Online Appendix to "Endogenous Belief Hazards"

This appendix collects the population extensions omitted from the main paper: the wedge between consensus and the representative individual (§OA.1), and the two economies—disaster pricing and overlapping generations—that instantiate the cross-section (§OA.2). Numbering is self-contained (results and equations carry the prefix IA); cross-references without that prefix point to the main paper. The main-paper results invoked here are proved in the main-paper appendices.

OA.1 Consensus is not the individual

Because revision-proneness is heterogeneous, the consensus belief is a selected, composition-weighted object that need not move like any investor in it. Two statistics of the cross-section m_t summarize it—where it *sits*, the survivor-weighted mean, and how far it *spreads*, the dispersion—and they are separate objects that can move in opposite directions. We take them in turn.

OA.1.0.1 Where the consensus sits

We take *monotonicity of the held belief in tolerance* as a maintained hypothesis—under a common newborn model and an ordered menu, more revision-prone types lag at more complacent (resp. fearful) rates under a complacent (resp. fearful) start—assumed rather than derived, since the exit-side direction of Proposition 4 holds only conditionally on the surviving path and an unconditional monotonicity is not available. Under it the composition shift signs the survivor-weighted *mean*.

Proposition OA.1 (Selection wedge)**.** *Order types by tolerance, narrower bands (larger α) being more revision-prone, and* assume the held belief is monotone in the tolerance label *(the maintained hypothesis above). Among the* unrevised survivors*—those still holding their adopted model—the ex-ante expected composition* $\mathbb{E}[G_a]$ *is selected toward wider-band, slower-hazard types, so this sub-population's mean belief differs from the average-tolerance investor's by a frailty term, pulled under the maintained monotonicity toward the beliefs those slower-hazard types hold. The ordering is by survival, not by the instantaneous rate—a frailty covariance between the held belief and survival across tolerance types, signed by their co-monotonicity. The full consensus* $\langle w, m_t \rangle$ *adds the revised entrants, who have re-selected through* J*, and so is* not *the survivor sub-population; the frailty ordering is a statement about that survivor component alone.*

This is a frailty effect in the demographic sense (Vaupel et al., 1979): the frail (revision-prone) exit first, the survivors are disproportionately robust, and the average belief drifts with composition rather than with any individual's learning. Wealth-weighted for pricing, the same logic runs through relative wealth—the market-selection channel (Sandroni, 2000; Blume and Easley, 2006)—which is why the selection term in (14) is wealth-weighted only.

OA.1.0.2 How far it spreads

The wedge concerns where the consensus *sits*; its *spread* is a separate object—and that spread, not the level, is where the mechanism makes direct contact with a documented business-cycle moment.

Proposition OA.2 (Disagreement dynamics)**.** *Let* $D_t := \mathrm{Var}_{m_t}(w)$ *be the cross-sectional dispersion of beliefs. Between common surprises* D_t *evolves continuously in the trend channel; in the crash channel it can additionally carry* predictable atoms *where a disaster-synchronized cohort reaches a shared floor-step at a tolerance-independent* *rank-swap date (Lemma D.1), which an atomless G does not smooth. Such a date moves D_t by an atom only when a positive G-mass of* cutoff-relevant *types crosses then* and *the re-selection changes the dispersion; a shared floor-step date alone does not imply a jump in D_t. When it does, the sign is set by the cohort's position relative to the consensus, as for the disaster jump below, not fixed downward. At a common disaster, under* (B+) *and an overshoot condition—the revising boundary mass lands* farther *from the stayers' mean than it sat before, net of any change in within-group variance—D_t jumps up: disagreement widens at that disaster. When the overshoot condition fails the same crossing narrows it, so widening is the conditional case, not a law. The dispersion D_t is distinct from—and can move opposite to—the consensus* level $\langle w, m_t \rangle$.

Disagreement is thus endogenous in the crash channel and *countercyclical under the overshoot condition*: when the revising boundary mass overshoots, a disaster does not only move average perceived risk, it *spreads* it across investors, because only those at the boundary revise and they overshoot the rest; when the overshoot fails the same disaster compresses dispersion. The overshoot is not a free parameter—in the crash channel it reduces to $\hat{\lambda}^{+} + \hat{\lambda}^{-} > 2\bar{w}^{\text{stay}}$, a condition on the menu rungs and the endogenous consensus (Appendix E, condition (U))—but can fail, depending on where the rungs sit relative to the consensus. Level and dispersion are separate statistics of m_t, not one: that discipline carries to the next sections.

Mankiw et al. (2003) document that forecaster disagreement is substantial, rises with the level and absolute change of inflation, and is not generated endogenously by most models. The mechanism generates it, subject to a channel caveat: inflation is a disputed-*trend* variable, whose revision flux is continuous and carries no boundary-mass atom (Corollary OA.3), so its comovement with $|\Delta\text{inflation}|$ is continuous rather than the discrete *spike* of Proposition OA.2, which we claim only for rare-event variables carrying boundary mass. We do not sign it: a larger $|\Delta\mu|$ compresses the crossing dates while widening the belief gap, and which dominates turns on the threshold distribution and re-selection destinations. Disagreement is then *sustained by demography* (Corollary OA.8)—replacement holds stationary dispersion at an interior level rather than letting it collapse to agreement—and the level effect rides the continuous channel higher in high-inflation regimes: the mechanism's contact with a documented business-cycle moment, with no dispersion in preferences or priors.

Disagreement, lumpy individuals against a sluggish consensus, and the cyclical spread of forecasts are not three mechanisms but one—the revision hazard and the cross-section it traces—taken at three aggregations, all belief-level, not price-level.

Proof of Proposition OA.1. Index types by tolerance and write $S_\xi(a)$ for the survival of type ξ. The survivor composition at age a is $G_a(d\xi) \propto S_\xi(a)\, G(d\xi)$, and the *survivor-component* mean is $\int w_a(\xi)\, G_a(d\xi)$—the mean over *unrevised survivors* only. The full consensus $\langle w, m_t \rangle$ adds the revised entrants and is a distinct object (Proposition OA.2); the wedge below orders the survivor component, not the full cross-section. This G_a is the *ex-ante* expected composition; under the common path the realized survivor set at age a is a deterministic subset of types, agreeing with G_a in expectation over paths (Proposition 2), and the ordering below is a statement about that expected composition. The survival ordering across tolerance types comes directly from the band nesting, with no claim about hazard *measures.* For $\alpha' > \alpha$ the bands are nested, $A_{\alpha'}(a) \subseteq A_\alpha(a)$ at every age (Lemma D.1), so along the common path the narrower-band type's record leaves its band no later than the wider-band type's: the first-exit times satisfy $\tau_{\xi_1} \leq \tau_{\xi_2}$ *pathwise*, hence $S_{\xi_1}(a) = \mathbb{P}(\tau_{\xi_1} > a) \leq \mathbb{P}(\tau_{\xi_2} > a) = S_{\xi_2}(a)$ for every a. The narrower-band, more revision-prone types thus have uniformly smaller survival, so G_a tilts toward the wider-band types ξ_2 relative to G in the first-order stochastic-dominance sense. When the held belief w is monotone in the tolerance type—the maintained condition of Proposition OA.1—this tilt of the survivor composition orders the survivor-weighted *mean*: the consensus is pulled toward the belief of the slower-hazard, wider-band types. This is the frailty composition of Vaupel et al. (1979), the revision-prone exiting first and the survivors being disproportionately robust; the ordering is by the *integrated* hazard that governs survival, not by the instantaneous rate, and the same holds in the survival tail. Consistently, a narrower band gives the *lighter* holding-time tail—the larger principal eigenvalue $\theta(\alpha)$ of Corollary 1(a), by domain monotonicity of the principal Dirichlet eigenvalue: the variational characterization $\theta_1 = \inf\{\int \psi'^2\, dm / \int \psi^2\, dm\}$ is an infimum over test functions vanishing at the boundary, and a smaller interval admits fewer test functions—so the more revision-prone types are selected out both early and in the tail. $\square$

Corollary OA.3 (Micro-jumpy, macro-sluggish)**.** *Suppose G is atomless. Then every individual belief path is pure-jump (the held model is piecewise constant, jumping at each* model-changing *revision; self-re-adoptions, which leave the held model unchanged, are not jumps), while the aggregate $\langle \varphi, m_t \rangle$ is continuous in t at every date at which no positive G-mass revises at once. In the trend channel revisions are a flux, spread across a continuum of types, so past the common probation onset no positive mass revises at once and the aggregate is continuous at every such date, though in general only* singular-*continuous and not Lipschitz: the crossing flux can carry a running-maximum, non-absolutely-continuous part, and continuity follows from injectivity (not continuity) of the type-to-exit-date map under atomless G and* interior initialization*—the cohort beginning inside its band; a synchronized onset otherwise leaves a single atom at the probation date (Appendix E). In the crash channel the* *aggregate is atomic by design: a common disaster makes a positive boundary mass cross simultaneously and the aggregate jumps—upward for a belief-monotone index that the re-selection raises, with the sign of* $\int(\varphi\circ J-\varphi)\,db_t$ *in general, nonzero precisely when that boundary mass is positive—and* between *disasters the cohorts that past disasters synchronized reach their scheduled floor-steps together, so the aggregate also carries* predictable atoms*—downward for a belief-monotone index, signed by* $\int(\varphi\circ J-\varphi)\,db_t$ *in general. Because the cohort's shared floor-steps fall at tolerance-independent rank-swap dates (Lemma D.1), atomless* G *does not smooth them where a positive sub-mass of cutoff-relevant tolerances shares the step; the crash-channel aggregate is then not continuous between disasters. The trend-channel continuity is what carries the micro-jumpy/macro-sluggish contrast.*

Proof of Corollary OA.3. The held model changes only at the countably many revision times, so each individual path is piecewise constant in $\hat{\theta}$, i.e. pure-jump. For the aggregate, by (E.2) the jump at t is $\int(\varphi\circ J-\varphi)\,db_t$, where b_t is the G-mass of types whose exit time equals exactly t. In the trend channel, along a continuous record R, the exit time $\tau_\xi=\inf\{a\geq a_0:|R_a|>z_{\alpha(\xi)}\}$ is a non-decreasing function of the band level $z_{\alpha(\xi)}$, and because R—hence its running maximum—is continuous, it has no flat stretch *after* the probation onset a_0: for $t>a_0$ the level set $\{\xi:\tau_\xi=t\}$ is a single band level (the map jumps at past running maxima but is otherwise strictly increasing), a G-null set when G is atomless and types are indexed monotonically by tolerance, so $b_t=0$. At the onset itself the record enters monitoring at $|R_{a_0}|$, so every type with band $z_{\alpha(\xi)}<|R_{a_0}|$ is already breached and exits at a_0 together—a flat stretch $\tau_\xi=a_0$ on $\{z_{\alpha(\xi)}<|R_{a_0}|\}$, of G-mass $G(\{z_\alpha<|R_{a_0}|\})$. Under *interior initialization*—the synchronized cohort begins inside its band, $|R_{a_0}|<z_{\alpha(\xi)}$ for G-a.e. type, or that onset-breached mass is excluded—this set is G-null and $\langle\varphi,m_t\rangle$ is con-

tinuous at every t; without it the aggregate carries a single atom at the onset a_0 and is continuous for all $t > a_0$. (The operative condition throughout is that the type-to-exit-date map push G forward to an atomless law; atomlessness of G alone does not suffice.) In the crash channel a disaster at t raises the count discontinuously, so every type whose ceiling lies below the new count exits at t simultaneously—an interval of band levels. That interval carries positive G-mass exactly under condition (17) (atomlessness alone does not give an interval positive measure), and then $b_t > 0$ and the aggregate jumps. Between disasters, moreover, the flux carries the predictable atoms of Corollary 2: a positive-G-mass cohort that a past disaster synchronized onto a common reset date shares the common disaster count, and by Lemma D.1 the floor of its Sterne band steps up at two kinds of age: *rank-swap* ages (D.1), fixed by the held rate and the integer pair alone and so tolerance-independent, and *cumulative-mass-threshold* ages, which vary with α. The sub-cohort whose Sterne floor the rank-swap moves—a G-mass set of *cutoff-relevant* tolerances, those for which the swapped ranks sit at the lower endpoint of the band—crosses together at that tolerance-independent date, so a *predictable* atom $b_t > 0$ forms whenever that set has positive G-mass, which atomless G does not disperse; the α-dependent threshold steps are spread across the tolerance continuum and feed the continuous part. The crash-channel aggregate is therefore not continuous between disasters. $\square$

Proof of Proposition OA.2. Between common surprises, $D_t = \langle w^2, m_t\rangle - \langle w, m_t\rangle^2$ is continuous in the trend channel, each index $\langle w^k, m_t\rangle$ being continuous there off common-surprise dates (Corollary OA.3); in the crash channel it inherits the predictable floor-step atoms of that proposition, between which it is again continuous. The jump calculation below treats a common disaster; the between-disaster floor-step atoms are signed the same way, by the crossing cohort's position relative to the consensus. At a common disaster u, split $m_{u-} = (1-\beta)\, m^{\text{stay}} + \beta\, \hat{b}_{u-}$, with $\beta = b_u(\mathcal{S}) > 0$ the crossing mass under (17) and $\hat{b}$ its normalization; the jump carries $\hat{b}_{u-} \mapsto \hat{b}_u = J\#\hat{b}_{u-}$. The mixture variance decomposes as

$$D = (1-\beta)\, V^{\text{stay}} + \beta\, V^{\text{cross}} + (1-\beta)\beta\, \big(\bar{w}^{\text{stay}} - \bar{w}^{\text{cross}}\big)^2, \tag{OA.1}$$

with V and $\bar{w}$ the within-group variances and means. The disaster changes only the crossing group, so

$$\Delta D_u = \beta\big(V_u^{\text{cross}} - V_{u-}^{\text{cross}}\big) + (1-\beta)\beta\Big[(\bar{w}^{\text{stay}} - \bar{w}_u^{\text{cross}})^2 - (\bar{w}^{\text{stay}} - \bar{w}_{u-}^{\text{cross}})^2\Big]. \tag{OA.2}$$

The overshoot condition is exactly that the second bracket is positive and dominates—the revising mass moves its mean past the consensus, widening the between-group gap, $|\bar{w}^{\text{stay}} - \bar{w}_u^{\text{cross}}| > |\bar{w}^{\text{stay}} - \bar{w}_{u-}^{\text{cross}}|$, by more than any contraction of the within-group variance—under which $\Delta D_u > 0$. A transparent primitive sufficient condition obtains in the crash channel: a disaster carries the complacent crossing mass from a common rung $\hat{\lambda}^- < \bar{w}^{\text{stay}}$ up to the *adjacent* menu rung $\hat{\lambda}^+$ (Proposition 4; a multi-rung re-selection gives the same inequality with $\hat{\lambda}^+$ the realized destination rung), leaving V^{cross} unchanged, so that by (OA.2)

$$\Delta D_u = (1-\beta)\beta\Big[(\bar{w}^{\text{stay}} - \hat{\lambda}^+)^2 - (\bar{w}^{\text{stay}} - \hat{\lambda}^-)^2\Big] > 0 \quad \Longleftrightarrow \quad \hat{\lambda}^+ + \hat{\lambda}^- > 2\,\bar{w}^{\text{stay}}, \tag{OA.3}$$

a condition on the menu rungs and the endogenous consensus: the upward re-selection step of the detectably-misspecified crossing mass (condition (U)) lands the group farther above consensus than it began below it. It is not a free parameter, but neither is it unconditional—it can fail when the rungs straddle the consensus the other way.

The level jump $\Delta\langle w, m_u\rangle = \beta(\bar{w}_u^{\mathrm{cross}} - \bar{w}_{u-}^{\mathrm{cross}})$ and the dispersion jump (OA.2) are separate functionals of the same crossing mass and need not share sign. $\square$

OA.2 The aggregate patterns and the two economies

OA.2.1 Reading the patterns

The aggregate dynamics above are channel-free; below we derive four documented patterns from the cross-section, formal statements and proofs in Appendices E–OA.2. The statements are *conditional*: each upward consensus move, fear jump, and price-of-risk imprint needs the crossing boundary mass to be *model-changing*—to re-select a different rate rather than self-re-adopt (conditions (B+), (U), and the menu condition of Corollary OA.9)—since the carried ranking (Lemma A.5) can keep an agent on her model after an exit; a primitive characterization is the companion paper's. All conditions are *environment-specific*, adding no behavioral parameter beyond α; the tolerance law G they involve is a primitive of the population, specified, not fitted.

OA.2.1.1 Lumpy individuals, sluggish consensus

Individual belief paths are pure-jump—the held model is piecewise constant, jumping only at a *model-changing* revision—while the consensus averaging them is *continuous* in the trend channel (a flux spread across a continuum of types) and jumps in the crash channel: up at a common disaster carrying positive boundary mass (condition (B+), upward for a belief-monotone index), down at the predictable scheduled floor-steps (Corollary OA.3).

This is the texture beneath two findings usually taken as opposed—Coibion and Gorodnichenko (2015) find consensus forecasts under-react, Bordalo et al. (2020) find individual forecasts over-react. The mechanism speaks to the *discreteness* they document, not the reaction sign: an unconsolidated individual belief is pure-jump, moving in a discrete revision where a smooth updater would inch (micro *lumpiness*), while the consensus, averaging staggered jumps across a continuum (atomless G; Corollary OA.3), trails the news (macro *sluggishness*). A discrete jump can under- or overshoot the efficient update, and the Coibion–Gorodnichenko sign is target-dependent, reversing without persistence—two textures, one phenomenon at two aggregations.

OA.2.1.2 The disaster economy: fear waves and scheduled relief

Place the population in a rare-disaster economy with the disaster intensity disputed (Appendix OA.2.2). The consensus fear F_t moves by two routes only (Proposition OA.4): it jumps *up* at each disaster carrying model-changing boundary mass (conditions (B+), (U)) and bleeds *down* only on the scheduled lower-exit dates between disasters—the disaster asymmetry of §2.5 at the population level, with the true intensity and every preference parameter held fixed. Under a monotone kernel the jump-risk premium inherits this saw-tooth (Corollary OA.5): a countercyclical price of risk with risk aversion and the true intensity both fixed, the documented spike-and-decay of option-implied tail fears (Bollerslev and Todorov, 2011), sustained over calendar time once turnover replenishes the complacent mass (Corollary OA.9). A disaster-only stand-in kernel supplies the monotonicity; the full heterogeneous-agent equilibrium is the companion paper's (Remark OA.6).

OA.2.1.3 The overlapping-generations economy: demographic disagreement

Switch on demographic turnover: investors are born and die at a constant hazard, each newborn inheriting the common prior and resetting her record (Assumption OA.7).

Cross-sectional disagreement is then *demographic*, not behavioral. Under a dispersed prior ($\hat{\theta}_0 \neq \theta^\circ$) distinct ages hold distinct rates and stationary disagreement $\bar{D}(p)$ is positive at moderate turnover while vanishing in both limits—no turnover consolidates the population, instantaneous turnover makes everyone a newborn (Corollary OA.8); the non-monotonicity is the testable signature a behavioral-dispersion account has no reason to produce. Turnover also sustains the fear saw-tooth—replenished young understaters give each disaster a stationary mass of complacency to correct (Corollary OA.9)—and delivers experience effects as a demographic survival weight, recovering Malmendier and Nagel (2011) and the fading memory of Nagel and Xu (2022) (Lemma OA.10, Corollary OA.11).

OA.2.1.4 An informational transition path

The same survival decay imprints a multidecade trend on any observable loading on the consensus, even with stationary fundamentals. A date-0 surprise sending a *model-changing* cohort across its boundary displaces the consensus; as that cohort dies off and revises, its contribution decays as $e^{-(p+\kappa_{\text{rev}})t}$, mortality p the robust floor (Corollary OA.11), so any observable loading on the consensus carries the same decaying drift (Appendix OA.2.4). With $p \approx 0.02$ and a quiet-spell revision rate $\kappa_{\text{rev}} \approx 0.03$ the half-life is near fourteen years, mortality alone near thirty-five—multidecade either way, appearing to an econometrician with a thirty-year sample as a secular trend. We claim only an *informational transition path* on a stationary process, not a theory of the natural rate: the mechanism supplies the *timescale*, the trend's sign and size the date-0 displacement and the loading.

OA.2.2 The disaster economy

Let the priced fundamental have a rare-disaster component,

$$d \log C_t = \mu_C \, dt + \sigma_C \, dW_t - b \, dN_t, \tag{OA.4}$$

with N a Poisson disaster process of true intensity $\lambda^\star$ and $b > 0$ the log loss in a disaster. The disputed object is the intensity: the menu is $\Theta = \{\lambda_1 < \cdots < \lambda_K\}$ of candidate rates, and each investor holds a perceived intensity $\hat{\lambda}_t$, monitored through the disaster count N_a of its record against the Sterne band (7)—the disputed-crash-rate channel of Sections 2–2.4. Write the consensus perceived intensity (the head-count fear)

$$F_t := \langle \hat{\lambda}, m_t \rangle = \int \hat{\lambda} \, dm_t. \tag{OA.5}$$

The population is fixed in this economy, so the turnover channel of (14) is off; generational turnover is the subject of §OA.2.1.3.

Proposition OA.4 (Fear waves and scheduled relief)**.** *In the disaster economy* (OA.4), *with the population fixed, on the surviving path the consensus fear F_t moves by two routes only.* (Exact.) (i) *At each disaster carrying a positive boundary mass, condition* (B+) *of* (17)*, F_t jumps by $\Delta F_u = \int (\hat{\lambda} \circ J - \hat{\lambda}) \, db_u \geq 0$, weakly* up*—the sign exact, because a disaster raises every count and the Poisson re-selection is weakly increasing in the count on an upper crossing, for* every *record-age composition of b_u; and* (ii) *between disasters F_t moves* down *only, through the predictable scheduled lower exits of the crash band (Lemma D.1), a continuous bleed under an atomless reset-age law and a predictable step when a positive cohort shares a reset date. There is no other route and no predictable upward move—this structure and the weak sign of each move hold exactly, for any population composition.* (Asymptotic; local misspecification, condition (U), conditional on survival to the detection scale.) *The upper exits of route (i) are then the deaths of understaters and the lower exits of route (ii) the deaths of overstaters (Proposition 4, in its conditional form), and the jump in (i) is* strictly *positive when a positive sub-mass of the boundary mass is detectably understating— condition* (U)*, so that its re-selection moves strictly up toward $\lambda^\star$—and is zero on the self-re-adopting sub-mass. In the fixed population the understating and overstating masses are depleted by consolidation, so both routes fade and the saw-tooth is* transient*; Corollary OA.9 shows that turnover with a complacent prior replenishes the understating mass—keeping its expected sub-mass stationary and positive—so that the saw-tooth is sustained over calendar time.*

Proof of Proposition OA.4. The held model is a rate $\hat{\lambda}$ and the record is the disaster count N_a, monitored against the Sterne band of Lemma D.1. The two *exact* claims rest only on the Poisson re-selection and the band structure, not on the directional map: Routes (i) and (ii) below establish the weak sign of each move and the two-route structure for any population composition. The *exit-side* identification of the asymptotic clause—that the route-(i) upper exits are the deaths of understaters and the route-(ii) lower exits those of overstaters, and that re-selection moves an understater *strictly* up under (U)—is Proposition 4 in its conditional form (on the surviving path, conditional on survival to the detection scale): a complacent $\hat{\lambda} < \lambda^\star$ then dies by an upper exit re-selecting weakly higher, strictly under (U); a fearful $\hat{\lambda} > \lambda^\star$ by a lower exit re-selecting weakly lower.

Route (i). At a disaster date u the count jumps by one for every investor at once (the common disaster), so every investor whose count sits at its ceiling exits simultaneously; the G-mass that crosses is b_u, and by Corollary 3 the consensus moves by $\Delta F_u = \int (\hat{\lambda} \circ J - \hat{\lambda})\, db_u$. We separate the *exact* sign from the asymptotic exit-side and strict-movement claims, because b_u in general mixes records of different ages. *Exact sign.* At a disaster the count jumps up by one, so any investor who crosses at u crosses the *upper* edge—an upper exit by construction, for every crossing record age. For an upper exit the Poisson re-ranking (8) re-selects the MAP rate of a record that has just risen: by the monotone-likelihood-ratio property of the Poisson family—for menu rates $\lambda' > \lambda$ and a record of length a, $L_{\lambda'}(n)/L_{\lambda}(n) = (\lambda'/\lambda)^n e^{-(\lambda'-\lambda)a}$ is strictly increasing in the count n—a higher count shifts the maximizer of $n \mapsto L_{\cdot}(n)\rho(\cdot)$ over the finite menu weakly up, so $\hat{\lambda} \circ J \geq \hat{\lambda}$ pointwise on the crossing mass. This is an exact monotonicity of the Poisson-MAP map on upper crossings, not an asymptotic statement, and it gives $\Delta F_u \geq 0$ regardless of the record-age composition of b_u. *Strict movement (asymptotic).* The inequality is strict, $\Delta F_u > 0$, when a positive sub-mass of b_u is detectably understating (condition (U), so that its re-selection moves strictly up toward $\lambda^\star$ in the long-record regime of Assumption 3), the self-re-adopting sub-mass contributing zero; the identification of *which* held rates exit by upper versus lower (Proposition 4) is likewise the asymptotic part. Thus a strict upward jump requires both (B+) ($b_u > 0$) and a positive understating sub-mass. The disaster times are inaccessible stopping times of the count's filtration.

Route (ii). Between disasters the count does not jump, so no upper exit can occur; the only revisions are the lower exits at the deterministic band-step dates of Lemma D.1, where the floor steps up past a stalled count. Each re-selects a weakly lower rate, strictly under (U) (Proposition 4), weakly lowering F_t, and each such date is predictable (a continuous bleed under an atomless reset-age law, a predictable step when a positive cohort shares a reset date). No other revision is possible in the crash channel, so F_t moves only through (i) and (ii); the upward moves are at inaccessible disaster times and the downward moves at predictable scheduled times, giving the saw-tooth with no predictable upward move.

Transience in the fixed population. The strict moves of both routes draw on detectably-misspecified mass: route (i) on understaters poised at the ceiling, route (ii) on overstaters at a floor-step. In a fixed population consolidation (Proposition 1(a)) depletes both as every held rate settles at the Kullback–Leibler projection θ° (the menu rate nearest $\lambda^\star$, equal to it only when the truth is on the menu): once a held rate is its own record's maximizer, condition (U) fails and the crossing self-re-adopts, so the boundary sub-mass satisfying (U) vanishes asymptotically and the saw-tooth flattens; the result is therefore transient absent replenishment. Section OA.2.1.3 restores a stationary understating mass through turnover, sustaining it. $\square$

The same shape imprints on the price of risk, in sign and timing if not in level. Let the economy's jump-risk premium be a function $\Pi(F_t)$ of the consensus fear, strictly increasing in perceived intensity, with risk aversion and the true intensity held fixed; the stand-in kernel of Remark OA.6 below supplies that monotonicity, and here we use only the monotonicity.

Corollary OA.5 (Price comovement, sign and timing)**.** *This is a corollary of Proposition OA.4 under monotone pricing, carrying no equilibrium of its own. Under any strictly increasing pricing* function Π*, the jump-risk premium* $\Pi(F_t)$ *inherits the saw-tooth of the consensus fear—it jumps weakly up at each disaster, strictly under* (U)*, declines on the scheduled dates between, and makes no predictable upward move—by direct composition of* Π *with the head-count fear* F_t*, the monotone transform passing sign and timing through unchanged. Not claimed here are the level, the size of the jumps, or that* F_t *is the pricing-relevant belief: the stand-in kernel that supplies* Π *(Remark OA.6) replaces the head-count* F_t *in genuine equilibrium with a* wealth-weighted *belief aggregate, and because wealth flows toward the more accurate model, selection on wealth could smooth or even invert the head-count saw-tooth—a contest* *decided only in the companion paper's equilibrium, where whether a predictable decline in a* traded *price is arbitrage-free is also settled.*

Proof of Corollary OA.5. By Proposition OA.4, F_t jumps by $\Delta F_u \geq 0$ at each disaster—strictly when (B+) holds and a positive understating sub-mass satisfies (U)—and declines through predictable scheduled lower exits between disasters. As Π is strictly increasing, $\Pi(F_u) \geq \Pi(F_{u^-})$ at each disaster, strictly under the same condition, and $\Pi(F_t)$ declines on the same scheduled dates, contemporaneously with F_t: the sign and timing of the comovement follow. The only predictable moves of F_t are declines, hence the only predictable moves of $\Pi(F_t)$ are declines, so the premium function has no predictable upward jump. These are statements about the non-traded objects F_t and the premium function $\Pi(F_t)$; whether a predictable decline in a *traded* price is arbitrage-free turns on the equilibrium kernel and gains process of the companion paper, which we do not invoke here. The level of Π and the size of $\Delta\Pi$ depend on the preferences and equilibrium of the stand-in kernel (Remark OA.6) and the companion paper; only sign and timing are claimed. $\square$

Remark OA.6 (Stand-in kernel). Take a representative agent with constant relative risk aversion $\gamma > 0$ whose consumption falls by a fixed log size $\zeta > 0$ at each disaster—the disaster log-consumption size, not the boundary mass b—disasters arriving at the predictable consensus intensity F_{t-}. In this disaster-only stand-in (continuous consumption risk suppressed) the state-price density $M_t \propto C_t^{-\gamma}$ is pure-jump, rising by $e^{\gamma\zeta}$ at each disaster; by the intensity tilting of Lemma A.1(ii) with constant multiplier $e^{\gamma\zeta}$, the risk-neutral disaster intensity is

$$\lambda_t^{\mathbb{Q}} = e^{\gamma\zeta} F_{t-}, \tag{OA.6}$$

strictly increasing in F. The jump-risk premium is the function $\Pi(F) = (e^{\gamma\zeta} - 1)F = \lambda^{\mathbb{Q}}(F) - F$ of the fear level, with $\lambda^{\mathbb{Q}}(F) = e^{\gamma\zeta}F$ the risk-neutral intensity as a function of fear and the predictable process $\lambda_t^{\mathbb{Q}} = e^{\gamma\zeta}F_{t-}$. In this stand-in the agent prices disasters as arriving at its own perceived intensity F_{t-}—its subjective physical rate—so $\Pi(F) = \lambda^{\mathbb{Q}}(F) - F$ is the wedge of the risk-neutral over the *perceived* intensity, not over the true constant $\lambda^\star$, with which F coincides only when the consensus fear is correct. The premium therefore inherits the saw-tooth of F_t in sign and timing (Corollary OA.5), countercyclical, with γ and $\lambda^\star$ both fixed. The *level* $e^{\gamma\zeta} - 1$ is a preference-and-disaster-size constant. The full heterogeneous-agent equilibrium—in which F_t gives way to a wealth-weighted belief aggregate, wealth is itself moved by the disasters, and the gains process is checked for arbitrage—is the companion paper's, and the deferral is not innocuous: wealth flows toward whoever holds the more accurate model (Sandroni, 2000; Blume and Easley, 2006), so the pricing-relevant belief is wealth-weighted and its drift could in principle smooth or invert the head-count saw-tooth—decided in the contest between selection, which retires the fragile complacent types, and the turnover that replenishes them (Corollary OA.9).

OA.2.3 The overlapping-generations economy

Assumption OA.7 (Demographic turnover: perpetual youth)**.** *A continuum of investors is born and dies at hazard rate $p > 0$ (Yaari, 1965; Blanchard, 1985; Gârleanu and Panageas, 2015), so lifespans are exponential with mean $1/p$ and, in the stationary population, ages ℓ have density $p\,e^{-p\ell}$. At birth an investor* (i) *inherits the common prior ρ_0—so her ranking, and its MAP $\hat{\theta}_0 = \arg\max_\theta \rho_0(\theta)$, are common across cohorts—and* (ii) *resets her test record to empty, beginning the per-test rule of Definition 1 on the data observed over her own since-birth window. She holds $\hat{\theta}_0$*

until her first exit and re-ranks and re-adopts by the MAP rule thereafter, exactly as an incumbent; the menu, tolerance α, recording resolution δ, and probation interval a_0 are common to all cohorts.

All living investors share the common path (Assumption 2); they differ only in age and hence in the record their held model has survived. Here the belief functional w of §3 is the perceived intensity $\hat{\lambda}$, so the cross-sectional disagreement $D_t = \mathrm{Var}_{m_t}(w)$ is the object of Proposition OA.2.

Whether turnover sustains disagreement turns on a single primitive feature of the prior: whether the *age-to-held-rate map* $\ell \mapsto \hat{\theta}(\ell)$—the rate a vintage of *demographic* age ℓ (time since birth) carries along the common path, determined by its full since-birth window and so indexed by ℓ rather than by the record age that resets at each revision—is nonconstant on a positive-age-measure set. The prior drives a systematic young-old rate gap when the prior-implied model $\hat{\theta}_0 = \arg\max_\theta \rho_0(\theta)$ differs from the consolidation destination θ° (Proposition 1(a)): young vintages hold $\hat{\theta}_0$, consolidated vintages hold θ°, and intermediate ages hold the menu rungs the record passes through en route. We call this the *dispersed-prior condition*, $\hat{\theta}_0 \neq \theta^\circ$. It is a spread in the support of the *one shared* prior relative to θ°, not cross-agent prior heterogeneity; and it sustains disagreement because distinct ages *hold* distinct rates, not because they re-select to distinct destinations—a population in which every revision selects the same θ° still disagrees while its young vintages have not yet revised.

Corollary OA.8 (Stationary disagreement under turnover)**.** *Under Assumptions OA.7 and 2, a stationary-increment DGP, and the dispersed-prior condition $\hat{\theta}_0 \neq \theta^\circ$, the cross-sectional disagreement $D_t = \mathrm{Var}_{m_t}(\hat{\lambda})$ is stationary in law, and its stationary mean $\bar{D}(p) := \mathbb{E}[D_t]$ is positive at intermediate turnover and vanishes at both extremes: $\bar{D}(p) \to 0$ as $p \to 0$ and as $p \to \infty$, while $\bar{D}(p) > 0$ for every $p \in (0, \infty)$,* *so $\bar{D}$ attains an interior maximum. Disagreement is sustained at moderate generational turnover and disappears in both the no-turnover and instantaneous-turnover limits. Absent the condition, and provided every post-exit re-ranking self-re-adopts θ° (the carried ranking keeping θ° the front-runner at all ages, so a false-alarm exit re-installs it rather than a transient rung), distinct ages share the held rate θ° and $\bar{D}(p) \equiv 0$. Without that self-re-adoption, fixed-size testing moves a flow of young vintages onto other rungs between exits, and $\bar{D}(p)$ is positive even at $\hat{\theta}_0 = \theta^\circ$—a testing-driven dispersion distinct from the prior-driven one.*

Proof of Corollary OA.8. Stationarity. Under a stationary-increment DGP the since-adoption *window* record of a vintage of age ℓ—$X_{(t-\ell,t]}$ in the trend channel, $N_t - N_{t-\ell}$ in the crash channel—has a law depending only on ℓ (by stationary increments, not on calendar time t); the stationary population has age density $p\,e^{-p\ell}$ (the exponential lifespan); so m_t, the pushforward (Proposition 2) of the joint age-and-window-record law under the common flow, is stationary in law. It is the window record, not the cumulative count N_t, that carries the stationarity. Hence $D_t = \mathrm{Var}_{m_t}(\hat{\lambda})$ is stationary in law and $\bar{D}(p) = \mathbb{E}[D_t]$ is well defined.

No within-age dispersion. Under Assumption 2 (common path), two investors of the same age ℓ were born at the same date and have observed the same window $[t-\ell, t]$, so they hold the identical model: $w_t(\ell)$, the held rate at age ℓ, is a single value. The disagreement is therefore the across-vintage variance,

$$D_t = \mathrm{Var}_{\ell \sim pe^{-p\ell}}\big(w_t(\ell)\big) = \int_0^\infty \big(w_t(\ell) - \bar{w}_t\big)^2 p\,e^{-p\ell}\,d\ell, \qquad \bar{w}_t = \int_0^\infty w_t(\ell)\,p\,e^{-p\ell}\,d\ell. \tag{OA.7}$$

Continuity. The held rates lie in the finite menu, so $|w_t(\ell)| \le \lambda_K$ and, writing $\Delta_\Theta := \lambda_K - \lambda_1$ for the menu diameter, the integrand of (OA.7) is bounded by Δ_Θ^2. For p in any compact $[\underline{p}, \bar{p}] \subset (0, \infty)$ the density obeys $pe^{-p\ell} \le \bar{p}\,e^{-\underline{p}\ell}$, integrable in ℓ, so dominated convergence makes $p \mapsto \bar{D}(p)$ continuous (indeed C^∞) on $(0, \infty)$.

Boundary limits. Fix a date t and a path ω. The cross-section ranges over cohorts of demographic age ℓ, the cohort of age ℓ carrying the window $[t-\ell, t]$; under Assumption 2 these windows are *nested* sub-paths of the one common path, each ending at t, so increasing ℓ extends the same record backward rather than sampling an independent cohort history. This nesting is what lets the single-vintage results act cross-sectionally. *(Young.)* A cohort holds $\hat{\theta}_0$ while its window has triggered no revision; the probation onset (Lemma A.7) makes the first-revision age positive, so there is a cross-cohort age $\ell_1(t, \omega) > 0$ with $w_t(\ell) = \hat{\theta}_0$ for all $\ell < \ell_1$. *(Old.)* For the older cohorts the relevant limit is *backward* in birth date, not forward in one cohort's age, so we argue from the latent log-likelihood ratio directly. The cohort of age ℓ holds, at t, the MAP of the ranking at its *last revision time* $\sigma(\ell) \in [t-\ell, t]$, the ranking being frozen since; by the chaining identity (Lemma A.5) that ranking carries the log-ratio of the window up to that revision, $\log \Lambda^{\theta/\theta^\circ}_{[t-\ell, \sigma(\ell)]}$, in pairwise form. Write $\Psi_\theta(r) := \log \Lambda^{\theta/\theta^\circ}_{[t-r,t]}$ for the backward-window log-ratio, a single one-parameter process, so the evaluated ratio is $\Psi_\theta(\ell) - \Psi_\theta(t - \sigma(\ell))$. Under $\mathbb{P}^\star$ the log-ratio has stationary increments (the data do), so by the strong law of large numbers applied to the *lengthening* window—the reversed increments $X_t - X_{t-\ell}$ are again stationary and independent, so the strong law of Proposition 1(a) applies to the backward window as it lengthens, $\ell \to \infty$ at fixed t—

$$\tfrac{1}{\ell} \log \Lambda^{\theta/\theta^\circ}_{[t-\ell,t]} \xrightarrow[\ell\to\infty]{\text{a.s.}} \kappa(\theta^\circ) - \kappa(\theta) \;<\; 0 \qquad (\theta \neq \theta^\circ),$$

so $\Psi_\theta(\ell) = -g_\theta\, \ell\, (1 + o(1))$ a.s. with $g_\theta := \kappa(\theta) - \kappa(\theta^\circ) > 0$. Two cases close the limit. If the final spell is short, $t - \sigma(\ell) \le \ell/2$, then $\Psi_\theta(\ell) - \Psi_\theta(t - \sigma(\ell)) \le -g_\theta \ell/2 + o(\ell) < 0$ eventually, uniformly over the finite menu, so the ranking at $\sigma(\ell)$—hence the held model at t—is θ°. If the final spell is long, $t - \sigma(\ell) > \ell/2$, the held model has survived a window of length exceeding $\ell/2$ ending at t; by the backward strong law the count over that window is $\lambda^\star$-typical, $N_{(\sigma(\ell),t]} = \lambda^\star(t - \sigma(\ell))(1 + o(1))$ a.s., while survival in the held band requires $|N_{(\sigma(\ell),t]} - \hat{\lambda}\,(t - \sigma(\ell))| \le z_\alpha \sqrt{\hat{\lambda}(t - \sigma(\ell))} + o(\sqrt{\ell})$ (Lemma D.1), which for ℓ beyond an a.s. finite threshold forces $\hat{\lambda} = \theta^\circ$ when the truth is on the menu and is impossible for any menu rate under misspecification—so long final spells eventually occur only at θ°, or not at all, returning the same conclusion through the short-spell case. Either way $w_t(\ell) = \theta^\circ$ for all ℓ beyond an a.s. finite $\ell_2(t,\omega)$. This is the backward-in-birth-date consolidation the cross-section needs—one threshold across cohorts, the nesting of the windows making the $\ell \to \infty$ limit the same lengthening record rather than a separate consolidation per birth date. Under misspecification θ° is the divergent-drift destination of Proposition 1(a), not a correct-model limit. Using $\mathrm{Var}(Y) \le \mathbb{E}[(Y - c)^2]$ for any constant c (the mean minimizes mean-square deviation), take $c = \hat{\theta}_0$ and then $c = \theta^\circ$: since $(w_t - \hat{\theta}_0)^2$ vanishes on $\ell < \ell_1$ and $(w_t - \theta^\circ)^2$ on $\ell > \ell_2$, each $\le \Delta_\Theta^2$ elsewhere,

$$
\begin{aligned}
\mathrm{Var}_{\ell \sim pe^{-p\ell}}\big(w_t(\ell)\big) &\le \Delta_\Theta^2 \int_{\ell_1}^{\infty} pe^{-p\ell}\, d\ell = \Delta_\Theta^2 e^{-p\ell_1}, \\
\mathrm{Var}_{\ell \sim pe^{-p\ell}}\big(w_t(\ell)\big) &\le \Delta_\Theta^2 \int_{0}^{\ell_2} pe^{-p\ell}\, d\ell = \Delta_\Theta^2 \big(1 - e^{-p\ell_2}\big).
\end{aligned}
\tag{OA.8}
$$

Taking expectations over the path, with $e^{-p\ell_1}$, $1 - e^{-p\ell_2} \le 1$ and $\ell_1 > 0$, $\ell_2 < \infty$ a.s., dominated convergence gives

$$
\bar{D}(p) \le \Delta_\Theta^2\, \mathbb{E}\big[e^{-p\ell_1}\big] \xrightarrow[p \to \infty]{} 0, \qquad \bar{D}(p) \le \Delta_\Theta^2\, \mathbb{E}\big[1 - e^{-p\ell_2}\big] \xrightarrow[p \to 0]{} 0. \tag{OA.9}
$$

Interior positivity. Under the dispersed-prior condition $\hat{\theta}_0 \ne \theta^\circ$, young vintages *hold* $\hat{\theta}_0$ on $\ell < \ell_1$ and consolidated vintages *hold* θ° on $\ell > \ell_2$, each a set of positive

$pe^{-p\ell}$-measure since the density is everywhere positive on $(0, \infty)$. The held rate $w_t(\ell)$ is therefore not $pe^{-p\ell}$-a.s. constant—it takes the two distinct values $\hat{\theta}_0 \neq \theta^\circ$ on sets of positive mass—so

$$\bar{D}(p) = \mathbb{E}\big[\mathrm{Var}_{\ell \sim pe^{-p\ell}}(w_t(\ell))\big] > 0 \qquad \text{for every } p \in (0, \infty). \tag{OA.10}$$

(Absent the condition, and if every post-exit re-ranking self-re-adopts θ°, distinct ages share θ° and $\bar{D}(p) \equiv 0$; otherwise testing excursions sustain a positive $\bar{D}$ even at $\hat{\theta}_0 = \theta^\circ$.)

Interior maximum. Collect the three facts: $\bar{D}$ is continuous on $(0, \infty)$, strictly positive there by (OA.10), and $\bar{D}(p) \to 0$ at both endpoints by (OA.9). Fix any $p_0 \in (0, \infty)$ and put $M := \bar{D}(p_0) > 0$, $\epsilon := M/2$. By (OA.9) there exist $0 < p_1 < p_0 < p_2 < \infty$ with

$$\bar{D}(p) < \epsilon \qquad \text{for all } p \in (0, p_1] \cup [p_2, \infty). \tag{OA.11}$$

On the compact $[p_1, p_2] \subset (0, \infty)$ the function $\bar{D}$ is continuous, so by the extreme-value theorem it attains its maximum over $[p_1, p_2]$ at some $p^\star$, with $\bar{D}(p^\star) \geq \bar{D}(p_0) = M > \epsilon$. The endpoints satisfy $\bar{D}(p_1), \bar{D}(p_2) < \epsilon < \bar{D}(p^\star)$ by (OA.11), so $p^\star \notin \{p_1, p_2\}$, i.e. $p^\star \in (p_1, p_2) \subset (0, \infty)$, an *interior* point. It is the *global* maximizer: on $[p_1, p_2]$ by construction, and on $(0, p_1] \cup [p_2, \infty)$ because there $\bar{D} < \epsilon < \bar{D}(p^\star)$ by (OA.11). Hence $\bar{D}$ attains a global maximum at an interior $p^\star \in (0, \infty)$. No single-peakedness is claimed: the argument yields existence of an interior maximizer, not uniqueness. □

Corollary OA.9 (Turnover sustains the fear saw-tooth)**.** *Take a* complacent *prior—the prior-implied rate* $\hat{\theta}_0 = \arg\max_\theta \rho_0(\theta)$ *strictly below the truth,* $\hat{\theta}_0 < \lambda^\star$*—and suppose the menu offers a rate above the ceiling understaters' held rate that their* *upper-exit record favors, so that condition (U) of* (U) *holds on a positive understating sub-mass (their disaster-driven upper exit re-selects strictly upward, not into self-re-adoption). Under Assumption OA.7, a stationary disaster DGP, and Assumption 3, the* understating boundary sub-mass*—the population mass detectably understating $\lambda^\star$ whose count sits at its ceiling—is then stationary in law with positive mean, $\mathbb{E}[b_t^{(U)}] \geq c(p, \rho_0, \lambda^\star, \alpha) > 0$, and the* expected *upward jump in consensus fear at a disaster, $\mathbb{E}[\Delta F_u \mid \text{disaster at } u]$, is strictly positive and constant in calendar time. The fixed-population transience of Proposition OA.4 is thereby overturned* in the mean*: under turnover with a complacent prior the fear saw-tooth is stationary in law, with a positive expected upward jump at disaster dates and scheduled relief between. The bound is on the mean sub-mass; it does not assert a positive realized jump at every disaster date.*

Proof of Corollary OA.9. Each newborn adopts the prior-implied rate $\hat{\theta}_0 < \lambda^\star$ and so enters *understating*. While a held rate $\hat{\lambda} < \lambda^\star$ is in force the count drifts up at rate $\lambda^\star - \hat{\lambda} > 0$ (the empirical rate $\lambda^\star$ exceeds the held rate), reaching the ceiling $u_\alpha(a)$ recurrently and sitting there until either the ceiling steps up or the next disaster arrives; a disaster that finds the count at the current ceiling forces an upper exit and, by Proposition 4 under Assumption 3 and (U), a strictly upward re-selection toward the consolidation destination. An investor leaves the *understating phase*—the transient period from birth during which her held rate sits below the consolidation destination θ° (the menu rate nearest $\lambda^\star$, equal to $\lambda^\star$ when the truth is on the menu and below it otherwise) and upper exits re-select strictly upward—on consolidating to θ° (Proposition 1(a) gives eventual dominance of θ°, not a single terminal date); the phase has positive expected duration, $\mathbb{E}[T_U] > 0$, since an understating model is detected only at a strictly positive holding time ($T_U > 0$ a.s.; the detection scale $a^\star = z_\alpha^2 \hat{\theta}_0 / (\lambda^\star - \hat{\theta}_0)^2$ sets its order in the local-misspecification regime but is not a pathwise lower bound, since early fluctuation-driven exits have positive probability).

By the stationary structure of the overlapping-generations population—constant birth inflow p and the stationary age law $p\,e^{-p\ell}$—Little's law (Little, 1961) gives the stationary *mean* mass in the understating phase as $\mathbb{E}[M_U] = p\,\mathbb{E}[\min(T_U, L)]$, the living-population sojourn truncated by the exponential lifespan $L \sim \mathrm{Exp}(p)$; since $T_U > 0$ a.s. this mean is strictly positive. The spell law in that conversion is *adoption-aligned*—the law of τ for a spell begun, the renewal cycle law—not the length-biased law of the spell in progress at a fixed calendar date, which reweights the same law by duration; the standing cross-section's age and residual-life densities are the corresponding renewal objects, and no length-biased moment is used. The realized cross-sectional mass is path-dependent and stationary in law; what the conclusion uses is the mean, not a realized floor at every date. Within the phase an investor is *poised*—count at the current ceiling, awaiting the disaster that pushes it over. Because $\lambda^\star > \hat{\lambda}$ the count overtakes the held model's rising ceiling near the detection scale $a^\star$ with positive probability, and once there it dwells until the earliest of three competing clocks fires: the next disaster (rate $\lambda^\star$, which forces the upper exit), the next scheduled ceiling step-up (which re-opens the gap), and death (rate p). Each wait has positive expected length, so the per-life time spent *poised and detectably misspecified*—count at the ceiling and condition (U) in force, $T_P = \int_0^L \mathbf{1}\{\text{count} = \text{ceiling},\ (U) \text{ holds}\}\, da$—satisfies $\mathbb{E}[T_P] > 0$ and $T_P > 0$ with positive probability, since a newborn at $\hat{\theta}_0$ that becomes poised has a long upper-exit record favoring a higher rung (Assumption 3), so (U) holds on a positive-probability, positive-duration set. Little's law applied to this poised, (U)-satisfying state, under the same stationary OLG structure, gives its stationary *mean* sub-mass $\mathbb{E}[b^{(U)}] = p\,\mathbb{E}[\min(T_P, L)] > 0$, strictly positive because $\min(T_P, L) > 0$ on $\{T_P > 0\}$ (the conclusion below needs the mean, not a pathwise floor).

At a disaster u the poised mass crosses—its count steps past the ceiling—so the crossing measure b_u contains the poised understating mass, of which $b^{(U)}_{u^-}$ is the *detectably-misspecified* sub-mass defined above: those poised understaters whose just-ended upper-exit record (U) prefers a strictly higher menu rate, so re-selection is strictly upward, $\hat{\lambda} \circ J - \hat{\lambda} > 0$ (Proposition 4). This is in general a *proper* sub-mass of the poised understaters: a poised agent already holding an intermediate underestimating rate that is her own record's maximizer is an understater who nonetheless self-re-adopts at the upper exit and contributes zero. Complacency alone does not force the strict move: it also requires the menu to admit a rate the upper-exit record prefers, failing which the crossing self-re-adopts and the jump is null. Restricting to the detectably-misspecified sub-mass, $\Delta F_u = \int (\hat{\lambda} \circ J - \hat{\lambda})\, db_u \geq \Delta_{\min}\, b^{(U)}_{u^-}$, with $\Delta_{\min} > 0$ the smallest upward rung step; taking expectations over the stationary law of $b^{(U)}$, $\mathbb{E}[\Delta F_u \mid \text{disaster at } u] \geq \Delta_{\min}\, \mathbb{E}[b^{(U)}] > 0$, constant in u. The upward jumps do not fade, so the saw-tooth is stationary. □

OA.2.4 Experience effects and the transition path

Lemma OA.10 (Demographic decay envelope)**.** *Let a common surprise at date t shift the held models of the cohort alive and at their boundary at t, and write $\mathcal{I}(h)$ for that cohort's* retained-imprint *contribution to the consensus belief $\langle w, m_{t+h}\rangle$ at lag $h \geq 0$, in magnitude—the part still carried by members holding the date-t shifted model (a model-changing reviser drops out of it; any later contribution she makes through a new model is excluded). In the overlapping-generations economy the imprint is bounded above by the cohort's survival,*

$$0 \leq \mathcal{I}(h) \leq e^{-ph}\, \mathcal{I}(0), \tag{OA.12}$$

because at most the fraction e^{-ph} *of the experiencing cohort is still alive at* $t+h$*, and any survivor who has since undergone a* model-changing *revision no longer carries the imprint (a self-re-adoption, which leaves the held model unchanged, does not erode it). Mortality alone therefore sets the* slowest *admissible decay—an upper envelope at the demographic rate* p*, attained with equality only if no survivor changes model—and ongoing model-changing revision makes the imprint decay* at least *as fast as* p*.*

Proof of Lemma OA.10. A common surprise at t shifts the held models of the cohort alive and at its boundary at t; write Δ_t for the shift this induces in the consensus $\langle w, m\rangle$. A member of that cohort still carries the shifted model at $t+h$ only if she is alive at $t+h$ *and* has not since revised. Mortality is a rate-p thinning independent of the record, so at most the fraction e^{-ph} of the cohort survives to $t+h$; setting aside any further revision, the cohort's contribution to $\langle w, m_{t+h}\rangle$ is exactly e^{-ph} times its date-t contribution. Further revision can only remove additional members from the shifted state, so the contribution is *bounded above* by the survival factor: $0 \leq \mathcal{I}(h) \leq e^{-ph}\mathcal{I}(0)$, with equality only if no survivor revises. Thus e^{-ph} is the slowest admissible decay—an upper envelope—and the imprint decays *at least* as fast as the demographic rate p. □

Corollary OA.11 (Experience impulse response)**.** *If, in addition, each surviving investor undergoes a* model-changing *revision at a constant hazard* κ_{rev} *(self-re-adoptions, which leave the held model unchanged, excluded; a constant-hazard idealization), the experience impulse response is*

$$\mathcal{I}(h) = e^{-(p+\kappa_{\mathrm{rev}})h}\,\mathcal{I}(0), \tag{OA.13}$$

mortality and revision eroding the memory of a past surprise additively in the exponent, so the imprint decays strictly faster than the mortality-only envelope. The *envelope* (OA.12) *is the* $\kappa_{\text{rev}} \to 0$ *limiting (slowest) case; the idealization is the only place a constant revision hazard is invoked, and it is flagged as such.*

Proof of Corollary OA.11. Add a constant *model-changing* revision hazard κ_{rev} (self-re-adoptions excluded). A surviving cohort member retains her date-t held model at $t+h$ only if she has neither died nor undergone a model-changing revision in $(t, t+h]$. Death arrives at rate p and model-changing revision at rate κ_{rev}, independently, so the retention probability is $e^{-ph}\, e^{-\kappa_{\text{rev}} h} = e^{-(p+\kappa_{\text{rev}})h}$, which is the impulse response. As $\kappa_{\text{rev}} \to 0$ it returns the floor (OA.12). The constant revision hazard is an idealization: the exact hazard is age- and record-dependent (Sections 2–2.4), and it is assumed constant only here. $\square$

The informational transition path. Let Y_t be a macroeconomic observable loading on the fundamentals X_t (the common path of §3) and the consensus $F_t = \langle w, m_t \rangle$, smooth in both; linearizing about a steady state,

$$Y_t \approx \eta_X X_t + \eta_F F_t, \tag{OA.14}$$

with η_F the loading on the consensus. A date-0 surprise that carries a positive boundary mass across its band—the positive-crossing condition $b_0(\partial) > 0$ in the sense of (B+)—displaces the consensus only if that mass re-selects to a different model: the displacement is the signed increment $\int (w \circ J - w)\, db_0$, and (B+) alone does not make it nonzero, since the crossing can self-re-adopt or carry offsetting re-selections that leave it zero. When the crossing is model-changing in one direction the displacement has *magnitude* $\mathcal{I}(0) > 0$ and *sign* $s_0 \in \{-1, +1\}$ fixed by the re-selection direction (the crossing cohort re-selecting above or below the prevailing consensus, so the signed increment is $s_0\, \mathcal{I}(0)$); absent such a crossing the path is degenerate. The retained-imprint magnitude decays, by Corollary OA.11, as

$$\mathcal{I}(t) = e^{-(p+\kappa_{\mathrm{rev}})t}\,\mathcal{I}(0) \le e^{-pt}\,\mathcal{I}(0), \tag{OA.15}$$

inside the survival envelope of Lemma OA.10, so the cohort's signed retained-imprint contribution to F_t is $s_0\,\mathcal{I}(t)$. By (OA.14) this enters the expected path as a single decaying term,

$$\mathbb{E}_0[Y_t] = \eta_X\,\mathbb{E}_0[X_t] + \eta_F\bar{F} + \underbrace{\eta_F\,s_0\,\mathcal{I}(0)\,e^{-(p+\kappa_{\mathrm{rev}})t}}_{\text{retained imprint}} + \mathcal{R}_t, \tag{OA.16}$$

with $\bar{F}$ the steady-state consensus and $\mathcal{R}_t$ a remainder defined below: the bracketed term is a deterministic, exponentially decaying drift in Y_t whose sign is $\operatorname{sign}(\eta_F s_0)$ and whose timescale is the demographic-revision rate. It is the *retained-imprint* contribution alone—the part carried by survivors still holding the date-0 model. The remainder $\mathcal{R}_t$ is an *indirect* effect the exponential omits: by shifting the records of surviving non-revisers relative to their boundaries it alters their future exit propensities, a channel that feeds back into F_t at later dates. We claim the retained-imprint *timescale*, not the full impulse response, which would net this indirect channel against the decay above.

References

Olivier J. Blanchard. Debt, deficits, and finite horizons. *Journal of Political Economy*, 93(2):223–247, 1985.

Lawrence Blume and David Easley. If you're so smart, why aren't you rich? belief selection in complete and incomplete markets. *Econometrica*, 74(4):929–966, 2006.

Tim Bollerslev and Viktor Todorov. Tails, fears, and risk premia. *Journal of Finance*, 66(6):2165–2211, 2011.

Pedro Bordalo, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782, 2020.

Olivier Coibion and Yuriy Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678, 2015.

Nicolae Gârleanu and Stavros Panageas. Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 123(3):670–685, 2015.

John D. C. Little. A proof for the queuing formula: $L = \lambda W$. *Operations Research*, 9(3):383–387, 1961.

Ulrike Malmendier and Stefan Nagel. Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, 126(1):373–416, 2011.

N. Gregory Mankiw, Ricardo Reis, and Justin Wolfers. Disagreement about inflation expectations. *NBER Macroeconomics Annual*, 18:209–248, 2003.

Stefan Nagel and Zhengyang Xu. Asset pricing with fading memory. *Review of Financial Studies*, 35(5):2190–2245, 2022.

Alvaro Sandroni. Do markets favor agents able to make accurate predictions? *Econometrica*, 68(6):1303–1341, 2000.

James W. Vaupel, Kenneth G. Manton, and Eric Stallard. The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography*, 16(3):439–454, 1979.

Menahem E. Yaari. Uncertain lifetime, life insurance, and the theory of the consumer. *Review of Economic Studies*, 32(2):137–150, 1965.