17); so dominated convergence justifies differentiation under the integral, and

$$\begin{aligned}
\frac{dU_t}{dt} &= \lambda^\star \int_{\mathcal{J}} e^{\alpha x}\, \alpha\, g_t(x)^{\alpha-1} \frac{dg_t(x)}{dt}\, f(x)\, dx \\
&= \lambda^\star \int_{\mathcal{J}} e^{\alpha x}\, \alpha\, g_t(x)^{\alpha-1} \Big(-\tfrac{1}{\alpha} \operatorname{Cov}_{S_t}(\lambda, r(\cdot,x)^{1/\alpha}) \Big) f(x)\, dx && \text{(F18)} \\
&= -\lambda^\star \int_{\mathcal{J}} e^{\alpha x} g_t(x)^{\alpha-1} \operatorname{Cov}_{S_t}(\lambda, r(\cdot,x)^{1/\alpha})\, f(x)\, dx, && \text{(F19)}
\end{aligned}$$

(F18) substituting (F17) and (F19) cancelling $\alpha \cdot \frac{1}{\alpha}$. Divide by $U_t = \lambda^\star \mathbb{E}_f[e^{\alpha J} g_t^\alpha]$ (chain rule $d \log U_t/dt = (dU_t/dt)/U_t$):

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

which is (??). In the intensity channel $\operatorname{Cov}_{S_t}(\lambda, r(\cdot,x)^{1/\alpha}) \geq 0$ for every x (F15), and $e^{\alpha J}, g_t^{\alpha-1}, g_t^\alpha > 0$, so both expectations are ≥ 0 and the ratio is ≥ 0; hence $d \log U_t/dt \leq 0$, strict iff $\operatorname{Var}_{S_t}(\lambda) > 0$. At a disaster of realised size x_0, the reweight $S_t = r(\cdot, x_0)^{1/\alpha} g_{t-}(x_0)^{-1} S_{t-}$ moves g to

$$\begin{aligned}
g_t(x) &= \mathbb{E}_{S_t}[r(\cdot,x)^{1/\alpha}] \\
&= \frac{\mathbb{E}_{S_{t-}}[r(\cdot,x)^{1/\alpha}\, r(\cdot,x_0)^{1/\alpha}]}{\mathbb{E}_{S_{t-}}[r(\cdot,x_0)^{1/\alpha}]}, && \text{(F20)}
\end{aligned}$$

(F20) substituting the reweight into $\mathbb{E}_{S_t}$. In the intensity channel this equals $\mathbb{E}_{S_t}[(\lambda/\lambda^\star)^{1/\alpha}]$, which rose: apply the association argument of (F15) with $Z := (\lambda/\lambda^\star)^{1/\alpha}$ in place of λ—both factors are then non-decreasing in λ, so the product inside (F15) is pointwise ≥ 0 and $\operatorname{Cov}_{S_{t-}}(Z,Z) = \operatorname{Var}_{S_{t-}}(Z) \geq 0$—so the post-reweight mean satisfies $\mathbb{E}_{S_{t-}}[Z^2]/\mathbb{E}_{S_{t-}}[Z] \geq \mathbb{E}_{S_{t-}}[Z]$, strict iff $\operatorname{Var}_{S_{t-}}(Z) > 0$. The reweighted aggregator therefore rises for every x, and since the integrand $e^{\alpha x} g^\alpha$ of (??) is increasing in g, the reweight-only fear factor $U(\widetilde{S}_t^{x_0})$ exceeds U_{t-}—the unconditional statement. The final post-disaster U_t adds the triggered relocation, which in the intensity channel changes the aggregator by $\widehat{s}_t^{\text{sw}}\big(Z(\phi_t^*) - \widehat{\mathbb{E}}_t^{\text{sw}}[Z]\big)$, with $\widehat{s}_t^{\text{sw}}$ and $\widehat{\mathbb{E}}_t^{\text{sw}}$ the post-reweight triggered-switcher mass and mean; the exact sign condition for the relocation is therefore $Z(\phi_t^*) \geq \widehat{\mathbb{E}}_t^{\mathrm{sw}}[Z]$, the likelihood-ratio form of the complacent-switcher condition. For $\alpha \geq 1$ the condition of (ii) implies it:

$$Z(\phi_t^*) = \Big(\tfrac{\lambda(\phi_t^*)}{\lambda^\star}\Big)^{1/\alpha} \;\geq\; \Big(\tfrac{\widehat{\mathbb{E}}_t^{\mathrm{sw}}[\lambda]}{\lambda^\star}\Big)^{1/\alpha} \;\geq\; \widehat{\mathbb{E}}_t^{\mathrm{sw}}\big[(\lambda/\lambda^\star)^{1/\alpha}\big],$$

the first inequality by the condition and monotonicity of $z \mapsto z^{1/\alpha}$, the second by Jensen, $z \mapsto z^{1/\alpha}$ concave for $\alpha \geq 1$. Hence the relocation weakly raises the aggregator and with it U_t, so $U_t \geq U(\widetilde{S}_t^{x_0}) > U_{t-}$: the jump holds under the complacent-switcher condition for $\alpha \geq 1$ (under its likelihood-ratio form for all α), sufficiency not equivalence, and the reweight component holds unconditionally. The disaster premium's U_t-loading component (Proposition **??**) inherits the direction; the full premium adds the valuation jump ψ_t, hence v (Proposition **??**), which this argument leaves unsigned.

(iv) Staggered recovery. The drifts of (i),(iii) hold on any disaster- and switch-free interval. By Proposition D.1 a calm spell ends at the deadline $\Delta^\star(\phi, \varepsilon)$ (D4), where the reviser re-selects the mode prevailing at her deadline; by Lemma D.1 the mode is non-increasing through a disaster-free spell, so under monotone re-selection her held intensity weakly falls—a downward step in $\bar{\lambda}^S$ and, in the intensity channel, in U_t. By Proposition D.2 the deadlines are weakly ordered by ε, so the steps are staggered, occurring at distinct times where the deadline map is non-constant; and by Proposition D.1 $\Delta^\star$ is finite and increases as $\varepsilon \to 0$ or the held intensity $\lambda(\phi) \to 0$ (the dependence on $\lambda^\star$ is indirect, through which held models survive). The jump of (ii) is instantaneous (fast in); the drift of (i),(iii) and the staggered steps act over the deadline horizon (slow out). □

Proof of Theorem **??**. *(i)* By Proposition **??** and Proposition **??**, $F_t(0) = \mathbb{E}[U_t \mid \mathcal{F}_t] = U_t = U^{\mathrm{RA}}\kappa_t$ is a functional of S_t.

(ii) The slope is the generator applied to U. The fear factor $U_t = \lambda^\star \mathbb{E}_f[e^{\alpha J} g_t^\alpha]$ is a bounded functional of the cross-section ($g_t \in [\underline{r}^{1/\alpha}, \overline{r}^{1/\alpha}]$ on the compact $\mathcal{J} \times \Phi$), hence a bounded

functional $U \circ \pi_\phi$ of the record-resolved state Ψ_t, which evolves as the piecewise-deterministic Markov process of Proposition **??**—deterministic calm flow, disaster jumps at rate $\lambda^\star$. By the extended-generator (Dynkin) identity for piecewise-deterministic Markov processes (Davis, 1984; Jacobsen, 2006), with $U \circ \pi_\phi$ in the domain (bounded, absolutely continuous along the calm flow),

$$\partial_s^+ F_t(0) = \frac{d}{ds}\mathbb{E}[U_{t+s} \mid \mathcal{F}_t]\Big|_{s=0+} = (\mathcal{A}_\Psi U)(\Psi_t) = \frac{dU_t}{dt}\Big|_{\text{calm}} + \lambda^\star \int_{\mathcal{J}} \big(U(\Psi_t^x) - U(\Psi_{t-})\big) f(x)\, dx, \tag{F21}$$

the generator splitting into the calm drift and the disaster compensator. The identity is evaluated at a predictable t (a.s. every fixed date: disasters are totally inaccessible), where $\Psi_{t-} = \Psi_t$; at a jump date the right slope is generated from the post-jump state.

Calm drift. Differentiate $U_t = \lambda^\star \int_{\mathcal{J}} e^{\alpha x} g_t(x)^\alpha f(x)\, dx$; the integrand's t-derivative $\alpha e^{\alpha x} g_t(x)^{\alpha-1}\, dg_t(x)/$ is bounded uniformly on the compact $\mathcal{J}$ (the aggregator $g_t(x)$ lies between the menu's extreme size ratios, uniformly in t), so dominated convergence justifies differentiating under the integral,

$$\frac{dU_t}{dt}\Big|_{\text{calm}} = \alpha\lambda^\star \int_{\mathcal{J}} e^{\alpha x} g_t(x)^{\alpha-1}\, \frac{dg_t(x)}{dt}\Big|_{\text{calm}} f(x)\, dx. \tag{F22}$$

The calm motion of $g_t(x) = \langle r(\cdot, x)^{1/\alpha}, S_t\rangle$ has two parts—the selection drift (**??**) and the calm switch flow (**??**) of Proposition **??**, at $h = r(\cdot, x)^{1/\alpha}$:

$$\frac{dg_t(x)}{dt}\Big|_{\text{sel}} = -\tfrac{1}{\alpha}\,\mathrm{Cov}_{S_t}\big(\lambda, r(\cdot, x)^{1/\alpha}\big), \qquad \frac{dg_t(x)}{dt}\Big|_{\text{sw}} = \int_\Phi \big(r(\phi_t^*, x)^{1/\alpha} - r(\phi, x)^{1/\alpha}\big)\, \Phi_t^{\text{sw}}(d\phi), \tag{F23}$$

the first by (F17), the second because a switcher relocating $\phi \to \phi_t^*$ changes her contribution to g from $r(\phi, x)^{1/\alpha}$ to $r(\phi_t^*, x)^{1/\alpha}$, at the wealth-weighted deadline rate Φ_t^{sw}. Substituting each part of (F23) into (F22),

$$\begin{aligned}
\frac{dU_t}{dt}\Big|_{\text{sel}} &= \alpha\lambda^\star \int_{\mathcal{J}} e^{\alpha x} g_t^{\alpha-1}\Big(-\tfrac{1}{\alpha}\operatorname{Cov}_{S_t}(\lambda, r(\cdot,x)^{1/\alpha})\Big) f\,dx = -\lambda^\star \mathbb{E}_f\big[e^{\alpha J} g_t^{\alpha-1}\operatorname{Cov}_{S_t}(\lambda, r(\cdot,J)^{1/\alpha})\big] \\
&= -U_t\, b_{\text{sel}}(S_t), \qquad\qquad (\text{F24}) \\
\frac{dU_t}{dt}\Big|_{\text{sw}} &= \alpha\lambda^\star \int_{\mathcal{J}} e^{\alpha x} g_t^{\alpha-1}\Big(\int_\Phi (r(\phi_t^*,x)^{1/\alpha} - r(\phi,x)^{1/\alpha})\Phi_t^{\text{sw}}(d\phi)\Big) f\,dx \\
&= U_t\, c_{\text{sw}}(\Phi_t^{\text{sw}}),
\end{aligned}$$

with $b_{\text{sel}}(S_t) = \mathbb{E}_f[e^{\alpha J} g_t^{\alpha-1}\operatorname{Cov}_{S_t}(\lambda, r(\cdot,J)^{1/\alpha})]/\mathbb{E}_f[e^{\alpha J} g_t^{\alpha}] \geq 0$ and c_{sw} the ratio of the statement. *The factor α cancels in the selection term* ($\alpha\cdot\frac{1}{\alpha} = 1$, the share drift (**??**) carrying the $\frac{1}{\alpha}$), recovering the calm log-drift (**??**); *it does not cancel in the switch term*, where the deadline flux (**??**) carries no $\frac{1}{\alpha}$—so c_{sw} retains the α, the asymmetry between the two calm channels. Under monotone re-selection the switch term is signed, $c_{\text{sw}} \leq 0$: on the *disaster-free* switch set scheduled relief gives $r(\phi_t^*,\cdot) \leq r(\phi,\cdot)$ (Proposition D.1); switchers whose records carry non-rejecting disasters need not satisfy it—exactly the condition's content, and the theorem leaves c_{sw} unsigned.

Disaster compensator. The post-disaster state Ψ_t^x applies the reweight (**??**) and then relocates the triggered investors (E21), so the jump in U splits as

$$\lambda^\star \int_{\mathcal{J}} \big(U(\Psi_t^x) - U(\Psi_{t-})\big) f\,dx = \underbrace{\lambda^\star\,\Theta_{\text{rw}}(S_t)}_{\text{reweight}} + \underbrace{\lambda^\star\,\Theta_{\text{sw}}(\Psi_t)}_{\text{relocation}}, \qquad (\text{F25})$$

$\Theta_{\text{rw}}(S_t)$ the change in $U = \lambda^\star\mathbb{E}_f[e^{\alpha J} g^\alpha]$ produced by the size-x reweight of g_{t-} (**??**)—a functional of S_{t-}, since the reweight $r(\cdot,x)^{1/\alpha}/g_{t-}(x)$ and the reweighted aggregator are functionals of S_{t-}; and $\Theta_{\text{sw}}(\Psi_t)$ the additional change from relocating to ϕ_t^* those investors whose updated record leaves the typical set, at the record-dependent rate $\lambda^\star\pi(\phi,\tau,\mathcal{R};\varepsilon)$ (E21)—a functional of Ψ_t, not of S_t, signed under the complacent-switcher condition (for $\alpha \geq 1$, or its likelihood-ratio form)—$\lambda(\phi_t^*)$ above the switchers' post-reweight mean—which the destination's count-monotonicity

(Proposition **??**(ii)) supports but does not imply; the theorem leaves Θ_{sw} unsigned.

Collecting. Substituting (F24)–(F25) into (F21),

$$\partial_s^+ F_t(0) = \underbrace{-U_t\, b_{\text{sel}}(S_t) + \lambda^\star \Theta_{\text{rw}}(S_t)}_{\mathrm{b}_S(S_t)} + U_t\, c_{\text{sw}}(\Phi_t^{\text{sw}}) \;+\; \lambda^\star \Theta_{\text{sw}}(\Psi_t), \tag{F26}$$

the first group a functional of S_t through g_t (Proposition **??**); the calm flow Φ_t^{sw} and the trigger propensity π are functionals of Ψ_t and not of S_t, since the deadline-hitting flow and the triggered set depend on the joint law of held models, ages, and records, which the held-model marginal does not fix. This is the decomposition of (ii).

(iii) For Ψ_t, Ψ_t' with $\pi_\phi \Psi_t = \pi_\phi \Psi_t' = S_t$: the levels coincide by (i), $F_t(0) = F_t'(0) = U_t$, and by (F26) the $\mathrm{b}_S(S_t)$ part of the slope coincides, leaving the two staggered-revision channels $U_t c_{\text{sw}} + \lambda^\star \Theta_{\text{sw}}$. Conditioning on no disaster over the horizon suppresses Θ_{sw} and leaves the calm channel $U_t c_{\text{sw}}$; under monotone re-selection $c_{\text{sw}} \le 0$, so the state with the larger imminent deadline flow Φ_t^{sw} carries the weakly steeper decline, $\partial_s^+ F_t^{0\prime}(0) \le \partial_s^+ F_t^{0}(0)$; under that condition and the complacent-switcher condition, as qualified above, the two channels are oppositely signed. 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 **??**).

(iv) Under heterogeneous Bayesian learning each posterior is a sufficient statistic for its own future, so the wealth-weighted cross-section is Markov in itself; there is no first-passage deadline and no rejection trigger, hence neither Φ_t^{sw} nor Θ_{sw}; the Bayesian generator carries its own posterior-drift terms in place of the lazy calm structure, so the slope is *some* S_t-measurable functional—the benchmark's own, which Section IA.9 computes—not the lazy formula's $\mathrm{b}_S(S_t)$. The forward curve is then pinned by the current cross-section at every horizon (Section IA.9); an affine latent-state model and a representative filtering economy give the same conclusion through their own Markov states. □

Proof of Proposition G.1. (i) Under the clock, every compensator that moves the wealth measure reads only the coordinates of S_t^{ck} and the public posterior: revisions arrive at rate $b(\phi,\tau)$ with the common destination $\mathfrak{m}(\rho_t)$, resetting $(\tau,\mathcal{R})$; the selection drift and the disaster reweight read the held model; demographic turnover reads cohort age; and ρ_t evolves on the public record alone. The pair $(S_t^{\mathrm{ck}},\rho_t)$ is therefore closed—a piecewise-deterministic Markov process by the argument of Proposition **??**, the since-adoption record entering no rate and no destination—and every pricing functional of Sections **??**–**??** reads Ψ_t through the held-model marginal and the revision flows, each $(S_t^{\mathrm{ck}},\rho_t)$-measurable here. The projection argument of Theorem **??**(iii) then applies with $(S_t^{\mathrm{ck}},\rho_t)$ in place of S_t: every forward object equals its own projection on the state, and the residual loading of any record functional is zero. *(ii)* Revisions carry no jump component, so the trigger set of (E21) is empty and the relocation term of Proposition **??**(iii) vanishes identically; the per-disaster mark-averaged relocation increment is zero, $\Theta_{\mathrm{sw}}\equiv 0$, and the disaster jump of U_t is the reweight component of Proposition **??**(ii)–(iii) alone, for every mark. *(iii)* Fix (ϕ,τ,ε) in the intensity channel, $\phi=(\lambda,\cdot)$, and two since-adoption counts $n\neq n'$. Between disasters the record is fixed and the deviance (**??**) drifts deterministically along the regime age Δ, $D_{\phi,\Delta}=n\log\frac{n}{\lambda\Delta}-n+\lambda\Delta$ (the size term vanishes in this channel). The difference across counts satisfies

$$\frac{\partial}{\partial\Delta}\big[D_{\phi,\Delta}(n)-D_{\phi,\Delta}(n')\big]\;=\;-\frac{n-n'}{\Delta},$$

of one strict sign, so the two deviance paths cross at most once. Each path meets the acceptance boundary at its scheduled calm deadline $\Delta^\star$ (Propositions D.1 and E.2); since the paths cross at most once, their first passages over the common boundary coincide only if the passage occurs exactly at that crossing age—at most one coincidence age. Off it the calm deadlines differ, the deadline flows Φ_t^{sw} differ on a set of positive mass, and by Theorem **??**(iii) the calm-channel slope difference is $\int_\Phi\big(h(\phi_t^*)-h(\phi)\big)\,d(\Phi-\Phi')_t^{\mathrm{sw}}$, nonzero whenever the differential mass carries a nonzero valuation gap. $\square$

Proof of Corollary **??**. Couple the two economies as in Theorem **??**(iii), on a common driving path with common public posterior. Write $D := e^{-\int_t^{t+s_1} r_u\,du}$; to first order in $s_2 - s_1$ the calm-conditional tranche is the window's expected discounted count at the prevailing risk-neutral intensity, $P_t^0 = \mathbb{E}_t^{\mathbb{Q},0}[D\,U_{t+s_1}]\,(s_2 - s_1) + o(s_2 - s_1)$ (Remark **??**; discounting inside the window is second order). Hence

$$P_t^0 - P_t'^0 = (s_2 - s_1)\Big\{\mathbb{E}_t^{\mathbb{Q},0}\big[D\,(U_{t+s_1} - U'_{t+s_1})\big] + \mathbb{E}_t^{\mathbb{Q},0}\big[(D - D')\,U'_{t+s_1}\big]\Big\} + o(s_2 - s_1).$$

The single $\mathbb{E}_t^{\mathbb{Q},0}$ is licensed by the coupling: conditional on no disaster over $[t, t + s_1]$, D, U_{t+s_1}, and their primed counterparts are deterministic functionals of the two initial record-resolved states on the common path, so the primed values are constants under either economy's measure and the subtraction is measure-free. The fear-path term: conditional on no disaster the wedge $U - U'$ is generated by the calm channel—Theorem **??**(iii) under the S_{t-}-measurable tilt of Remark **??**, which moves only the S-loadings—so it is measurable in Ψ_t and vanishes under a Markov-in-S_t benchmark, whose calm U-path is an autonomous S-functional shared by the pair. The discount-path term: $D - D' = -D\int_t^{t+s_1}(r_u - r'_u)\,du\,(1 + o(1))$, and $r_u - r'_u = r(S_u) - r(S'_u)$ with r a bounded-Lipschitz functional of S in total variation ((**??**) is affine in $\bar{\lambda}^S$ and U, both bounded Lipschitz on the compact menu), while along the calm coupling $\|S_u - S'_u\|_{\mathrm{TV}}$ is generated by the differing switch flows; so this term is likewise measurable in Ψ_t, and under the Markov benchmark the shared S-path forces $r_u \equiv r'_u$ and it vanishes. For the unconditional tranche, decompose on the first disaster in $(t, t+s_1]$: its compensator adds the disaster-triggered relocation channel through the post-jump state, entering the difference at first order through Θ_{sw}. □

Appendix G. Record-independent clocks

Proposition G.1 (Record-independent clocks): *Modify the economy of Section* **??** *in one respect: revisions arrive with predictable intensity* $b(\phi_i, \tau_i)$ *for a bounded measurable* $b \geq 0$, *re-selection is* *at the public mode $\mathfrak{m}(\rho_t)$, the since-adoption record and regime age reset at adoption, and no revision is disaster-triggered. Write S_t^{ck} for the wealth measure over (held model ϕ, regime age τ, cohort age a). Then:* (i) *$(S_t^{\mathrm{ck}}, \rho_t)$ is a Markov pricing state: every forward object is a functional of it, and conditional on it every since-adoption record functional has zero forward loading—$\mathcal{H}_t \equiv 0$ in the sense of Theorem* **??***(iii) with $(S_t^{\mathrm{ck}}, \rho_t)$ in place of S_t.* (ii) *The disaster-triggered channel is absent: $\Theta_{\mathrm{sw}} \equiv 0$, and at a disaster the fear factor moves by the reweight alone, for every mark.* (iii) *Under the testing rule, in the intensity channel, $(S_t^{\mathrm{ck}}, \rho_t)$ is not sufficient: at fixed $(\phi, \tau, \varepsilon)$, records with different since-adoption counts reach the acceptance boundary at different calm deadlines (coinciding at most at one age), so two testing economies sharing $(S_t^{\mathrm{ck}}, \rho_t)$ and the conditional tolerance composition, with different count compositions, have different deadline flows, and their forward slopes differ whenever the differential mass carries a nonzero valuation gap $h(\phi_t^*) - h(\phi)$.*

Duration dependence alone is conceded—ages are a priced state under any staggered clock. What the clock cannot carry is the record: the deadline that reads how unlikely the surviving evidence is under the held model, and the trigger that fires on it.

Appendix H. What history prices, and what it does not

The state that prices disasters is not what the market believes but a pair: what it believes, and how long it has believed it without rechecking. The first coordinate is the cross-section S_t. The second is the cross-section of *fragility*: how each holding's since-adoption record stands against the model it holds (Teguia, 2026), hence how near it is to revision. Two markets can share their beliefs and their near-dated tail and still differ here—one has just re-anchored, records fresh, deadlines far; the other has held the same fear through a long calm, records aged, deadlines near. They price the forward tail apart. Belief fragility is priced, and current belief does not recover it (Proposition **??**).

The disaster law $(\lambda^\star, F_J)$ is fixed—known to the modeler, disputed by investors—so all the record does here is govern revision. It prices the tail—it moves the forward fear factor and the valuation $v(\Psi_t)$—yet forecasts nothing about how often or how hard disasters strike. That is the content of the non-Markov bridge: a market-memory variable orthogonal to every fundamental forecast, yet priced. No model whose priced state is a belief about fundamentals—filtering, Bayesian learning, or an affine latent-state filter among them—can do this: there, history moves prices only by moving that belief, so the two move together; here they part.

The mechanism also meets data where a latent disaster process cannot: there the intensity's persistence is identified only jointly with the belief that reads it, while here the coordinate that moves the forward tail is the *survival record*. Time since the last disaster is read off the path alone; accumulated calm surprise and distance to a rejection boundary are model-based reconstructions—functions of the path jointly with the held model, adoption age, and tolerance under the maintained rule—proxies an econometrician builds, not raw readings. That partial visibility is what makes the over-identifying restriction of Theorem **??** a test.

Three things follow. Through a continuing calm the same near-dated tail relaxes faster going forward where the fear has aged, nearer the deadline that revises it down (Theorem **??**). A disaster need not reprice two markets of identical beliefs alike: the reprice runs through the holdings it tips into revision at once, read from the records and not from S_t (Proposition **??**), so whether more revision deepens or cushions the fall is a cross-sectional condition, not a sign the model fixes. And the effect reaches the level, not the slope alone: $v(\Psi_t)$ and the premium read the whole record.

A renewal process is Markov once one tracks age, so one might ask whether "(beliefs, age)" is a Markov state in disguise. It is not. The record resets at every rejection, and both the revision hazard and its destination depend on the full record and the public posterior, not a scalar clock. These are claims of different strength, and the paper rests on the weaker one. That the forward surface reads the record Ψ_t and not the cross-section S_t alone holds in *either* channel: it follows from the non-autonomy of the law of motion, whose switch flow and trigger read the within-cohort record distribution rather than the held-model marginal (Proposition **??**). The stronger statement—that no *finite* statistic of the state is sufficient—is a size-channel result (Proposition IA.1) the intensity-channel calibration never invokes: there the instantaneous tail is the single scalar G_t (Corollary E.1), and only the *forward* objects carry the record. We claim the dimensionality only where the mathematics forces it and the economics where it bites—a few measures of fragility: time since the last disaster, accumulated calm surprise, distance to a deadline.

One boundary. Fragility is priced; it does not feed back on the rule that makes it. Whether a long quiet breeds a more violent revision—fragility itself building in good times, not merely its price—is a separate question, left to other work. This is the pricing face of one fact about the rule: because rejection never stops—the calm deadline is finite even at the truth, so a correct-on-average holding is revised forever (Lemma IA.2)—fragility keeps accumulating and never collapses into the belief—priced, though, only through surviving dispersion or turnover-refreshed disagreement: a self-re-adoption resets the clock and moves nothing, so a degenerate economy at the pseudo-true model re-tests forever without a forward wedge.

Appendix I. Quantitative development for Section ??

Two consequences. First, a cross-section whose wealth-weighted mean belief equals the truth, $\bar{\lambda}_t^S = \lambda^\star$, still prices the tail *below* the representative-agent benchmark—dispersion alone is complacency: at an equal-mass two-point spread $\{\lambda^\star - s, \lambda^\star + s\}$ the multiplier is 0.96 at $s = 1\%/\text{yr}$ and 0.82 at $s = 2\%/\text{yr}$. Second, fear ($U_t > U^{\text{RA}}$) requires $\bar{\lambda}_t^S$ to exceed $\lambda^\star$ by enough to overcome this Jensen discount, which a sufficiently severe disaster delivers by reweighting wealth toward and concentrating the cross-section on the fearful models (Proposition **??**). Equation (**??**) is the analytic core of the chronic-complacency/disaster-spike asymmetry: the kernel prices the power mean of beliefs, which sits below the wealth-weighted arithmetic mean.

Sensitivity. Three dials govern the fear cycle, each through a displayed formula. The newborn dispersion sets the crisis amplitude through the jump covariance (F10): at a given size the jump is $\mathrm{Cov}_{S_t}\big(\lambda, (\lambda/\lambda^\star)^{1/\alpha}\big)$ scaled by the mean transformed intensity, so dispersion moves it one-for-one holding that mean fixed. The turnover rate sets the demographic component of calm relaxation one-for-one: the pull in (J2) has half-life $\ln 2/\Omega$—thirty-five years at $\Omega = 0.02$—so halving Ω doubles it. And the menu coarseness sets the revision step: a dense menu replaces the two-point relocations by many small grid steps, so the crisis spike is milder and the post-spike plateau smooths into near-continuous decay, all other blocks unchanged.

One model-free finding runs the other way. Gormsen and Jensen (2025) estimate the *physical* return distribution from option prices and find it more left-skewed and fat-tailed in good times—a *procyclical* conditional tail risk they read as hard to square with disaster models in which the physical law is elevated in bad times. This economy is not such a model: the physical law $(\lambda^\star, F_J)$ is fixed, and what is countercyclical is the *priced* multiplier κ_t, not the world's disaster risk. The link is tighter than coexistence—the complacency that holds $\kappa_t < 1$ through calm *is* an under-pricing of the tail in precisely the spells where, on their evidence, the conditional distribution is fattest—so the gap between the thin priced tail and the fat physical tail is the belief wedge (time variation in the physical *return* tail, as opposed to the consumption tail, runs here through the state-dependent valuation jump $e^{-x}\psi_t(x)$ of Proposition **??**—the channel a procyclical return-tail estimate would load on—while the consumption disaster law never moves; its sign we leave to the valuation operator), bracketed from opposite ends: its priced side the modest option-implied disasters of Backus et al. (2011) the calibration matches, its physical side the procyclical return tail Gormsen and Jensen measure. The mechanism turns this into a discriminating prediction—the physical-minus-priced gap should widen through calm and collapse at disasters, developed in the supplementary part below—so Gormsen and Jensen's fact is a second empirical handle on the price of disaster fear, not an embarrassment to it.

The forward-slope ensemble (Section **??***).* One hundred fifty independent histories are drawn at the calibration of Table ??: disasters Poisson at $\lambda^\star$, the full rule dynamics from the common-prior start with demographic turnover, evaluated at $t^* = 60$ (a finite-history ensemble with the public-mode path alive, not draws from the locked-destination invariant law; that regime is reported below); each terminal state is then transported calm-conditionally for twelve years, once as the path left it and once with every since-adoption record reset (the reset-record counterfactual; the constructed witness of the figure retains its fresh/aged naming), at the converged configuration of the Figure ?? protocol. The pipeline reproduces the constructed witness to 16.17% against its converged 15.8–16.2 band. Regressions are linear in S^{hi} within each band; the interquartile time-since effect after disasters is +5.3 percentage points of $F^0_{t^*}(12)/U_{t^*}$ at fixed S^{hi}. Within-band regressions of the twelve-year slope on the spot state and the record proxies give: after disasters ($S^{hi} \geq 0.40$, $n = 75$), R^2 of 0.641 on S^{hi} alone, 0.922 adding time since the last disaster (+0.0028 per year), unchanged adding deadline-proximity mass; in deep calm ($S^{hi} < 0.10$, $n = 73$), 0.869 on S^{hi} alone and 0.993 with the proxies, the proximity coefficient −0.654, a −0.10 percentage-point effect across the calm interquartile range of proximity mass. Among near-matched pairs in the comparison window ($|\Delta S^{hi}| < 0.02$, 103 pairs) the twelve-year slope gaps have median 2.1, ninetieth percentile 5.0, and maximum 5.8 percentage points. Sampling is at the sixty-year horizon from the common-prior start; the transient middle band is visited in passing rather than at the fixed sampling date, which is why the constructed witness, not the ensemble, marks the channel's ceiling. The exercise is a finite-history event study from the calibrated start: the stationary economy of Proposition K.1 guarantees a well-posed long run but is not the sampling frame—under the misspecified menu its locked destination is the fearful point, while the finite-history public mode governs the scheduled calm relief here. These calm-conditional, physical-measure paths are mechanism illustrations; the traded object remains the discounted risk-neutral calendar spread of Corollary ?? (Remark ??). Ensemble quantities are Monte Carlo statistics at $M = 150$; the transport itself is a deterministic numerical scheme—step $dt = 0.01$, the two-block tolerance grid, the calm-branch deadline lookup, newborn batches at half-year spacing, low-weight atom compression—with witness convergence reported across configurations; seed and M stability is reported, with the exact-criterion and dominance-slack checks, in the supplementary part below (Section IA.10).

The short rate along the cycle. The safe rate is the S_t-functional (**??**); at the calibration it spans 0.4–9.6% between the belief vertices, and across the ensemble terminal states it has mean 7.4%, standard deviation 237 basis points, and interdecile range 3.7–9.6%—the familiar cost of moving a large fear factor under CRRA, disclosed rather than smoothed. Corollary **??** prices the induced rate path explicitly: the rate term is part of the same survival-record residual and vanishes across the benchmark class, so the calendar-spread test is not contaminated by it; level and volatility realism of the safe rate are outside the claim set (Section **??**).

The locked-destination regime. The stationarity proposition's invariant regime locks the re-selection destination at the menu's Kullback–Leibler projection—here the fearful point λ_{hi}. We simulate it directly: sixty independent histories of one hundred twenty years (master seed 11; the chunked driver ships with the code), destination locked throughout, terminal states read as approximate invariant draws. The cross-section shifts fearful and disperses—S^{hi} deciles and median 0.24/0.50/0.87—because every deadline and every trigger now re-anchors at λ_{hi} and only the newborn inflow replenishes complacency. The record's incremental content survives: within the low band ($S^{hi} < 0.40$, $n = 22$) the twelve-year-slope R^2 rises from 0.948 on S^{hi} alone to 0.987 adding time since the last disaster, and within the high band ($n = 38$) from 0.898 to 0.968. The reset wedge, by contrast, changes sign structure: it is negative essentially everywhere (quantiles $-8.0/-3.8/-0.2$ overall; $-8.8/-7.8/-5.1$ in the low band), the mirror of the baseline mechanism—with the destination locked fearful, the operative calm-deadline flow is the complacent holders' delayed re-anchoring upward, so resetting records postpones it and lowers the reset curve, where the baseline's long-calm destination made the same channel push downward and the calm wedge positive. What is invariant across the two regimes is that records organize the calm-conditional forward slope beyond the spot state; the wedge's sign is a destination statement, exactly the deadline channel of Remark **??**.

Appendix J. Demographic turnover and the overlapping-generations economy

Assumption J.1 (Demographic turnover: perpetual youth): *We extend Assumption* **??** *with overlapping generations in the manner of Yaari (1965) and Blanchard (1985), as used for heterogeneous agents by Gârleanu and Panageas (2015):*

(i) *each living investor faces a constant death hazard $\Omega > 0$—an independent Poisson clock, the same for all and independent of the aggregate path (aggregating the continuum of independent clocks into the deterministic flows of (ii)–(iii) is the exact law of large numbers (Judd, 1985; Sun, 2006); equivalently, and as (iii) states, the share-replacement law is the primitive);*

(ii) *newborns arrive at mass Ω per unit time (population mass constant), each drawn independently with type $(\varepsilon, \phi_0) \sim \Pi$, a fixed law on $(0,1) \times \Phi$ with ϕ_0-marginal Π_ϕ—the surprise threshold $\varepsilon \in (0,1)$ as in Assumption* **??**, *so the typical-set test of Definition* **??** *is well defined for every newborn type;*

(iii) (allocation of cohorts) *the symmetric allocation of Assumption* **??***(iv) applies to every cohort, so the* effective *cohort multiplier is common (no preference- or endowment-dispersion is priced) and each newborn enters with no since-birth record ($M = 1$ at birth). Demographic turnover is then a primitive share-replacement rule: the newborn cohort carries aggregate consumption share $\Omega\, dt$, distributed over held models by the type law Π_ϕ. If formally decentralized (Appendix J), each entering cohort receives the* born-at-par *Negishi weight $\widetilde{\psi}_i^{-1/\alpha} = \Xi_{b_i}$—the entrant's weight equals the living population's current average—which is exactly what makes its consumption share $\Omega\, dt$ and the replacement law (J1) exact. The symmetry of Assumption* **??***(iv) acts* within *each cohort—equal weights across contemporaneous entrants—while the cohort-level multiplier Ξ_{b_i} is common within a birth date and varies across birth dates by construction; no fixed economy-wide Pareto weight*

is asserted for entrants. Throughout the turnover economy, M^i denotes the belief density restarted at birth, $M^i_{b_i} = 1$—the conditional post-birth likelihood on the investor's own record—with all $\mathcal{F}_{b_i}$-measurable factors absorbed into $\widetilde{\psi}_i$; the par-entry cancellation and the first-order conditions below are exact under this normalization. We take the share-replacement law as the primitive*; no result below requires its decentralization, which Appendix J supplies only to exhibit the supporting weights. The budget behind them uses two separate instruments: fair annuities insure mortality among the living—estates fund the survivor credit Ω (Yaari, 1965)—while entry is financed by diluting incumbents' tree claims at rate Ω, each entrant receiving new shares worth her birth endowment—type-dependent under heterogeneous beliefs, the cohort total matching the diluted flow; the two Ω-flows offset in every survivor's return, so the Euler equation and the short rate are untouched. Explicitly, a survivor holding the (diluted) tree earns $\frac{C_t}{P_t}\,dt + \frac{dP_t}{P_{t-}} - \Omega\,dt + \Omega\,dt$—dividend yield and capital gain, less the dilution drag, plus the fair annuity credit—so Ω cancels from her return; on the goods side the dying cohorts' consumption share is $\Omega\,dt$ (death is wealth-proportional across the living), exactly the entrants' assigned share, so the resource constraint is undisturbed. Death is actuarially fairly annuitized (Yaari, 1965), so Ω does not enter any survivor's effective discount; the turnover acts on the share simplex and is aggregate-resource-neutral by construction, $\langle 1, S_t\rangle \equiv 1$.*

(iv) (newborn information). *A newborn inherits the common second-order posterior ρ_t prevailing at her birth date—the public object of Lemma D.1, a deterministic functional of the aggregate record, shared by every living investor—and enters holding her drawn model ϕ_0 with a reset test state: regime age $\tau = 0$ and empty since-adoption record ($n = 0$). Inheriting ρ_t keeps the re-selection destination $\mathfrak{m}(\rho_t)$ common across cohorts (Lemma D.1, Proposition D.2); resetting the held model and test clock is what refreshes the cross-section, since young cohorts hold ϕ_0 and have not yet revised. The birth holding is a* boundary condition *distinct from the post-rejection adoption rule of Definition* **??***: at birth the held model is* drawn *from Π_ϕ, not set to the posterior mode $\mathfrak{m}(\rho_t)$, and the MAP rule governs*

only subsequent rejections. This birth-specific endowment—not unconcentrated newborn posteriors—is what sustains the stationary dispersion. Disagreement is thus carried by the held models, not by posterior dispersion. Economically, ρ_t is a public statistic*—a deterministic function of the public record, on the same footing as an equilibrium price—so this clause is the standard common-information convention: every investor, newborn or incumbent, conditions on the same public posterior, while the only datum that resets at birth is the agent's private adoption clock. The primitives of Assumptions* **??**–**??** *do not by themselves fix a newborn's posterior; this selects the common-information reading, a specification choice within the existing information structure rather than a new economic primitive.*

We now specialize to a finite *menu $\Phi = \{\phi_1, \ldots, \phi_K\}$ (the setting of Corollary E.2), an additional restriction for the stationary and quantitative exercise—in general Φ is a compact continuum (Assumption* **??***), not a finite set.*

Antecedent and difference. Perpetual youth is Blanchard's (1985) continuous-time overlapping-generations device, built on Yaari (1965); Gârleanu and Panageas (2015) use it precisely to keep heterogeneity from washing out and to obtain a nondegenerate stationary equilibrium, with the dead's wealth redistributed to the living. We inherit that structure and its purpose, and differ in the dimension of heterogeneity (the surprise threshold ε and the initial held model ϕ_0, not risk aversion) and the environment (CRRA with disaster jumps and lazy revision, not recursive-preference diffusion). The convention in (iii) is the Gârleanu and Panageas reduced form, adopted to avoid the annuity bookkeeping of the full Yaari (1965) market; it makes the demographic effect on the cross-section a single replacement term.

Under Assumption J.1(iii) the law of motion of the pricing state (Proposition **??**) gains exactly one term. Over $[t, t + dt]$ the population mass is held at one by balanced birth and death: a fraction $\Omega\, dt$ dies (proportionally in wealth, so the redistribution to survivors leaves their relative shares—hence the normalized measure—unchanged), while newborn mass $\Omega\, dt$ enters with type law Π_ϕ. The post-turnover normalized cross-section is therefore the mixture of the surviving S_t (mass $1 - \Omega\, dt$) and the newborn Π_ϕ (mass $\Omega\, dt$):

$$S_{t+dt}\big|_{\text{demog}} = (1-\Omega\, dt)\, S_t + \Omega\, dt\, \Pi_\phi, \qquad \text{hence} \qquad dS_t\big|_{\text{demog}} = S_{t+dt}\big|_{\text{demog}} - S_t = \Omega\,(\Pi_\phi - S_t)\, dt. \tag{J1}$$

Pairing the measure identity (J1) with any bounded h,

$$\frac{d}{dt}\langle h, S_t\rangle\Big|_{\text{demog}} = \Omega\big(\langle h, \Pi_\phi\rangle - \langle h, S_t\rangle\big), \tag{J2}$$

added to the selection, switch, and disaster channels of (??)–(??): a mean-reversion of the cross-section toward the newborn law Π_ϕ at rate Ω.

The replacement is resource-neutral by construction: it operates in consumption-share space, where the inflow $\Omega\, \Pi_\phi\, dt$ and outflow $\Omega\, S_t\, dt$ of (J1) preserve $\int_{\text{living}} s_t^i\, di = 1$, so goods continue to clear, $\int_{\text{living}} c_t^i\, di = C_t$, with no second claim on the Lucas tree. The share-endowment birth rule of Assumption J.1(iii) is the Gârleanu and Panageas (2015) reduced form for a Pareto-weight (human-capital) convention—a newborn receives a consumption share, not an additional financial claim funded out of the dead's annuitized wealth—which is what makes the demographic effect a single share-replacement term rather than a wealth-accounting transfer between the dead, the survivors, and the newborns. Formally the demographic turnover is a primitive sharing rule on the simplex; if decentralized, it is the Negishi construction of Appendix J with a *state-dependent* cohort weight ζ_t that implements the $\Omega\, dt$ share assignment, priced under fair annuities, so that the death hazard cancels from every survivor's Euler equation (IA.1.1) and enters the equilibrium only through the composition of the living. The share law (J1) is in turn invariant to the death-share bookkeeping—whether the dead's wealth-proportional share is annuitized to survivors and the block then diluted by the newborn inflow, or the dead's share is replaced directly by newborns, both give $(1 - \Omega\, dt)S_t + \Omega\, dt\, \Pi_\phi$ to first order in dt—which is what licenses writing the demographic effect as the single replacement term (J2) without committing to either ledger.

Proposition J.1 (The kernel under demographic turnover): *Under Assumptions* **??**–**??** *and J.1, an overlapping-generations equilibrium (Appendix J) exists, and its state-price density retains the form of Proposition* **??**,

$$\xi_t = e^{-\delta t}\, C_t^{-\alpha}\, \Xi_t^{\alpha}, \qquad \Xi_t = \int_{\text{living}} \widetilde{\psi}_i^{-1/\alpha} \big(M_t^i\big)^{1/\alpha}\, di, \tag{J3}$$

now the $\frac{1}{\alpha}$*-power mean of beliefs over the* living *population, with* born-at-par *cohort weights* $\widetilde{\psi}_i^{-1/\alpha} = \Xi_{b_i}$*. Under par entry the inflow and outflow of weight offset exactly, so* Ξ_t *carries* no demographic term*: it solves the selection–jump dynamics of Proposition E.1 applied to the living cross-section—calm drift* $-\frac{1}{\alpha}(\bar{\lambda}_t^S - \lambda^\star)$*, disaster factor* $g_t(x)$*—and* (**??**) *holds verbatim in the turnover economy. The death hazard* Ω *does not enter the time preference of the kernel—the fair annuity of Assumption J.1(iii) offsets mortality in each survivor's Euler equation—and enters the equilibrium only through the composition of the living, i.e. through the demographic channel* (J2) *of the cross-section. The bridge relation (Proposition* **??***) and the selection drift* (**??**)*—instantaneous relations of the kernel to the cross-section—hold verbatim in the turnover economy; the cross-section dynamics are those of Proposition* **??** *augmented by* (J2)*, so the calm erosion of Proposition* **??***(i) acquires a mean-reversion of* $\bar{\lambda}_t^S$ *toward the newborn mean at rate* Ω*, while the disaster jump and staggered recovery are unchanged in form.*

Three readings of the proposition. Demographic turnover is a primitive share-replacement law on the simplex (Assumption J.1(iii), eq. (J2)): newborns enter carrying an aggregate consumption share $\Omega\, dt$ drawn from Π_ϕ, the dead's wealth-proportional share is released, and $\langle 1, S_t\rangle \equiv 1$ is preserved. The born-at-par weight is what reconciles the two requirements: a *common* underlying weight would give the entering cohort the share $\zeta^{-1/\alpha}/\Xi_b$, fluctuating with the aggregator and growing with the birth date, incompatible with the constant-$\Omega\, dt$ inflow the demographic law requires; entry at par delivers the $\Omega\, dt$ share identically. Only the ratios and the drift of Ξ_t enter prices, so the explicit $e^{-\delta t}$ in (J3) is untouched and the riskless-rate *level* keeps its $+\delta$ term. The *growth* of Ξ_t is stationary under the stationary cross-section; its level is a multiplicative selection functional and need not be, and no priced object reads it.[1] The entrant thus arrives at the living population's current average weight—at par—and the aggregator's composition turns over while its dynamics do not.

Finite lives change *who* sits in the power mean. With fair annuities the mortality risk is fully insured, so each survivor optimises as if infinitely-lived at rate δ (the annuity offsets the death hazard Ω in the Euler equation); clearing on the exogenous Lucas endowment returns the kernel $\xi_t = e^{-\delta t} C_t^{-\alpha} \Xi_t^{\alpha}$ over the living cross-section, which turnover keeps non-degenerate. Under par entry the aggregate Ξ_t carries no demographic term: the entering weight $\Omega\, \Xi_t\, dt$ and the dying weight $\Omega\, \Xi_t\, dt$ offset exactly, so Ξ_t evolves by the within-life terms alone—the calm drift $d\Xi_t/\Xi_t = -\frac{1}{\alpha}(\bar{\lambda}_t^S - \lambda^\star)\, dt$ from each cohort's $d(M^i)^{1/\alpha}$, and the disaster factor $g_t(x)$—Proposition E.1 applied to the living cross-section. Hence $\Xi_t > 0$ and finite a.s. (bounded drift, finitely many bounded jump factors on compacts), its *growth* is a stationary functional of the cross-section, and its level is a multiplicative selection functional that no *normalized* priced object reads: conditional prices and returns use Ξ-ratios and the local drift alone, while the state-price density carries Ξ_t^{α}, so date-0 valuations of dated payoffs read the accumulated level only relative to the normalization date. The whole apparatus of Sections **??**–**??**, including the finite-menu fear factor (E17), therefore transfers to the stationary economy of Proposition K.1 without re-derivation; turnover supplies the stationary S_t these results take as input. The level formula (**??**) is unchanged: with the *growth* of Ξ_t stationary and its level unread by prices, the explicit $e^{-\delta t}$ survives, so r_t^f keeps its $+\delta$ term, and turnover enters the rate only through the demographic mean-reversion of $\bar{\lambda}_t^S$ and U_t in (J2), which smooths its realized path—the stationary safe rate of the perpetual-youth economy, its fluctuations smoothed—not removed—by the demographic mean-reversion (Gârleanu and Panageas, 2015).

Primitives. Under Assumption J.1 an investor born at b with type (ε, ϕ_0) and belief $\mathbb{P}^i$ (change-

[1] Par entry precludes the overlapping-generations δ-reduction in the rate *level*; the low safe rate in calm rests entirely on the fear-factor mechanism $\bar{\lambda}_t^S - U_t$, not on a demographic accounting offset.

of-measure M^i) chooses consumption to solve

$$\max\ \mathbb{E}^i\Big[\int_b^\infty e^{-\delta(s-b)}\mathbf{1}\{T_i > s\}\, u(c_s^i)\, ds\Big] = \max\ \mathbb{E}^i\Big[\int_b^\infty e^{-(\delta+\Omega)(s-b)}\, u(c_s^i)\, ds\Big], \qquad u(c) = \tfrac{c^{1-\alpha}}{1-\alpha}, \tag{J4}$$

where the residual lifetime satisfies $T_i - b_i \sim \text{Exp}(\Omega)$ (memoryless: $\mathbb{P}(T_i > s \mid T_i > b_i) = e^{-\Omega(s-b_i)}$, Assumption J.1(i)) and the second equality takes the expectation over the death time T_i, independent of the path. Markets are complete in the aggregate risks (Assumption **??**) with state-price density ξ; under the fair annuity of Assumption J.1(iii) a survival-contingent unit of consumption at s costs $\xi_s\, e^{-\Omega(s-b)}$ (the market price under $\mathbb{P}$ times the survival probability), so the market value of the consumption plan—the budget—is taken under the pricing measure $\mathbb{P}$,

$$\mathbb{E}\Big[\int_b^\infty \frac{\xi_s}{\xi_b}\, e^{-\Omega(s-b)}\, c_s^i\, ds\,\Big|\,\mathcal{F}_b\Big] \le W_b^i, \tag{J5}$$

with birth wealth W_b^i the value the birth endowment of Assumption J.1(iii) implies at the birth date b (type-dependent under heterogeneous beliefs, not per-capita), in date-b consumption units and $\mathcal{F}_b$-measurable (Assumption J.1(ii)–(iii)); ξ_b is its date-0 price, and the normalization by ξ_b is absorbed into the cohort multiplier $\widetilde{\psi}_i$ in the proof, so the dating affects no equilibrium object. An *overlapping-generations equilibrium* is an SPD ξ and allocations $\{c^i\}$ such that every living investor solves (J4)–(J5) and goods clear, $\int_{\text{living}} c_t^i\, di = C_t$.

Proofs for this section are in Section IA.1.

Appendix K. The stationary cross-section

Proposition K.1 (Stationary cross-section under demographic turnover): *Under Assumptions* **??**–**??** *and the finite-menu demographic turnover of Assumption J.1:*

(i) (Age structure.) *For $\Omega > 0$ the age distribution of the living is stationary with density $\Omega e^{-\Omega a}$ on $a \ge 0$ (at $\Omega = 0$ there is no turnover and no such stationary age law).*

(ii) (No refresh $\Rightarrow$ concentration.) *If* $\Omega = 0$*, the second-order posterior is consistent,* $\mathrm{MAP}(\rho_t^i) \to \phi^\star$ *a.s. (Berk, 1966); every re-selection* (**??**) *then returns* $\phi^\star$ *and the held-model cross-section converges to the single vertex* $\delta_{\phi^\star}$*.*

(iii) (Refresh $\Rightarrow$ unique nondegenerate stationary cross-section.) *If* $\Omega > 0$*, the cohort-resolved cross-section* Ψ_t*—the wealth measure over (held model* $\phi \in \Phi$*, since-adoption record* $\mathcal{R}$*, regime age* τ*, t cohort age* $a \geq 0$ *carried alongside for the demographics—is, in the limiting regime where the public posterior has concentrated (*$\mathrm{MAP}(\rho_t) = \phi^\star$*, so the re-selection destination is calendar-invariant), a time-homogeneous measurable functional of the marked disaster point process* N *on the* two-sided *stationary line and the stationary demographic-and-type structure (Proposition D.3 with Assumption J.1). On that two-sided process* Ψ_t *is a stationary, ergodic process with a unique stationary law* Ψ^∞*, to which the economy from any finite prior converges, geometrically at a rate proportional to* Ω *under the turnover-dominance condition* $\Omega > \kappa_\Omega$*, an explicit constant of the menu and preferences displayed in the proof—a crude sufficient bound (the two-sided construction, the stationary laws, and every time-average moment below hold without this condition)—in the bounded-Lipschitz (Wasserstein-*1*) metric on laws over* $\mathcal{P}(\Sigma)$*, where* Σ *carries the bounded product metric* $d_\Sigma\big((\phi, \mathcal{R}, \tau, \varepsilon), (\phi', \mathcal{R}', \tau', \varepsilon')\big) = |\lambda - \lambda'| + |\theta - \theta'| + \big(\mathbf{1}\{n \neq n'\} + \mathbf{1}\{n = n'\}\big(\sum_{k \leq n} |x_k - x_k'|\big) \wedge 1\big) + |\tau - \tau'| \wedge 1 + |\varepsilon - \varepsilon'|$*—each coordinate bounded (*Φ *compact, marks in* $\mathcal{J}$*), so* d_Σ *is a bounded metric—and* not *in total variation (Step 5)—as* $\mathrm{MAP}(\rho_t) \to \phi^\star$*; the pricing state* $S_t \in \Delta^{K-1}$ *(the* ϕ*-marginal of* Ψ_t*) and its instantaneous functionals (*U_t*,* $\bar{\lambda}_t^S$*), together with the price–dividend ratio* $v(\Psi_t)$*, the jump ratio* ψ_t*, and the disaster premium* DRP_t*—measurable and integrable functionals of the* full Ψ_t*—inherit unique stationary laws, with unconditional moments equal to long-run time averages (Birkhoff).* Ψ^∞ *places mass on at least two models whenever* Π_ϕ *does; the dependence of* Ψ_t *on the remote past (and on any initial cross-section) decays geometrically, at a rate proportional to* Ω*.*

Four notes on the statement. In the selection-only reduction with re-selection suppressed every vertex of Δ^{K-1} is absorbing; under the full re-selecting dynamics, by contrast, only $\delta_{\phi^\star}$ is invariant, since any other vertex is left at the next rejection—Remark E.1. The record coordinates of d_Σ are compared entrywise at equal counts and at distance one otherwise; the coupling bounds of the proof are insensitive to this choice up to constants, and the turnover-dominance constant κ_Ω is a crude sufficient bound. The rate is the demographic one for the head-count component; for the wealth-weighted state the coupling delivers convergence with the rate qualified by the selection and reweight perturbations (a Foster–Lyapunov strengthening in the sense of Costa and Dufour 2008 would sharpen it). The jump ratio reads the full state through $\psi_t(x) = v(\Psi_t^x)/v(\Psi_{t-})$ (Proposition **??**), and integrability rather than continuity is what the Birkhoff step needs. A fifth note, on the destination: in this locked regime the re-selection destination is the menu's Kullback–Leibler projection—the fearful point at the binary menu—so the calm switch flow is non-negative there and the stationary economy's calm relaxation rests on selection and turnover; the scheduled-relief reading of the calm channel is a finite-history statement (Proposition **??**).

The turnover does exactly one thing to the mechanism and one thing to the mathematics. To the mechanism (eq. J2): it adds a constant pull of the cross-section back toward the newborn dispersion Π_ϕ, at rate Ω, which is precisely the force that offsets the selection-and-consistency concentration of Remark E.1—young investors, whose *held models* have not yet moved to $\phi^\star$—they inherit the already-concentrated common posterior but enter with a drawn held model—are continually reborn. To the mathematics: with the menu finite and the inflow of fresh types constant, the cross-section becomes a time-homogeneous functional of a stationary, ergodic disaster clock, so it is itself stationary and ergodic, and the unconditional moments the calibration needs are well-defined and computable as time averages. The parameter that governs the stationary *dispersion*—and hence, through Proposition **??**, the amplitude of the synchronized revision wave—is the newborn joint law Π over birth held models *and* thresholds (its model marginal is Π_ϕ), together with the turnover rate Ω: more dispersed births leave more disagreement, the threshold law sets how long birth beliefs survive before revision (through the rejection and deadline times), and Ω governs how quickly fresh cohorts replace the selected cross-section. This is the sense in which, as anticipated, stationarity rests on the assumed joint law of threshold and intensity at birth.

Two beliefs, two clocks. The refresh and the synchronization live on different objects, and the stationary economy keeps them apart. The common second-order posterior ρ_t (over the menu, the source of the re-selection destination) is a public functional of the record, inherited by every cohort (Assumption J.1(iv)); by part (ii) it has concentrated on the fixed pseudo-true model, $\mathfrak{m}(\rho_t) \equiv \phi^\star$, so in the stationary regime the common destination is constant and carries no remote-past dependence. The per-agent state $(\phi, a, \mathcal{R})$ is reset at birth and its dependence on the past decays geometrically, at a rate proportional to Ω (Step 5 of the proof). Non-degeneracy is therefore a held-model phenomenon: a stationary age-mixture of ϕ_0-holding young and $\phi^\star$-absorbed old. The fast-in jump at a disaster is correspondingly carried by the reweighting channel—wealth flowing to the higher-intensity holders (Proposition **??**(ii))—and by relocation to the fixed $\phi^\star$, not by a jumping destination; the destination's disaster-time movement of Lemma D.1 is a transient, finite-record feature. The *direction* of that jump is a restriction, not an identity: relocation raises mean intensity only under the complacent-switcher condition of Proposition **??**—$\lambda(\phi^\star)$ above the *post-reweight mean of the triggered switchers*, a cross-sectional condition on the stationary state. A complacent newborn law makes it plausible, but Π_ϕ bears on the switcher distribution only through the stationary composition; we state the condition rather than derive it.

Two boundaries, kept sharp. First, stationarity delivers the invariant *distribution* Ψ^∞; it does not by itself close the price level. Even under the finite menu the price–dividend ratio $v(\Psi_t)$ solves the Feynman–Kac operator equation (IA.7.1) on the record-resolved state Ψ_t (Corollary IA.1), collapsing to a finite-dimensional-state operator equation in the shares only in the selection-only specialization with no record-dependent switching; with a continuous menu v likewise solves the analogous infinite-dimensional Feynman–Kac valuation problem of Proposition **??** (the displayed operator equation (IA.7.1) is the finite-menu form of Corollary IA.1).

So the invariant-distribution statement and the price-level statement are distinct. Second, the turnover modifies the equilibrium kernel only through this cross-section: Proposition J.1 showed the kernel form is preserved, with Ω entering the dynamics of Ξ_t only through the composition of the living—under par entry the demographic inflow and outflow offset, leaving no direct mortality drift—and hence r_t^f only through the stationary composition of S_t, never the time preference, so the pricing of Sections **??**–**??** transfers to the stationary economy: the realized record-resolved state Ψ_t is then *distributed* according to the invariant law Ψ^∞ of Proposition K.1, and the state-contingent pricing formulas are evaluated at the realized Ψ_t (with marginal S_t), not at the invariant law itself.

The proof is in Section IA.2.

Supplementary Part

The remaining sections collect the proofs for the demographic and stationarity constructions, secondary results and their proofs, the benchmark comparisons, and the numerical stability checks. References of the form IA.x point within this part; all other numbered references point to the main text.

Appendix IA.1. Proofs for the overlapping-generations construction

Proof of Proposition 18. Individual optimum. Markets are complete in the aggregate risks (Assumption 3), so investor i's dynamic problem reduces to the static budget problem under the unique state-price density ξ (the martingale method, valid for complete markets and a concave, increasing utility; Cox and Huang, 1989; Karatzas et al., 1990). Writing the objective (F4) under $\mathbb{P}$ via $\mathbb{E}^i[\,\cdot\,] = \mathbb{E}[M_s^i\,\cdot\,]$ and the budget (F5) under $\mathbb{P}$, the Lagrangian

$$\mathbb{E}\Big[\int_b^\infty \Big(e^{-(\delta+\Omega)(s-b)}u(c_s^i)\,M_s^i - \psi_i\,\xi_s\,e^{-\Omega(s-b)}\,c_s^i\Big)ds\Big]$$

is concave in c^i (u concave), so its pointwise (in (s,ω)) stationarity condition is necessary and sufficient: The date-b problem conditions on $\mathcal{F}_b$: the density and the price enter through the ratios M_s^i/M_b^i and ξ_s/ξ_b, the $\mathcal{F}_b$-measurable factors M_b^i and ξ_b being absorbed into the cohort multiplier ψ_i; with the absorbed multiplier the pointwise condition reads

$$e^{-(\delta+\Omega)(s-b)}\,u'(c_s^i)\,M_s^i = \psi_i\,\xi_s\,e^{-\Omega(s-b)},$$

$$e^{-\delta(s-b)}\,u'(c_s^i)\,M_s^i = \psi_i\,\xi_s, \tag{IA.1.1}$$

$$e^{-\delta(s-b)}\,(c_s^i)^{-\alpha}\,M_s^i = \psi_i\,\xi_s, \tag{IA.1.2}$$

line (IA.1.1) cancelling the common mortality factor $e^{-\Omega(s-b)}$ (the Yaari (1965) insurance result: the survival weight in the objective and the annuity price offset), and (IA.1.2) substituting $u'(c) = c^{-\alpha}$. Solve (IA.1.2) for c_s^i in one step each:

$$
\begin{aligned}
(c_s^i)^{-\alpha} &= \psi_i\, \xi_s\, e^{\delta(s-b)}/M_s^i \\
c_s^i &= \psi_i^{-1/\alpha}\big(\xi_s e^{\delta(s-b)}\big)^{-1/\alpha}(M_s^i)^{1/\alpha} && \text{(IA.1.3)} \\
&= \widetilde{\psi}_i^{-1/\alpha}\left(\xi_s e^{\delta s}\right)^{-1/\alpha}(M_s^i)^{1/\alpha}, \qquad \widetilde{\psi}_i := \psi_i\, e^{-\delta b}, && \text{(IA.1.4)}
\end{aligned}
$$

The born-at-par assignment sets the cohort weight at entry equal to the living average,

$$
\widetilde{\psi}_i^{-1/\alpha} \;=\; \Xi_{b_i}, \tag{IA.1.5}
$$

so the entering cohort's consumption share is, from (IA.1.3) at $s = b$ with $M_b^i = 1$, $c_b^i/C_b = \widetilde{\psi}_i^{-1/\alpha}/\Xi_b = 1$ per unit of entering mass—exactly the $\Omega\, dt$ aggregate share of (F1), with no residual dependence on the state at birth.

line (IA.1.4) writing $e^{\delta(s-b)} = e^{\delta s}e^{-\delta b}$ and absorbing the birth-time factor $e^{-\delta b}$ into $\widetilde{\psi}_i$ (so $(e^{-\delta b})^{-1/\alpha}$ joins $\psi_i^{-1/\alpha}$). Assumption 5(iii) supplies exactly these weights, (IA.1.5); the inflow and outflow of weight then offset, and Ξ carries no demographic term.

Market clearing fixes the kernel. Impose $\int_{\text{living}} c_s^i\, di = C_s$ with $\Xi_s := \int_{\text{living}} \widetilde{\psi}_i^{-1/\alpha}(M_s^i)^{1/\alpha}\, di$; the factor $(\xi_s e^{\delta s})^{-1/\alpha}$ is constant in i, so

$$
\begin{aligned}
\left(\xi_s e^{\delta s}\right)^{-1/\alpha}\int_{\text{living}} \widetilde{\psi}_i^{-1/\alpha}(M_s^i)^{1/\alpha} di &= C_s \\
\left(\xi_s e^{\delta s}\right)^{-1/\alpha}\Xi_s &= C_s && \text{(IA.1.6)} \\
\left(\xi_s e^{\delta s}\right)^{-1/\alpha} &= C_s/\Xi_s && \text{(IA.1.7)} \\
\xi_s\, e^{\delta s} &= \left(C_s/\Xi_s\right)^{-\alpha} = C_s^{-\alpha}\Xi_s^{\alpha} && \text{(IA.1.8)} \\
\xi_s &= e^{-\delta s}\, C_s^{-\alpha}\, \Xi_s^{\alpha}, && \text{(IA.1.9)}
\end{aligned}
$$

(IA.1.6) substituting the definition of Ξ_s, (IA.1.7) dividing by Ξ_s (> 0 as in Step 1 of Proposition 1, the bounds (C2) holding cohort by cohort), (IA.1.8) raising to the $-\alpha$, and (IA.1.9) multiplying by $e^{-\delta s}$. This is (F3), identical in form to (12) with the integral over the living. The death hazard appears nowhere in (IA.1.9): it cancelled at (IA.1.1) and enters the equilibrium only through which investors are alive at s—the demographic channel (F2)—and through the cohort multipliers $\widetilde{\psi}_i$ fixed at birth. Existence and uniqueness follow as in Proposition 1 (strict concavity of u and the complete-market budget), applied cohort by cohort. Since (F3) is the kernel of Proposition 1 evaluated at the living cross-section, the instantaneous relations of the kernel to the cross-section—the bridge (Proposition 2) and the selection drift (21)—hold verbatim with S_t the living cross-section, whose dynamics are those of Proposition 3 augmented by the demographic term (F2). □

Appendix IA.2. Proof of the stationary cross-section

Proof of Proposition 19. (i) Age structure. Population mass is constant: births of mass Ω per unit time balance deaths at hazard Ω. A living agent of age a was born a units ago and has survived the death clock with probability $e^{-\Omega a}$. The birth flow a units ago was Ω, so the mass of agents currently of age in $[a, a + da]$ is the birth flow times the survival probability,

$$\Omega\, e^{-\Omega a}\, da, \tag{IA.2.1}$$

and the total living mass is

$$\int_0^\infty \Omega e^{-\Omega a}\, da = \left[-e^{-\Omega a}\right]_{a=0}^{a=\infty} = 0 - (-1) = 1,$$

consistent with unit population. The age density $\Omega e^{-\Omega a}$ does not depend on t: it is stationary (Blanchard, 1985).

(ii) No refresh. By Assumption 4 the family is identifiable and, by Lemma 3, the Kullback–Leibler projection $\phi^\star$ of the true law onto Φ exists and is unique (Assumption 4(iv)). The second-order posterior ρ_t^i updates by Bayes on the full history (6). Under Assumption 1 the homogeneous Poisson process has stationary *independent* increments and the marks and diffusion increments are i.i.d., so the record splits into the i.i.d. sequence of unit-interval blocks $X_k =$ (the marked-Poisson record and diffusion increment on $(k-1,k]$), $k = 1, 2, \dots$; the full-history log-likelihood (6) is the sum of the block log-likelihoods, and the posterior at integer $t = n$ is the Bayes posterior from n i.i.d. observations $X_1, \dots, X_n$. Berk (1966)'s theorem applies to exactly this i.i.d. setting: for an identifiable dominated family observed i.i.d., with finite Kullback–Leibler divergence to the family, the posterior concentrates on the divergence-minimising parameter. Its three hypotheses hold block by block: identifiability (Assumption 4), i.i.d. blocks (stationary independent increments, Assumption 1), and finite divergence—the per-block divergence is the unit-time rate $\mathcal{K}(\phi)$ of (A2), finite with unique minimiser $\phi^\star$ (Lemma 3, (A3)). Hence $\rho_n^i \Rightarrow \delta_{\phi^\star}$ along integers. For non-integer t, the log-likelihood ratio of any $\phi \neq \phi^\star$ against $\phi^\star$ is the integer-time sum $-n\big(\mathcal{K}(\phi) - \mathcal{K}(\phi^\star)\big) + o(n)$ ($n = \lfloor t \rfloor$, $\mathcal{K}(\phi) > \mathcal{K}(\phi^\star)$ strictly by the unique minimiser) plus the contribution of the partial block $(\lfloor t \rfloor, t]$, bounded in absolute value by $C(1+N_{(\lfloor t \rfloor, \lfloor t \rfloor+1]})$ with r bounded (Assumption 1); since $\mathbb{P}(N_{(n,n+1]} > 2e \log n)$ is summable by the Poisson tail, Borel–Cantelli gives a partial-block contribution of order $O(\log n)$ uniformly over $t \in (n, n+1]$, almost surely, which the linear drift dominates; the ratio still diverges to $-\infty$ and $\rho_t^i \Rightarrow \delta_{\phi^\star}$ and $\mathrm{MAP}(\rho_t^i) \to \phi^\star$ for all t; on the finite menu the MAP is an argmax over finitely many vertices, and the divergence is uniform across them: each of the $K-1$ rivals has log-posterior-ratio to $\phi^\star$ tending to $-\infty$ a.s., so their maximum does too, and there is an a.s.-finite *lock time* τ_0 with $\mathrm{MAP}(\rho_t) = \phi^\star$ *exactly* for all $t \geq \tau_0$ (under a continuum of models finite-time equality fails, and mode convergence would need an argmax-consistency argument we do not make; the stationary exercise is finite-menu)—the *public-posterior lock time*, a property of the common record shared by the no-refresh dynamics here and the refresh economy of part (iii). Once an investor's MAP is $\phi^\star$, the common re-selection destination (11)

(Lemma 4) is $\phi^\star$, so every rejection returns $\phi^\star$ and her held model is absorbed there. With $\Omega = 0$ the demographic term (F2) is absent, so no other model receives inflow. Convergence to $\delta_{\phi^\star}$ is driven by re-selection and disaster reweighting, *not* by the calm selection drift (21): between disasters selection favors below-average-intensity (complacent) models and can move wealth *away* from $\phi^\star$, but every rejection returns the held model to $\phi^\star$, so each holder is re-selected there at its next rejection and $S_t \to \delta_{\phi^\star}$. The unique invariant law is therefore $\delta_{\phi^\star}$ itself: the other vertices of Δ^{K-1} are not invariant, since any holder there is re-selected to $\phi^\star$ at its next rejection (the selection-only contrast above)—the formalisation promised in Remark 6.

(iii) Refresh. Step 1 (functional representation). Proposition 10 established, for a *single* cohort sharing one path, that each investor's state is the deterministic measurable image $\mathsf{T}_t(\varepsilon)$ of her type and the cross-section the pushforward $\mathsf{T}_{t\#}G$—no law of large numbers, there being no idiosyncratic randomness to average. That argument applies cohort by cohort here, with two extensions it accommodates without change to its logic: the type is now the *joint* birth draw $(\varepsilon, \phi_0) \sim \Pi$ rather than a common ϕ_0—the revision times (B4) and the path-determined re-selection destinations (Lemma 4) are deterministic measurable functions of (ε, ϕ_0) exactly as they were of ε alone, so each cohort's contribution is a pushforward $\mathsf{T}^a_{t\#}\Pi$ over its since-birth segment $N|_{[t-a,t]}$—and the cross-section carries a *biological-age* coordinate (constant, hence suppressed, in the single-cohort case) recording survival and wealth. Under Assumption 5 the population is the wealth-weighted age-mixture of such cohorts—age density $\Omega e^{-\Omega a}$ by (IA.2.1), birth type Π. An agent of age a alive at t was born at $t-a$ with $(\varepsilon, \phi_0) \sim \Pi$ and has evolved her held model, record, and wealth share under the marked segment $N|_{[t-a,t]} = \{(T_k, J_k) : t-a < T_k \leq t\}$ by the revision rule (Proposition 12) and the share dynamics (Proposition 11); the marks J_k enter through the likelihood ratios and the share jumps (C), so it is the marked process, not the bare count, that drives her trajectory. The *destination* of any revision she makes is the common mode $\mathfrak{m}(\rho_t)$, inherited identically by every living cohort (Assumption 5(iv), Lemma 4); by part (ii) it has concentrated on the fixed $\phi^\star$, so it contributes no separate since-birth dependence. Her belief change-of-measure M^i, hence her share $s^i = \zeta_i^{-1/\alpha}(M^i)^{1/\alpha}/\Xi$ (13), is

then a functional of the marked disaster history alone (the diffusion is undisputed), and the birth multiplier ζ_i is common across cohorts because the equal-weight normalization of Assumption 3 assigns every cohort the common effective cohort multiplier ζ (Assumption 5(iii) with the share-implementing birth endowment), not because birth wealth is equal, cancelling from s^i. Aggregating over the stationary age-and-type structure—each age-a cohort contributing its Π-distributed, $N|_{[t-a,t]}$-evolved states with its *wealth* weight $\Omega e^{-\Omega a}\,(M_a)^{1/\alpha}/\Xi_t$ (the demographic mass $\Omega e^{-\Omega a}$ times the belief share of Proposition 18, not the bare age density)—gives

$$\Psi_t = \mathfrak{F}\big(N|_{(-\infty,t]}\big).$$

The aggregation over $a \in [0,\infty)$ converges: each cohort contributes a probability measure on the finite Φ weighted by its wealth share in $[0,1]$, and the age weight integrates to one, the wealth-weighted normalizer Ξ_t positive and finite a.s. (Proposition 18) ($\int_0^\infty \Omega e^{-\Omega a}\,da = 1$ by (IA.2.1)), so Ψ_t is a well-defined probability measure and $\mathfrak{F}$ is measurable and independent of calendar time (rule, DGP, and demographics are time-homogeneous).

Step 2 (stationarity and ergodicity). We establish stationarity and ergodicity by the *factor* method—representing Ψ_t as a measurable functional of the stationary ergodic disaster process—rather than by exhibiting an infinitesimal generator for the measure-valued birth–death dynamics and verifying a Doeblin/minorisation condition; the factor route is the natural one for a pricing state that is itself non-Markov (Proposition 3), and the stationary moments the calibration needs are time averages, for which ergodicity alone suffices. N is a homogeneous Poisson process; its law is invariant under the time shift $\theta_h : N|_{(-\infty,t]} \mapsto N|_{(-\infty,t+h]}$, and the shift is ergodic—a process with stationary independent increments is mixing, hence ergodic, under time translation. We construct the invariant law on the *two-sided* stationary marked Poisson measure $N|_{(-\infty,t]}$: on it we *define* the re-selection destination as the locked limiting regime: MAP $\equiv \phi^\star$ at every date, the stationary counterpart of the a.s.-finite lock of part (ii)—no posterior is computed on $(-\infty,t]$, and no infinite-past likelihood enters; the functional $\mathfrak{F}$ is then genuinely time-homogeneous. A time-homogeneous measurable functional of the stationary ergodic two-sided process is stationary and ergodic (a measurable factor of an ergodic shift is ergodic), so $\Psi_t = \mathfrak{F}(\theta_t N)$ is stationary and ergodic with law Ψ^∞. The functional is moreover wealth-summable: remote cohorts carry vanishing *wealth* weight, not only vanishing demographic mass—the old-cohort tail obeys (IA.2.3) of Step 5, whose pathwise product estimate applies verbatim on the two-sided line, so likelihood selection cannot offset demographic forgetting under the turnover-dominance condition. *Uniqueness* and convergence from an arbitrary one-sided initialisation rest not on factor inheritance alone but on the demographic coupling, and are proved in Step 5 by an explicit synchronous coupling: the cohort weight $\Omega e^{-\Omega a}$ forgets the remote past at rate Ω, and the one-sided economy couples to the two-sided regime as $\mathrm{MAP}(\rho_t) \to \phi^\star$ (Berk; on the finite menu each false model's unit-window log-likelihood-ratio increments against $\phi^\star$ are i.i.d. with negative mean and finite exponential moments—the exponential-family structure, Küchler and Sørensen 1989—so a Chernoff bound with Borel–Cantelli drives the posterior odds to the projection geometrically; the fractional block over $(\lfloor t \rfloor, t]$ is $O(1 + N_{(\lfloor t \rfloor, \lfloor t \rfloor+1]})$ with exponential moments, hence $o(t)$ a.s.)).

Step 3 (inheritance and moments). $S_t = \pi_\phi(\Psi_t)$, the ϕ-marginal, is a continuous linear functional of Ψ_t; the instantaneous objects U_t, $f_t^{\mathbb{Q}}$, and $\bar{\lambda}_t^S$ are functionals of S_t (Proposition 1, Proposition 2), finite-dimensional on Δ^{K-1} under the finite menu, while the price–dividend ratio $v(\Psi_t)$ *and* the disaster-premium *level* DRP_t are functionals of the full record-resolved state Ψ_t (Propositions 13, 14, 15), bounded and measurable in Ψ_t—continuous off the deadline and trigger thresholds of the revision rule, where they may jump, but not assumed continuous. Each is thus a bounded measurable functional of Ψ_t and inherits a unique stationary law—the premium as a functional of the stationary Ψ_t, not as a finite-dimensional function of S_t alone. These functionals are bounded under (T): U_t and $\bar{\lambda}_t^S$ are continuous on the compact Δ^{K-1}; $0 < v(\Psi_t) \le (\delta - \sup m)^{-1}$ by (E25), so by (E.A) the jump ratio $\psi_t = v(\Psi_t^J)/v(\Psi_{t-})$ is bounded above and below away from zero; and DRP_t (E27) is then bounded, loading on the bounded U_t and on the bounded disaster return $\Delta R_t^e = e^{-J}\psi_t - 1$. All are integrable, so the Birkhoff ergodic theorem gives, for any such g,

$$\frac{1}{T}\int_0^T g(\Psi_t)\,dt \xrightarrow{\text{a.s.}} \mathbb{E}_{\Psi^\infty}[g] \qquad (T \to \infty),$$

so the unconditional, calendar-time moments are well-defined and equal these limits. Disaster-conditional *event-study* averages—means taken over disaster epochs—are not calendar-time averages; for the stationary, ergodic marked process N they are the Palm expectation $\mathbb{E}^0_{\Psi^\infty}[g]$, and the Palm–Campbell ergodic theorem gives the corresponding $\mathbb{P}^0$-a.s. convergence of disaster-epoch averages to it. The two limits coincide only for functionals that do not condition on an arrival.

Step 4 (non-degeneracy). Fix $\eta > 0$ and let $A_t = t - T_{N_t}$ be the current calm age; set $w := \eta \wedge A_t$, strictly positive a.s. at every fixed t (disasters are a.s. not at t). Cohorts born in $(t-w, t]$ carry an *empty* disaster record, so none is disaster-triggered; the only revisions among them are calm deadlines inside the window, of type mass $r(w) := \Pi(\{\varepsilon : \Delta^\star(\phi_0, \varepsilon) \le w\})$, and since $\Delta^\star > 0$ pointwise on $(0,1)$ these sets decrease to $\varnothing$, so $r(\eta) \downarrow 0$; fix η with $r(\eta) \le \frac{1}{2}\min_{\phi\in\operatorname{supp}\Pi_\phi} \Pi_\phi(\{\phi\})$. Each such cohort's weight is its par-entry level Ξ_{t-a} times its own calm factor $(M^i_{t-a\to t})^{1/\alpha} \ge e^{-Ca}$, $C := \lambda_{hi}/\alpha$; and over the jump-free window the aggregate moves only by its bounded calm drift, $\Xi_{t-a}/\Xi_t = \exp\{\frac{1}{\alpha}\int_{t-a}^t (\bar\lambda^S_u - \lambda^\star)\,du\} \ge e^{-Ca}$. Hence, pathwise, for every $\phi \in \operatorname{supp}\Pi_\phi$,

$$S_t(\{\phi\}) \;\ge\; e^{-2C\eta}\big(1 - e^{-\Omega w}\big)\,\tfrac{1}{2}\,\Pi_\phi(\{\phi\}) \;>\; 0 \qquad \text{a.s.},$$

the Ξ-level cancelling between the cohort weight and the normaliser, and the same bound holds for Ψ^∞ with $A_0 \sim \operatorname{Exp}(\lambda^\star)$ strictly positive a.s. If $|\operatorname{supp}\Pi_\phi| \ge 2$, at least two models carry positive stationary mass a.s.: Ψ^∞ is nondegenerate, in contrast to (ii). Step 5 (convergence from an arbitrary init We prove that the law of the one-sided economy started from any initial cross-section Ψ_0 converges to Ψ^∞, and we identify the mode. Couple the two economies on a common driving path: run the one-sided economy from Ψ_0 at $t = 0$ and the stationary two-sided economy

$\Psi_t^\infty = \mathfrak{F}(\theta_t N)$ on the *same* marked Poisson process N, sharing the entire post-zero birth stream—birth epochs, types $(\varepsilon, \phi_0) \sim \Pi$, and death clocks. The economies then differ only in their pre-zero composition: the one-sided economy carries the initial cohorts encoded in Ψ_0, the two-sided economy the cohorts born in $(-\infty, 0]$. No likelihood comparison over that infinite pre-history is made—it carries infinitely many arrivals a.s., so no bounded pre-zero product exists; the argument runs through the one-sided lock alone.

Destination lock. The public re-selection destination is the posterior MAP $\mathfrak{m}(\rho_t)$ (11), which by part (ii) *locks* on the finite menu at the a.s.-finite public-posterior lock time: each rival vertex's log-posterior-ratio to $\phi^\star$ diverges to $-\infty$ a.s., so with finitely many vertices the MAP equals $\phi^\star$ exactly for all t beyond that time. The stationary economy's destination is $\phi^\star$ at every t by construction of $\mathfrak{F}$—its re-selection uses the locked two-sided regime of part (ii); the one-sided posterior from ρ_0 locks on its own, by the finite-menu Chernoff argument above, at an a.s.-finite time τ_0. For $t \geq \tau_0$ both destinations equal $\phi^\star$, and every revision in either economy sends the reviser there.

Common cohorts coincide. A cohort born at $b \geq \tau_0$ has, in both economies, the same birth type, the same since-birth segment $N|_{[b,t]}$, and the same destination $\phi^\star$ at every revision; its held model, since-adoption record, age, and unnormalised belief weight $(M^i)^{1/\alpha}$ are therefore *identical* across the two economies. Cohorts born in $[0, \tau_0)$ may have revised to differing destinations before τ_0 and can differ between the two economies in records, ages, and adoption dates until their deaths; no coupling of their states is claimed, and their joint contribution is the W_t^{pre} mass bounded above.

Distance bound. The two cross-sections share the post-τ_0 cohorts' *path-only* factors—types, records, destinations, and hence the $(M^i)^{1/\alpha}$ terms coincide cohort by cohort—while the par-entry levels Ξ_b need not: entry is at the economy's own aggregate, and the aggregates differ through the coupled paths. The states therefore differ through three wealth groups and one distortion channel—the pre-zero initial cohorts, mass W_t^{init}; the pre-zero infinite-past cohorts, mass W_t^∞; and the $[0,\tau_0)$-born cohorts that revised before τ_0 and may differ across the two economies until their deaths, mass W_t^{pre}—and the renormalisation these induce in Ξ_t (their belief-weight factors dominated uniformly by Lemma 7, so the age-mass envelope $Ce^{-\Omega(t-\tau_0)}$ bounds the conditional and unconditional contribution alike). For any test function h on the per-investor state space that is 1-Lipschitz with $\|h\|_\infty \le 1$, the common part cancels in $\langle h, \Psi_t\rangle - \langle h, \Psi_t^\infty\rangle$ up to the entry-level distortion and the common renormalisation, and each offending group contributes at most twice its wealth share—the pre-zero groups through their own mass and the Ξ_t renormalisation, the pre-lock group through its differing copy in each economy (the pre-lock mass cancels in $\Xi_t - \Xi_t^\infty$). The distortion tilts the common cohorts' relative weights: writing $\rho_b := \log\bigl(\Xi_b/\Xi_b^\infty\bigr)$ for the entry-level gap, $\delta_u := \|S_u - S_u^\infty\|_{\text{TV}}$, and $d\Lambda_u := \bigl[\frac{\lambda_{hi}-\lambda_{lo}}{\alpha}\,du + \frac{q^{-1}-1}{2}\,dN_u\bigr]\delta_{u-}$, the coupled aggregate dynamics give $|\rho_b - \rho_{b'}| \le \Lambda(b \wedge b', b \vee b']$—the aggregate drift and the jump normalisations are Lipschitz in the cross-section (the Lipschitz constants of the main text's total-variation stability lemma)—so the two normalised common parts differ in total variation by at most $2e\int_{\tau_0}^t A_{t,r}\,d\Lambda_r$ whenever $\Lambda(\tau_0, t] \le 1$, and by 2 always, with $A_{t,r}$ the Ψ^∞-share of common cohorts born before r. Hence

$$d_{\text{BL}}(\Psi_t, \Psi_t^\infty) = \sup_h \bigl|\langle h, \Psi_t\rangle - \langle h, \Psi_t^\infty\rangle\bigr| \;\le\; 2\bigl(W_t^{\text{init}} + W_t^\infty + W_t^{\text{pre}}\bigr) \;+\; 2\min\Bigl(1,\; e\int_{\tau_0}^t A_{t,r}\,d\Lambda_r\Bigr). \tag{IA.2.2}$$

Decay of the bound. W_t^∞ is the wealth share of two-sided cohorts older than t—the old-cohort tail

$$W_{>a}(t) = \int_a^\infty \Omega e^{-\Omega a'}\,\frac{\Xi_{t-a'}}{\Xi_t}\,\overline{(M^{a'})^{1/\alpha}}\,da' \tag{IA.2.3}$$

at $a = t$, with $\overline{(M^{a'})^{1/\alpha}}$ the age-a' cohort average of $(M^i)^{1/\alpha}$ and $\Xi_{t-a'}$ the cohort's par-entry level (IA.1.5); by that identity the integrand over $[0,\infty)$ integrates to one pathwise (shares sum to one). Two facts bound its tail, and they must be taken *pathwise together*: the cohort's belief factor and the aggregate ratio are driven by the same post-birth disasters, so their expectations do not factor. Along the window, each investor's contribution to the product $(\Xi_{t-a'}/\Xi_t)\,\overline{(M^{a'})^{1/\alpha}}$ multiplies a calm factor $\exp\{\frac{1}{\alpha}\int_{t-a'}^t(\bar\lambda_u^S - \lambda_{\text{held},u})\,du\} \le e^{(\lambda_{hi}/\alpha)a'}$ by a jump factor $\prod_{\text{jumps}} r(\phi_{\text{held}}, x)^{1/\alpha}\, g_u(x)^{-1} \leq q^{-N_{(t-a',t]}}$, with $q := (\underline{r}/\overline{r})^{1/\alpha}$ the contraction ratio of the main text's total-variation stability lemma, since $g_u \geq \underline{r}^{1/\alpha}$ and $r^{1/\alpha} \leq \overline{r}^{1/\alpha}$ on the compact menu; the cohort average obeys the same pathwise bound, and one application of the Poisson moment generating function gives $\mathbb{E}[q^{-N_{a'}}] = e^{\lambda^\star(q^{-1}-1)a'}$. Combining, $\mathbb{E}_{\mathbb{P}}[W_{>a}(t)] \leq \int_a^\infty \Omega e^{-\Omega a'} e^{\kappa_\Omega a'}\, da' = \frac{\Omega}{\Omega-\kappa_\Omega}\, e^{-(\Omega-\kappa_\Omega)a}$ under the turnover-dominance condition $\Omega > \kappa_\Omega$, with

$$\kappa_\Omega := \lambda_{hi}/\alpha + \lambda^\star(q^{-1}-1) + 2e\Big[\tfrac{\lambda_{hi}-\lambda_{lo}}{\alpha} + \tfrac{1}{2}\lambda^\star(q^{-1}-1)\Big]$$

the crude menu constant—its first two terms bound this tail; the bracketed load, scaled by the tilt factor $2e$, serves the distortion closure of Step 5—, so $\mathbb{E}_{\mathbb{P}}[W_t^\infty] \leq Ce^{-c\Omega t}$ with $c := 1 - \kappa_\Omega/\Omega \in (0,1]$. The one-sided group W_t^{init} has the Ψ_0-dependent age profile of the initial cohorts, which need *not* be the stationary $\Omega e^{-\Omega a'}$; but those cohorts survive from 0 to t only with probability $e^{-\Omega t}$, their aggregate par-normalized weight at 0 is one, and the same pathwise product bound and one Poisson expectation give $\mathbb{E}_{\mathbb{P}}[W_t^{\text{init}}] \leq e^{-\Omega t} e^{\kappa_\Omega t} = e^{-c\Omega t}$ as well. The third group decays unconditionally, by a split at $t/2$: on $\{\tau_0 \leq t/2\}$ every $[0,\tau_0)$ cohort still alive at t has deterministic age at least $t/2$, so its wealth share is bounded by the old-cohort tail $W^{\text{pre}}_{>t/2}(t) \leq Ce^{-c\Omega t/2}$ of (IA.2.3), an estimate at a *fixed* age requiring no conditioning on τ_0; and $\mathbb{P}(\tau_0 > t/2) \leq C'e^{-c't}$, since for each of the finitely many rivals the log-posterior-ratio at u has mean $-\mathcal{K}\, u$ and finite exponential moments, so, writing X_u for the ratio and $\psi(\theta)$ for its increment cumulant—finite for every θ by the bounded likelihood ratios, with $\psi'(0) = -\mathcal{K} < 0$—there is $\theta^\star > 0$ with $\psi(\theta^\star) < 0$; $e^{\theta^\star X_u - u\psi(\theta^\star)}$ is then a positive martingale, on $\{X_u \geq 0\}$ with $u \geq t/2$ it exceeds $e^{|\psi(\theta^\star)|t/2}$, and Doob's maximal inequality with a union over the finitely many rivals gives $\mathbb{P}(\sup_{u\geq t/2} X_u \geq 0) \leq Ce^{-|\psi(\theta^\star)|t/2}$. Hence $\mathbb{E}[W_t^{\text{pre}}] \leq Ce^{-c\Omega t/2} + C'e^{-c't}$, with no conditioning on the last-passage time. The lock time is the last-passage time above zero of the negative-drift log-posterior-ratio walk, whose increments are light-tailed (bounded r, Poisson marks), so $\mathbb{E}_{\mathbb{P}}[e^{c\Omega\tau_0}] < \infty$ after possibly shrinking c; the outer expectation then gives $\mathbb{E}_{\mathbb{P}}[W_t^{\text{pre}}] \leq C'e^{-c\Omega t}$. (A $[0,\tau_0)$-born cohort whose *first* revision falls after τ_0 re-selects $\phi^\star$ in both economies and never leaves the common part; a cohort that revised before τ_0 stays in this group to its death, and the age bound above needs nothing more.) *Closing the distortion.* The kernel $A_{t,r}$ is the common-part wealth of cohorts older than $t-r$, so the pathwise product estimate above gives, along almost every path, $A_{t,r} \leq C(\omega)e^{-(\Omega-\tilde{\kappa})(t-r)}$ with $\tilde{\kappa} := \lambda_{hi}/\alpha + \lambda^\star \log(1/q) < \kappa_\Omega$—the law-of-large-numbers rate of the jump count, with $C(\omega) < \infty$ a.s. by Borel–Cantelli over unit blocks. Dominating δ_u by the realised total-variation distance $\mathcal{D}_u := \|\Psi_u - \Psi_u^\infty\|_{\mathrm{TV}}$, which obeys the same right-hand side, and weighting by $e^{\beta t}$ turns (IA.2.2) into the pathwise renewal inequality $\sup_{s \leq t} e^{\beta s}\mathcal{D}_s \leq C'(\omega) + \theta \sup_{s\leq t} e^{\beta s}\mathcal{D}_s$, with $\theta = 2e\big(\frac{\lambda_{hi}-\lambda_{lo}}{\alpha} + \lambda^\star \frac{q^{-1}-1}{2}\big)/(\Omega - \tilde{\kappa} - \beta)$ and the jump sum handled by the same unit-block argument; $\theta < 1$ for $\beta \leq \frac{1}{2}(\Omega - \kappa_\Omega)$: by construction $\kappa_\Omega - \tilde{\kappa} \geq 2e\big[(\lambda_{hi} - \lambda_{lo})/\alpha + \frac{1}{2}\lambda^\star(q^{-1}-1)\big]$, so $\Omega - \tilde{\kappa} - \beta$ exceeds θ's numerator by at least $\frac{1}{2}(\Omega - \kappa_\Omega)$. Hence $\mathcal{D}_t \leq C''(\omega)e^{-\beta t}$ almost surely; since $\mathcal{D}_t \leq 2$, dominated convergence gives $\mathbb{E}_{\mathbb{P}}[d_{\mathrm{BL}}(\Psi_t, \Psi_t^\infty)] \to 0$ geometrically at a rate proportional to Ω, and Markov's inequality gives convergence in probability.

Mode and uniqueness. The coupled measures themselves converge: because the post-τ_0 part enters both as a common measure, the bound (IA.2.2) holds a fortiori in total variation, $\|\Psi_t - \Psi_t^\infty\|_{\mathrm{TV}}$ bounded by the same right-hand side and $\to 0$, with $d_{\mathrm{BL}} \leq \|\cdot\|_{\mathrm{TV}}$ (this is the total-variation distance between the coupled *realisations* Ψ_t, Ψ_t^∞ on a common probability space, not between their laws). Through $\mathrm{W}_1\big(\mathrm{Law}(\Psi_t), \Psi^\infty\big) \leq \mathbb{E}_{\mathbb{P}}[d_{\mathrm{BL}}(\Psi_t, \Psi_t^\infty)] \leq C''e^{-c\Omega t}$ this gives convergence of the *law* of Ψ_t to Ψ^∞ in the Wasserstein-1 (bounded-Lipschitz) metric on laws over $\mathcal{P}(\Phi \times \mathbb{R}_+ \times \mathcal{R})$, geometrically at rate $c\Omega$—the asymptotic-coupling route to Wasserstein ergodicity of Hairer et al. (2011). We do *not* claim total-variation convergence of the law in the Meyn and Tweedie (2009) sense: the synchronous coupling never sets $\Psi_t = \Psi_t^\infty$ (a positive, if vanishing, wealth of distinct-atom cohorts always differs), so it is not a successful coupling and yields no minorisation, and the measure-valued birth–death dynamics are not shown to satisfy a Doeblin condition in total variation—which the stationary moments do not require, being long-run *time averages* delivered by the Birkhoff theorem of Step 3 under integrability alone, a route that needs no continuity and no control of the law beyond ergodicity. Because the two-sided regime Ψ^∞ is independent of Ψ_0, every initialisation converges to the *same* limit, so the stationary law is unique; the dependence of Ψ_t, and of every functional of Step 3, on the initial cross-section and the remote past decays geometrically, at a rate proportional to Ω. $\square$

Appendix IA.3. Identification of the physical disaster law

The lemma below records the asymmetry motivated in the main text: the cross-section is observable in principle while the physical law is not.

Lemma IA.1 (Identification asymmetry): *Under Assumptions 1 and 4, observe the path of C on $[0,t]$.*

(i) (Diffusion: pathwise.) *The continuous quadratic variation of $\log C$ satisfies, for every $t \geq 0$ and $h > 0$,*

$$[\log C]^c_{t+h} - [\log C]^c_t = \sigma^2 h \qquad \textit{exactly,} \tag{IA.3.1}$$

and is recovered from a discrete record on any fixed $[s, s+h] \subseteq [0,t]$ as the probability limit, as the mesh shrinks, of the truncated realized variation. A held diffusion coefficient $\neq \sigma$ is contradicted on every interval.

(ii) (Realized disasters: pathwise.) *The disaster times $\{T_k\}$ and sizes $J_k = -\Delta \log C_{T_k}$ are the discontinuities of $\log C$ and their magnitudes; from a discrete record they are recovered consistently, at rate $\sqrt{n}$ in the mesh (n the number of high-frequency observation intervals, not the disaster count N_t). The recovered $\{J_k\}_{k:T_k \leq t}$ is an i.i.d. sample from F_J of size $N_t \sim \mathrm{Poisson}(\lambda^\star t)$.*

(iii) (The disaster law: statistical, and slow in both coordinates.) *Neither the intensity nor the size law is a functional of the record on $[0,t]$; both are identified only through the realized disaster record—for the intensity, the Poisson count over exposure time, so that the*

informative zero-count spells enter as well, not only the arrivals—and the information is block-diagonal between them. Intensity: *the count N_t is a sufficient statistic for λ in the Poisson observation on $[0,t]$, with Fisher information t/λ; the maximum-likelihood estimator $\hat{\lambda}_t = N_t/t$ has $\operatorname{Var}(\hat{\lambda}_t) = \lambda^\star/t$, so its* relative *error is $(\hat{\lambda}_t - \lambda^\star)/\lambda^\star = O_{\mathbb{P}}\big((\lambda^\star t)^{-1/2}\big)$ (the standard-deviation-normalized error is $O_{\mathbb{P}}(1)$).* Size law: *let $\hat{\theta}_t$ be the size (pseudo-)maximum-likelihood estimator on the realized sizes; under the regularity of Assumption 4(vi) it is asymptotically normal at rate $n^{-1/2}$ conditional on $N_t = n$ (Appendix A of the main text), and with $N_t \sim \operatorname{Poisson}(\lambda^\star t)$ this pins the rate at*

$$\hat{\theta}_t - \theta^\star = O_{\mathbb{P}}\big((\lambda^\star t)^{-1/2}\big). \tag{IA.3.2}$$

Both errors vanish only as $t \to \infty$, and slowly when $\lambda^\star$ is small, since both are governed by the disaster count $N_t \sim \operatorname{Poisson}(\lambda^\star t)$. By full support $f_\theta(J_k) > 0$ at every realized size, so a held model is never contradicted by an individual disaster—only by the accumulated record.

Proof of Lemma IA.1. (i). By (4), $\log C_t = \log C_0 + \int_0^t (\mu_0 - \frac{1}{2}\sigma^2)\, ds + \sigma W_t - \sum_{k:T_k \le t} J_k$. The drift $\int_0^\cdot (\mu_0 - \frac{1}{2}\sigma^2)\, ds$ is of finite variation, hence contributes nothing to the quadratic variation; the disaster sum $\sum_{k:T_k \le \cdot} J_k$ is a finite-activity pure-jump process, hence contributes nothing to the *continuous* quadratic variation. The only continuous component of $\log C$ is σW, so its continuous quadratic variation is

$$[\log C]^c_t = [\sigma W]_t = \sigma^2 [W]_t = \sigma^2 t,$$

the last equality using $[W]_t = t$. Therefore

$$[\log C]^c_{t+h} - [\log C]^c_t = \sigma^2 (t+h) - \sigma^2 t = \sigma^2 h,$$

which is (IA.3.1). Mancini's truncated realized variation is a *fixed-horizon, shrinking-mesh* object, so we fix any interval $[s, s+\delta] \subseteq [0, t]$ of positive length δ and partition it with mesh $\Delta \to 0$ (*not* $\delta \to 0$). With the hypotheses verified above, (A) applied on $[s, s+\delta]$ with mesh-dependent threshold $r(\Delta)$ gives

$$\operatorname*{plim}_{\Delta\to 0} \sum_{i:\,[s_{i-1},s_i]\subseteq[s,s+\delta]} (\Delta_i \log C)^2\, \mathbf{1}\{(\Delta_i \log C)^2 \le r(\Delta)\} = [\log C]^c_{s+\delta} - [\log C]^c_s = \sigma^2\delta,$$

the interval length held fixed while the mesh vanishes. Hence $\sigma^2 = (\sigma^2\delta)/\delta$ is recovered on every fixed interval, so σ is a pathwise functional of the continuous record and two investors who disagree about the diffusion coefficient are contradicted on every interval.

(ii). By Assumption 1, N is Poisson with finite intensity, so on $[0, t]$ it has finitely many jumps a.s.; by (4) the only discontinuities of $\log C$ are at the T_k, where $\Delta \log C_{T_k} = -J_k$, so the times and sizes are read off the path as its discontinuities and their magnitudes. From a discrete record, the jump-time estimator $\widehat{\Delta_i N} = \mathbf{1}\{(\Delta_i \log C)^2 > r(h)\}$ recovers the jump intervals (Mancini Thm. 3.1) and the size estimator $\hat{\gamma}^{(i)}$ recovers each realized size: in this finite-activity setting each jump is eventually isolated in its own mesh interval, so the truncated increment over that interval recovers the size with a Brownian error that vanishes as the mesh shrinks; the rate-$\sqrt{n}$ mixed-normal limit (A7) is the joint fluctuation of these componentwise recoveries, not a substitute for them. The marks are i.i.d. F_J and independent of N by Assumption 1, and $N_t \sim \text{Poisson}(\lambda^\star t)$, so $\{J_k\}_{k:T_k\le t}$ is an i.i.d. F_J-sample of size $N_t \sim \text{Poisson}(\lambda^\star t)$.

(iii). By (i)–(ii), the record on $[0, t]$ determines σ and the realized sizes exactly in the high-frequency limit, but the size *law* enters only through the i.i.d. sample $\{J_k\}_{k\le N_t}$; the diffusion and drift carry no information about F_J. Condition on $N_t = n$ and let $\hat{\theta}_t \in \arg\max_{\theta\in\Theta} \sum_{k\le n} \log f_\theta(J_k)$ be the size maximum-likelihood estimator.

Achieved rate. Imported result (asymptotic normality of the possibly-misspecified maximum-likelihood estimator; White, 1982, Thm. 3.2; van der Vaart, 1998, Thm. 5.39). For i.i.d. $J_1, \ldots, J_n \sim F_J$ with pseudo-true value $\theta^\star = \arg\max_{\theta\in\Theta} M(\theta)$, $M(\theta) := \mathbb{E}_{F_J}[\log f_\theta(J)]$, interior and unique, and under twice-continuous differentiability of $\theta \mapsto \log f_\theta$ with first and second derivatives dominated by an F_J-square-integrable envelope and nonsingular pseudo-Hessian $H := -M''(\theta^\star)$,

$$\sqrt{n}\,\big(\hat\theta_t - \theta^\star\big) \;\xrightarrow{\;d\;}\; N\big(0,\, H^{-1}BH^{-1}\big), \qquad B := \mathbb{E}_{F_J}\big[\dot\ell_{\theta^\star}(J)^2\big]. \tag{IA.3.3}$$

Hypotheses in our model. The pseudo-true $\theta^\star$ is the unique maximiser of M. By the definition of the Kullback–Leibler divergence and linearity of $\mathbb{E}_{F_J}$,

$$\mathrm{KL}(F_J\|F_\theta) = \mathbb{E}_{F_J}\Big[\log \tfrac{f(J)}{f_\theta(J)}\Big] = \mathbb{E}_{F_J}[\log f(J)] - \mathbb{E}_{F_J}[\log f_\theta(J)] = \mathbb{E}_{F_J}[\log f(J)] - M(\theta),$$

and $\mathbb{E}_{F_J}[\log f(J)]$ does not depend on θ, so minimising $\mathrm{KL}(F_J\|F_\theta)$ in θ is the same as maximising $M(\theta)$: $\arg\max_\theta M = \arg\min_\theta \mathrm{KL}(F_J\|F_\theta) = \{\theta^\star\}$ by Assumption 4(iv); interiority of $\theta^\star$, the C^2 and domination conditions, $H > 0$, and $B < \infty$ are Assumption 4(vi). All hypotheses hold, so (IA.3.3) applies and, given $N_t = n$, $\hat\theta_t - \theta^\star = O_{\mathbb{P}}(n^{-1/2})$. When the size family contains F_J (correct specification) $H = B = I(\theta^\star)$ and the asymptotic variance is $I(\theta^\star)^{-1}$; the order $n^{-1/2}$ is the same either way. The lemma asserts only this rate—an upper bound on the estimation error—so no efficiency or lower-bound claim is made or required.

From count to calendar time. Since $N_t \sim \mathrm{Poisson}(\lambda^\star t)$ has $\mathbb{E}[N_t] = \lambda^\star t$ and $N_t/(\lambda^\star t) \to 1$ a.s. as $t \to \infty$ (strong law for the Poisson process), substituting $n = N_t$ gives

$$\sqrt{N_t}\,\big(\hat\theta_t - \theta^\star\big) = O_{\mathbb{P}}(1), \qquad \text{hence} \qquad \hat\theta_t - \theta^\star = O_{\mathbb{P}}\big((\lambda^\star t)^{-1/2}\big),$$

which is (IA.3.2); the rate is governed by the realized disaster count $\lambda^\star t$, so it vanishes only as $t \to \infty$ and slowly when $\lambda^\star$ is small. For the intensity, the count N_t is a sufficient statistic for λ in the Poisson observation on $[0, t]$, with Fisher information t/λ. The maximum-likelihood estimator is $\hat\lambda_t = N_t/t$, and since $N_t \sim \mathrm{Poisson}(\lambda^\star t)$ has $\mathrm{Var}(N_t) = \lambda^\star t$,

$$\mathrm{Var}(\hat\lambda_t) = \mathrm{Var}\Big(\frac{N_t}{t}\Big) = \frac{\mathrm{Var}(N_t)}{t^2} = \frac{\lambda^\star t}{t^2} = \frac{\lambda^\star}{t}, \qquad \mathrm{Var}\Big(\frac{\hat\lambda_t - \lambda^\star}{\lambda^\star}\Big) = \frac{\mathrm{Var}(\hat\lambda_t)}{(\lambda^\star)^2} = \frac{1}{\lambda^\star t},$$

so the *relative* error is $(\hat{\lambda}_t - \lambda^\star)/\lambda^\star = O_{\mathbb{P}}\big((\lambda^\star t)^{-1/2}\big)$, the same disaster-count rate. The count and the marks are independent (Assumption 1), so the information is block-diagonal between λ and θ and the two are learned in parallel. Finally, by Assumption 4(i), $f_{\theta_i}(x) > 0$ for all $x \in \mathcal{J}$, so every realized $J_k \in \mathcal{J}$ has strictly positive density under any held model and is never a zero-probability event: no single mark *falsifies* a held model by impossibility. Test rejection is a separate matter—the deviance can cross its threshold at a disaster epoch, including the first, when the updated record leaves the typical set—so full support rules out falsification by impossibility, not rejection by the test. □

Appendix IA.4. Secondary belief-dynamics results

A long-run property of the rule. The pricing-state dynamics of Section 4 and the fragility reading of Section 5.D rest on one fact about the rule: revision never stops, even for a correct model. We record it here, self-contained.

Lemma IA.2 (Recurrence of revision): *Fix a tolerance $\varepsilon \in (0,1)$ and run the rule of Definition 1 under the physical law $\mathbb{P}$ (Assumption 1). For any held model $\phi = (\lambda, \theta) \in \Phi$:* (i) *the holding time τ^{rev} is finite $\mathbb{P}$-almost surely; and* (ii) *re-adopting the held model after each rejection, it is rejected infinitely often, $\mathbb{P}$-almost surely. Consequently, when the menu contains the data-generating model—$\lambda^\star \in \operatorname{int}\Lambda$ and $F_J = F_{\theta^\star}$ for some $\theta^\star \in \Theta$—a correct-on-average model $\phi = (\lambda^\star, \theta^\star)$ is, on a finite menu (or when the second-order prior has an atom at the truth), rejected and re-adopted forever: past a finite (a.s.) date every regime begins at the correct model and is nonetheless abandoned again.*

Proof. Hold $\phi = (\lambda, \theta)$ adopted at age 0; write $n_\Delta = N_{a+\Delta} - N_a$ for the since-adoption count over a window of length Δ, $M_\Delta := n_\Delta - \lambda^\star\Delta$ for the count centered at the physical rate, and $D_\Delta := D_{\phi,\Delta}(\mathcal{R})$ for the deviance (9). By Definition 1 the model is rejected at the first age at which its surprise $\bar{\varepsilon}(\mathcal{R}; \phi, \Delta)$ (10)—the predictive deviance tail under the held model's own law $Q_{\phi,\Delta}$—falls to ε or below; it therefore suffices to exhibit a finite age at which the realized deviance clears the bulk of that predictive law—the surprise then falls to ε or below at that age; subsequential clearing suffices for the first-passage rejection time, and no convergence of $\bar{\varepsilon}$ over calendar time is claimed (nor holds at the truth, where the deviance recurrently revisits typical values). The count channel does this: the size deviance in (9) is nonnegative, so $D_\Delta \geq D_\Delta^{\mathrm{c}} := \Delta\, d_P(n_\Delta/\Delta \| \lambda)$. The acceptance threshold it must clear is bounded in Δ: under $Q_{\phi,\Delta}$ the deviance D_Δ converges in distribution to a fixed chi-square-type law—its count part to a scaled χ_1^2 by the central limit theorem, its size part to a scaled χ_1^2 by Wilks's theorem—so the $(1-\varepsilon)$-quantile $r_{\phi,\Delta}(\varepsilon)$ stays bounded.

Step 1 (a finite exit, $\mathbb{P}$-a.s.). Under $\mathbb{P}$ the disaster count is a Poisson process of rate $\lambda^\star$, so the empirical rate $n_\Delta/\Delta \to \lambda^\star$ almost surely. If the held intensity is wrong, $\lambda \neq \lambda^\star$, then $D_\Delta^{\mathrm{c}} = \Delta\, d_P(n_\Delta/\Delta\|\lambda) \sim \Delta\, d_P(\lambda^\star\|\lambda)$ grows linearly, since $d_P(\lambda^\star\|\lambda) > 0$, and clears the bounded threshold in finite time. If the held intensity is correct, $\lambda = \lambda^\star$, the centered count $M_\Delta = n_\Delta - \lambda^\star\Delta$ has i.i.d. unit-age increments of mean 0 and variance $\lambda^\star$ and so obeys the law of the iterated logarithm (Hartman and Wintner, 1941), $\limsup_{\Delta\to\infty} M_\Delta/\sqrt{2\lambda^\star\Delta\log\log\Delta} = 1$ $\mathbb{P}$-a.s. Expanding d_P about $\lambda^\star$ ($d_P(\lambda^\star\|\lambda^\star) = 0$, $\partial_a d_P|_{\lambda^\star} = 0$, $\partial_a^2 d_P = 1/a$) gives $D_\Delta^{\mathrm{c}} \sim M_\Delta^2/(2\lambda^\star\Delta)$—the third-order remainder is $O(|M_\Delta/\Delta|^3)\,\Delta = O(M_\Delta^3/\Delta^2)$, of order $(\log\log\Delta)^{3/2}/\sqrt{\Delta} \to 0$ on the iterated-logarithm scale, since $\partial_a^3 d_P = -1/a^2$ is bounded near $\lambda^\star$ and $n_\Delta/\Delta \to \lambda^\star$ a.s.—, hence $\limsup_{\Delta\to\infty} D_\Delta^{\mathrm{c}}/\log\log\Delta = 1$, and the count deviance again clears the bounded threshold at a finite age—a fixed band crossed in finite time even at the truth, the boundary-crossing behind the foregone conclusion of fixed-level sequential testing (Robbins and Siegmund, 1970); the disaster-free record is the transparent instance, rejected at the deterministic calm deadline $\Delta^\star(\phi,\varepsilon)$ (B4), where $\Delta\, d_P(0\|\lambda) = \lambda\Delta$ crosses $r_{\phi,\Delta}(\varepsilon)$. In either case the realized record leaves the typical set—the surprise falls below the tolerance—at a finite age, so $\tau^{\mathrm{rev}} < \infty$ $\mathbb{P}$-a.s. This is (i).

Step 2 (infinitely often). Each rejection time is an $(\mathcal{F}_t)$-stopping time, so by the strong Markov property the model re-adopted at it monitors a fresh since-adoption record, and the

k-th holding time depends only on the model adopted at the k-th rejection; by Step 1 each is finite $\mathbb{P}$-a.s., correct model or not. The menu permits re-selection in either coordinate direction (Definition 1), so no model is absorbing and each rejection is followed by a further regime. Rejections arrive through two channels, and any bounded interval contains finitely many of each. A *calm* rejection on a holding whose since-adoption record is disaster-free cannot occur before the fresh-record deadline, bounded below uniformly—$\Delta^\star(\phi, \varepsilon) \geq \delta(\varepsilon) > 0$, continuous in ϕ on the compact menu with the zero-count surprise equal to one at age zero—and successive such rejections are separated by at least $\delta(\varepsilon)$, so an interval of length L holds at most $\lfloor L/\delta(\varepsilon) \rfloor + 1$ of them. A calm rejection on a holding whose record *contains* a disaster can come arbitrarily soon—a nonrejecting disaster can leave the record next to the boundary—but each such rejection is charged to the most recent disaster of its own holding, and successive holdings charge distinct disasters, so their number is bounded by the interval's disaster count. A *disaster* rejection consumes a disaster, and the disaster process is non-explosive, so an interval holds no more disaster rejections than its almost-surely finite disaster count. The holding time itself is *not* bounded below—an early disaster can leave the typical set at once—but the crash channel is rate-limited by non-explosiveness, not by a floor; the channels together place finitely many rejections in any bounded interval. The rejection times therefore increase to infinity: the held model is rejected infinitely often, $\mathbb{P}$-a.s. This is (ii).

Step 3 (correct-on-average). When $\lambda^\star \in \operatorname{int} \Lambda$ and $F_J = F_{\theta^\star}$ with $\theta^\star \in \Theta$, the model $(\lambda^\star, \theta^\star)$ lies in Φ, and the second-order posterior ρ_t (6)—the Bayesian update on the full marked-Poisson record over the identified menu (Assumption 4)—concentrates on it under $\mathbb{P}$—the re-selection posterior merges with the truth (Blackwell and Dubins, 1962)—so on a finite menu whose support contains the truth with positive prior mass, its mode $\mathfrak{m}(\rho_t)$ equals $(\lambda^\star, \theta^\star)$ for all t past a finite, a.s. date; on a continuum menu the posterior concentrates but finite-time mode equality fails, and the statement is the finite-menu case (a prior atom at the truth on a continuum menu would need a mode convention privileging atoms, which we do not adopt). From that date every re-selection (11) adopts the correct model, which by the correct-intensity case of Step 1 is rejected and re-adopted again. A correct-on-average model is thus revised forever. □

Appendix IA.5. Further properties of the belief multiplier

Remark IA.1 (The size channel can lift the binary magnitude ceiling (general case, not exercised)). This remark concerns the size channel, which the quantitative sections do not exercise (Scope, above). Under a two-point *size* menu $\{\theta_L, \theta_H\}$ at fixed intensity, $g_t(x)$ is a two-point mixture of the $1/\alpha$ powers $r(\phi, x)^{1/\alpha}$ of the two likelihood ratios, so the attainable risk-neutral tilt is bounded by the more fearful of the two. With a full-support size family $\{F_\theta\}$ whose held indices realize on an interval of Θ, $g_t(x)$ is a continuum-weighted aggregator that can tilt $f_t^{\mathbb{Q}}$ continuously across $\mathcal{J}$. Two caveats. First, this is a statement about the size indices θ_i; the continuum-support mechanism of Remark 4 concerns the calm-time *intensity* mode, not θ, so it does not by itself populate an interval of Θ—a size-dispersed newborn law would. Second, a continuum size family widens the *range* of attainable tilts but need not raise the supremum if the two-point menu already contains the most fearful model; the gain is continuity and dispersion-sensitivity—a comparison across designs, not an unboundedness claim (compactness of $\mathcal{J}$ and Θ keeps $f_t^{\mathbb{Q}}$ a proper density at all times).

Remark IA.2 (Which belief dynamics can move the fear factor). Two features of the bridge (18) bear on what kind of belief process can generate a time-varying fear factor.

(a) With constant fundamentals the fear factor is a pure belief object (proved). Volatility σ, the physical law $(\lambda^\star, F_J)$, and risk aversion α are constant (Assumptions 1 and 3: $\sigma, \lambda^\star, F_J$ from the former, α from the latter), so U^{RA} is constant and, by the bridge relation (19), $d \log U_t = d \log \kappa_t$: all movement in priced fear is belief movement. A representative agent, or any population holding correct beliefs, has $\kappa_t \equiv 1$ and a flat U_t. This fear factor is the

structural counterpart of the unspanned left-tail factor of Andersen et al. (2015), which moves the risk-neutral tail with no matching move in volatility or realized tail risk; here that property is exact and entirely belief-driven. What must be stripped out for a constant factor is *moving* beliefs, not merely heterogeneity: a correct-belief or fully merged population has κ_t constant ($\equiv 1$ when correct), but a single *learning* Bayesian—no heterogeneity at all—still moves the factor through her evolving posterior (Appendix IA.9). Cross-sectional dispersion is what keeps κ_t from merging to a constant; it is not the only belief process that can.

(b) Beliefs that merge make it flat (standard theory and Remark 6). If agents update on the common public filtration, their predictive laws merge (Blackwell and Dubins, 1962) and, under the identifiable family of Assumption 4, concentrate on the Kullback–Leibler model (Berk, 1966); the held likelihood ratios then share a common limit, κ_t converges to a constant, and the fear factor loses its time variation. This applies to a textbook single Bayesian and, as Remark 6 notes, to an infinitely-lived lazy agent alike: a time-varying fear factor requires beliefs that have not merged.

What is and is not claimed. We do *not* claim only lazy revision can move the factor. A Bayesian overlapping-generations population of as-yet-unconverged learners would also keep dispersion alive and update upward at disasters. The distinction is qualitative, not a theorem: lazy beliefs stay fixed within a tolerance and then jump discretely at rejection (Proposition 12), with scheduled relief in calm (Proposition 8), giving abrupt recurring spikes and persistent calm wedges—the regime-shift shape of the Bollerslev and Todorov (2011) fears series. The contrast with Bayesian learning is at the level of the individual belief path—fixed within tolerance, then a discrete jump, versus smooth between-disaster updating—not at disaster arrivals, where a rare-event Bayesian fear factor also jumps; and the persistence of the aggregate wedge owes as much to demographic turnover as to the learning rule. Non-representability here is relative to the named public filtration and the single-prior class (Section 4), not to Bayesians on enriched states. Whether the multiplier sits below one on average ($\kappa_t < 1$, the Backus et al. (2011) moderation—the inverted intensity $\lambda^\star \kappa_t$ below the physical $\lambda^\star$) and any quantitative match are calibration questions (Section 6).

Appendix IA.6. The pricing state beyond finite moments

Proposition IA.1 (Strengthening of Proposition 3: no continuous finite-dimensional sufficient statistic): *Fix t and the primitives. No* continuous *finite-dimensional statistic of the pricing state is sufficient for the priced disaster law: for every $n \geq 1$ and every continuous map $T : \mathcal{P}(\Phi) \to \mathbb{R}^n$ ($\mathcal{P}(\Phi)$ in the total-variation topology) there are admissible pricing states S, S' with $T(S) = T(S')$ yet $(U_t, f_t^{\mathbb{Q}}) \neq (U_t', f_t^{\mathbb{Q}\prime})$; equivalently, the instantaneous map $S \mapsto (U_t, f_t^{\mathbb{Q}})$ factors through no continuous $\mathbb{R}^n$-valued statistic. The claim is a* size-disagreement *statement on the full-support class $\mathcal{S}$ used in the proof (densities bounded below on Φ): dynamically reachable cross-sections include Dirac and finitely supported states outside $\mathcal{S}$, and the separation is proved on $\mathcal{S}$, asserting nothing beyond it. In the intensity channel the aggregator is the scalar G_t, and on a finite menu the instantaneous price objects are functions of the simplex shares—finite-dimensional; the impossibility concerns the infinite-dimensional size-mixture map. The continuity hypothesis is essential: under the weak topology $\mathcal{P}(\Phi)$ is Polish, hence a standard Borel space, so a measurable bijection with a subset of $\mathbb{R}$ furnishes a (discontinuous, non-constructive) measurable sufficient statistic of dimension one—the measurable structure here being the weak-Borel one, decoupled from the total-variation topology in which the continuity hypothesis is stated.*

Proof. Write $g_t = T_{\text{lin}} S$ for the linear map $T_{\text{lin}} : S \mapsto \int_\Phi H(\cdot, \phi)\, S(d\phi)$ with $H(x, \phi) = (\lambda/\lambda^\star)^{1/\alpha}(f_\theta(x)/f(x))^{1/\alpha}$, as in the proof of Proposition 3. The price determines the aggregator: if $f_t^{\mathbb{Q}\prime} = f_t^{\mathbb{Q}}$ then $g_t'^{\,\alpha} = c\, g_t^{\alpha}$ f-a.e. with $c = \mathbb{E}_f[e^{\alpha J} g_t'^{\,\alpha}]/\mathbb{E}_f[e^{\alpha J} g_t^{\alpha}]$, and $U_t' = U_t$ forces $c = 1$, whence $g_t' = g_t$ f-a.e. (both positive); so $(U_t, f_t^{\mathbb{Q}})$ and $g_t(\cdot)$ determine one another, and a statistic is sufficient for the price iff g_t factors through it.

Suppose some continuous $T : \mathcal{P}(\Phi) \to \mathbb{R}^n$ were sufficient, and put $m = n + 1$. By the spanning condition (Assumption 4(v)) choose $\phi_0, \ldots, \phi_m$ with $H(\cdot, \phi_0), \ldots, H(\cdot, \phi_m)$ linearly independent in $L^1(\mathcal{J}, f)$; the m differences $H(\cdot, \phi_k) - H(\cdot, \phi_0)$, $k = 1, \ldots, m$, are then linearly independent. Let S_0 have a continuous density bounded below by $c_0 > 0$ on the compact Φ (an admissible state), let β_j be a bounded-density probability measure concentrated near ϕ_j, and set $\eta_k := \beta_k - \beta_0$ (zero total mass). By continuity of H on the compact $\mathcal{J} \times \Phi$ (Assumption 4(i)), $T_{\mathrm{lin}}\eta_k \to H(\cdot, \phi_k) - H(\cdot, \phi_0)$ in $L^1(\mathcal{J}, f)$ as the β_j concentrate, so for tight enough β_j the $\{T_{\mathrm{lin}}\eta_k\}_{k=1}^m$ are linearly independent. On the open box $B = \{\varepsilon \in \mathbb{R}^m : \|\varepsilon\|_\infty < c_0/\sum_k \|\eta_k\|_\infty\}$ the affine map $\varepsilon \mapsto S_\varepsilon := S_0 + \sum_{k=1}^m \varepsilon_k \eta_k$ takes values in $\mathcal{P}(\Phi)$—its density is $\geq c_0 - \sum_k |\varepsilon_k| \, \|\eta_k\|_\infty \geq c_0 - \|\varepsilon\|_\infty \sum_k \|\eta_k\|_\infty > 0$ (the $\sum_k$ bound, not a per-term one, is what keeps this nonnegative for $m > 1$) and its mass is 1 (the η_k have zero mass)—and is continuous in total variation. Along it $g_t(S_\varepsilon) = g_t(S_0) + \sum_k \varepsilon_k \, T_{\mathrm{lin}}\eta_k$ is injective in ε, the $T_{\mathrm{lin}}\eta_k$ being independent. Sufficiency then gives, for $\varepsilon, \varepsilon' \in B$,

$$T(S_\varepsilon) = T(S_{\varepsilon'}) \implies (U_t, f_t^{\mathbb{Q}}) \text{ coincide} \implies g_t(S_\varepsilon) = g_t(S_{\varepsilon'}) \implies \varepsilon = \varepsilon',$$

so $\varepsilon \mapsto T(S_\varepsilon)$ is a continuous injection of the open set $B \subseteq \mathbb{R}^m$, $m = n + 1$, into $\mathbb{R}^n$. This contradicts Brouwer's invariance of domain, which forbids a continuous injection of a nonempty open subset of $\mathbb{R}^m$ into $\mathbb{R}^n$ when $m > n$. Hence no continuous finite-dimensional sufficient statistic exists.

For the final clause, equip $\mathcal{P}(\Phi)$ with the *weak* topology, under which—Φ being compact metric—it is Polish, hence a standard Borel space; the total-variation topology of the main argument is finer and, for uncountable Φ, not separable, so it is this weak-Borel structure that is standard. By the Borel isomorphism theorem $\mathcal{P}(\Phi)$ is then Borel-isomorphic to a subset of $\mathbb{R}$, and g_t is weak-Borel measurable—indeed weak-continuous, being integration of the bounded continuous kernel $H(\cdot, \phi)$ against the measure—so composing it with such an isomorphism writes it as a measurable function of one real coordinate, a measurable, onedimensional sufficient statistic. It is discontinuous and non-constructive, which is why the continuity hypothesis is not cosmetic. □

Appendix IA.7. Finite-menu reduction and higher-order expansions

Corollary IA.1 (Finite menu: instantaneous reduction, valuation on the record-resolved state): *Under the finite menu of Corollary 3, the instantaneous growth rate is a function of the finite share vector, $m_t = m(S_t)$ with $S_t \in \Delta^{K-1}$, and the price–dividend ratio solves the Feynman–Kac equation*

$$[\,\delta - m(S_t)\,]\, v(\Psi_t) = 1 + (\mathcal{A}_\Psi v)(\Psi_t), \tag{IA.7.1}$$

with $\mathcal{A}_\Psi$ the pricing *generator of the record-resolved state Ψ_t—the state generator of Proposition 3 with the disaster term reweighted by the* payoff-growth *jump of $Y_t = C_t^{1-\alpha}\Xi_t^{\alpha}$, namely $e^{-(1-\alpha)x} g_t(x)^{\alpha}$ (not the state-price-density jump $e^{\alpha x} g_t(x)^{\alpha}$), so that it carries the Y-weighted disaster-jump contribution $\lambda^\star \mathbb{E}_f[e^{-(1-\alpha)J} g_t(J)^{\alpha}\,(v(\Psi_t^J) - v(\Psi_{t-}))]$ rather than the bare state jump (proof in Appendix E of the main text). Even with K finite, Ψ_t is infinite-dimensional—the since-adoption records and ages are continuous—so (IA.7.1) does* not *reduce to a finite-dimensional-state equation in $S \in \Delta^{K-1}$ unless the cross-section is Markov in the shares alone, i.e. the revision/switch flow is autonomous in S. The lazy-revision rule violates this: the switch flow depends on the within-class record distribution (Proposition 12). The linear Feynman–Kac operator equation $[\delta - m(S)]v(S) = 1 + (\mathcal{A}_S v)(S)$ on the continuous simplex Δ^{K-1} holds only in the selection-only specialization with no record-dependent switching, where Ψ_t collapses to S_t; it is an operator equation on a continuous state space, not a finite system of algebraic equations—for $K = 2$, a scalar-state equation for a function of S^{hi} over a continuum.*

Proof of Corollary IA.1. Under the finite menu (Corollary 3), $\bar{\lambda}_t^S$ and g_t are functions of the share vector (eqs. (D)–(D18)), so $m_t = m(S_t)$ with $S_t \in \Delta^{K-1}$. Because ξ prices the consumption claim, $e^{-\delta t} Y_t\, v(\Psi_t) + \int_0^t e^{-\delta s} Y_s\, ds$ is a martingale, where $Y_t = C_t^{1-\alpha} \Xi_t^{\alpha}$ has local growth $m(S_t)$ (E14). Setting the drift of $e^{-\delta t} Y_t\, v(\Psi_t)$ plus the running term $e^{-\delta t} Y_t\, dt$ to zero and dividing by $e^{-\delta t} Y_t$,

$$-\delta\, v(\Psi_t) + m(S_t)\, v(\Psi_t) + (\mathcal{A}_\Psi v)(\Psi_t) + 1 = 0.$$

Here $\mathcal{A}_\Psi$ is the *pricing* (Y-weighted) generator of Ψ_t. Writing $\mathcal{A}_\Psi = \mathcal{A}_\Psi^{c} + \mathcal{A}_\Psi^{\mathrm{dis}}$ with $\mathcal{A}_\Psi^{c}$ the finite-variation selection/switch drift and the disaster part

$$\mathcal{A}_\Psi^{\mathrm{dis}} v = \lambda^\star \int_{\mathcal{J}} e^{-(1-\alpha)x} g_t(x)^\alpha \big(v(\Psi_t^x) - v(\Psi_{t-})\big) f(x)\, dx,$$

the kernel jump $Y_t/Y_{t-} = e^{-(1-\alpha)x} g_t(x)^\alpha$ weights the post-jump valuation, so the simultaneous Y–$v(\Psi)$ disaster covariation in the Itô product rule is carried *inside* $\mathcal{A}_\Psi$. Equivalently, relative to the physical state generator $\mathcal{A}_\Psi^P v = \mathcal{A}_\Psi^c v + \lambda^\star \int_{\mathcal{J}} (v(\Psi_t^x) - v(\Psi_{t-})) f\, dx$, the pricing generator adds the covariation correction $\lambda^\star \int_{\mathcal{J}} (e^{-(1-\alpha)x} g_t(x)^\alpha - 1)(v(\Psi_t^x) - v(\Psi_{t-})) f\, dx$. Multiplying through by -1 and collecting the two terms in $v(\Psi_t)$,

$$\delta\, v(\Psi_t) - m(S_t)\, v(\Psi_t) = 1 + (\mathcal{A}_\Psi v)(\Psi_t), \qquad \text{i.e.} \qquad [\,\delta - m(S_t)\,]\, v(\Psi_t) = 1 + (\mathcal{A}_\Psi v)(\Psi_t),$$

which is the Feynman–Kac equation (IA.7.1). The generator $\mathcal{A}_\Psi$ acts on functions of Ψ_t: the selection drift and the disaster reweighting are autonomous in S_t, but the switch term carries the within-class record distribution (Proposition 3(ii) through the hazard of Proposition 12), so $(\mathcal{A}_\Psi v)$ is in general not a function of S alone; moreover Ψ_t—with its continuous since-adoption records and ages—is infinite-dimensional even for finite K. If the switch flow is suppressed (selection-only), Ψ_t collapses to $S_t \in \Delta^{K-1}$, $\mathcal{A}_\Psi = \mathcal{A}_S$ is the finite-dimensional share generator read off the selection drift and disaster reweighting of Proposition 3, and (IA.7.1) becomes a finite-dimensional-state operator equation on the simplex, $[\delta - m(S)]v(S) = 1 + (\mathcal{A}_S v)(S)$ for $S \in \Delta^{K-1}$—an integro-differential equation in the $K-1$ share coordinates, not a finite linear system, since v remains a function on Δ^{K-1}—solvable under (T); for $K = 2$ a scalar integro-differential equation in S^{hi}. $\square$

Corollary IA.2 (First-order belief-heterogeneity expansion around the homogeneous economy): *Consider the homogeneous benchmark in which all investors hold the common model $\phi_0 = (\lambda_0, \theta_0)$—because at t they share the common held model ϕ_0, either as the common posterior mode under Bayesian updating from a shared prior (no tolerance threshold) or under a shared tolerance with synchronized switching (one threshold)—so the instantaneous cross-section is the point mass $S_t = \delta_{\phi_0}$ and Ψ_t is degenerate. Perturb it by a small wealth-weighted dispersion, $S_t = \delta_{\phi_0} + \varepsilon\, \eta$, with η an* admissible dispersion*—$\eta = \nu - \mu\, \delta_{\phi_0}$ for a finite positive measure ν on Φ of mass $\mu = \nu(\Phi)$ (weight shifted off the benchmark)—so $\int_\Phi \eta = 0$ and $S_t = (1-\varepsilon\mu)\, \delta_{\phi_0} + \varepsilon\, \nu$ is a probability measure for $0 \le \varepsilon \le 1/\mu$. Write $\langle \psi \rangle_\eta := \int_\Phi \psi\, d\eta$, $g_0(x) := r(\phi_0, x)^{1/\alpha}$, and $g_1(x) := \langle r(\cdot, x)^{1/\alpha} \rangle_\eta$. Then, to first order in ε,*

$$
\begin{aligned}
g_t(x) &= g_0(x) + \varepsilon\, g_1(x) && (\textit{exact}), && \text{(IA.7.2)}\\
\bar{\lambda}_t^S &= \lambda_0 + \varepsilon\, \langle \lambda \rangle_\eta && (\textit{exact}), && \text{(IA.7.3)}\\
U_t &= U_0 + \varepsilon\, U_1 + O(\varepsilon^2), && U_0 = \lambda_0 \mathbb{E}_{\theta_0}[e^{\alpha J}], \quad U_1 = \alpha \lambda^\star \mathbb{E}_f\big[e^{\alpha J} g_0^{\alpha-1} g_1\big], &&\\
r_t &= r_0 + \varepsilon\, (\langle \lambda \rangle_\eta - U_1) + O(\varepsilon^2), && r_0 = \delta + \alpha\mu_0 - \tfrac{1}{2}\alpha(1+\alpha)\sigma^2 + \lambda_0 - U_0, && \text{(IA.7.4)}
\end{aligned}
$$

where $\mathbb{E}_{\theta_0}$ is under f_{θ_0}. The price–dividend ratio satisfies $v_t = v_0 + O(\varepsilon)$ with the homogeneous value $v_0 = (\delta - m_0)^{-1}$ and $m_0 = (1-\alpha)\mu_0 - \frac{1}{2}\alpha(1-\alpha)\sigma^2 - (\lambda_0 - \lambda^\star) + \lambda^\star(\mathbb{E}_f[e^{-(1-\alpha)J} g_0^{\alpha}] - 1)$. The perturbation is an $O(\varepsilon)$ move of the record-resolved *state Ψ_t—the held-model marginal moving as displayed, the perturbed mass carrying fresh records—and the remainder is $O(\varepsilon)$ by Lemma 8; the displayed first-order coefficient is computed at the homogeneous benchmark, where it is closed form. The Bayesian and synchronized-switching readings fix the date-t state only. The zeroth order is the representative-agent rare-disaster economy under the belief* frozen *at ϕ_0—under readings in which the common belief moves, v_0 would inherit its law, so the closed form is the frozen-ϕ_0 benchmark's; the first-order terms are the leading correction from belief* *heterogeneity, linear in the dispersion* η.

Proof of Corollary IA.2. Instantaneous objects. The aggregator g_t and the mean intensity $\bar{\lambda}_t^S$ are linear in S_t (Definitions 2, 3). Substituting $S_t = \delta_{\phi_0} + \varepsilon\eta$,

$$\begin{aligned}
g_t(x) &= \int_\Phi r(\phi, x)^{1/\alpha}\big(\delta_{\phi_0} + \varepsilon\eta\big)(d\phi) \\
&= r(\phi_0, x)^{1/\alpha} + \varepsilon \int_\Phi r(\phi, x)^{1/\alpha}\eta(d\phi) = g_0(x) + \varepsilon\, g_1(x), \\
\bar{\lambda}_t^S &= \int_\Phi \lambda\big(\delta_{\phi_0} + \varepsilon\eta\big)(d\phi) = \lambda_0 + \varepsilon\, \langle\lambda\rangle_\eta,
\end{aligned} \tag{IA.7.5}$$

both exact (the functionals are linear, and δ_{ϕ_0} evaluates at ϕ_0); these are (IA.7.2)–(IA.7.3). For $U_t = \lambda^\star \mathbb{E}_f[e^{\alpha J} g_t(J)^\alpha]$ (17), expand the power by Taylor's theorem ($g_0 > 0$ and continuous on the compact $\mathcal{J}$, g_1 bounded), uniformly in x:

$$\begin{aligned}
g_t(x)^\alpha &= \big(g_0(x) + \varepsilon g_1(x)\big)^\alpha \\
&= g_0(x)^\alpha + \alpha\, g_0(x)^{\alpha-1}\, \varepsilon\, g_1(x) + O(\varepsilon^2).
\end{aligned}$$

Integrating against $\lambda^\star e^{\alpha x} f(x)\, dx$,

$$\begin{aligned}
U_t &= \lambda^\star \mathbb{E}_f[e^{\alpha J} g_0^\alpha] + \varepsilon\, \alpha \lambda^\star \mathbb{E}_f[e^{\alpha J} g_0^{\alpha-1} g_1] + O(\varepsilon^2) \\
&= U_0 + \varepsilon\, U_1 + O(\varepsilon^2),
\end{aligned} \tag{IA.7.6}$$

and the zeroth-order term simplifies with $g_0^\alpha = r(\phi_0, \cdot)$ and $r(\phi_0, x) = \lambda_0 f_{\theta_0}(x)/(\lambda^\star f(x))$:

$$\begin{aligned}
U_0 &= \lambda^\star \mathbb{E}_f[e^{\alpha J} r(\phi_0, J)] = \lambda^\star \int_{\mathcal{J}} e^{\alpha x} \frac{\lambda_0 f_{\theta_0}(x)}{\lambda^\star f(x)} f(x)\, dx \\
&= \lambda_0 \int_{\mathcal{J}} e^{\alpha x} f_{\theta_0}(x)\, dx = \lambda_0 \mathbb{E}_{\theta_0}[e^{\alpha J}].
\end{aligned}$$

Substituting (IA.7.5) and (IA.7.6) into the short rate (E1),

$$
\begin{aligned}
r_t &= \delta + \alpha\mu_0 - \tfrac{1}{2}\alpha(1+\alpha)\sigma^2 + \big(\lambda_0 + \varepsilon\langle\lambda\rangle_\eta\big) - \big(U_0 + \varepsilon U_1\big) + O(\varepsilon^2) \\
&= r_0 + \varepsilon\big(\langle\lambda\rangle_\eta - U_1\big) + O(\varepsilon^2),
\end{aligned}
$$

with $r_0 = \delta + \alpha\mu_0 - \frac{1}{2}\alpha(1+\alpha)\sigma^2 + \lambda_0 - U_0$; this is (IA.7.4).

Valuation. At the frozen homogeneous belief ($\varepsilon = 0$, ϕ_0 held fixed) $g_t = g_0$ is constant, so by (E14) the growth rate is the constant $m_0 = (1-\alpha)\mu_0 - \frac{1}{2}\alpha(1-\alpha)\sigma^2 - (\lambda_0 - \lambda^\star) + \lambda^\star(\mathbb{E}_f[e^{-(1-\alpha)J}g_0^\alpha] - 1)$. Because the primitive increments are i.i.d. (Assumption 1) and the jump factor $e^{-(1-\alpha)x}g_0(x)^\alpha$ is time-invariant and independent of Y—a fixed function of the size x, not a constant in x—its compensator contribution $\lambda^\star(\mathbb{E}_f[e^{-(1-\alpha)J}g_0^\alpha] - 1)$ is constant, so the increment dY_s/Y_s has conditional mean $m_0\,ds$ independent of Y_s/Y_t, so $\mathbb{E}_t[Y_s/Y_t]$ (conditioning on $\mathcal{F}_t$) solves

$$
\frac{d}{ds}\mathbb{E}_t[Y_s/Y_t] = m_0\,\mathbb{E}_t[Y_s/Y_t], \qquad \mathbb{E}_t[Y_t/Y_t] = 1,
$$

whose solution is

$$
\mathbb{E}_t[Y_s/Y_t] = e^{m_0(s-t)}. \tag{IA.7.7}
$$

Since $(C_s/C_t)^{1-\alpha}(\Xi_s/\Xi_t)^\alpha = Y_s/Y_t \geq 0$, Tonelli lets $\mathbb{E}_t$ pass inside the ds-integral of (E13), and

$$
\begin{aligned}
v_0 &= \int_t^\infty e^{-\delta(s-t)}\,\mathbb{E}_t[Y_s/Y_t]\,ds \\
&= \int_t^\infty e^{-\delta(s-t)}e^{m_0(s-t)}\,ds && \text{(IA.7.8)} \\
&= \int_0^\infty e^{-(\delta-m_0)u}\,du && \text{(IA.7.9)} \\
&= \frac{1}{\delta - m_0}, && \text{(IA.7.10)}
\end{aligned}
$$

line (IA.7.8) by (IA.7.7), (IA.7.9) substituting $u = s-t$ and combining $e^{-\delta u}e^{m_0 u} = e^{-(\delta-m_0)u}$, and (IA.7.10) evaluating the integral, which converges because $\delta - m_0 > 0$ ($m_0 \leq \sup m < \delta$ under (T)). For $\varepsilon > 0$ the perturbed path of the *full* record-resolved state $\Psi_t(\varepsilon)$ is an ε-perturbation of the benchmark—ε here the displaced *wealth mass*, the total-variation size of the state move, not an investor tolerance— (Lemma 8, under its margin $\delta > \kappa_L$ for the valuation step), and all three of its drivers are stable in ε: the selection drift (21) and the disaster reweighting are Lipschitz in the cross-section (Proposition 3), g_t enters through the bounded factor (15), and—the piece the marginal S_t alone does not control—the record-dependent *switch flow* carries only the $O(\varepsilon)$ mass that the perturbation displaces off the benchmark cohort. However the relocation *times* of that mass differ between the perturbed and benchmark economies, it is the same $O(\varepsilon)$ total mass that relocates, so its contribution to $\|\Psi_t(\varepsilon) - \Psi_t(0)\|_{\mathrm{TV}}$ is $O(\varepsilon)$ at every t, with no continuity of the calm-deadline map in the tolerance required (the count atoms make $\Delta^\star(\phi, \cdot)$ a step function of the tolerance, but the displaced mass is bounded whatever its switch times). With all three drivers thus $O(\varepsilon)$-stable, the total-variation rate is $\|\Psi_t(\varepsilon) - \Psi_t(0)\|_{\mathrm{TV}} = O(\varepsilon)$ uniformly on compact time intervals. The valuation is Lipschitz in the state in total variation by Lemma 8(iii), under its margin $\delta > \kappa_L$: the coupled economies share all but an $O(\|\Psi - \Psi'\|_{\mathrm{TV}})$ mass of investors pathwise, the shared part's payoff contributions cancel exactly, and the unshared part's contribution is integrated against the domination envelope of Lemma 7, giving $|v(\Psi) - v(\Psi')| \leq L\,\|\Psi - \Psi'\|_{\mathrm{TV}}$ for a finite L. Composing the two Lipschitz transfers—with (E25) dominating the time integrand for dominated convergence—gives $v_t = v_0 + O(\varepsilon)$. The $O(\varepsilon)$ correction solves the Feynman–Kac equation (IA.7.1) linearized at $\varepsilon = 0$, which is not closed form in general (Proposition 3). $\square$

Appendix IA.8. The option surface against its benchmarks

The benchmark decomposition of Proposition 5 sorts the disaster-option literature by which side of the surface each model occupies: every documented feature reads the current cross-section, the level any Markov pricing state reproduces, while the term-structure slope reads the survival records that no belief-sufficient model carries; among clock-bearing alternatives, the two signs of Remark 1 attribute the residual to testing.

Table I: What the option surface shows, and what reproduces it.

Object on the surface (anchor work)	What it reads	Reproduced by
Thin far tail on average (Backus et al., 2011)	the current cross-section: complacency $\kappa_t < 1$	any Markov or Bayesian benchmark
Investor-fears wedge, spiking after crises (Bollerslev and Todorov, 2011; Andersen et al., 2015)	a current-state factor	a Markov latent state
Fast-in, slow-out jump premium (self-exciting models)	a post-crisis spike, slow reversion	a clustering intensity
Leverage effect, volatility dynamics (Dew-Becker et al., 2025)	a smoothly filtered intensity	its own Markov pricing state
Raw path dependence of implied volatility, *weakening* with maturity (Guyon and Lekeufack, 2023; Andrès et al., 2024)	the recent path, already in the cross-section	any belief-sufficient model: its current state spans it
Slope of the tail term structure, entering at first order beyond the short end (this paper)	**the records under which beliefs have survived**	**no belief-sufficient model**
Slope residual under exogenous staggered clocks (Abel et al., 2013; Bacchetta and van Wincoop, 2010)	time since last adjustment, whatever the held model	escapes belief-sufficiency as well; lacks held-model-dependent deadlines and the disaster-fired upward wave (made exact by the main text's record-independent-clock proposition, Proposition 6 there)

Appendix IA.9. A representative-agent Bayesian benchmark

This appendix records the benchmark against which the lazy economy is read: a *single* representative agent who updates a common prior by Bayes' rule and prices under the resulting predictive law, with no surprise threshold, no lazy holding, and—being one agent—no cross-section. It isolates what the main mechanism adds by removing all three. All objects are derived here from the same primitives (Assumptions 1, 3, and 4); the verbal comparison with the body is made only after each benchmark object is derived.

Remark IA.3 (Heterogeneous Bayesian benchmark). With finitely many prior types $\rho_0^{(1)}, \dots, \rho_0^{(J)}$ held by wealth shares $w_t^{(j)}$, each type's posterior $\rho_t^{(j)}$ is a sufficient statistic for its own predictive law, and the pair $(\rho_t^{(j)}, w_t^{(j)})_{j \le J}$ is Markov: posteriors update by Bayes on the common record, and shares move by the replicator in the type predictive densities, both functionals of the current vector and the increment. Every forward price object is therefore measurable in the current belief state, the survival-history residual is identically zero, and the comparisons of Theorem 1(iv) extend to this heterogeneous benchmark verbatim.

The benchmark economy

Endow the economy of Assumption 1 with a single representative agent with common CRRA utility $u(c) = c^{1-\alpha}/(1-\alpha)$ and discount δ, who consumes the aggregate endowment C_t. For a model $\phi = (\lambda, \theta) \in \Phi$ let $M_t^\phi := \frac{d\mathbb{P}^\phi}{d\mathbb{P}}\big|_{\mathcal{F}_t}$ be its belief density (C1), a strictly positive $(\mathcal{F}_t, \mathbb{P})$-martingale with $\mathbb{E}[M_t^\phi] = 1$ (Appendix C of the main text, verbatim, with ϕ in place of the held model). The agent holds the common second-order prior ρ_0 of Assumption 2 and prices under the *Bayesian predictive measure*

$$\mathbb{P}^B := \int_\Phi \mathbb{P}^\phi \, \rho_0(d\phi), \qquad Z_t^B := \frac{d\mathbb{P}^B}{d\mathbb{P}}\Big|_{\mathcal{F}_t} = \int_\Phi M_t^\phi \, \rho_0(d\phi). \tag{IA.9.1}$$

Lemma IA.3 (The Bayesian belief density): *Z^B is a strictly positive $(\mathcal{F}_t, \mathbb{P})$-martingale with $Z_0^B = 1$ and $\mathbb{E}[Z_t^B] = 1$, and the implied posterior is the common posterior of (6): writing $\rho_t^B(d\phi) := M_t^\phi \rho_0(d\phi)/Z_t^B$,*

$$\rho_t^B = \rho_t \text{ (of (6))}, \qquad \left.\frac{Z_t^B}{Z_{t-}^B}\right|_{J=x} = \frac{\int_\Phi M_{t-}^\phi \, r(\phi, x) \, \rho_0(d\phi)}{\int_\Phi M_{t-}^\phi \, \rho_0(d\phi)} = \mathbb{E}_{\rho_{t-}^B}\big[r(\cdot, x)\big], \tag{IA.9.2}$$

the posterior-mean disaster likelihood ratio, with $r(\phi, x) = \lambda f_\theta(x)/(\lambda^\star f(x))$ as in (23).

Proof. Each M^ϕ is a positive $\mathbb{P}$-martingale with $\mathbb{E}[M_t^\phi] = 1$ (Appendix C of the main text); by Tonelli the mixture $Z_t^B = \int_\Phi M_t^\phi \rho_0(d\phi) \geq 0$ has, for $s < t$, $\mathbb{E}[Z_t^B \mid \mathcal{F}_s] = \int_\Phi \mathbb{E}[M_t^\phi \mid \mathcal{F}_s]\rho_0(d\phi) = \int_\Phi M_s^\phi \rho_0(d\phi) = Z_s^B$ (the interchange justified by positivity and $\mathbb{E}\int M_t^\phi \rho_0(d\phi) = \int \mathbb{E}[M_t^\phi]\rho_0(d\phi) = 1 < \infty$), so Z^B is a martingale with $\mathbb{E}[Z_t^B] = 1$; $Z_0^B = \int M_0^\phi \rho_0 = 1$, and strict positivity holds because each $M_t^\phi > 0$ (Appendix C of the main text) and ρ_0 has full support (Assumption 2). The normalized measure $\rho_t^B(d\phi) = M_t^\phi \rho_0(d\phi)/Z_t^B$ is, by (C1) and Bayes' rule, proportional to $\rho_0(\phi) e^{-\lambda t} \lambda^{N_t} \prod_k f_\theta(J_k)$, which is (6); hence $\rho_t^B = \rho_t$. At a disaster of size x, $M_t^\phi = M_{t-}^\phi r(\phi, x)$ (C1), so $Z_t^B / Z_{t-}^B = \int M_{t-}^\phi r(\phi, x) \rho_0 / \int M_{t-}^\phi \rho_0 = \mathbb{E}_{\rho_{t-}^B}[r(\cdot, x)]$, which is (IA.9.2). □

The benchmark kernel and fear factor

Proposition IA.2 (Benchmark state-price density and fear factor): *In the representative-agent Bayesian economy the unique state-price density is*

$$\xi_t^B = e^{-\delta t} C_t^{-\alpha} Z_t^B, \tag{IA.9.3}$$

and the benchmark risk-neutral disaster law tilts the physical law by risk aversion and by the predictable (pre-jump) posterior-mean likelihood ratio: with $\zeta_t^B(x) := \mathbb{E}_{\rho_{t-}^B}[r(\cdot, x)]$*,*

$$\nu^{\mathbb{Q},B}(dt, dx) = \lambda^\star e^{\alpha x}\zeta_t^B(x)\,dt\,F_J(dx), \qquad U_t^B = U^{\mathrm{RA}}\kappa_t^B, \qquad \kappa_t^B = \frac{\mathbb{E}_f[e^{\alpha J}\zeta_t^B(J)]}{\mathbb{E}_f[e^{\alpha J}]}. \quad \text{(IA.9.4)}$$

Proof. The agent values an $\mathcal{F}_t$-payoff X at $\mathbb{E}^{\mathbb{P}}[\,e^{-\delta t}u'(C_t)\,(d\mathbb{P}^B/d\mathbb{P})X\,] = \mathbb{E}^{\mathbb{P}^B}[e^{-\delta t}C_t^{-\alpha}X]$ (change of measure by (IA.9.1)), so the state-price density under $\mathbb{P}$ is $\xi_t^B = e^{-\delta t}u'(C_t)Z_t^B = e^{-\delta t}C_t^{-\alpha}Z_t^B$; market clearing is automatic (one agent consumes C_t) and uniqueness is the strict concavity of u. This is the single-agent case of (12): with one belief the aggregator $\Xi_t = (Z_t^B)^{1/\alpha}$ (the multiplier cancels), so $\Xi_t^\alpha = Z_t^B$ and the $\frac{1}{\alpha}$- and α-powers compose to the identity—the power-mean distortion of (12) is absent. The ξ^B jump ratio at a disaster of size x is, by (IA.9.3), $C_t^{-\alpha}/C_{t-}^{-\alpha}$ times Z_t^B/Z_{t-}^B, i.e. $e^{\alpha x}\zeta_t^B(x)$ by (IA.9.2); the change-of-measure theorem (Appendix C of the main text), applied to the risk-neutral density $\xi_t^B e^{\int_0^t r_s^B\,ds}/\xi_0^B$—a positive $\mathbb{P}$-martingale by the definition of r^B, with disaster-jump ratio $e^{\alpha x}\zeta_t^B(x)$ by (IA.9.2), ζ_t^B predictable, strictly positive, and bounded by $\overline{r}$ via (C1)—gives $\nu^{\mathbb{Q},B} = e^{\alpha x}\zeta_t^B(x)\nu$, the first equation of (IA.9.4). Integrating the size out, $U_t^B = \lambda^\star\mathbb{E}_f[e^{\alpha J}\zeta_t^B(J)] = U^{\mathrm{RA}}\kappa_t^B$ with κ_t^B as displayed, exactly as in Proposition 2 but with $\zeta_t^B = \mathbb{E}_{\rho_{t-}^B}[r(\cdot,\cdot)]$ in place of g_t^α. □

Proposition IA.3 (Merging: the benchmark wedge has no persistent variation): *Under Assumptions 1 and 4, with the prior* ρ_0 *assigning positive mass to every Kullback–Leibler neighborhood of* $\phi^\star$*, as* $t \to \infty$ *the posterior concentrates on the Kullback–Leibler model,* $\rho_t^B \Rightarrow \delta_{\phi^\star}$ $\mathbb{P}$*-a.s. (with* $\phi^\star = (\lambda^\star, \theta^\star)$ *of (A3)), so*

$$\zeta_t^B(x) \longrightarrow r(\phi^\star, x) = \frac{f_{\theta^\star}(x)}{f(x)} \quad \textit{for each } x, \qquad \kappa_t^B \longrightarrow \kappa_\infty^B := \frac{\mathbb{E}_f[e^{\alpha J}f_{\theta^\star}(J)/f(J)]}{\mathbb{E}_f[e^{\alpha J}]}, \quad \text{(IA.9.5)}$$

a constant; if the size family is correctly specified ($f_{\theta^\star} = f$) *then* $\zeta_t^B \to 1$*,* $\kappa_\infty^B = 1$*, and* $U_t^B \to U^{\mathrm{RA}}$*. The approach is smooth and shrinking: the predictive moments converge at the rare-disaster rate,* $\zeta_t^B(x) - r(\phi^\star, x) = O_{\mathbb{P}}\big((\lambda^\star t)^{-1/2}\big)$*, so the increments of* $\log U_t^B$*—its calm drift* *and its disaster-jump size—vanish in probability at that rate,* $d \log U_t^B \to 0$.

Proof. Concentration follows from Berk's theorem applied directly to the single posterior $\rho_t^B = \rho_t$ (Lemma IA.3). Three facts hold. The record splits into i.i.d. unit-interval blocks by the stationary independent increments of Assumption 1; the family is identifiable with a unique Kullback–Leibler minimiser $\phi^\star$ (Assumption 4(iii)–(iv), Lemma 3); and the per-block divergence $\mathcal{K}(\phi)$ is finite, with that unique minimiser (Lemma 3). These are exactly the hypotheses of Berk (1966), whose theorem—for an identifiable dominated family observed i.i.d. with finite divergence to the family—gives posterior concentration on the divergence-minimising parameter, $\rho_t^B \Rightarrow \delta_{\phi^\star}$ $\mathbb{P}$-a.s. (The same three facts drive the wealth-weighted concentration of Proposition 19(ii); here they are applied to the single posterior, requiring none of the demographic machinery.) Since $x \mapsto r(\phi, x)$ is continuous and bounded on the compact $\mathcal{J} \times \Phi$ (Assumption 4(i)), weak convergence of ρ_t^B (hence of its left limit ρ_{t-}^B) to $\delta_{\phi^\star}$ gives $\zeta_t^B(x) = \mathbb{E}_{\rho_{t-}^B}[r(\cdot, x)] \to r(\phi^\star, x)$ for each x, and bounded convergence gives $\kappa_t^B \to \kappa_\infty^B$ of (IA.9.5); $r(\phi^\star, x) = (\lambda^\star/\lambda^\star)(f_{\theta^\star}(x)/f(x)) = f_{\theta^\star}(x)/f(x)$, which is 1 when $f_{\theta^\star} = f$. The rate is the posterior-contraction (Bernstein–von Mises) counterpart of Lemma IA.1(iii): under the smoothness of Assumption 4(vi) the posterior contracts on the Kullback–Leibler point $\phi^\star$ at the maximum-likelihood rate—by the misspecified Bernstein–von Mises theorem (Kleijn and van der Vaart, 2012) when $f_{\theta^\star} \neq f$, and the classical one (van der Vaart, 1998) when $f_{\theta^\star} = f$—itself governed by the disaster count $N_t \sim \text{Poisson}(\lambda^\star t)$, so the posterior-predictive mean ζ_t^B concentrates at order $(\lambda^\star t)^{-1/2}$—this requires the contraction rate, not merely the maximum-likelihood rate that Lemma IA.1(iii) supplies; hence $d \log U_t^B = d \log \kappa_t^B \to 0$ at that rate—a loose upper bound; under the regular expansion the calm drift and one-disaster update are typically $O_{\mathbb{P}}((\lambda^\star t)^{-1})$, the level rate sufficing here—the disaster-jump increment is bounded by pre- and post-update contraction, and the calm-drift claim uses only this upper bound. $\square$

What the lazy economy adds

With the benchmark objects derived, the contrast is exact and lives in four places.

(i) No endogenous disagreement. The benchmark has one agent and no surprise threshold, so there is no cross-section and nothing to disperse; $\mathrm{Var}(\cdot) = 0$ identically. The lazy economy manufactures a non-degenerate cross-section from an identical start through the threshold heterogeneity and the revision rule (Proposition 4)—disagreement that is a *result*, not a primitive, and that the single Bayesian cannot have.

(ii) Merging kills the wedge; lazy revision need not. Here κ_t^B converges to the constant κ_∞^B (Proposition IA.3): the belief-driven *variation* of the fear factor dies out as the predictive laws merge (Blackwell and Dubins, 1962; Kalai and Lehrer, 1993), and with correct specification the fear factor relaxes to the frictionless benchmark U^{RA}. This is the single-agent face of Remark IA.2(b). The lazy economy escapes it only through demographic turnover, which keeps the cross-section non-degenerate and sustains a *stationary, fluctuating* wedge (Proposition 19); an infinitely-lived lazy agent alone merges just as the Bayesian does (Remark 6).

(iii) Shrinking-and-merging versus threshold dynamics. The benchmark Bayesian belief is itself càdlàg—it jumps at each disaster and moves continuously between—but its variation *shrinks and merges*, vanishing at rate $(\lambda^\star t)^{-1/2}$ (Proposition IA.3). The lazy economy differs not by carrying jumps the benchmark lacks, but in two structural ways: *individual* held beliefs stay flat within a tolerance and then jump at a threshold-triggered rejection, and the *aggregate* fear factor does not vanish—it drifts in calm through wealth-reallocating selection and, under turnover, settles at a non-degenerate stationary cross-section with persistent wedges (Propositions 6, 17). The contrast is shrinking-and-merging versus threshold holding/switching with a non-vanishing stationary spread, not continuity versus jumps.

(iv) Arithmetic posterior mean versus wealth-weighted power mean. The benchmark prices the posterior *arithmetic* mean $\zeta_t^B(x) = \mathbb{E}_{\rho_{t-}^B}[r(\cdot, x)]$—mark-dependent in the joint channel, a scalar in the intensity channel— (Proposition IA.2); the lazy economy prices the wealth-

weighted $\frac{1}{\alpha}$-*power* mean $g_t(x)^\alpha = \big(\mathbb{E}_{S_{t-}}[r(\cdot,x)^{1/\alpha}]\big)^\alpha$ (Proposition 1, Proposition 2). The power-mean gap below the wealth-weighted arithmetic mean—generic under disagreement, and the source of complacency in calm spells—is intrinsically a heterogeneity phenomenon: it closes (power mean equals arithmetic mean) for the single agent and is therefore invisible in the benchmark.

Appendix IA.10. Numerical stability of the quantitative appendix

This section reports three further exercises for the quantitative appendix of the main text: the slack in the weighted dominance condition of the witness valuation step of Appendix B of the main text, the effect of replacing the transport scheme's split-branch testing criterion by the exact two-sided criterion, and the seed and Monte Carlo stability of the ensemble headline statistics. Every number is generated by the companion script `wp_stability.py` from the model laws at the calibration; nothing is tuned to the printed values.

The weighted dominance at the calibration. The valuation step concludes that the two witness states carry two values whenever the forcing kernel

$$K(u) \;=\; -\,v(\Psi_u)\,\delta\bar{\lambda}_u \;+\; \lambda^\star\,\alpha\,G_u^{\alpha-1}\,\mathbb{E}_f\big[e^{(\alpha-1)J}\big]\,v(\Psi_u^\diamond)\,\delta G_u$$

keeps one sign on the swept range, met with slack at the calibration. We verify the assertion on the aged witness at $t^\star = 44$—a reduced tolerance grid, every third atom with block weights renormalized, transport step 0.02, landing at $S^{hi} = 0.200$ against 0.207 at full resolution—sweeping the sixteen-year calm-conditional path along which S^{hi} declines from 0.200 to 0.068. The age-advance derivatives are finite differences at advances $a \in \{0.02, 0.01\}$; the raw levels are step functions of the atomic transport, but the channel ratio $\delta G_u/\delta\bar{\lambda}_u$ agrees across the two advances to numerical zero on all thirteen substantive grid points. The values come from the renewal representation, truncated at twenty-two years with the frozen-state tail closure and iterated twice; the second iteration moves the level by at most 11.1%, and the conservative-rung ratio below takes the smaller post-jump value across iteration rungs against the larger current value.

u	S_u^{hi}	$10^3\delta\bar{\lambda}_u$	$10^3\delta G_u$	$v(\Psi_u)$	$v(\Psi_u^\diamond)$	weighted ratio	unweighted
0	0.200	0.0000	0.0000	16.56	27.65	—	—
3	0.151	0.0040	0.0389	16.16	26.82	1.968	1.186
5	0.127	0.0064	0.0622	15.93	26.38	1.895	1.144
8	0.089	0.0043	0.0416	15.64	25.62	1.766	1.079
10	0.086	0.0040	0.0389	15.46	25.10	1.740	1.072
12	0.076	0.0034	0.0333	15.29	24.71	1.708	1.057
15	0.068	0.0028	0.0275	15.02	23.92	1.659	1.042

Table II: The witness kernel along the calm-conditional sweep. The weighted ratio is the G-channel term over the drift term of $K(u)$; the unweighted column drops the valuation weights. The origin row is the theorem's exact zero.

The kernel is positive at every substantive point. The minimum weighted ratio is 1.659, against an unweighted minimum of 1.042: the post-disaster value runs some 65% above the current value, so the valuation weights multiply the thin unweighted margin roughly fivefold. The conservative-rung ratio bottoms at 1.250, and the discounted integrated forcing is $+3.06\times10^{-4}$ with a wrong-sign share of exactly zero. The main text's "met with slack at our calibration" is confirmed as written.

The exact two-sided criterion against the split-branch scheme. The transport evaluates calm deadlines by the lower-tail branch of the count criterion and disaster triggers by the upper tail; the criterion of the main text is the full two-sided deviance tail. We implement the exact criterion from precomputed crossing tables—deviance-ordered Poisson tail sums on a 0.02-year grid to two hundred years, since-adoption counts to thirty—and rerun the deterministic witness and a paired ensemble of sixty histories on common disaster paths. Under the split-branch control, $S^{hi}(t^\star) = 0.2067$ with forward-tail wedges 13.55% and 16.17% at horizons six and twelve, the landmarks of the main text reproduced exactly. Under the exact criterion the deadlines lengthen, the scheduled relief slows, more fearful mass survives to $t^\star$ ($S^{hi} = 0.2378$), and the wedges rise to 14.48% and 18.76%. The paired ensemble tells the same story: the time-since loading rises (+0.00357 against +0.00267 per year), the paired twelve-year wedge difference averages +1.22 percentage points, the largest matched-pair slope gap jumps from 4.0 to 19.4 percentage points, and the post-disaster R^2 on the spot state alone falls from 0.65 to 0.24, reaching 0.50 with all record proxies. Signs and orderings are preserved throughout; the record channel is stronger and less summarizable by the coarse proxies under the exact criterion, so the scheme-tagged headline numbers of the main text are conservative for its claim.

Seed and Monte Carlo stability. We rerun the ensemble under the scheme of the main text at seeds 11, 23, and 47, one hundred fifty histories each, with nested seventy-five-path subsamples and a pooled set of four hundred fifty.

seed	M	$R^2(S)$	$R^2(S, t_{\text{since}})$	$b_{t_{\text{since}}}$	IQR effect	pairs $[n, \text{med}, q_{90}, \max]$
11	75	0.574	0.925	+0.00293	+3.6pp	[35, 1.8, 5.5, 5.8]
11	150	0.641	0.922	+0.00280	+5.3pp	[103, 2.1, 5.0, 5.8]
23	75	0.594	0.888	+0.00312	+5.5pp	[33, 1.5, 4.2, 5.5]
23	150	0.644	0.917	+0.00287	+5.1pp	[122, 1.4, 4.5, 6.4]
47	75	0.576	0.926	+0.00299	+3.2pp	[24, 2.2, 5.1, 6.6]
47	150	0.638	0.909	+0.00306	+4.5pp	[106, 2.2, 4.8, 6.6]
pooled	450	0.642	0.915	+0.00290	+4.9pp	[1060, 2.0, 4.7, 7.5]

Table III: Headline stability across seeds and M. Pairs are all near-matched pairs ($|\Delta S^{hi}| < 0.02$) in the comparison window, gaps in percentage points of $F(12)/U$.

The post-disaster ladder $0.64 \to 0.92$ of the main text holds at every seed to within ± 0.003 and ± 0.013; the time-since loading varies by about ten percent; the interquartile effect spans 4.5–5.3 percentage points at $M = 150$, with the printed +5.3 the seed-11 draw. The near-matched-pairs quadruple printed in the main text, [103, 2.1, 5.0, 5.8], is likewise the seed-11 realization, with cross-seed ranges as tabulated; the deep-calm wedge median is 1.08% at every seed. The seventy-five-path rows wobble visibly, confirming $M = 150$ as adequate for the reported precision. The pooled short-rate statistics are mean 7.48%, standard deviation 257 basis points, range [0.72, 9.57]% against the printed 7.4%, 237 basis points, [0.44, 9.60]%; the path grid sampled here is two years, which misses the extreme trough a finer grid records—a sampling-grid effect, not a model one.

Scope. The kernel exercise runs on the reduced tolerance grid at transport step 0.02 with the truncated, twice-iterated renewal valuation; its conclusions are stated with the conservative-rung bound. The criterion comparison holds the disaster paths fixed across schemes. The seed grid covers three seeds at the reported M; the subsample rows are nested, not independent. The code and raw outputs are available from the author.

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.

Aït-Sahalia, Yacine, and Jean Jacod, 2009, Testing for jumps in a discretely observed process, *The Annals of Statistics* 37, 184–222.

Andersen, Torben G., Nicola Fusari, and Viktor Todorov, 2015, 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.

Anscombe, F. J., 1954, Fixed-sample-size analysis of sequential observations, *Biometrics* 10, 89–100.

Armitage, Peter, C. K. McPherson, and B. C. Rowe, 1969, Repeated significance tests on accumulating data, *Journal of the Royal Statistical Society. Series A (General)* 132, 235–244.

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.

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

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

Blackwell, David, and Lester Dubins, 1962, Merging of opinions with increasing information, *Annals of Mathematical Statistics* 33, 882–886.

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.

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

Cho, In-Koo, and Kenneth Kasa, 2008, Learning dynamics and endogenous currency crises, *Macroeconomic Dynamics* 12, 257–285.

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

Costa, Oswaldo L. V., and François Dufour, 2008, Stability and ergodicity of piecewise deterministic Markov processes, *SIAM Journal on Control and Optimization* 47, 1053–1077.

Cover, Thomas M., and Joy A. Thomas, 2006, *Elements of Information Theory*, second edition (Wiley-Interscience, Hoboken, NJ).

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.

Davis, M. H. A., 1984, Piecewise-deterministic markov processes: A general class of non-diffusion stochastic models, *Journal of the Royal Statistical Society, Series B* 46, 353–388.

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

Dembo, Amir, and Ofer Zeitouni, 1998, *Large Deviations Techniques and Applications*, second edition (Springer, New York).

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.

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.

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.

Hairer, Martin, Jonathan C. Mattingly, and Michael Scheutzow, 2011, Asymptotic coupling and a general form of Harris' theorem with applications to stochastic delay equations, *Probability Theory and Related Fields* 149, 223–259.

Hartman, Philip, and Aurel Wintner, 1941, On the law of the iterated logarithm, *American Journal of Mathematics* 63, 169–176.

Jacobsen, Martin, 2006, *Point Process Theory and Applications: Marked Point and Piecewise Deterministic Processes* (Birkhäuser, Boston).

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

Judd, Kenneth L., 1985, The law of large numbers with a continuum of IID random variables, *Journal of Economic Theory* 35, 19–25.

Kadane, Joseph B., Mark J. Schervish, and Teddy Seidenfeld, 1996, Reasoning to a foregone conclusion, *Journal of the American Statistical Association* 91, 1228–1235.

Kalai, Ehud, and Ehud Lehrer, 1993, Rational learning leads to Nash equilibrium, *Econometrica* 61, 1019–1045.

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.

Kasa, Kenneth, 2004, Learning, large deviations, and recurrent currency crises, *International Economic Review* 45, 141–173.

Kleijn, Bas J. K., and Aad W. van der Vaart, 2012, The Bernstein–von Mises theorem under misspecification, *Electronic Journal of Statistics* 6, 354–381.

Küchler, Uwe, and Steffen L. Lauritzen, 1989, Exponential families, extreme point models and minimal space-time invariant functions for stochastic processes with stationary and independent increments, *Scandinavian Journal of Statistics* 16, 237–261.

Küchler, Uwe, and Michael Sørensen, 1989, Exponential families of stochastic processes: A unifying semimartingale approach, *International Statistical Review* 57, 123–144.

Küchler, Uwe, and Michael Sørensen, 1998, On exponential families of Markov processes, *Journal of Statistical Planning and Inference* 66, 3–19.

Mancini, Cecilia, 2009, Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps, *Scandinavian Journal of Statistics* 36, 270–296.

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

Meyn, Sean, and Richard L. Tweedie, 2009, *Markov Chains and Stochastic Stability*, second edition (Cambridge University Press).

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

Robbins, Herbert, 1970, Statistical methods related to the law of the iterated logarithm, *Annals of Mathematical Statistics* 41, 1397–1409.

Robbins, Herbert, and David Siegmund, 1970, Boundary crossing probabilities for the Wiener process and sample sums, *Annals of Mathematical Statistics* 41, 1410–1429.

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

Sun, Yeneng, 2006, The exact law of large numbers via Fubini extension and characterization of insurable risks, *Journal of Economic Theory* 126, 31–69.

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

van der Vaart, Aad W., 1998, *Asymptotic Statistics* (Cambridge University Press, Cambridge).

Wald, Abraham, 1949, Note on the consistency of the maximum likelihood estimate, *The Annals of Mathematical Statistics* 20, 595–601.

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

White, Halbert, 1982, Maximum likelihood estimation of misspecified models, *Econometrica* 50, 1–25.

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.