

Rare States and Long-Run Pricing

Ye Lu^a and John Stachurski^b

^a University of Sydney

^b National Graduate Institute for Policy Studies

September 2026

ABSTRACT: Macroeconomic crises are rare, yet they can have large effects on asset prices. Because the probability of entering a crisis from normal times is small relative to the probability of recovery, the dynamics are nearly reducible: setting the crisis-entry probability to zero makes crisis states unreachable from normal states. We develop a graph-theoretic and spectral analysis of long-run pricing under reducible and nearly reducible dynamics. We show that, at the reducible limit, physical and pricing dynamics induce the same classes of economic states and accessibility relations, but rank their long-run importance differently: by recurrence and by class-specific pricing rates. Moreover, a claim's price depends only on its *pricing corridor*—the classes of states lying on directed paths from the initial state to states where the payoff is positive—and the highest class-specific pricing rate within this corridor determines the claim's long-run pricing rate. When the reducible-limit corridor excludes the globally pricing-dominant class, we characterize how restoring rare transitions introduces a contribution from that class with a weight that vanishes as the transitions become rarer. This contribution becomes dominant only beyond a *crossover maturity*, which increases with the rarity of the connecting paths and decreases with the dominant class's pricing-rate advantage. Finally, a consumption-based disaster-recovery application shows how preferences, consumption risk, cash-flow exposure, and recovery timing jointly shape the long-run pricing dominance of disaster states and crossover maturities.

KEYWORDS: Long-run asset pricing; Reducible Markov chains; Rare transitions; Pricing corridors; Positive pricing semigroups; Crossover horizons; Rare disasters.

1. INTRODUCTION

Severe macroeconomic disasters, such as those associated with wars, depressions, sovereign-debt crises, and banking or financial crises, are rare, but the resulting contractions can last for years before recovery begins.¹ These episodes also coincide with marked changes in cash-flow prospects, uncertainty, and risk compensation, creating pricing conditions distinct from those prevailing in normal times.² The pricing importance of adverse states therefore cannot be inferred from their physical frequency alone. It depends jointly on how rarely they are entered, how long they persist once reached, and how they affect cash flows and discounting, as well as on the payoff horizon.

More generally, economic dynamics may feature regions of the state space associated with distinct economic conditions for which entry intensities are much lower than exit intensities. For an economy initially outside such a region, ignoring rare entry may have little effect on the pricing of near-term payoffs, yet lead to substantial pricing errors for more distant payoffs. Two questions naturally arise. Under what conditions can a region that the economy rarely enters and eventually leaves nevertheless govern how a claim’s price grows or declines as its payoff horizon increases? When this occurs, what determines the horizon at which the region’s influence becomes quantitatively important, and how does that horizon change as entry becomes rarer?³

Two strands of literature address complementary aspects of these questions. Rare-disaster models show how low-probability contractions can have large effects on asset prices (Barro, 2006; Gourio, 2012; Gabaix, 2012; Wachter, 2013), while models with disaster recovery highlight how post-disaster dynamics can reshape term structures (Hasler and Marfè, 2016; Wu, 2025). Furthermore, Martin (2012) shows that exceptionally adverse aggregate outcomes can drive the valuation of a broad class of long-dated assets, even without an explicit disaster mechanism.

Separately, the operator-semigroup approach provides a general framework for studying asset prices across maturities and characterizing their long-run behavior. For each maturity, a pricing operator maps a future payoff function into its current price as a function of the initial economic state. In a time-homogeneous continuous-time Markov setting, these operators form a *semigroup* (Garman, 1985; Hansen and Scheinkman, 2009). The pricing semigroup combines physical state transitions with stochastic discounting and can also incorporate cash-flow growth, thereby providing a unified framework for analyzing the pricing of risk across investment horizons (Hansen and Scheinkman, 2012; Hansen, 2012). In particular, Hansen and Scheinkman (2009) factorize the stochastic discount

¹Using twentieth-century international data, Barro (2006) estimates an annual disaster probability of 1.5–2% and documents peak-to-trough declines in per capita GDP of 15–64%. Using a 10% threshold and data extending back to 1870, Barro and Ursúa (2008) estimate an average peak-to-trough duration of about 3.5 years. In an estimated model of consumption disasters, Nakamura et al. (2013) find that the average disaster reaches its trough after about 6 years and that roughly half of the decline is subsequently reversed.

²Nakamura et al. (2013) find sharply elevated uncertainty about future consumption growth during consumption disasters, with important consequences for asset prices. Muir (2017) finds that risk premia rise substantially during financial crises, while Krishnamurthy and Muir (2025) document a sharp repricing of credit at crisis onset.

³The empirical literature on equity term structures also underscores the importance of the maturity dimension. Evidence from dividend strips and estimated equity term structures documents maturity-dependent risk premia and discount rates (van Binsbergen et al., 2012, 2013; Giglio et al., 2024).

factor (SDF) into an exponential term, a positive martingale, and an eigenfunction ratio. Under suitable recurrence and integrability conditions, the exponential term determines the long-run pricing rate; see also [Qin and Linetsky \(2016\)](#).⁴

In standard finite-state formulations of the operator-semigroup approach, the state process is commonly assumed to be *irreducible*, meaning that every state is accessible from every other. Given a strictly positive SDF process, Perron–Frobenius theory then yields strictly positive left and right principal eigenvectors common to the pricing operators. Consequently, every nonzero nonnegative terminal payoff has a positive loading on the dominant component from every initial state, and all such payoffs share the same long-run pricing rate.⁵

The questions raised above lie at the intersection of these literatures. Rare-disaster and recovery models illuminate the roles of arrival, persistence, recovery, and risk compensation within particular economic specifications. The operator-semigroup approach provides a general framework for pricing across horizons, but its standard irreducible formulations do not by themselves characterize how a region’s contribution to a particular price changes as access to that region becomes rarer. An economy can remain irreducible while approaching a *reducible limit*: letting the intensities of entering into a rare region vanish while retaining transitions out of the region gives a natural reducible approximation. At that limit, the dominant eigenvectors need no longer be strictly positive, and the dominant pricing rate need not govern a given claim’s price from a given initial state. Work on nearly decomposable systems ([Simon and Ando, 1961](#); [Ando and Fisher, 1963](#); [Courtois, 1975](#); [Meyer, 1989](#)) provides related tools for studying weakly connected dynamics, but its primary focus is on aggregation and, in stochastic settings, transition dynamics and stationary distributions rather than the pricing of state-contingent claims.

This paper provides a unified characterization of how rare access to economically distinct states shapes asset prices across initial states, claims, and payoff horizons. We model the economic state as a finite-state, time-homogeneous continuous-time Markov chain, allowing the dynamics to be reducible, and represent prices by a positive pricing semigroup. Using a graph-theoretic approach, our analysis begins with a structural observation: the physical and pricing dynamics partition the state space into the same *communicating classes* (or simply *classes*)—maximal sets of mutually accessible states—and induce the same accessibility relations among them. Yet the two dynamics need not attach the same importance to these classes. Physical dynamics distinguish recurrent from transient classes, whereas pricing dynamics assign each class its own exponential long-run pricing rate. Physical transience therefore need not imply pricing irrelevance: a class with vanishing long-run physical probability (i.e., a class of rare states) can nevertheless have the highest long-run pricing rate.

⁴Related contributions extend long-term factorization to general semimartingale environments ([Qin and Linetsky, 2017](#)), develop recovery results linking state prices to physical probabilities and discounting under additional restrictions ([Ross, 2015](#)), and provide nonparametric methods for estimating the eigenfunction decomposition ([Christensen, 2017](#)). For spectral conditions for the existence and uniqueness of infinite-horizon asset prices, see [Borovička and Stachurski \(2021\)](#).

⁵See, for example, [Hansen and Scheinkman \(2009, Example 6.1 and Section 7.1\)](#), [Hansen \(2012, Section 5\)](#), and [Borovička et al. \(2016, Section I.B\)](#). Appendix C reviews this benchmark in the finite-state irreducible setting.

Our analysis under reducible dynamics yields exact representations of claim prices and their long-maturity asymptotics. In particular, we show that the class-specific pricing rates relevant for a given claim depend jointly on the initial state and the payoff, and we summarize this dependence through an object we call the *pricing corridor*: the collection of classes of states lying on directed paths from the class containing the initial state to classes on which the payoff is positive. Restricting the pricing dynamics to this corridor leaves the claim price unchanged at every maturity. Importantly, when the pricing corridor is nonempty, the highest class-specific pricing rate within the corridor, which we call the *corridor rate*, determines the long-run exponential rate of the claim price. If several classes attain the corridor rate, their accessibility relations determine the degree of the accompanying polynomial factor.

We next consider rare transitions that restore connections absent at the reducible limit. Suppose the limiting economy has a uniquely dominant pricing class that lies outside the claim's pricing corridor and has a higher pricing rate than the corridor rate. Rare transitions can create paths from the initial state through this class to the payoff support. We show that the contribution of the dominant class then enters the claim price with a small coefficient but eventually dominates at sufficiently long maturities.

More specifically, let ε index the strength of the rare connections, with $\varepsilon = 0$ corresponding to the reducible limit and sufficiently small $\varepsilon > 0$ restoring irreducibility. Using a weighted-path analysis related to minimum-resistance methods (Young, 1993; Kandori et al., 1993), we show that the pricing component associated with the dominant class has a coefficient of order ε^d , where $d > 0$ measures the rarity of the connecting paths. Under an additional two-term spectral separation condition, we define the *crossover maturity*, t_ε^c , as the maturity at which this component equals the component associated with the reducible-limit corridor rate. Beyond this maturity, the contribution associated with the dominant class exceeds the corridor-rate component and ultimately governs the claim's price. We derive an explicit expression for t_ε^c and show that

$$t_\varepsilon^c \sim \frac{d}{\Delta} \log \frac{1}{\varepsilon}, \quad \varepsilon \downarrow 0,$$

where $\Delta > 0$ is the difference between the dominant class's pricing rate and the corridor rate. Therefore, holding other quantities fixed, rarer access delays the crossover, whereas a larger pricing-rate advantage brings it forward.

We illustrate these results in a consumption-based disaster-recovery economy with CRRA preferences. Without disaster entry, the disaster shock does not affect prices from normal initial states, but it can govern long-maturity prices starting in disaster even when the claim pays only after recovery. The dominance condition connects risk aversion, consumption dynamics, cash-flow exposure, and recovery intensity. Changing cash-flow growth exposure can change which class is pricing dominant, holding physical dynamics and preferences fixed. Restoring rare entry gives exact two-state prices and crossover maturities from the normal state.

Under an illustrative parameterization with empirically motivated entry frequencies, disaster durations, and consumption growth characteristics, we show that the disaster-associated dominant component accounts for roughly 95% of the price of a fixed real payment due in 25 years, starting

from the normal state. Greater disaster persistence can nevertheless delay crossover: it widens the pricing-rate gap but reduces the dominant-component weight.

The rest of the paper is organized as follows. Section 2 introduces the pricing environment and framework. Sections 3–4 establish the global pricing rate and class-based semigroup asymptotics, and Section 5 derives their state- and payoff-specific counterparts. Section 6 develops the disaster-recovery illustration and its consumption-based specification. Section 7 studies rare transitions, crossover maturities, and the continuation of the application with rare disaster entry. The appendices provide the semimartingale foundation, additional comparative statics, and proofs.

2. ENVIRONMENT AND PRICING SEMIGROUP

This section introduces the finite-state continuous-time environment and the pricing objects used throughout the paper. The economic state follows a finite-state Markov process, while the underlying information flow may include additional shocks to stochastic discounting and cash-flow growth. We model the SDF and cash-flow growth processes as strictly positive multiplicative functionals whose product generates the pricing semigroup. Appendix A provides a semimartingale foundation for this formulation and derives the pricing generator. We use operator and matrix representations interchangeably; Appendix B.1 records their equivalence and collects the associated notation. Throughout, as a terminological convention, we use *pricing* for payoffs delivered at specified maturities and *valuation* for infinite-horizon cash-flow streams.

2.1. Markov environment. Let $\mathsf{X} = \{x_1, \dots, x_n\}$ be a finite state space. Consider a filtered measurable space $(\Omega, \mathcal{F}, (\mathcal{F}_t^\circ)_{t \geq 0})$ supporting a càdlàg state process $X = (X_t)_{t \geq 0}$ that is adapted to $(\mathcal{F}_t^\circ)$ and takes values in X . For each $x \in \mathsf{X}$, let $\mathbb{P}_x$ be a probability measure on $(\Omega, \mathcal{F})$ such that $\mathbb{P}_x(X_0 = x) = 1$, and write $\mathbb{E}_x$ for expectation under $\mathbb{P}_x$. Fix a full-support probability distribution ϖ on X , and define $\mathbb{P} = \sum_x \varpi(x) \mathbb{P}_x$. We denote by $(\mathcal{F}_t)_{t \geq 0}$ the usual right-continuous $\mathbb{P}$ -augmentation of the raw filtration $(\mathcal{F}_t^\circ)_{t \geq 0}$. Let $\mathcal{F}_\infty^\circ := \sigma(\bigcup_{t \geq 0} \mathcal{F}_t^\circ)$, the smallest σ -algebra containing $\mathcal{F}_t^\circ$ for every $t \geq 0$. Let $(\theta_t)_{t \geq 0}$ be a family of measurable shift maps $\theta_t : \Omega \rightarrow \Omega$ satisfying $\theta_0 = \text{id}$, $\theta_{t+u} = \theta_u \circ \theta_t$, and $X_u \circ \theta_t = X_{t+u}$ for all $t, u \geq 0$. We impose the following Markov condition on the economic environment.

Assumption 2.1 (Time-homogeneous Markov environment). *For every $x \in \mathsf{X}$, $t \geq 0$, and bounded $\mathcal{F}_\infty^\circ$ -measurable random variable Y ,*

$$\mathbb{E}_x(Y \circ \theta_t \mid \mathcal{F}_t) = \mathbb{E}_{X_t} Y, \quad \mathbb{P}_x\text{-a.s.} \quad (2.1)$$

Assumption 2.1 is stronger than requiring X alone to be a time-homogeneous Markov process. In particular, taking $Y = f(X_u)$ in (2.1) and using $X_u \circ \theta_t = X_{t+u}$ shows that, under each $\mathbb{P}_x$, X is a time-homogeneous continuous-time Markov chain with respect to $(\mathcal{F}_t)_{t \geq 0}$.

Denote by $\mathbb{R}^\mathsf{X}$ the vector space of real-valued functions on X and by $\mathbb{R}_+^\mathsf{X}$ its cone of nonnegative functions. Equip $\mathbb{R}^\mathsf{X}$ with the supremum norm $\|f\|_\infty := \max_{x \in \mathsf{X}} |f(x)|$, and use $\|\cdot\|_\infty$ also for the

induced operator norm. For $f \in \mathbb{R}^X$, the *transition operators* $(P_t)_{t \geq 0}$ of X are defined by

$$(P_t f)(x) := \mathbb{E}_x f(X_t), \quad t \geq 0. \quad (2.2)$$

Identifying $\mathbb{R}^X$ with $\mathbb{R}^n$, each P_t is represented by an $n \times n$ row-stochastic transition matrix. The family $(P_t)_{t \geq 0}$ satisfies $P_0 = I$ and $P_{t+u} = P_t P_u$ for $t, u \geq 0$. Since X is finite-state and càdlàg, $P_t \rightarrow I$ as $t \rightarrow 0$; hence $(P_t)_{t \geq 0}$ is a continuous matrix semigroup. Its infinitesimal generator, which we call the *Markov generator* of X , is

$$L_X := \lim_{t \downarrow 0} \frac{P_t - I}{t}. \quad (2.3)$$

In particular, L_X is a Metzler matrix (i.e., a matrix with nonnegative off-diagonal entries) with zero row sums, and $P_t = \exp(tL_X)$ for every $t \geq 0$.⁶ We interpret $(\mathbb{P}_x)_{x \in X}$ as the physical probability measures and accordingly also refer to L_X as the *physical generator*, to distinguish it from the pricing generator introduced below.

2.2. Stochastic discounting and cash-flow growth. Let $S = (S_t)_{t \geq 0}$ denote the stochastic discount factor (SDF) process, and let $G = (G_t)_{t \geq 0}$ denote a cumulative stochastic cash-flow growth process. The process S prices payoffs: under $\mathbb{P}_x$, the time-zero price of an $\mathcal{F}_t$ -measurable payoff Z_t delivered at time t is $\mathbb{E}_x(S_t Z_t)$, provided that $\mathbb{E}_x|S_t Z_t| < \infty$.

To represent stochastic discounting and cash-flow growth in a time-homogeneous manner, we use multiplicative functionals, which are standard in operator formulations of long-run valuation; see Hansen and Scheinkman (2009, 2012). A jointly measurable, $(\mathcal{F}_t^\circ)$ -adapted process $M = (M_t)_{t \geq 0}$ is called a *multiplicative functional* if $M_0 = 1$ and

$$M_{t+u} = M_t(M_u \circ \theta_t), \quad t, u \geq 0, \quad \mathbb{P}\text{-a.s.}$$

It is *strictly positive* if $M_t > 0$, $\mathbb{P}$ -almost surely, for every $t \geq 0$. Appendix A.2 shows how exponential additive functionals driven by diffusion, state-transition, and non-transition jump shocks generate such processes.

Assumption 2.2 (Discounting and growth). *The SDF process S and the cash-flow growth process G are strictly positive multiplicative functionals with càdlàg sample paths. Moreover, $\mathbb{E}_x(S_t G_t) < \infty$ for $x \in X$ and $t \geq 0$.*

Appendix A.2 gives a semimartingale parameterization of S and G in terms of continuous shocks, state-transition jumps, and non-transition jumps. Such parameterizations can accommodate SDF and cash-flow processes derived from more specific structural models.⁷

For $0 \leq s \leq t$, multiplicativity and strict positivity imply $S_t/S_s = S_{t-s} \circ \theta_s$ and $G_t/G_s = G_{t-s} \circ \theta_s$, $\mathbb{P}$ -almost surely. For S , this identity expresses intertemporal consistency in dynamic valuation (Garman, 1985; Duffie and Garman, 1986); for G , it is the corresponding compounding rule for

⁶For background on Markov transition semigroups and their infinitesimal generators on general state spaces, see, for example, Applebaum (2009, Sections 3.1–3.2). For continuous-time finite-state Markov chains, the properties of the generator matrix and the transition matrix semigroup are provided in Norris (1997, Sections 2.1 and 2.8).

⁷See Hansen and Scheinkman (2009, Sections 3.2–3.5) for related constructions and examples.

cumulative cash-flow growth. Since the product of two multiplicative functionals is itself multiplicative, the *growth-adjusted SDF* process defined by

$$S_t^G := S_t G_t, \quad t \geq 0,$$

is also a strictly positive multiplicative functional. The moment condition in Assumption 2.2 can therefore be written as $\mathbb{E}_x S_t^G < \infty$ for $x \in \mathbf{X}$ and $t \geq 0$.

2.3. Pricing semigroup and pricing generator. A function $g \in \mathbb{R}^{\mathbf{X}}$ specifies a state-contingent payoff profile: $g(x)$ determines the terminal payoff associated with state x . Thus $G_t g(X_t)$ is the payoff delivered at time t , with G_t capturing cumulative cash-flow growth and $g(X_t)$ capturing its dependence on the economic state at maturity.

For each $t \geq 0$, we call the linear operator $\Pi_t : \mathbb{R}^{\mathbf{X}} \rightarrow \mathbb{R}^{\mathbf{X}}$, defined by

$$(\Pi_t g)(x) := \mathbb{E}_x S_t^G g(X_t) = \mathbb{E}_x S_t G_t g(X_t), \quad x \in \mathbf{X}, \quad g \in \mathbb{R}^{\mathbf{X}}, \quad (2.4)$$

the *pricing operator* for maturity- t payoffs. Thus $(\Pi_t g)(x)$ is the time-zero price of $G_t g(X_t)$, conditional on $X_0 = x$. Assumption 2.2 makes Π_t well defined, since $|(\Pi_t g)(x)| \leq \|g\|_{\infty} \mathbb{E}_x S_t^G < \infty$. Moreover, $S_0^G = 1$ and $X_0 = x$, $\mathbb{P}_x$ -almost surely, imply $\Pi_0 = I$. Since $S_t^G > 0$, we also have $\Pi_t g \geq 0$ whenever $g \geq 0$. Hence Π_t is a positive linear operator. Identifying $\mathbb{R}^{\mathbf{X}}$ with $\mathbb{R}^n$, its representing matrix is entrywise nonnegative. Moreover, Assumptions 2.1 and 2.2 imply the *semigroup* property:

$$\Pi_{t+u} = \Pi_t \Pi_u, \quad t, u \geq 0. \quad (2.5)$$

Thus $\Pi = (\Pi_t)_{t \geq 0}$ is a positive matrix semigroup, which we call the *pricing semigroup*.⁸ To obtain the infinitesimal generator of the pricing semigroup, we impose the following continuity condition.

Assumption 2.3. For every $x \in \mathbf{X}$, $\lim_{t \rightarrow 0} \mathbb{E}_x S_t G_t = 1$.

Together with the strict positivity of S and G , Assumption 2.3 and the càdlàg paths imply that $\Pi_t \rightarrow I$ in operator norm as $t \rightarrow 0$; Appendix A.1 gives the short proof. Hence Π is a continuous matrix semigroup. Its infinitesimal generator, which we call the *pricing generator*, is

$$L_{\Pi} := \lim_{t \downarrow 0} \frac{\Pi_t - I}{t}, \quad (2.6)$$

where the limit is taken in operator norm. Finite-dimensional semigroup theory then gives

$$\Pi_t = e^{tL_{\Pi}}, \quad t \geq 0.$$

Because Π is a positive matrix semigroup, its generator L_{Π} is a Metzler matrix.⁹ Unlike the physical generator L_X , however, L_{Π} need not have zero row sums, because it incorporates stochastic

⁸This semigroup is the counterpart of the cash-flow-growth semigroup Q in Hansen and Scheinkman (2009, Section 4.4 and Table I). Hansen and Scheinkman describe this object as valuation with stochastic growth. Under our terminology, we call it a pricing semigroup because Π_t assigns a time-zero price to a payoff delivered at maturity t . Appendix A.1 verifies the semigroup property under our assumptions.

⁹For the definition of the infinitesimal generator and the exponential representation of a continuous matrix semigroup, see Bátkai et al. (2017, Proposition 9.4 and Definition 9.14). The equivalence between positivity of the matrix semigroup and the Metzler property of its generator is given in their Theorem 7.1.

discounting and cash-flow growth in addition to physical state transitions. Appendix A.3 derives the entries of L_Π from the semimartingale specification and gives its Feynman–Kac decomposition into pricing-adjusted transition dynamics and state-specific effective discounting.

2.4. Strip prices, yields, and infinite-horizon values. For $t \geq 0$ and $g \in \mathbb{R}^X$, define the *maturity- t strip price*

$$q_t(x, g) := (\Pi_t g)(x) = (e^{tL_\Pi} g)(x), \quad x \in X. \quad (2.7)$$

In light of (2.4), $q_t(x, g)$ is the time-zero price of the payoff $G_t g(X_t)$, delivered at time t , conditional on $X_0 = x$. For $t > 0$, whenever $q_t(x, g) > 0$, define the corresponding continuously compounded strip yield by

$$y_t(x, g) := -\frac{1}{t} \log q_t(x, g), \quad x \in X.$$

Let $\mathbf{1} \in \mathbb{R}^X$ denote the function identically equal to one. The special case $g = \mathbf{1}$ gives

$$z_t(x) := q_t(x, \mathbf{1}) = (\Pi_t \mathbf{1})(x), \quad t \geq 0, \quad y_t(x) := -\frac{1}{t} \log z_t(x), \quad t > 0. \quad (2.8)$$

Here $z_t(x)$ is the time-zero price of the payoff G_t delivered at date t . We call z_t the *growth-adjusted zero-coupon price* and y_t its continuously compounded yield. Since $S_t^G > 0$, $z_t(x) > 0$ for every $x \in X$ and $t \geq 0$.

To pass from single-maturity payoffs to infinite-horizon assets¹⁰, fix $g \in \mathbb{R}_+^X$ and consider an asset paying the instantaneous cash-flow rate $G_t g(X_t)$ for $t \geq 0$. Its cash flow over an interval $[a, b]$ is $\int_a^b G_t g(X_t) dt$. We define the *valuation operator* V for such assets by

$$(Vg)(x) := \mathbb{E}_x \int_0^\infty S_t G_t g(X_t) dt = \int_0^\infty q_t(x, g) dt \in [0, \infty], \quad x \in X. \quad (2.9)$$

The equality follows from Tonelli’s theorem. At this stage, V may take the value ∞ ; under the global valuation stability condition introduced in the next section, it extends to a finite positive linear operator.

3. GLOBAL LONG-RUN PRICING RATE AND VALUATION STABILITY

In this section, we identify the global long-run pricing rate and relate its sign to the finiteness of infinite-horizon values. The starting point is the *spectral bound* of the pricing generator. For a square matrix A , let $\sigma(A)$ denote its spectrum, and $s(A) := \max\{\Re \lambda : \lambda \in \sigma(A)\}$ denote its spectral bound. Define

$$\lambda^* := s(L_\Pi).$$

Since L_Π is Metzler, Lemma B.1 implies that

$$\lambda^* \in \sigma(L_\Pi) \cap \mathbb{R}, \quad \Re \mu < \lambda^* \text{ for every } \mu \in \sigma(L_\Pi) \setminus \{\lambda^*\}. \quad (3.1)$$

That is, λ^* is a real eigenvalue of L_Π whose real part strictly exceeds that of every distinct eigenvalue.

¹⁰Examples include common equity, modeled as a claim to an indefinite dividend stream, and perpetuities such as consols, which pay coupons without a fixed maturity.

If L_Π is irreducible, λ^* is the algebraically simple principal eigenvalue associated with the strictly positive eigenfunction/eigenvector used in the Hansen–Scheinkman factorization of the growth-adjusted SDF; Appendix C reviews this case and its implications for long-run pricing. Under reducibility, λ^* need not be algebraically simple or admit strictly positive eigenvectors, and hence it need not govern every individual price.

To identify the global pricing interpretation of λ^* without imposing irreducibility, define the *state-wise zero-coupon price envelope* and its associated *envelope yield* by

$$z_t^{\text{env}} := \max_{x \in \mathbf{X}} z_t(x), \quad t \geq 0, \quad y_t^{\text{env}} := -\frac{1}{t} \log z_t^{\text{env}} = \min_{x \in \mathbf{X}} y_t(x), \quad t > 0.$$

The following proposition relates their long-maturity behavior to λ^* .

Proposition 3.1. *Under Assumptions 2.1–2.3,*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log z_t^{\text{env}} = \lambda^*, \quad \lim_{t \rightarrow \infty} y_t^{\text{env}} = -\lambda^*.$$

Proposition 3.1 identifies λ^* as the *global long-run pricing rate*: it is the asymptotic exponential rate of the statewise price envelope, while $-\lambda^*$ is the corresponding long-run envelope yield.¹¹ In fact, the sign of this rate also determines whether every modeled nonnegative infinite-horizon cash-flow stream has finite value from every initial state. We call this property *global valuation stability*.

Definition 1. We say that *global valuation stability* holds if

$$(Vg)(x) < \infty \quad \text{for every } g \in \mathbb{R}_+^{\mathbf{X}} \text{ and } x \in \mathbf{X}. \quad (3.2)$$

Because $0 \leq g \leq \|g\|_\infty \mathbf{1}$ for every $g \in \mathbb{R}_+^{\mathbf{X}}$ and Π_t is positive, $0 \leq q_t(x, g) \leq \|g\|_\infty z_t(x)$. Consequently, global valuation stability holds if and only if $(V\mathbf{1})(x) = \int_0^\infty z_t(x) dt < \infty$ for every $x \in \mathbf{X}$. Thus this finiteness condition can be tested using the constant function $g = \mathbf{1}$ alone. The following theorem links this condition to λ^* , the price envelope, and the yield envelope.

Theorem 3.2 (Global valuation stability). *Under Assumptions 2.1–2.3, the following statements are equivalent:*

- (G1) *Global valuation stability holds.*
- (G2) *The global long-run pricing rate is negative: $\lambda^* < 0$.*
- (G3) *The long-run envelope yield is positive: $\lim_{t \rightarrow \infty} y_t^{\text{env}} > 0$.*
- (G4) *The price envelope is below one at some maturity: $z_t^{\text{env}} < 1$ for some $t > 0$.*
- (G5) *The price envelope vanishes at long maturities: $\lim_{t \rightarrow \infty} z_t^{\text{env}} = 0$.*

When these conditions hold, the infinite-horizon valuation operator V extends uniquely to a positive linear operator on $\mathbb{R}^{\mathbf{X}}$, and the operator-valued integral $\int_0^\infty \Pi_t dt$ converges in norm. Moreover, L_Π is invertible and V admits the representation

$$V = \int_0^\infty \Pi_t dt = -L_\Pi^{-1}.$$

¹¹The same rate also governs ex ante zero-coupon prices under any full-support distribution of the initial state; see Corollary D.1 in Appendix D.1.

Theorem 3.2 provides equivalent spectral, yield-based, and price-based characterizations of global valuation stability. Specifically, global valuation stability holds if and only if the global long-run pricing rate is negative, or equivalently, the long-run envelope yield is positive. It also holds if and only if the price envelope falls below one at some finite maturity, in which case the envelope vanishes at long maturities.

The sign condition $\lambda^* = s(L_\Pi) < 0$ is stronger than invertibility of L_Π . For example, in a one-state economy with $L_\Pi = a > 0$, the equation $L_\Pi v = -1$ has the solution $v = -1/a$, but this negative number cannot be the value of a nonnegative cash-flow stream, whose infinite-horizon value is infinite. Thus algebraic solvability of $L_\Pi v = -1$ does not ensure that $-L_\Pi^{-1}$ is a positive valuation operator; stability, rather than invertibility alone, justifies the inverse-generator representation.

The global rate λ^* and the valuation stability criterion provide an economy-wide characterization, but do not identify which parts of the state space determine that rate or which prices inherit it. Addressing these questions requires examining which states can reach one another and how this accessibility structure interacts with pricing.

4. CLASS STRUCTURE AND PRICING SEMIGROUP ASYMPTOTICS

We organize the state space into communicating classes and show how their class-specific pricing rates and accessibility relations determine the pricing semigroup's long-run behavior. In particular, the physical and pricing dynamics share this class structure, but physical recurrence need not coincide with pricing dominance. The matrix and graph-theoretic preliminaries are collected in Appendix B.

4.1. Communicating classes and pricing rates. We call $\mathcal{G}(L_X)$ and $\mathcal{G}(L_\Pi)$ the *physical graph* and *pricing graph*, respectively. In each graph, directed edges correspond to positive off-diagonal entries of the generator; see Appendix B.3 for the graph conventions.

Strict positivity of the growth-adjusted SDF ensures that the two graphs have the same accessibility relations. Specifically, by (2.2) and (2.4), for any states x_i, x_j , $(P_t)_{ij} = \mathbb{P}_{x_i}(X_t = x_j)$ and $(\Pi_t)_{ij} = \mathbb{E}_{x_i}[S_t^G 1_{\{X_t = x_j\}}]$. Since $S_t^G > 0$ almost surely, $(P_t)_{ij} > 0$ if and only if $(\Pi_t)_{ij} > 0$ for all $t > 0$. Lemma B.2 therefore implies that the physical and pricing graphs induce the same accessibility relation on $\mathbf{X}$. Because communication is mutual accessibility, they also partition $\mathbf{X}$ into the same communicating classes. Denote these classes by $\mathcal{C}_1, \dots, \mathcal{C}_m$, and write

$$\mathcal{C} := \{\mathcal{C}_1, \dots, \mathcal{C}_m\}.$$

We call the economy *irreducible* if all states belong to a single class ($m = 1$), and *reducible* otherwise. Clearly, this is equivalent to irreducibility or reducibility of either L_X or L_Π .

Collapsing each class into a single vertex gives the *reduced physical graph* $\mathcal{G}_r(L_X)$ and the *reduced pricing graph* $\mathcal{G}_r(L_\Pi)$. These graphs share the vertex set $\mathcal{C}$ and the same class-level accessibility relation. We write $x_i \rightsquigarrow x_j$ and $\mathcal{C}_k \rightsquigarrow \mathcal{C}_\ell$ for the common state- and class-level accessibility

relations, respectively. Paths of length zero are allowed, so every state and class is accessible from itself. Unless stated otherwise, $\mathcal{C}_k \rightarrow \mathcal{C}_\ell$ denotes a directed edge of the reduced pricing graph.

A class is *closed* if it has no access to any other class. The common accessibility relation therefore gives the physical and pricing graphs the same closed classes. Under the finite-state physical Markov dynamics, closed classes are *recurrent*, while all other classes are *transient*. From every initial state, the probability of occupying a transient class tends to zero as time increases.

The reduced graphs are acyclic, so we may order the classes such that $\mathcal{C}_k \rightsquigarrow \mathcal{C}_\ell$ with $k \neq \ell$ implies $k > \ell$. We place all closed classes first, in arbitrary order, and order the remaining classes subject to this requirement. The corresponding ordering of states puts the pricing generator L_Π in block lower-triangular Frobenius form (Schneider, 1986, Section 2):

$$L_\Pi = \begin{pmatrix} L_1 & 0 & 0 & \cdots & 0 \\ L_{21} & L_2 & 0 & \cdots & 0 \\ L_{31} & L_{32} & L_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ L_{m1} & L_{m2} & L_{m3} & \cdots & L_m \end{pmatrix}. \quad (4.1)$$

Each diagonal block L_k is the principal submatrix of L_Π indexed by the states in $\mathcal{C}_k$. Since each $\mathcal{C}_k$ is a class of the pricing graph, L_k is an irreducible Metzler matrix. For $k > \ell$, the off-diagonal block $L_{k\ell}$ is nonzero if and only if the reduced pricing graph contains the directed edge $\mathcal{C}_k \rightarrow \mathcal{C}_\ell$.

For each class, define its *class long-run pricing rate* (or simply *class pricing rate*) by

$$\lambda_k := s(L_k), \quad k = 1, \dots, m.$$

By Perron–Frobenius theory for irreducible Metzler matrices, λ_k is an algebraically simple real eigenvalue of L_k that is strictly dominant in real part over all other eigenvalues. Since L_Π is block triangular, $\sigma(L_\Pi) = \bigcup_{k=1}^m \sigma(L_k)$, and therefore

$$\lambda^* = s(L_\Pi) = \max_{1 \leq k \leq m} \lambda_k. \quad (4.2)$$

We say a class $\mathcal{C}_k \in \mathcal{C}$ is *globally pricing-dominant* when $\lambda_k = \lambda^*$. Let

$$\mathcal{C}^* := \{\mathcal{C}_k \in \mathcal{C} : \lambda_k = \lambda^*\}$$

denote the set of globally pricing-dominant classes. Together with (4.2), Theorem 3.2 implies that global valuation stability holds if and only if $\lambda_k < 0$ for every $k = 1, \dots, m$.

The Frobenius form (4.1) separates class-level accessibility from class pricing rates. Given the class decomposition, the zero pattern of the off-diagonal blocks determines the reduced pricing graph and, through the common accessibility relation, the recurrent–transient classification under the physical dynamics. The diagonal blocks, by contrast, determine the class pricing rates and hence the global rate λ^* . Thus, global pricing dominance is a diagonal-block property.

Figure 1 illustrates a pricing graph and its reduced graph in a three-class disaster-recovery economy comprising a disaster class $\mathcal{D}$, a recovery class $\mathcal{R}$, and a normal class $\mathcal{N}$. Communication within each class, together with the interclass edges shown in panel (a), gives every disaster state access




FIGURE 1. The pricing graph and its reduced graph in a disaster-recovery economy. Panel (b) collapses the classes of the pricing graph in panel (a). The double border identifies $\mathcal{N}$ as the unique closed class and hence the recurrent class supporting the long-run physical distribution.

to recovery and every recovery state access to normal. Collapsing the classes therefore yields the reduced graph $\mathcal{D} \rightarrow \mathcal{R} \rightarrow \mathcal{N}$, which implies the class-level accessibility relations $\mathcal{D} \rightsquigarrow \mathcal{R}$, $\mathcal{R} \rightsquigarrow \mathcal{N}$, and $\mathcal{D} \rightsquigarrow \mathcal{N}$. Because the physical and pricing graphs induce the same class-level accessibility relation, $\mathcal{N}$ is the unique closed class in both reduced graphs and therefore the unique recurrent class under the physical dynamics. The nonclosed classes $\mathcal{D}$ and $\mathcal{R}$ are transient. Consequently, the state process eventually enters the normal class with probability one, and the probability of occupying that class tends to one.

Pricing dominance need not coincide with physical recurrence. Once discounting, risk adjustment, and cash-flow growth are incorporated into the class blocks of L_Π , the recurrent normal class need not have the largest pricing rate. If the disaster block has the largest spectral bound, the physically transient class $\mathcal{D}$ determines λ^* even though its physical probability vanishes asymptotically. If this rate is nonnegative, global valuation stability fails despite eventual physical recovery. Thus, a physically transient class can determine the global long-run pricing rate.

4.2. Unique pricing dominance. Suppose that a single class attains the global pricing rate, $\mathcal{C}^* = \{\mathcal{C}^*\}$. This condition allows the economy to be reducible. The following theorem shows that the exponentially normalized pricing semigroup converges to a rank-one spectral projection, whose right and left eigenvectors have supports determined by accessibility to and from $\mathcal{C}^*$, respectively.

Theorem 4.1. *Under Assumptions 2.1–2.3, suppose that $\mathcal{C}^* = \{\mathcal{C}^*\}$. Then λ^* is an algebraically simple eigenvalue of L_Π . There exist unique nonzero nonnegative right and left eigenvectors φ and ψ , normalized by $\mathbf{1}^\top \varphi = 1$ and $\psi^\top \varphi = 1$, such that $L_\Pi \varphi = \lambda^* \varphi$ and $L_\Pi^\top \psi = \lambda^* \psi$. Moreover, $\varphi \psi^\top$ is the rank-one spectral projection of L_Π associated with λ^* , and*

$$\lim_{t \rightarrow \infty} e^{-\lambda^* t} \Pi_t = \varphi \psi^\top. \quad (4.3)$$

The supports of these eigenvectors are determined by class-level accessibility. For each $\mathcal{C} \in \mathcal{C}$, let $\varphi_{\mathcal{C}}$ and $\psi_{\mathcal{C}}$ denote the subvectors of φ and ψ indexed by the states in $\mathcal{C}$. Then,

$$\varphi_{\mathcal{C}} \begin{cases} \gg 0, & \text{if } \mathcal{C} \rightsquigarrow \mathcal{C}^*, \\ = 0, & \text{otherwise,} \end{cases} \quad \psi_{\mathcal{C}} \begin{cases} \gg 0, & \text{if } \mathcal{C}^* \rightsquigarrow \mathcal{C}, \\ = 0, & \text{otherwise.} \end{cases} \quad (4.4)$$

For every $x \in \mathbf{X}$ and nonzero $g \in \mathbb{R}_+^{\mathbf{X}}$, (4.3) implies that

$$e^{-\lambda^* t} q_t(x, g) \rightarrow \varphi(x) \psi^\top g, \quad \text{as } t \rightarrow \infty.$$

Let $\mathcal{C}$ be the class containing x . By (4.4), the limit is positive precisely when $\mathcal{C} \rightsquigarrow \mathcal{C}^* \rightsquigarrow \mathcal{C}'$ for some class $\mathcal{C}'$ on which g is nonzero. Thus the global pricing rate governs the price's leading asymptotic behavior exactly when the initial state can reach the dominant class and that class can reach a state where the payoff is positive.

The dominant eigenvectors also give a Markov interpretation of long-run pricing. Under reducibility, the right eigenvector may vanish, so the Hansen–Scheinkman change of measure is available only from states in its support. Let φ_+ be the restriction of φ to $\text{supp}(\varphi)$ and $L_\Pi^{\varphi_+}$ the corresponding principal submatrix of L_Π . The *right Doob chain* has generator

$$L_\Pi^{\rightarrow} := \text{diag}(\varphi_+)^{-1} (L_\Pi^{\varphi_+} - \lambda^* I) \text{diag}(\varphi_+).$$

This is the Doob h -transform with $h = \varphi_+$ (see, e.g., Doob, 1957). Under irreducibility, it reduces to the whole-state Hansen–Scheinkman transform reviewed in Appendix C.

The right Doob chain has $\mathcal{C}^*$ as its unique recurrent class, and its distribution converges from every state in $\text{supp}(\varphi)$ to $\pi^{\text{LR}} := \varphi \odot \psi$, extended by zero outside that support. We call π^{LR} the *long-run pricing law*. It is supported on $\mathcal{C}^*$, so a physically transient class can carry all long-run pricing mass.¹² Appendix D.3 establishes these claims and the change-of-measure construction. It also constructs the left Doob chain from the adjoint pricing semigroup. The left chain has the same long-run pricing law, but its state space can differ.

Returning to Figure 1, suppose the disaster class $\mathcal{D}$ is uniquely pricing-dominant. Then $\text{supp}(\varphi) = \mathcal{D}$ and $\text{supp}(\psi) = \mathbf{X}$. Thus the disaster pricing rate governs every nonzero nonnegative payoff priced from a disaster state, including payoffs delivered only after recovery, but does not govern prices starting in recovery or normal states. The right Doob chain remains within $\mathcal{D}$, and its long-run pricing law is concentrated there, although the physical state process eventually enters $\mathcal{N}$.

4.3. Tied pricing dominance. When several classes attain the global pricing rate λ^* , their accessibility relations determine whether polynomial factors accompany the leading exponential term. The relevant quantity is the largest number of globally dominant classes that can be visited along a single directed path. We call this number the *global critical depth* and define

$$\nu^* := \max_{\substack{\gamma \text{ is a chain of classes} \\ \gamma \subseteq \mathcal{C}}} \#\{\mathcal{C}_k \in \gamma : \lambda_k = \lambda^*\}. \quad (4.5)$$

Here *chains* are defined by accessibility in the reduced pricing graph $\mathcal{G}_r(L_\Pi)$; see Appendix B.3.¹³

¹²The same weights determine the sensitivity of the long-run envelope yield to state-specific effective discount rates; see Appendix A.4.

¹³Upon a scalar shift that makes L_Π nonnegative, the globally dominant classes in $\mathcal{C}^*$ are precisely the *basic classes* of Rothblum (1975). In that terminology, the *length* of a chain is the number of basic classes it contains; ν^* is the maximum such length over all chains.

The following theorem describes the long-maturity asymptotics of the pricing semigroup, allowing for multiple globally dominant classes. Let $P^* := P_{\lambda^*}(L_\Pi)$ be the spectral projection onto the generalized eigenspace associated with λ^* , and define $B^* := [(L_\Pi - \lambda^* I)P^*]^{\nu^*-1}P^*/(\nu^* - 1)!$.

Theorem 4.2. *Under Assumptions 2.1–2.3, the largest Jordan block of L_Π associated with λ^* has size ν^* . Moreover,*

$$\lim_{t \rightarrow \infty} t^{1-\nu^*} e^{-\lambda^* t} \Pi_t = B^*, \quad (4.6)$$

where B^* is a nonzero, nonnegative matrix satisfying $L_\Pi B^* = B^* L_\Pi = \lambda^* B^*$. In particular, λ^* is semisimple if and only if no chain of classes contains more than one class in $\mathcal{C}^*$; in that case $\nu^* = 1$ and $B^* = P^*$.

Theorem 4.2 is the continuous-time Metzler formulation of classical reduced-graph index results (Rothblum, 1975; Friedland and Schneider, 1980; Schneider, 1986). It covers (4.3) as a special case: when there is a unique pricing dominant class, $\nu^* = 1$ and $B^* = P^* = \varphi\psi^\top$. If $\nu^* > 1$, then the long-run behavior of the pricing operators includes an additional polynomial factor.¹⁴

For a given initial state x and payoff g , the price term associated with the global exponential rate and polynomial order has coefficient $(B^*g)(x)$, which may be zero. A positive price may then have a lower exponential rate, or retain the global exponential rate with a lower polynomial order. In the next section, we refine the analysis and identify the leading rate and order by restricting the pricing generator to classes lying on paths from the initial state to states where the payoff is positive.

5. PRICING CORRIDORS AND STATE- AND PAYOFF-SPECIFIC ASYMPTOTICS

For a given initial state and nonnegative payoff, restricting the pricing generator and payoff to the pricing corridor defined below preserves the price at every maturity. The restricted generator determines the price's asymptotic exponential rate and polynomial order, as well as whether the corresponding perpetual cash-flow stream has finite value.

5.1. Pricing corridors and restricted pricing generators. For $g \in \mathbb{R}_+^\mathbf{X}$, define its state-level support by $\text{supp}(g) := \{x \in \mathbf{X} : g(x) > 0\}$. Writing $g_\mathcal{C}$ for its restriction to $\mathcal{C} \in \mathcal{C}$, define its class-level support by $\mathcal{S}(g) := \{\mathcal{C} \in \mathcal{C} : g_\mathcal{C} \neq 0\}$. For an initial class $\mathcal{C} \in \mathcal{C}$, we call the following collection of classes the *pricing corridor*:

$$\mathcal{K}(\mathcal{C}, g) := \{\mathcal{C}' \in \mathcal{C} : \mathcal{C} \rightsquigarrow \mathcal{C}' \rightsquigarrow \mathcal{C}'' \text{ for some } \mathcal{C}'' \in \mathcal{S}(g)\}. \quad (5.1)$$

Thus, $\mathcal{K}(\mathcal{C}, g)$ contains exactly the classes lying on at least one directed class path from the current class $\mathcal{C}$ to a class on which the payoff g is nonzero.

Figure 2 extends the disaster–recovery example by adding liquidation and fiscal-stress classes. With initial class $\mathcal{D}$ and payoff support $\mathcal{S}(g) = \{\mathcal{D}, \mathcal{N}\}$, the pricing corridor is $\mathcal{K}(\mathcal{D}, g) = \{\mathcal{D}, \mathcal{R}, \mathcal{N}\}$. Liquidation is accessible from disaster but cannot reach the payoff support; fiscal stress can reach

¹⁴Tied class rates alone do not produce polynomial factors: if no globally dominant class can reach another, $\nu^* = 1$ and exponential normalization suffices.




FIGURE 2. The pricing corridor in an extension of the reduced pricing graph in Figure 1(b). For current class $\mathcal{D}$ and payoff support $\mathcal{S}(g) = \{\mathcal{D}, \mathcal{N}\}$, the shaded classes form the corridor $\mathcal{K}(\mathcal{D}, g) = \{\mathcal{D}, \mathcal{R}, \mathcal{N}\}$: each is accessible from disaster and can reach a class where the payoff is positive. Liquidation cannot reach the payoff support, while fiscal stress is inaccessible from disaster.

the payoff support but is inaccessible from disaster. For the same payoff, $\mathcal{K}(\mathcal{N}, g) = \{\mathcal{N}\}$ and $\mathcal{K}(\mathcal{L}, g) = \emptyset$. Thus, holding the pricing graph and payoff fixed, the corridor may contain several classes, one class, or no class, depending on the initial state.

Restricting the pricing generator and payoff to the corridor leaves the price unchanged at every maturity. To state this precisely, for a nonempty corridor $\mathcal{K} = \mathcal{K}(\mathcal{C}, g)$, let

$$\mathbf{X}_{\mathcal{K}} := \bigcup_{\mathcal{C}' \in \mathcal{K}} \mathcal{C}', \quad L_{\mathcal{K}} := L_{\Pi}(\mathbf{X}_{\mathcal{K}}, \mathbf{X}_{\mathcal{K}}), \quad g_{\mathcal{K}} := g_{\mathbf{X}_{\mathcal{K}}}.$$

The *restricted* pricing generator $L_{\mathcal{K}}$ is the principal submatrix of L_{Π} indexed by the corridor states.

Proposition 5.1. *Under Assumptions 2.1–2.3, let $g \in \mathbb{R}_+^{\mathbf{X}}$ and fix $\mathcal{C} \in \mathcal{C}$.*

- (i) *If $\mathcal{K}(\mathcal{C}, g) = \emptyset$, then $q_t(x, g) = 0$ for every $x \in \mathcal{C}$ and $t \geq 0$.*
- (ii) *If $\mathcal{K}(\mathcal{C}, g) \neq \emptyset$, then $\mathcal{C} \in \mathcal{K}(\mathcal{C}, g)$, $\mathcal{S}(g) \cap \mathcal{K}(\mathcal{C}, g) \neq \emptyset$, and $g_{\mathcal{K}} \neq 0$. The reduced graph of $L_{\mathcal{K}}$ is the subgraph of the reduced pricing graph induced by $\mathcal{K}(\mathcal{C}, g)$, and*

$$q_t(x, g) = (e^{tL_{\Pi}}g)(x) = (e^{tL_{\mathcal{K}}}g_{\mathcal{K}})(x), \quad x \in \mathcal{C}, \quad t \geq 0. \quad (5.2)$$

Moreover, $q_t(x, g) > 0$ for every $x \in \mathcal{C}$ and $t > 0$.

Proposition 5.1 distinguishes two cases. If the corridor is empty, no state with a positive payoff is accessible from the current class, so the price is zero at every maturity. If the corridor is nonempty, the price is strictly positive at every positive maturity and can be computed entirely from the restricted generator and payoff. Every class path from $\mathcal{C}$ to $\mathcal{S}(g)$ lies entirely within the corridor, and the restricted generator retains every such path. The restriction is therefore exact at every maturity, not merely asymptotic.

5.2. Corridor-specific price and yield asymptotics. The exact corridor restriction allows us to apply Theorem 4.2 to the restricted pricing generator and thereby characterize the price of a given maturity- t payoff from a given initial state. The price's asymptotic exponential rate and

polynomial order are determined by the highest class pricing rate within the corridor and the largest number of classes attaining that rate along a single directed path within it.

Specifically, for a nonempty pricing corridor $\mathcal{K}(\mathcal{C}, g)$, define the *corridor long-run pricing rate* (or simply *corridor pricing rate*) by

$$\lambda(\mathcal{C}, g) := \max_{\mathcal{C}_k \in \mathcal{K}(\mathcal{C}, g)} \lambda_k = s(L_{\mathcal{K}}). \quad (5.3)$$

We call a class $\mathcal{C}_k \in \mathcal{K}(\mathcal{C}, g)$ *corridor-critical* when $\lambda_k = \lambda(\mathcal{C}, g)$. The *critical depth* is the largest number of corridor-critical classes that can be visited along a single directed path:

$$\nu(\mathcal{C}, g) := \max_{\substack{\gamma \text{ is a chain of classes} \\ \gamma \subseteq \mathcal{K}(\mathcal{C}, g)}} \#\{\mathcal{C}_k \in \gamma : \lambda_k = \lambda(\mathcal{C}, g)\}.$$

In Figure 2, the corridor from $\mathcal{D}$ satisfies $\mathcal{D} \rightsquigarrow \mathcal{R} \rightsquigarrow \mathcal{N}$ and hence forms a chain of classes. Its critical depth is one if only one of these classes is critical, two if two are critical, and three if all three are critical.

The corridor-critical classes are precisely the globally dominant classes of the restricted pricing generator $L_{\mathcal{K}}$. Theorem 4.2, applied to $L_{\mathcal{K}}$, therefore identifies $\nu(\mathcal{C}, g)$ as the size of its largest Jordan block at $\lambda(\mathcal{C}, g)$. To state the coefficient of the leading price term, abbreviate $\lambda = \lambda(\mathcal{C}, g)$ and $\nu = \nu(\mathcal{C}, g)$, let $P_{\mathcal{K}} := P_{\lambda}(L_{\mathcal{K}})$ be the spectral projection associated with λ , and define

$$B_{\mathcal{K}} := [(L_{\mathcal{K}} - \lambda I)P_{\mathcal{K}}]^{\nu-1}P_{\mathcal{K}}/(\nu-1)!. \quad (5.4)$$

An equivalent eigenvector representation of $B_{\mathcal{K}}$ is given in Appendix D.4.

Theorem 5.2. *Under Assumptions 2.1–2.3, let $g \in \mathbb{R}_+^{\mathbf{X}}$ be nonzero and fix $\mathcal{C} \in \mathcal{C}$. Suppose that $\mathcal{K}(\mathcal{C}, g) \neq \emptyset$. Then, for every $x \in \mathcal{C}$,*

$$\lim_{t \rightarrow \infty} t^{1-\nu(\mathcal{C}, g)} e^{-\lambda(\mathcal{C}, g)t} q_t(x, g) = (B_{\mathcal{K}} g_{\mathcal{K}})(x) =: c_x(g) > 0. \quad (5.5)$$

The corresponding yield satisfies

$$y_t(x, g) = -\lambda(\mathcal{C}, g) - [\nu(\mathcal{C}, g) - 1] \frac{\log t}{t} - \frac{1}{t} \log c_x(g) + o(t^{-1}), \quad \text{as } t \rightarrow \infty. \quad (5.6)$$

Unlike the global coefficient $(B^*g)(x)$, which may vanish, the corridor coefficient $c_x(g) = (B_{\mathcal{K}}g_{\mathcal{K}})(x)$ in (5.5) is strictly positive. Thus $\lambda(\mathcal{C}, g)$ and $\nu(\mathcal{C}, g) - 1$ are the actual asymptotic exponential rate and polynomial order of the price. For a fixed pricing generator, these depend only on the initial class and payoff support: changing payoff magnitudes without changing the support, or moving the initial state within its class, can change $c_x(g)$ but not the rate or order.

In an irreducible economy, $\mathcal{C} = \{\mathbf{X}\}$, and every nonzero $g \in \mathbb{R}_+^{\mathbf{X}}$ satisfies $\lambda(\mathbf{X}, g) = \lambda^*$, $\nu(\mathbf{X}, g) = 1$. Hence, for every initial state $x \in \mathbf{X}$,

$$q_t(x, g) \sim c_x(g) e^{\lambda^* t}, \quad y_t(x, g) = -\lambda^* - t^{-1} \log c_x(g) + o(t^{-1}), \quad \text{as } t \rightarrow \infty.$$

Reducibility allows not only the leading coefficient $c_x(g)$ but also the asymptotic exponential rate and polynomial order to depend on the initial class and payoff support.

If the corridor has a unique critical class, then $\nu(\mathcal{C}, g) = 1$, so the leading price term is $c_x(g)e^{\lambda(\mathcal{C}, g)t}$, and (5.4) gives $B_{\mathcal{K}} = P_{\mathcal{K}}$. Moreover, Theorem 4.1, applied to $L_{\mathcal{K}}$, implies that $\lambda(\mathcal{C}, g)$ is an algebraically simple eigenvalue of $L_{\mathcal{K}}$. Let $\varphi_{\mathcal{K}}$ and $\psi_{\mathcal{K}}$ be the corresponding nonzero nonnegative right and left eigenvectors, normalized by $\mathbf{1}^\top \varphi_{\mathcal{K}} = 1$ and $\psi_{\mathcal{K}}^\top \varphi_{\mathcal{K}} = 1$. The spectral projection and leading price coefficient therefore satisfy $P_{\mathcal{K}} = \varphi_{\mathcal{K}} \psi_{\mathcal{K}}^\top$ and $c_x(g) = \varphi_{\mathcal{K}}(x) \psi_{\mathcal{K}}^\top g_{\mathcal{K}} > 0$. With multiple critical classes, $\nu(\mathcal{C}, g) = 1$ still holds if no directed path within the corridor contains two of them. Otherwise, the leading price term includes the polynomial factor $t^{\nu(\mathcal{C}, g)-1}$.

5.3. Perpetual values and state-dependent yields. In a reducible economy, a perpetual cash-flow stream can have finite value from a given initial state even when global valuation stability fails. The positive leading coefficient in Theorem 5.2 implies that, for a nonempty corridor, the price is integrable over maturities exactly when the corridor pricing rate is negative. Together with the zero-price conclusion of Proposition 5.1 for an empty corridor, this gives the following corollary.

Corollary 5.3. *Under Assumptions 2.1–2.3, let $g \in \mathbb{R}_+^X$ and $x \in \mathcal{C}$. If $\mathcal{K}(\mathcal{C}, g) = \emptyset$, then $(Vg)(x) = 0$. Otherwise, $(Vg)(x) < \infty$ if and only if $\lambda(\mathcal{C}, g) < 0$.*

For a nonempty corridor, finite value requires all class pricing rates within the corridor to be negative, not all class rates in the economy. Classes with nonnegative pricing rates therefore do not prevent finite value when they lie outside the relevant corridor.

For $g = \mathbf{1}$, every class supports the payoff, so $\mathcal{K}(\mathcal{C}, \mathbf{1}) = \{\mathcal{C}_k \in \mathcal{C} : \mathcal{C} \rightsquigarrow \mathcal{C}_k\}$ consists of all classes accessible from $\mathcal{C}$. Consequently, every modeled nonnegative perpetual cash-flow stream has finite value from every state in an initial class $\mathcal{C}$ if and only if

$$\max_{\mathcal{C}_k: \mathcal{C} \rightsquigarrow \mathcal{C}_k} \lambda_k < 0.$$

Requiring this condition for every initial class recovers global valuation stability. Applying the yield expansion (5.6) in Theorem 5.2 to $g = \mathbf{1}$ gives the corresponding long-run zero-coupon yield:

Corollary 5.4 (State-dependent long-run zero-coupon yields). *Under Assumptions 2.1–2.3, for every $\mathcal{C} \in \mathcal{C}$ and $x \in \mathcal{C}$,*

$$\lim_{t \rightarrow \infty} y_t(x) = - \max_{\mathcal{C}_k: \mathcal{C} \rightsquigarrow \mathcal{C}_k} \lambda_k.$$

Since each class is accessible from itself, $\lambda^* = \max_{\mathcal{C} \in \mathcal{C}} \lambda(\mathcal{C}, \mathbf{1})$. Thus the global long-run pricing rate is the maximum of the state-dependent zero-coupon pricing rates, and the long-run envelope yield is the minimum of their limiting yields.

6. DISASTER RECOVERY AND LONG-RUN PRICING

We use disaster recovery to illustrate how a physically transient class can govern long-horizon prices and how this possibility depends on preferences and cash-flow exposure. Suppose the pricing graph has two communicating classes, a normal class $\mathcal{N}$ and a disaster class $\mathcal{D}$, with the single interclass edge $\mathcal{D} \rightarrow \mathcal{N}$. We call this a *one-way disaster-recovery economy*. The common physical and pricing

accessibility relations imply that $\mathcal{N}$ is physically recurrent and $\mathcal{D}$ is transient: recovery eventually occurs, and the physical probability of disaster vanishes. Nevertheless, the disaster class may have the higher long-run pricing rate.

Section 6.1 develops this structure allowing multiple states within each class, specializing the general results of Sections 4 and 5. Section 6.2 then studies a two-state consumption-based specification in which the pricing-rate gap is determined by risk aversion, consumption dynamics, cash-flow exposure, and recovery intensity.

6.1. One-way recovery and long-run prices. Ordering normal states before disaster states gives the pricing generator

$$L_{\Pi} = \begin{pmatrix} L_{\mathcal{N}} & 0 \\ L_{\mathcal{D}\mathcal{N}} & L_{\mathcal{D}} \end{pmatrix}, \quad (6.1)$$

where $L_{\mathcal{N}}$ and $L_{\mathcal{D}}$ are the irreducible Metzler blocks associated with the normal and disaster classes, respectively, and $L_{\mathcal{D}\mathcal{N}} \geq 0$ is the nonzero recovery block. Partition the payoff accordingly:

$$g = \begin{pmatrix} g_{\mathcal{N}} \\ g_{\mathcal{D}} \end{pmatrix}, \quad (6.2)$$

where $g_{\mathcal{N}}$ and $g_{\mathcal{D}}$ are the restrictions of g to $\mathcal{N}$ and $\mathcal{D}$. Define the class pricing rates

$$\lambda_{\mathcal{N}} := s(L_{\mathcal{N}}), \quad \lambda_{\mathcal{D}} := s(L_{\mathcal{D}}).$$

Then $\lambda^* = s(L_{\Pi}) = \max\{\lambda_{\mathcal{N}}, \lambda_{\mathcal{D}}\}$. The next proposition gives the exact pricing semigroup and the limiting strip yields implied by Theorem 5.2.

Proposition 6.1. *Under Assumptions 2.1–2.3, suppose the pricing generator L_{Π} has the block form (6.1), where $L_{\mathcal{N}}$ and $L_{\mathcal{D}}$ are irreducible Metzler matrices and $L_{\mathcal{D}\mathcal{N}} \geq 0$ is nonzero. Then,*

$$\Pi_t = \begin{pmatrix} e^{tL_{\mathcal{N}}} & 0 \\ H_t & e^{tL_{\mathcal{D}}} \end{pmatrix}, \quad H_t = \int_0^t e^{sL_{\mathcal{D}}} L_{\mathcal{D}\mathcal{N}} e^{(t-s)L_{\mathcal{N}}} ds, \quad t \geq 0. \quad (6.3)$$

Here H_t is the recovery block. For every nonzero $g \in \mathbb{R}_+^X$:

- (i) If $x \in \mathcal{N}$ and $g_{\mathcal{N}} \neq 0$, then $\lim_{t \rightarrow \infty} y_t(x, g) = -\lambda_{\mathcal{N}}$; if $g_{\mathcal{N}} = 0$, then $q_t(x, g) = 0$ for $t \geq 0$.
- (ii) If $x \in \mathcal{D}$, then

$$\lim_{t \rightarrow \infty} y_t(x, g) = \begin{cases} -\lambda_{\mathcal{D}}, & \text{if } g_{\mathcal{N}} = 0, \\ -\max\{\lambda_{\mathcal{N}}, \lambda_{\mathcal{D}}\}, & \text{if } g_{\mathcal{N}} \neq 0. \end{cases}$$

The block representation of Π_t in (6.3), together with (6.2), implies that

$$\Pi_t g = \begin{pmatrix} e^{tL_{\mathcal{N}}} g_{\mathcal{N}} \\ e^{tL_{\mathcal{D}}} g_{\mathcal{D}} + H_t g_{\mathcal{N}} \end{pmatrix}, \quad t \geq 0. \quad (6.4)$$

From a normal initial state, the upper block in (6.4) shows that the price depends only on $L_{\mathcal{N}}$ and $g_{\mathcal{N}}$ at every maturity. When $g_{\mathcal{N}} \neq 0$, the pricing corridor is $\{\mathcal{N}\}$, so the disaster block $L_{\mathcal{D}}$ cannot

affect the price.¹⁵ For a disaster initial state $X_0 = x \in \mathcal{D}$, denote the recovery stopping time by

$$\tau := \inf\{s \geq 0 : X_s \in \mathcal{N}\}. \quad (6.5)$$

Since the economy cannot return to disaster after recovery, the lower block in (6.4) decomposes the price $(\Pi_t g)(x)$ according to whether recovery has occurred by maturity. In particular,

$$(e^{tL_{\mathcal{D}}} g_{\mathcal{D}})(x) = \mathbb{E}_x S_t G_t g(X_t) 1_{\{\tau > t\}}, \quad (H_t g_{\mathcal{N}})(x) = \mathbb{E}_x S_t G_t g(X_t) 1_{\{\tau \leq t\}}.$$

These two terms average the same discounted payoff over sample paths *without* and *with recovery* by maturity, respectively. In the integral representation of H_t in (6.3), s is a possible recovery time, leaving $t - s$ units of time in the normal class before maturity.

A case of particular interest is when the disaster class has the higher pricing rate:

$$\lambda_{\mathcal{D}} > \lambda_{\mathcal{N}}. \quad (6.6)$$

Let $\varphi_{\mathcal{N}}, \psi_{\mathcal{N}} \gg 0$ be right and left eigenvectors of $L_{\mathcal{N}}$ associated with $\lambda_{\mathcal{N}}$, and let $\varphi_{\mathcal{D}}, \psi_{\mathcal{D}} \gg 0$ be the corresponding eigenvectors of $L_{\mathcal{D}}$ associated with $\lambda_{\mathcal{D}}$, normalized by $\mathbf{1}^{\top} \varphi_{\mathcal{N}} = \mathbf{1}^{\top} \varphi_{\mathcal{D}} = 1$ and $\psi_{\mathcal{N}}^{\top} \varphi_{\mathcal{N}} = \psi_{\mathcal{D}}^{\top} \varphi_{\mathcal{D}} = 1$. The following corollary gives the long-run asymptotics of strip prices for both initial classes.

Corollary 6.2. *Under the hypotheses of Proposition 6.1 and (6.6), let $g \in \mathbb{R}_+^{\mathcal{X}}$ be nonzero.*

(i) *For every $x \in \mathcal{N}$,*

$$\lim_{t \rightarrow \infty} e^{-\lambda_{\mathcal{N}} t} q_t(x, g) = \varphi_{\mathcal{N}}(x) \psi_{\mathcal{N}}^{\top} g_{\mathcal{N}},$$

where the limit is positive for $g_{\mathcal{N}} \neq 0$.

(ii) *For every $x \in \mathcal{D}$,*

$$\lim_{t \rightarrow \infty} e^{-\lambda_{\mathcal{D}} t} q_t(x, g) = \varphi_{\mathcal{D}}(x) \psi_{\mathcal{D}}^{\top} \left[g_{\mathcal{D}} + \int_0^{\infty} e^{-\lambda_{\mathcal{D}} v} L_{\mathcal{D}\mathcal{N}} e^{v L_{\mathcal{N}}} g_{\mathcal{N}} dv \right] > 0. \quad (6.7)$$

For a normal initial state and $g_{\mathcal{N}} \neq 0$, part (i) of the corollary gives $q_t(x, g) \sim e^{\lambda_{\mathcal{N}} t} \varphi_{\mathcal{N}}(x) \psi_{\mathcal{N}}^{\top} g_{\mathcal{N}}$, with $\lambda_{\mathcal{N}} < \lambda_{\mathcal{D}}$. For a disaster initial state, the two terms in brackets in (6.7), premultiplied by $e^{\lambda_{\mathcal{D}} t} \varphi_{\mathcal{D}}(x) \psi_{\mathcal{D}}^{\top}$, give the leading contributions to the price from sample paths without and with recovery, respectively. The recovery contribution is positive whenever $g_{\mathcal{N}} \neq 0$, so even a payoff supported entirely on normal states inherits the disaster pricing rate.

When $g_{\mathcal{N}} \neq 0$, the integrand in (6.7) decays exponentially at rate $\lambda_{\mathcal{D}} - \lambda_{\mathcal{N}} > 0$ as $v \rightarrow \infty$, where v represents the time between recovery and maturity. The asymptotics imply that, at sufficiently long maturities, most of the recovery contribution to the price comes from sample paths that remain in disaster for almost the entire horizon and recover near maturity. The pricing importance of these sample paths, however, does not imply that such late recovery is physically likely.

¹⁵Allowing rare entry into disaster removes the block-triangular structure of Π_t , so prices from normal initial states also reflect paths through disaster. Their contribution vanishes as entry becomes rare at each fixed maturity, but Section 7 shows that it can eventually dominate at long maturities when the disaster class pricing rate is higher.

6.2. Consumption risk and pricing dominance. We now specify economic primitives and characterize the conditions under which (6.6) holds. Specialize the one-way recovery economy to two states, $X = \{N, D\}$, with $\mathcal{N} = \{N\}$ and $\mathcal{D} = \{D\}$. Recovery occurs at physical intensity $\kappa > 0$, and the normal state is absorbing, so the physical generator is

$$L_X = \begin{pmatrix} 0 & 0 \\ \kappa & -\kappa \end{pmatrix}.$$

Conditional on $X_0 = D$, the recovery stopping time τ defined in (6.5) is exponentially distributed with mean $1/\kappa$, with density $f_\tau(s) := \kappa e^{-\kappa s}$ for $s > 0$. Therefore, a smaller κ means a longer expected disaster duration and hence higher disaster persistence.

6.2.1. A consumption-based pricing specification. We consider a consumption-based asset-pricing model with time-separable CRRA preferences, as in Hansen and Scheinkman (2009, Example 3.8). Here, the state process X is the two-state continuous-time Markov chain with physical generator L_X specified above.

Let $C_t > 0$ denote aggregate consumption at time t , with $C_0 > 0$. Between state transitions, log consumption has drift $\mu_x \in \mathbb{R}$ and volatility $\sigma_x \geq 0$ in state $x \in \{N, D\}$. At recovery, consumption is multiplied by $\chi_{DN} > 0$, so log consumption jumps by $\log \chi_{DN}$. Specifically, we assume

$$d \log C_t = \mu_{X_{t-}} dt + \sigma_{X_{t-}} dW_t + \log \chi_{DN} dN_t^X, \quad (6.8)$$

where W is a standard Brownian motion independent of X , and N_t^X counts the jumps of X over $(0, t]$. Since recovery is the only possible transition, $N_t^X = \mathbf{1}_{\{\tau \leq t\}}$ when $X_0 = D$.

Let $\alpha > 0$ be the relative risk-aversion parameter of the CRRA utility function u , and let $\beta > 0$ be the subjective discount rate. Since $u'(c) = c^{-\alpha}$, the SDF in this model takes the form

$$S_t = e^{-\beta t} \frac{u'(C_t)}{u'(C_0)} = e^{-\beta t} (C_t/C_0)^{-\alpha}. \quad (6.9)$$

We assume that cumulative cash-flow growth is a power of normalized consumption:

$$G_t = (C_t/C_0)^\vartheta, \quad \vartheta \geq 0, \quad (6.10)$$

where ϑ is the elasticity of cash-flow growth with respect to consumption.¹⁶ For $g = \mathbf{1}$, $\vartheta = 0$ gives a fixed real payment and $\vartheta = 1$ gives a payment proportional to aggregate consumption.

6.2.2. The pricing generator and pricing-rate gap. By (6.8)–(6.10), we have

$$d \log(S_t G_t) = (-\beta - (\alpha - \vartheta)\mu_{X_{t-}}) dt - (\alpha - \vartheta)\sigma_{X_{t-}} dW_t - (\alpha - \vartheta) \log \chi_{DN} dN_t^X.$$

At recovery, $\log S_t G_t$ jump by $(\vartheta - \alpha) \log \chi_{DN}$, so the growth-adjusted SDF is multiplied by $\chi_{DN}^{\vartheta - \alpha}$. Proposition A.1 then gives the pricing generator

$$L_\Pi = \begin{pmatrix} -\varrho_N & 0 \\ \kappa \chi_{DN}^{\vartheta - \alpha} & -\varrho_D - \kappa \end{pmatrix}, \quad (6.11)$$

¹⁶Appendix A.3.2 relaxes this power specification by allowing cash-flow growth to have its own drift, Brownian exposure, and transition multipliers.

where

$$\varrho_x := \beta + (\alpha - \vartheta)\mu_x - \frac{1}{2}(\alpha - \vartheta)^2\sigma_x^2, \quad x \in \{N, D\}.$$

Conditional on remaining in the initial state x , we have $\mathbb{E}_x(S_t G_t \mid N_t^X = 0) = \exp(-\varrho_x t)$.¹⁷ Thus ϱ_x is the within-state growth-adjusted discount rate, excluding transition effects. Its adjustment to subjective discounting β depends on $\alpha - \vartheta$, the excess of marginal-utility sensitivity over cash-flow sensitivity to consumption. When $\alpha > \vartheta$, higher expected log-consumption growth raises ϱ_x , whereas greater volatility lowers it, holding other parameters fixed.

If $\alpha > \vartheta$, $\mu_N > \mu_D$, and $\sigma_D^2 > \sigma_N^2$, then $\varrho_D < \varrho_N$. The difference $\varrho_N - \varrho_D$ measures how much lower the within-state discount rate is in disaster than in the normal state, before accounting for recovery. In particular,

$$\varrho_N - \varrho_D = (\alpha - \vartheta)(\mu_N - \mu_D) + \frac{1}{2}(\alpha - \vartheta)^2(\sigma_D^2 - \sigma_N^2). \quad (6.12)$$

Under the above conditions, greater risk aversion (i.e., a larger α) widens the discount-rate difference $\varrho_N - \varrho_D$, whereas greater cash-flow sensitivity to consumption (i.e., a larger ϑ) narrows it.¹⁸

Since each class contains a single state, its pricing rate is the corresponding diagonal entry of L_Π :

$$\lambda_N = -\varrho_N, \quad \lambda_D = -\varrho_D - \kappa. \quad (6.13)$$

Hence the pricing-rate gap is

$$\Delta := \lambda_D - \lambda_N = \varrho_N - \varrho_D - \kappa. \quad (6.14)$$

Disaster is pricing dominant exactly when its within-state discount rate is lower than the normal-state rate by more than the recovery intensity: $\varrho_N - \varrho_D > \kappa$. Holding other parameters fixed, greater disaster persistence—a smaller κ —raises Δ .

6.2.3. Exact prices and recovery timing. Applying Proposition 6.1 to the pricing generator in (6.11) gives the pricing semigroup

$$\Pi_t = \begin{pmatrix} e^{t\lambda_N} & 0 \\ H_t & e^{t\lambda_D} \end{pmatrix}, \quad t \geq 0,$$

where λ_N and λ_D are given by (6.13), and

$$H_t = \int_0^t h_t(s) ds, \quad h_t(s) := \kappa \chi_{DN}^{\vartheta-\alpha} e^{\lambda_D s} e^{\lambda_N(t-s)} > 0, \quad 0 < s < t. \quad (6.15)$$

As in Section 6.1, here s represents a possible value of the recovery stopping time τ , and $t - s$ the time spent in the normal state before maturity.

For a payoff g , write $g_N = g(N)$ and $g_D = g(D)$. The corresponding maturity- t strip prices are

$$q_t(N, g) = e^{\lambda_N t} g_N, \quad q_t(D, g) = e^{\lambda_D t} g_D + H_t g_N.$$

¹⁷On $\{N_t^X = 0\}$, starting from x , $S_t G_t = \exp\{(-\beta - (\alpha - \vartheta)\mu_x)t - (\alpha - \vartheta)\sigma_x W_t\}$. Independence of W and X yields $W_t \sim N(0, t)$ under this conditioning. The Gaussian exponential-moment formula then gives $\mathbb{E}_x(S_t G_t \mid N_t^X = 0) = \exp\{(-\beta - (\alpha - \vartheta)\mu_x + \frac{1}{2}(\alpha - \vartheta)^2\sigma_x^2)t\} = \exp(-\varrho_x t)$. At fixed log-consumption drift, greater volatility raises this expectation through the convexity of the exponential, lowering the corresponding discount rate when $\alpha \neq \vartheta$.

¹⁸Under these conditions, $\partial(\varrho_N - \varrho_D)/\partial\alpha = -\partial(\varrho_N - \varrho_D)/\partial\vartheta = (\mu_N - \mu_D) + (\alpha - \vartheta)(\sigma_D^2 - \sigma_N^2) > 0$.

To make the late-recovery interpretation at the end of Section 6.1 explicit, assume $g_N > 0$ and consider the *recovery price-contribution* $H_t g_N$ to $q_t(D, g)$. By (6.15), $g_N h_t(s)$ is the *density* of this recovery contribution: integrating it over any subinterval of $(0, t)$ gives the price contribution from paths that recover during that interval.

Using (6.13), we can factor $h_t(s)$ defined in (6.15) as

$$h_t(s) = f_\tau(s) \mathbb{E}_D(S_t G_t \mid \tau = s), \quad 0 < s < t, \quad (6.16)$$

where $f_\tau(s) = \kappa e^{-\kappa s}$ is the physical density of the recovery time τ under $\mathbb{P}_D$, and

$$\mathbb{E}_D(S_t G_t \mid \tau = s) = \chi_{DN}^{\vartheta - \alpha} e^{-\varrho_D s - \varrho_N(t-s)}$$

is the expected growth-adjusted SDF conditional on recovery at time $s \in (0, t)$. Thus the second factor in (6.16) weights the physical recovery-time density to obtain the recovery price-contribution density per unit of normal-state payoff. Since

$$\frac{d}{ds} \log f_\tau(s) = -\kappa, \quad \frac{d}{ds} \log \mathbb{E}_D(S_t G_t \mid \tau = s) = \varrho_N - \varrho_D,$$

we obtain

$$\frac{d}{ds} \log h_t(s) = \varrho_N - \varrho_D - \kappa = \Delta.$$

Therefore, $h_t(s)$ is strictly increasing in s exactly when $\Delta > 0$, or, equivalently, when the conditional expected growth-adjusted SDF increases with recovery time at a faster exponential rate than the physical density decreases.

To describe how the recovery contribution is distributed across recovery times, define the normalized recovery price-contribution density

$$\tilde{h}_t(s) := \frac{h_t(s)}{H_t} = \frac{h_t(s)}{\int_0^t h_t(r) dr}, \quad 0 < s < t. \quad (6.17)$$

For comparison, conditioning the recovery time τ on recovery by maturity t yields the density

$$f_\tau(s \mid \tau \leq t) = \frac{f_\tau(s)}{\mathbb{P}_D\{\tau \leq t\}} = \frac{\kappa e^{-\kappa s}}{1 - e^{-\kappa t}}, \quad 0 < s < t.$$

Both densities integrate to one over $(0, t)$, but they have distinct interpretations.

When $\Delta > 0$, (6.15) and (6.17) give $\tilde{h}_t(s) = \Delta e^{\Delta s} / (e^{\Delta t} - 1)$ for $s \in (0, t)$. For fixed $\Delta > 0$ and $v > 0$, and $t > v$, the share of the recovery contribution attributable to recovery during $[t - v, t]$ is

$$\int_{t-v}^t \tilde{h}_t(s) ds = \frac{1 - e^{-\Delta v}}{1 - e^{-\Delta t}} \rightarrow 1 - e^{-\Delta v} > 0, \quad \text{as } t \rightarrow \infty.$$

Thus, for any $\varepsilon \in (0, 1)$, choosing $v \geq \Delta^{-1} \log(1/\varepsilon)$ ensures that recovery during $[t - v, t]$ accounts for at least $1 - \varepsilon$ of the recovery price-contribution. For any such fixed v , however,

$$\int_{t-v}^t f_\tau(s \mid \tau \leq t) ds = \frac{e^{\kappa v} - 1}{e^{\kappa t} - 1} \rightarrow 0 \quad \text{as } t \rightarrow \infty.$$

This quantifies the distinction between pricing importance and physical likelihood.




FIGURE 3. Pricing-rate gap and physical versus price-contribution densities. Parameters follow (6.18) with $\vartheta = 0$. Left: contours of $\Delta = \lambda_D - \lambda_N$ (% per year), varying $\alpha \in [1, 6]$ and $\kappa \in [0.10, 0.24]$ per year. The bold $\Delta = 0$ boundary separates disaster (blue) from normal-state (gray) pricing dominance; the dot marks $(\alpha, \kappa) = (4, 0.18)$. Right: the conditional physical density $f_\tau(s | \tau \leq t)$ (blue) and normalized recovery price-contribution density $\tilde{h}_t(s) = h_t(s)/H_t$ (orange dashed), for $X_0 = D$ and $t = 50$ years at the baseline.

6.2.4. *A quantitative illustration.* We measure time in years and set

$$\beta = 0.02, \quad \mu_N = 0.025, \quad \mu_D = 0, \quad \sigma_N = 0.02, \quad \sigma_D = \sqrt{0.02^2 + 0.121^2}, \quad \kappa = 0.18. \quad (6.18)$$

These inputs are illustrative, not a calibration. The disaster-growth, excess-volatility, and recovery magnitudes are informed by Nakamura et al. (2013, Table 1).¹⁹

The recovery intensity κ has a direct duration interpretation: the mean disaster spell is $1/\kappa$, and its probability of lasting more than one year is $e^{-\kappa}$. Matching the annual disaster-continuation probability of 0.835 reported by Nakamura et al. (2013, Table 1) gives $\kappa = -\log(0.835) \approx 0.1803$. Our baseline therefore implies a mean disaster spell of 5.56 years, close in magnitude to the roughly six-year mean in their annual model. The range $\kappa \in [0.10, 0.24]$ in Figure 3 spans mean durations of approximately 4.2–10 years. Recovery restores normal consumption-growth dynamics but need not return consumption to its pre-disaster level.

Under (6.18), (6.12) and (6.14) imply that $\Delta = 0.025(\alpha - \vartheta) + 0.00732(\alpha - \vartheta)^2 - 0.18$. For $\alpha \geq \vartheta$, disaster dominance requires $\alpha - \vartheta > 3.5369$. At $\alpha = 4$, a fixed payoff ($\vartheta = 0$) gives $\lambda_N = -0.1168$, $\lambda_D = -0.0797$, and $\Delta = 0.0371$.²⁰ By contrast, consumption-proportional cash flows ($\vartheta = 1$) give $\Delta = -0.0391$; so the same physical economy and preferences do not imply disaster dominance for every cash-flow family.

¹⁹The growth and excess-volatility inputs draw on their permanent disaster-shock estimates. These are not simply contemporaneous consumption-growth shocks, and their benchmark preferences are Epstein–Zin rather than CRRA. Thus (6.18) is not a replication of their consumption process or equilibrium prices.

²⁰The implied normal-state real yield, $-\lambda_N = 11.68\%$, belongs to the counterfactual economy with disaster entry ruled out. It is not a calibrated prediction for yields observed during normal times, when investors still price the possibility of future disasters.

Figure 3 fixes $\vartheta = 0$. The left panel varies risk aversion α and recovery intensity κ , holding the remaining parameters at the values in (6.18). The pricing-rate gap Δ increases with α and decreases with κ , so greater risk aversion and more persistent disasters favor disaster dominance. The blue and gray regions correspond to disaster and normal-state dominance, respectively. The upward-sloping zero-gap boundary shows that faster recovery requires greater risk aversion for disaster dominance.

The right panel of Figure 3 considers the baseline setting $(\alpha, \kappa) = (4, 0.18)$ and maturity $t = 50$ years, starting from disaster. The conditional physical density places most probability on recovery within the first ten years, whereas the normalized recovery price-contribution density increases toward maturity. For example, recovery during the final ten years accounts for approximately 36.8% of the recovery price-contribution, but its physical probability conditional on recovery by maturity is only 0.062%. Thus physically unlikely late recovery can carry substantial pricing importance.

7. RARE TRANSITIONS AND CROSSOVER MATURITIES

The discussion following (6.4) in Section 6.1 shows that the disaster class can have the highest pricing rate yet remain absent from prices starting in the normal class. We now study how rare transitions bring an otherwise excluded pricing-dominant class into the pricing corridor. We characterize how the rarity of the connecting paths determines its vanishing price loading and how this loading, together with the pricing-rate gap, determines the crossover to long-maturity dominance.

7.1. Path rarity and dominant-component loadings. Consider Metzler pricing generators $L_\Pi(\varepsilon)$, $\varepsilon \geq 0$, such that $L_\Pi(\varepsilon)$ is irreducible for every sufficiently small $\varepsilon > 0$ and converges to a reducible generator $L_\Pi(0)$ as $\varepsilon \downarrow 0$. For $x \in \mathbf{X}$ and $g \in \mathbb{R}^{\mathbf{X}}$, write

$$\Pi_t(\varepsilon) := e^{tL_\Pi(\varepsilon)}, \quad q_t(x, g; \varepsilon) := (\Pi_t(\varepsilon)g)(x), \quad t \geq 0.$$

Throughout Section 7, we write $L_\Pi := L_\Pi(0)$ for the limiting pricing generator, and maintain Assumptions 2.1–2.3 for each sufficiently small $\varepsilon \geq 0$, on the same finite state space $\mathbf{X}$.

Let $\mathcal{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_m\}$, $m \geq 2$, be the class partition of the *limiting* graph $\mathcal{G}(L_\Pi(0))$. We retain these *limiting classes* as a fixed partition of $\mathbf{X}$ for $\varepsilon > 0$, even though irreducibility makes $\mathcal{G}(L_\Pi(\varepsilon))$ a single class. Partition

$$L_\Pi(\varepsilon) = [L_{k\ell}(\varepsilon)], \quad k, \ell = 1, \dots, m,$$

conformably with (4.1). Thus $L_{k\ell}(\varepsilon)$ has rows in $\mathcal{C}_k$ and columns in $\mathcal{C}_\ell$; when $k \neq \ell$, a nonzero block represents a direct edge $\mathcal{C}_k \rightarrow \mathcal{C}_\ell$ in $\mathcal{G}(L_\Pi(\varepsilon))$.

Set $L_k := L_{kk}(0)$ and $\lambda_k := s(L_k)$. Assume that the limiting network has a unique globally pricing-dominant class $\mathcal{C}^*$:

$$\lambda^* := s(L_\Pi(0)) = s(L_{k^*}) > \max_{\mathcal{C}_k \neq \mathcal{C}^*} \lambda_k, \quad (7.1)$$

where L_{k^*} is the block indexed by $\mathcal{C}^*$.

The next assumption is standard in singular perturbations of nearly decomposable systems; see, for example, Schweitzer (1968); Meyer (1989); Yin and Zhang (2012) and Avrachenkov et al. (2013).

Assumption 7.1 (Rare-edge perturbation). *For every sufficiently small $\varepsilon > 0$, $L_\Pi(\varepsilon)$ is Metzler and irreducible. There is a positive function $\eta(\varepsilon) \rightarrow 0$ such that $L_{kk}(\varepsilon) = L_k + O(\eta(\varepsilon))$, for every k . For each $k \neq \ell$, there is an edge-rarity exponent $w_{k\ell} \in [0, +\infty]$ such that, as $\varepsilon \downarrow 0$,*

$$L_{k\ell}(\varepsilon) = \begin{cases} \varepsilon^{w_{k\ell}} (K_{k\ell} + O(\eta(\varepsilon))), & w_{k\ell} < +\infty, \\ 0, & w_{k\ell} = +\infty, \end{cases}$$

where, in the first case, $K_{k\ell} \geq 0$ and $K_{k\ell} \neq 0$; in the second, the equality holds for every $\varepsilon \geq 0$.

Assumption 7.1 is inherited from a regular perturbation of the physical Markov dynamics. For example, if $r(\varepsilon) \rightarrow 0$, $\bar{Q}_{k\ell} \geq 0$ is nonzero, $\bar{\Gamma}_{k\ell} \gg 0$, and, for $k \neq \ell$ with $w_{k\ell} < +\infty$,

$$[L_X(\varepsilon)]_{k\ell} = \varepsilon^{w_{k\ell}} (\bar{Q}_{k\ell} + O(r(\varepsilon))), \quad \Gamma_{k\ell}(\varepsilon) = \bar{\Gamma}_{k\ell} + O(r(\varepsilon)),$$

then Proposition A.1 gives

$$L_{k\ell}(\varepsilon) = [L_X(\varepsilon)]_{k\ell} \odot \Gamma_{k\ell}(\varepsilon) = \varepsilon^{w_{k\ell}} (\bar{Q}_{k\ell} \odot \bar{\Gamma}_{k\ell} + O(r(\varepsilon))).$$

Thus regular risk adjustment changes the leading coefficient but preserves the edge exponent and zero pattern. If within-class errors are uniformly $O(r(\varepsilon))$, they and the diagonal effects of new exit intensities can be absorbed by taking, for example, $\eta(\varepsilon) = r(\varepsilon) + \varepsilon^{w_{\min}}$, where $w_{\min}$ is the smallest positive edge exponent. Positive limiting jump multipliers imply that physical irreducibility yields pricing irreducibility; singular or vanishing risk adjustments may instead change the effective exponents.

Define the weighted directed graph $\mathcal{G}_w$ on $\mathcal{C}$ by including the edge $\mathcal{C}_k \rightarrow \mathcal{C}_\ell$ exactly when $w_{k\ell} < +\infty$ and assigning it weight $w_{k\ell}$. It records which interclass edges can open and their asymptotic orders. An edge of weight zero is *non-rare* and is already present at the limit; one of positive finite weight is *rare* and opens at order $\varepsilon^{w_{k\ell}}$. An infinite weight means that the direct edge remains absent, although an indirect path may exist. The zero-weight subgraph of $\mathcal{G}_w$ is precisely the reduced pricing graph $\mathcal{G}_r(L_\Pi(0))$, while irreducibility for $\varepsilon > 0$ makes $\mathcal{G}_w$ strongly connected.

For $\mathcal{C}, \mathcal{C}' \in \mathcal{C}$, let $\mathcal{D}_w(\mathcal{C}, \mathcal{C}')$ be the set of directed paths from $\mathcal{C}$ to $\mathcal{C}'$ in $\mathcal{G}_w$. Here a path $\gamma = (\mathcal{C}_{i_0}, \dots, \mathcal{C}_{i_p})$ is a sequence of distinct vertices with $\mathcal{C}_{i_0} = \mathcal{C}$, $\mathcal{C}_{i_p} = \mathcal{C}'$, and a directed edge from each vertex to the next. Its weight is $W(\gamma) := \sum_{r=1}^p w_{i_{r-1}i_r}$. When $\mathcal{C} = \mathcal{C}'$, include the length-zero path with weight zero. Define

$$d_{\mathcal{C}\mathcal{C}^*} := \min_{\gamma \in \mathcal{D}_w(\mathcal{C}, \mathcal{C}^*)} W(\gamma), \quad d_{\mathcal{C}^*\mathcal{C}} := \min_{\gamma \in \mathcal{D}_w(\mathcal{C}^*, \mathcal{C})} W(\gamma). \quad (7.2)$$

The subscripts record direction, and $d_{\mathcal{C}^*\mathcal{C}^*} = 0$.

For a nonzero $g \geq 0$ with class-level support $\mathcal{S}(g)$, define

$$d(\mathcal{C}, g) := d_{\mathcal{C}\mathcal{C}^*} + \min_{\mathcal{C}' \in \mathcal{S}(g)} d_{\mathcal{C}^*\mathcal{C}'}. \quad (7.3)$$

It is the minimum combined rarity exponent of two directed paths: one from the initial class $\mathcal{C}$ to $\mathcal{C}^*$ and the other from $\mathcal{C}^*$ to the payoff support. Since zero-weight paths are exactly the paths in




FIGURE 4. A stylized distress–crisis–recovery weighted class graph. The payoff is supported on both the globally dominant crisis class $\mathcal{C}^*$ and the recovery class $\mathcal{C}'$. Solid and dashed arrows denote zero- and positive-weight edges, respectively. The rare relapse edge makes $\mathcal{G}_w$ strongly connected.

the reduced graph of $L_\Pi(0)$,

$$d(\mathcal{C}, g) = 0 \iff \mathcal{C}^* \in \mathcal{K}(\mathcal{C}, g), \quad (7.4)$$

where $\mathcal{K}(\mathcal{C}, g)$ is the limiting pricing corridor. Thus $d(\mathcal{C}, g) > 0$ exactly when at least one path in every such pair contains a rare edge.

Figure 4 may be read as a stylized distress–crisis–recovery network. Let $\mathcal{S}(g) = \{\mathcal{C}^*, \mathcal{C}'\}$, so the payoff is supported in both the globally dominant crisis class and the recovery class. Starting from the distress class $\mathcal{C}$, direct recovery is order one, whereas escalation to $\mathcal{C}^*$ occurs at order ε^2 . Recovery from crisis is again order one, while relapse from $\mathcal{C}'$ to $\mathcal{C}$ occurs at order ε . This economic reading is illustrative; the results below require only a rarely accessed globally dominant class, not a crisis interpretation.

The least-weight initial-state leg has weight 2. Because $\mathcal{C}^* \in \mathcal{S}(g)$, the length-zero path at $\mathcal{C}^*$ makes the least-weight payoff leg equal to zero; the solid recovery edge also gives $d_{\mathcal{C}^* \mathcal{C}'} = 0$. Hence

$$d_{\mathcal{C}\mathcal{C}^*} = 2, \quad \min_{\mathcal{C}'' \in \mathcal{S}(g)} d_{\mathcal{C}^* \mathcal{C}''} = 0, \quad d(\mathcal{C}, g) = 2.$$

The zero-weight graph gives the limiting corridor $\mathcal{K}(\mathcal{C}, g) = \{\mathcal{C}, \mathcal{C}'\}$, which excludes $\mathcal{C}^*$. Thus ε^2 is the order at which the dominant component enters the price, even though direct recovery to a payoff class remains order one. The positive rarity exponent comes entirely from the initial-state leg; neither singleton payoff support nor a rare payoff-side edge is required.

For sufficiently small $\varepsilon > 0$, let $\lambda^*(\varepsilon) = s(L_\Pi(\varepsilon))$ and let $P_*(\varepsilon)$ be its rank-one Riesz spectral projection. Write $P_{*, \mathcal{C}\mathcal{C}'}(\varepsilon)$ for the block with rows in $\mathcal{C}$ and columns in $\mathcal{C}'$, and define the long-run pricing law by $\pi_\varepsilon^{\text{LR}}(x) := P_*(\varepsilon)(x, x)$.

Theorem 7.2 (Perron projection and pricing-mass asymptotics). *Under Assumption 7.1 and the unique-dominant-class condition (7.1), as $\varepsilon \downarrow 0$, $\lambda^*(\varepsilon) \rightarrow \lambda^*$ and, for every $\mathcal{C}, \mathcal{C}' \in \mathcal{C}$, there is a matrix $\bar{P}_{\mathcal{C}\mathcal{C}'} \gg 0$ such that*

$$\varepsilon^{-(d_{\mathcal{C}\mathcal{C}^*} + d_{\mathcal{C}^* \mathcal{C}'})} P_{*, \mathcal{C}\mathcal{C}'}(\varepsilon) \rightarrow \bar{P}_{\mathcal{C}\mathcal{C}'}. \quad (7.5)$$

Moreover,

$$\varepsilon^{-(d_{CC^*} + d_{C^*C})} \sum_{x \in \mathcal{C}} \pi_\varepsilon^{\text{LR}}(x) \rightarrow \text{tr } \bar{P}_{CC} > 0. \quad (7.6)$$

In particular, $\sum_{x \in \mathcal{C}^*} \pi_\varepsilon^{\text{LR}}(x) \rightarrow 1$ and $\sum_{x \in \mathcal{C}} \pi_\varepsilon^{\text{LR}}(x) \rightarrow 0$ for every $\mathcal{C} \neq \mathcal{C}^*$.

For fixed sufficiently small $\varepsilon > 0$, Perron–Frobenius theory gives $e^{-\lambda^*(\varepsilon)t} \Pi_t(\varepsilon) \rightarrow P_*(\varepsilon)$ and hence, for $x \in \mathcal{C}$ and nonzero $g \geq 0$,

$$q_t(x, g; \varepsilon) \sim e^{\lambda^*(\varepsilon)t} (P_*(\varepsilon)g)(x), \quad t \rightarrow \infty.$$

For such a payoff, let $\mathcal{S}_{\min}(g) := \arg \min_{\mathcal{C} \in \mathcal{S}(g)} d_{C^*C}$. Combining this representation with (7.5) yields, for every $\mathcal{C} \in \mathcal{C}$,

$$\varepsilon^{-d(\mathcal{C}, g)} (P_*(\varepsilon)g)_c \rightarrow \sum_{\mathcal{C}' \in \mathcal{S}_{\min}(g)} \bar{P}_{c\mathcal{C}'} g_{\mathcal{C}'} \gg 0. \quad (7.7)$$

Only the minimizing payoff classes survive in this limit. Hence $d(\mathcal{C}, g)$ is the loading exponent of the globally dominant component, while the two directed exponents in (7.5) govern initial-state and terminal-payoff exposure. Equations (7.4) and (7.7) also separate the order-one and vanishing-loading cases. If $d(\mathcal{C}, g) = 0$, then $\mathcal{C}^* \in \mathcal{K}(\mathcal{C}, g)$, $\lambda(\mathcal{C}, g) = \lambda^*$, and the global loading has a strictly positive order-one limit; both iterated long-run yield limits equal $-\lambda^*$. If $d(\mathcal{C}, g) > 0$, then $\mathcal{C}^* \notin \mathcal{K}(\mathcal{C}, g)$ and the global loading vanishes at order $\varepsilon^{d(\mathcal{C}, g)}$. Only the latter case can generate a diverging local-to-global crossover driven by a vanishing loading on $\mathcal{C}^*$.

7.2. Crossover to global pricing dominance. Fix $x \in \mathcal{C}$ and nonzero $g \geq 0$. Throughout this subsection, suppose the limiting pricing corridor $\mathcal{K} = \mathcal{K}(\mathcal{C}, g)$ is nonempty and $d(\mathcal{C}, g) > 0$. Let $L_{\mathcal{K}}$ be its restricted pricing generator, as defined in Section 5, evaluated at $L_{\Pi} = L_{\Pi}(0)$, and abbreviate $\lambda_{\mathcal{K}} := \lambda(\mathcal{C}, g) = s(L_{\mathcal{K}})$.

By Proposition 5.1, $L_{\mathcal{K}}$ gives an exact representation of the limiting price. Because $d(\mathcal{C}, g) > 0$ is equivalent to $\mathcal{C}^* \notin \mathcal{K}$ by (7.4), we have $\lambda_{\mathcal{K}} < \lambda^*$. Theorem 7.2 nevertheless gives the global mode a loading of order ε^d . The key question is when its higher exponential rate overcomes this vanishing loading, which requires a spectral expansion uniform over maturities that diverge as $\varepsilon \downarrow 0$.

For every sufficiently small $\varepsilon > 0$, let $P_*(\varepsilon)$ be the rank-one Riesz projection of $L_{\Pi}(\varepsilon)$ associated with $\lambda^*(\varepsilon)$, and set $P_*(0) := P_*$, the corresponding projection of L_{Π} at λ^* . Unique global dominance makes λ^* algebraically simple, while $\ker P_*$ is invariant under L_{Π} and is the direct sum of the generalized eigenspaces associated with all eigenvalues other than λ^* . Define

$$s_{\circ} := s(L_{\Pi}|_{\ker P_*}), \quad \Sigma_{\circ} := \{z \in \sigma(L_{\Pi}) : \Re z = s_{\circ}\} \subset \sigma(L_{\Pi}).$$

Thus $\Sigma_{\circ}$ consists of the eigenvalues other than λ^* with the largest real part. Although $s_{\circ}$ is real by definition, it need not itself be an eigenvalue. Since $\sigma(L_{\Pi}|_{\ker P_*}) = \sigma(L_{\Pi}) \setminus \{\lambda^*\}$ and $\lambda_{\mathcal{K}}$ is a distinct eigenvalue of L_{Π} , we have $\lambda_{\mathcal{K}} \leq s_{\circ} < \lambda^*$.

For each $\lambda \in \sigma(L_{\Pi})$, let P_{λ} be its Riesz projection and define

$$P_{\circ} := \sum_{\lambda \in \Sigma_{\circ}} P_{\lambda}.$$

Choose a bounded open neighborhood $U_\circ$ of $\Sigma_\circ$ such that $\overline{U_\circ} \cap \sigma(L_\Pi) = \Sigma_\circ$. For all sufficiently small $\varepsilon \geq 0$, define

$$\Sigma_\circ(\varepsilon) := \sigma(L_\Pi(\varepsilon)) \cap U_\circ, \quad (7.8)$$

and let $P_\circ(\varepsilon)$ be its Riesz projection. Thus $\Sigma_\circ(0) = \Sigma_\circ$ and $P_\circ(0) = P_\circ$.

Also let $w_{\min} := \min\{w_{k\ell} : 0 < w_{k\ell} < +\infty\} > 0$. Such a positive-weight edge exists because the zero-weight subgraph is reducible whereas $\mathcal{G}_w$ is strongly connected. Fix a submultiplicative matrix norm. Assumption 7.1 gives

$$r_L(\varepsilon) := \|L_\Pi(\varepsilon) - L_\Pi\| = O(\eta(\varepsilon) + \varepsilon^{w_{\min}}) = o(1). \quad (7.9)$$

This is the generator-level perturbation rate used below. In the physical construction following Assumption 7.1, $r_L(\varepsilon) = O(r(\varepsilon) + \varepsilon^{w_{\min}})$, although the two error rates need not coincide.

For sufficiently small $\varepsilon \geq 0$, define the global Perron loading and the contribution associated with $\Sigma_\circ(\varepsilon)$ by

$$a_*(x, g; \varepsilon) := (P_*(\varepsilon)g)(x), \quad b_t(x, g; \varepsilon) := (e^{tL_\Pi(\varepsilon)}P_\circ(\varepsilon)g)(x), \quad t \geq 0.$$

The following result separates the vanishing global loading from the contribution associated with $\Sigma_\circ(\varepsilon)$ and controls all remaining spectral contributions uniformly over maturity.

Theorem 7.3. *Under Assumption 7.1 and the unique-dominant-class condition (7.1), fix $\mathcal{C} \in \mathcal{C}$, $x \in \mathcal{C}$, and a nonzero payoff $g \geq 0$. Suppose that $\mathcal{K} = \mathcal{K}(\mathcal{C}, g) \neq \emptyset$ and $d = d(\mathcal{C}, g) > 0$. There exist constants $\varepsilon_0, C, \delta > 0$ such that, for every $\varepsilon \in (0, \varepsilon_0]$ and $t \geq 0$,*

$$q_t(x, g; \varepsilon) = e^{\lambda^*(\varepsilon)t} a_*(x, g; \varepsilon) + b_t(x, g; \varepsilon) + R_{\varepsilon, t}, \quad |R_{\varepsilon, t}| \leq C e^{(s_\circ - \delta)t}. \quad (7.10)$$

Moreover, as $\varepsilon \downarrow 0$,

$$\lambda^*(\varepsilon) = \lambda^* + O(r_L(\varepsilon)), \quad a_*(x, g; \varepsilon) = a \varepsilon^d (1 + o(1)), \quad \text{for some } a > 0. \quad (7.11)$$

For every fixed $T < \infty$, $b_t(x, g; \varepsilon) = b_t(x, g; 0) + O(r_L(\varepsilon))$ uniformly over $0 \leq t \leq T$. At $\varepsilon = 0$, this contribution is

$$b_t(x, g; 0) = \sum_{\lambda \in \Sigma_\circ} e^{\lambda t} \sum_{j=0}^{\nu_\lambda - 1} \frac{t^j}{j!} ((L_\Pi - \lambda I)^j P_\lambda g)(x), \quad (7.12)$$

where ν_λ is the size of the largest Jordan block at λ . If $\lambda_{\mathcal{K}} < s_\circ$, then $b_t(x, g; 0) \equiv 0$: the spectral contribution associated with $\Sigma_\circ$ is absent from this limiting corridor price.

For bounded maturities, Theorem 7.3 makes the global contribution vanish at order ε^d , while $b_t(x, g; \varepsilon)$ converges to its reducible-limit counterpart. For every fixed positive perturbation, however, the higher Perron rate eventually governs the long run. The resulting order-of-limits discontinuity is immediate.

Corollary 7.4 (Noncommuting long-run limits). *Under the conditions of Theorem 7.3,*

$$\lim_{\varepsilon \downarrow 0} \lim_{t \rightarrow \infty} \left[-\frac{1}{t} \log q_t(x, g; \varepsilon) \right] = -\lambda^*, \quad \lim_{t \rightarrow \infty} \lim_{\varepsilon \downarrow 0} \left[-\frac{1}{t} \log q_t(x, g; \varepsilon) \right] = -\lambda_{\mathcal{K}} = -\lambda(\mathcal{C}, g). \quad (7.13)$$

Thus taking maturity first selects the global rate, whereas taking the reducible limit first restores the corridor rate. This conclusion does not require the local contribution to be a single exponential

term: by (7.12), it may combine several eigenvalues and Jordan terms. Moreover, eigenvalues other than λ^* with real parts above $\lambda_{\mathcal{K}}$ may have zero coefficients in the limiting corridor price but acquire nonzero loadings when rare edges open, producing intermediate crossovers. A two-term exponential formula therefore requires the following spectral condition.

Assumption 7.5 (Single simple eigenvalue at the corridor rate). *The corridor rate is the unique eigenvalue of L_{Π} , other than λ^* , with maximal real part:*

$$\lambda_{\mathcal{K}} = s_{\circ}, \quad \Sigma_{\circ} = \{\lambda_{\mathcal{K}}\}. \quad (7.14)$$

Moreover, $\lambda_{\mathcal{K}}$ is algebraically simple as an eigenvalue of L_{Π} .

Under Assumption 7.5, set $\lambda_{\circ} := s_{\circ} = \lambda_{\mathcal{K}} = \lambda(\mathcal{C}, g)$. The assumption implies that $P_{\circ} = P_{\lambda_{\circ}}$ has rank one, that the Jordan block at $\lambda_{\circ}$ has size one, and that every eigenvalue other than λ^* and $\lambda_{\circ}$ has real part strictly smaller than $\lambda_{\circ}$. By the projection estimate (D.31) in Lemma D.6 and rank stability, for every sufficiently small $\varepsilon \geq 0$ the set $\Sigma_{\circ}(\varepsilon)$ consists of a single real algebraically simple eigenvalue $\lambda_{\circ}(\varepsilon)$, with $\lambda_{\circ}(0) = \lambda_{\circ}$. Define $b_{\circ}(x, g; \varepsilon) := (P_{\circ}(\varepsilon)g)(x)$. Then

$$b_t(x, g; \varepsilon) = e^{\lambda_{\circ}(\varepsilon)t} b_{\circ}(x, g; \varepsilon), \quad t \geq 0. \quad (7.15)$$

The same projection estimate and the rank-one trace identity give, as $\varepsilon \downarrow 0$,

$$\lambda_{\circ}(\varepsilon) = \lambda_{\circ} + O(r_L(\varepsilon)), \quad b_{\circ}(x, g; \varepsilon) = b_{\circ}(x, g; 0) + O(r_L(\varepsilon)). \quad (7.16)$$

Write $\Delta(\varepsilon) := \lambda^*(\varepsilon) - \lambda_{\circ}(\varepsilon)$ and $\Delta := \Delta(0) = \lambda^* - \lambda_{\mathcal{K}} > 0$.

Theorem 7.6. *Under the conditions of Theorem 7.3 and Assumption 7.5, $b_{\circ}(x, g; 0) > 0$, and the following hold.*

(i) *There exist $\varepsilon_0, C, \delta > 0$ such that, for every $\varepsilon \in (0, \varepsilon_0]$ and $t \geq 0$,*

$$q_t(x, g; \varepsilon) = a_*(x, g; \varepsilon)e^{\lambda^*(\varepsilon)t} + b_{\circ}(x, g; \varepsilon)e^{\lambda_{\circ}(\varepsilon)t} + R_{\varepsilon, t}, \quad |R_{\varepsilon, t}| \leq Ce^{(\lambda_{\circ} - \delta)t}. \quad (7.17)$$

(ii) *For all sufficiently small $\varepsilon > 0$, the two exponential terms in (7.17) are equal at the unique positive maturity*

$$t_{\varepsilon}^c = \frac{\log[b_{\circ}(x, g; \varepsilon)/a_*(x, g; \varepsilon)]}{\Delta(\varepsilon)}. \quad (7.18)$$

With $a > 0$ as in (7.11), as $\varepsilon \downarrow 0$,

$$t_{\varepsilon}^c = \frac{d \log(1/\varepsilon) + \log[b_{\circ}(x, g; 0)/a]}{\Delta(\varepsilon)} + o(1). \quad (7.19)$$

(iii) *As $\varepsilon \downarrow 0$,*

$$\frac{a_*(x, g; \varepsilon)e^{\lambda^*(\varepsilon)(t_{\varepsilon}^c + s)}}{q_{t_{\varepsilon}^c + s}(x, g; \varepsilon)} \rightarrow \frac{e^{\Delta s}}{1 + e^{\Delta s}} \quad \text{uniformly for } s \text{ in compact subsets of } \mathbb{R}. \quad (7.20)$$

If the generator additionally satisfies

$$r_L(\varepsilon) \log(1/\varepsilon) \rightarrow 0, \quad \varepsilon \downarrow 0, \quad (7.21)$$

then (7.19) simplifies to

$$t_\varepsilon^c = \frac{d}{\Delta} \log(1/\varepsilon) + \frac{1}{\Delta} \log \frac{b_o(x, g; 0)}{a} + o(1). \quad (7.22)$$

The condition holds whenever $r_L(\varepsilon) = O(\varepsilon^p)$ for some $p > 0$ and ensures that replacing $\Delta(\varepsilon)$ by Δ changes the approximation by only $o(1)$. Holding d and Δ fixed, a larger global loading coefficient a relative to $b_o(x, g; 0)$ advances the crossover through the constant correction.

Under the theorem's assumptions, (7.19) implies the leading scale $t_\varepsilon^c \sim (d/\Delta) \log(1/\varepsilon)$. Holding the other quantity fixed, a larger rarity exponent d delays the crossover on this scale, whereas a larger pricing-rate gap Δ brings it forward. This is a spectral peso problem: a rarely accessed class can be nearly invisible at ordinary maturities yet govern the long end because its initially small price contribution has a higher exponential rate. In (7.20), the global term's limiting share of the full price passes through one half at the crossover. A larger Δ makes this transition sharper as maturity moves through t_ε^c .

7.3. Rare disaster entry with fixed recovery. In Section 6, we set normal-to-disaster entry equal to zero. We now restore entry at order ε , holding within-class dynamics and recovery fixed, and apply Sections 7.1–7.2 to determine its price loading and crossover horizon.

7.3.1. Pricing concentration and crossover in the two-class model. Let F_ε collect the physical normal-to-disaster transition intensities, set $D_\varepsilon := \text{diag}(F_\varepsilon \mathbf{1}_\mathcal{D})$, and let $\Gamma \gg 0$ collect the corresponding growth-adjusted SDF jump multipliers. With normal states ordered first, the physical and pricing generators are

$$L_X(\varepsilon) = \begin{pmatrix} Q - D_\varepsilon & F_\varepsilon \\ H & U \end{pmatrix}, \quad L_\Pi(\varepsilon) = \begin{pmatrix} L_\mathcal{N} - D_\varepsilon & E_\varepsilon \\ L_{\mathcal{D}\mathcal{N}} & L_\mathcal{D} \end{pmatrix}, \quad E_\varepsilon := F_\varepsilon \odot \Gamma. \quad (7.23)$$

Here Q and U govern the within-normal and within-disaster physical dynamics, respectively, H contains the recovery rates, and the pricing blocks are those introduced in Section 6. The diagonal adjustment $-D_\varepsilon$ records the additional exit intensity from the normal class. Suppose

$$F_\varepsilon = \varepsilon F + O(\varepsilon^{1+\eta_F}), \quad E_\varepsilon = \varepsilon J + O(\varepsilon^{1+\eta_F}), \quad J := F \odot \Gamma, \quad (7.24)$$

so regular risk adjustment changes the leading entry coefficient, through Γ , but not its rarity exponent. The following assumption collects the primitive regularity conditions; Appendix A.5 derives the pricing generator from those primitives.

Assumption 7.7 (Rare entry with fixed recovery). *The matrices Q and U are irreducible, Q is a Markov generator, U is a transient subgenerator, $H \geq 0$, and $H\mathbf{1}_\mathcal{N} + U\mathbf{1}_\mathcal{D} = 0$. Equation (7.24) holds for some $\eta_F > 0$ and some nonzero $F \geq 0$, and $F_\varepsilon \geq 0$ for every sufficiently small $\varepsilon > 0$. The matrices $L_\mathcal{N}$ and $L_\mathcal{D}$ are irreducible Metzler matrices, $L_{\mathcal{D}\mathcal{N}} \geq 0$ is nonzero, and Γ is fixed and strictly positive.*

These conditions imply that the physical and pricing generators in (7.23) are irreducible for every sufficiently small $\varepsilon > 0$: the entry and recovery blocks connect the two irreducible within-class graphs in opposite directions.

Maintain disaster pricing dominance, $\lambda_D > \lambda_N$. Let π_ε^P denote the stationary distribution of the physical generator $L_X(\varepsilon)$, and let $\pi_\varepsilon^{\text{LR}}$ be the long-run pricing law defined in Section 7.1. The next result makes their contrasting limits transparent.

Corollary 7.8 (Physical rarity and pricing concentration). *Under Assumption 7.7 and $\lambda_D > \lambda_N$, there exist constants $c_P, c_{\text{LR}} > 0$ such that, as $\varepsilon \downarrow 0$,*

$$\pi_\varepsilon^P(\mathcal{D}) = c_P \varepsilon + o(\varepsilon), \quad \pi_\varepsilon^{\text{LR}}(\mathcal{D}) = 1 - c_{\text{LR}} \varepsilon + o(\varepsilon).$$

Physical frequency and long-run pricing importance thus have opposite limits: the disaster class becomes physically negligible while the long-run pricing law concentrates on it. This does not determine when its contribution becomes important in a price starting from the normal class.

For the crossover, fix $x \in \mathcal{N}$ and the growth-adjusted zero-coupon payoff $g = \mathbf{1}$. A single rare entry edge reaches disaster, where g is positive, so $d(\mathcal{N}, \mathbf{1}) = 1$; the limiting corridor is $\{\mathcal{N}\}$ with rate λ_N . Theorem 7.6 requires λ_N to exceed the real part of every other eigenvalue except λ_D . Irreducibility of L_N therefore leaves only the condition

$$\sup_{z \in \sigma(L_D) \setminus \{\lambda_D\}} \Re z < \lambda_N, \quad (7.25)$$

with the convention $\sup \emptyset = -\infty$. Under this condition, Assumption 7.5 holds with $\lambda_{\mathcal{N}} = \lambda_o = \lambda_N$.

Corollary 7.9 (Rare-entry crossover from the normal class). *Under Assumption 7.7, $\lambda_D > \lambda_N$, and (7.25), fix $x \in \mathcal{N}$. Theorem 7.6 applies to $g = \mathbf{1}$ with $d = 1$ and $\Delta = \lambda_D - \lambda_N$. The coefficients*

$$a_x := \lim_{\varepsilon \downarrow 0} \varepsilon^{-1} a_*(x, \mathbf{1}; \varepsilon), \quad b_x := b_o(x, \mathbf{1}; 0)$$

exist and are strictly positive. As $\varepsilon \downarrow 0$,

$$t_\varepsilon^c(x) = \frac{\log(1/\varepsilon) + \log(b_x/a_x)}{\lambda_D - \lambda_N} + o(1). \quad (7.26)$$

With recovery held fixed, $r_L(\varepsilon) = O(\varepsilon)$, so (7.21) holds and yields the expansion (7.26). Corollary 7.4 also shows that the order of limits matters: letting maturity tend to infinity before sending entry intensity to zero gives the limiting yield $-\lambda_D$, whereas reversing these limits gives $-\lambda_N$.

7.3.2. A two-state consumption-based application. We retain the within-state consumption dynamics, preferences, and cash-flow specification of Section 6.2, together with the recovery intensity κ and consumption multiplier χ_{DN} , and vary only the physical disaster-entry intensity $\varepsilon\ell$. Here $\ell > 0$ is a fixed reference intensity, and $\varepsilon > 0$ controls the rarity of disaster entry. At entry to disaster, consumption is multiplied by $\chi_{ND} > 0$, so the growth-adjusted SDF is multiplied by $\chi_{ND}^{\vartheta-\alpha}$.

With time measured in years, an economy starting in the normal state N enters disaster within one year with probability $1 - e^{-\varepsilon\ell}$, and its mean waiting time to first entry is $1/(\varepsilon\ell)$. The recovery intensity κ , by contrast, determines the mean disaster duration $1/\kappa$.

The physical and pricing generators are, respectively,

$$L_X(\varepsilon) = \begin{pmatrix} -\varepsilon\ell & \varepsilon\ell \\ \kappa & -\kappa \end{pmatrix}, \quad L_\Pi(\varepsilon) = \begin{pmatrix} -\varrho_N - \varepsilon\ell & \varepsilon\ell\chi_{ND}^{\vartheta-\alpha} \\ \kappa\chi_{DN}^{\vartheta-\alpha} & -\varrho_D - \kappa \end{pmatrix}. \quad (7.27)$$

Here ϱ_N and ϱ_D are the within-state growth-adjusted discount rates from Section 6.2. Both generators are irreducible for every $\varepsilon > 0$. We retain $\lambda_N = -\varrho_N$ and $\lambda_D = -\varrho_D - \kappa$ for the zero-entry class pricing rates and assume

$$\Delta = \lambda_D - \lambda_N > 0.$$

The spectral separation condition (7.25) holds automatically as both limiting class blocks are scalar.

The two eigenvalues of the pricing generator are given by

$$\lambda_{\pm}(\varepsilon) = \frac{1}{2}[\lambda_N - \varepsilon\ell + \lambda_D \pm \Delta(\varepsilon)], \quad \Delta(\varepsilon) := \sqrt{(\Delta + \varepsilon\ell)^2 + 4\kappa\varepsilon\ell(\chi_{ND}\chi_{DN})^{\vartheta-\alpha}} > 0. \quad (7.28)$$

Here $\lambda_+(\varepsilon)$ is the global long-run pricing rate and $\Delta(\varepsilon)$ is the eigenvalue gap; as $\varepsilon \downarrow 0$, $\lambda_+(\varepsilon) \rightarrow \lambda_D$, $\lambda_-(\varepsilon) \rightarrow \lambda_N$, and $\Delta(\varepsilon) \rightarrow \Delta$. Direct calculation gives the physical stationary mass and long-run pricing mass of the disaster state as

$$\pi_{\varepsilon}^P(D) = \frac{\varepsilon\ell}{\kappa + \varepsilon\ell}, \quad \pi_{\varepsilon}^{\text{LR}}(D) = \frac{1}{2} \left(1 + \frac{\Delta + \varepsilon\ell}{\Delta(\varepsilon)} \right). \quad (7.29)$$

Their first-order expansions recover Corollary 7.8; Appendix A.5 derives these identities, their expansion constants, and the price formulas below.

As in Corollary 7.9, take $g = \mathbf{1}$ and start in the normal state N . The maturity- t payment is G_t ; it is a fixed real payment only when $\vartheta = 0$. Write $z_t(N; \varepsilon) = q_t(N, \mathbf{1}; \varepsilon)$. We have

$$z_t(N; \varepsilon) = \omega_{\varepsilon} e^{\lambda_+(\varepsilon)t} + (1 - \omega_{\varepsilon}) e^{\lambda_-(\varepsilon)t}, \quad \omega_{\varepsilon} := \frac{\lambda_+(\varepsilon) - \lambda_D + \varepsilon\ell\chi_{ND}^{\vartheta-\alpha}}{\Delta(\varepsilon)}. \quad (7.30)$$

The loading ω_{ε} on the dominant exponential is strictly positive, with²¹

$$\omega_{\varepsilon} = \varepsilon\omega + O(\varepsilon^2), \quad \omega := \frac{\ell\chi_{ND}^{\vartheta-\alpha}}{\Delta} \left(1 + \frac{\kappa\chi_{DN}^{\vartheta-\alpha}}{\Delta} \right) > 0. \quad (7.31)$$

For all sufficiently small ε , $0 < \omega_{\varepsilon} < 1/2$. The two positive price contributions in (7.30) then become equal at the strictly positive maturity

$$t_{\varepsilon}^c(N) = \frac{\log[(1 - \omega_{\varepsilon})/\omega_{\varepsilon}]}{\Delta(\varepsilon)} = \frac{\log(1/\varepsilon)}{\Delta} - \frac{\log \omega}{\Delta} + o(1), \quad \varepsilon \downarrow 0. \quad (7.32)$$

This specializes Corollary 7.9 with $a_N = \omega$ and $b_N = 1$.

At any ε for which $0 < \omega_{\varepsilon} < 1$, we have

$$\frac{\omega_{\varepsilon} e^{\lambda_+(\varepsilon)t}}{z_t(N; \varepsilon)} = \left[1 + \omega_{\varepsilon}^{-1}(1 - \omega_{\varepsilon})e^{-\Delta(\varepsilon)t} \right]^{-1}.$$

This share reaches $1/2$ at $t_{\varepsilon}^c(N)$ when the crossover is positive. If $1/2 \leq \omega_{\varepsilon} < 1$, the dominant contribution already accounts for at least half of the price at maturity zero, so there is no strictly positive crossover.²²

²¹The term in the brackets is the coefficient of $e^{\lambda_D t}$ in the long-run price $q_t(D, \mathbf{1}; 0)$ derived in Section 6.2.3. Its two terms arise from sample paths without and with recovery by maturity, respectively.

²²Outside the small-entry region, $1 - \omega_{\varepsilon}$ can be nonpositive. The two-positive-contribution interpretation then fails, although (7.30) remains an exact identity.




FIGURE 5. Exact crossover maturities and dominant-component price shares from the normal state. Both panels use $\alpha = 4$ and $\vartheta = 0$; remaining fixed parameters are given in the text. Left: crossover contours (years) where $\Delta > 0$, $0 < \omega_\varepsilon < 1/2$, and $\lambda_+(\varepsilon) < 0$; the blue band marks annual entry probabilities of 1.5–2%. Light gray denotes $\omega_\varepsilon \geq 1/2$, dark gray $\Delta \leq 0$, and tan $\lambda_+(\varepsilon) \geq 0$; the dashed line marks $\lambda_+(\varepsilon) = 0$. Right: $\omega_\varepsilon e^{\lambda_+(\varepsilon)t} / z_t(N; \varepsilon)$ at annual entry probability 1.75% and $\kappa = 0.14, 0.16, 0.18$ per year (all with $\Delta > 0$). Matching filled points identify these cases on the left; open circles mark positive crossovers at share 1/2 on the right.

Figure 5 uses the within-state consumption dynamics and subjective discounting specified in (6.18). Both panels set $\alpha = 4$, $\vartheta = 0$, $\ell = 1$ per year, $\chi_{ND} = e^{-0.111}$, $\chi_{DN} = e^{0.111}$, and $g = 1$. With time measured in years, the zero-entry pricing-rate gap in (6.14) becomes

$$\Delta = \lambda_D - \lambda_N = 4(0.025) + \frac{1}{2} 4^2 (0.121^2) - \kappa = 0.217128 - \kappa.$$

The contour plot on the left panel of Figure 5 shows how entry and recovery intensities affect crossover maturity from the normal state, holding the remaining parameters fixed. The blue band marks annual entry probabilities of 1.5–2%, using the disaster probabilities reported by Barro (2006) as a benchmark. With $\ell = 1$ per year, these probabilities correspond to $\varepsilon \approx 0.01511$ – 0.02020 and mean waiting times to entry of approximately 50–66 years from the normal state.²³ The recovery-intensity range matches that in Figure 3, corresponding to mean disaster durations of approximately 4.2–10 years. Within the region displaying crossover contours, higher entry intensity brings crossover forward. Higher recovery intensity also brings crossover forward in the current setting. In particular, more persistent disasters do not necessarily produce earlier crossover, because recovery intensity affects both the pricing-rate gap and the dominant-component loading.

The right panel of Figure 5 plots the disaster-associated dominant component's share of the price from the normal state. The annual disaster entry probability is fixed at 1.75%, implying $\varepsilon =$

²³These long waiting times reflect the rarity of severe disasters, rather than ordinary business-cycle downturns. These estimates provide an entry-frequency benchmark, not a full calibration of the consumption-based model.

$-\log(1 - 0.0175) \approx 0.01765$ and a mean waiting time to entry of 56.6 years. Recovery intensities $\kappa = 0.14, 0.16, 0.18$ correspond to mean disaster durations of 7.14, 6.25, 5.56 years, respectively. For the first two cases, the share reaches one half at maturities of 4.34 and 2.36 years. In the third case, $\omega_\varepsilon \approx 0.501$, so the dominant component already accounts for slightly more than half the price at maturity zero. All three shares increase with maturity, reaching approximately 94.5%, 95.0%, and 95.8% at 25 years.

The three recovery intensities give $\Delta \approx 0.0771, 0.0571, 0.0371$ per year, respectively, so all three cases satisfy $\Delta > 0$. We call the dominant component “disaster-associated” because $\lambda_+(\varepsilon) \rightarrow \lambda_{\mathcal{D}}$ as $\varepsilon \downarrow 0$; for every $\varepsilon > 0$, however, it is a global spectral component involving both states. Its price share is neither a disaster probability nor the price share of paths entering disaster. At the baseline $\kappa = 0.18$, for example, the physical stationary disaster probability is 8.9%, whereas the disaster mass under the long-run pricing law is 71.9%, by (7.29). These stationary state masses measure different objects from the maturity-dependent price share plotted in the right panel.

Both panels of Figure 5 use exact finite-entry formulas. The left-panel contours are computed from $\log[(1 - \omega_\varepsilon)/\omega_\varepsilon]/\Delta(\varepsilon)$ in (7.32), while the right-panel price shares are computed from the decomposition in (7.30). At the benchmark entry probabilities and parameter values considered here, the first-order loading approximation in (7.31) and the asymptotic crossover expansion in (7.32) need not provide close approximations to the exact quantities. Appendix A.5 illustrates this approximation error.

8. CONCLUSION

We show that physical recurrence and long-run pricing dominance are distinct. Although the physical and pricing generators share the same communicating classes and accessibility relations, stochastic discounting and cash-flow growth determine the ordering of class pricing rates. A physically transient class can therefore govern long-maturity prices.

The pricing corridor identifies when such dominance matters for a particular initial state and nonnegative payoff. Restricting the pricing generator to this corridor preserves the price at every maturity. For a nonempty corridor, its highest class rate determines the long-run yield, while chains of classes tied at that rate determine the polynomial order of the price. The corresponding perpetual cash-flow stream has finite value exactly when the corridor rate is negative. Global valuation instability can thus coexist with finite values for particular states and claims.

We characterize how rare transitions change these conclusions when the limiting corridor excludes the uniquely pricing-dominant class. Minimum-rarity paths from the initial class to the dominant class and onward to the payoff support determine the asymptotic order of its price loading. Under two-term spectral separation, this loading and the perturbed pricing-rate gap determine the crossover maturity, which diverges logarithmically as the connecting transitions vanish. The reducible and long-maturity limits need not commute: neglecting a rare transition can give an accurate fixed-maturity price but a different long-run yield.

The consumption-based disaster-recovery application illustrates the economic significance of this mechanism. With empirically motivated entry frequencies and disaster durations, the disaster-associated dominant component can account for most of a claim's price from a normal initial state at economically relevant maturities. Quantitative assessment nevertheless requires jointly disciplining physical transitions, cash-flow exposures, and stochastic discounting. Rarity alone does not justify neglecting a transition: its pricing importance depends on the claim and maturity.

APPENDIX

APPENDIX A. MULTIPLICATIVE FUNCTIONALS AND THE PRIMITIVE PRICING GENERATOR

Section 2 constructs the pricing semigroup abstractly from the SDF and cash-flow growth processes. This appendix provides its semimartingale foundation: it verifies the semigroup and continuity properties, parameterizes the two processes as exponential additive functionals, and derives the pricing generator from their primitive shock exposures. The construction builds on and extends the specification in [Hansen and Scheinkman \(2009\)](#) by explicitly allowing the information flow to contain continuous and jump shocks that affect stochastic discounting and cash-flow growth without changing the finite-state economic state. Related continuous-time formulations include [Hansen and Scheinkman \(2012\)](#), [Borovička et al. \(2011\)](#), and [Duffie et al. \(2000\)](#).

A.1. From multiplicative functionals to the pricing semigroup. Let $g \in \mathbb{R}_+^X$. Multiplicativity of S^G , iterated expectations, and the Markov property (2.1)²⁴ give

$$\begin{aligned} (\Pi_{t+u}g)(x) &= \mathbb{E}_x[S_t^G(S_u^G g(X_u)) \circ \theta_t] = \mathbb{E}_x\{S_t^G \mathbb{E}[(S_u^G g(X_u)) \circ \theta_t \mid \mathcal{F}_t]\} \\ &= \mathbb{E}_x[S_t^G \mathbb{E}_{X_t}(S_u^G g(X_u))] = (\Pi_t \Pi_u g)(x). \end{aligned}$$

Linearity extends this identity to every $g \in \mathbb{R}^X$, proving the semigroup property (2.5).

It remains to justify continuity at zero. The càdlàg paths and the normalization $S_0^G = 1$ imply $S_t^G \rightarrow 1$, $\mathbb{P}_x$ -almost surely, as $t \rightarrow 0$. Assumption 2.3 and positivity then imply $\mathbb{E}_x|S_t^G - 1| \rightarrow 0$ by Scheffé's lemma. For $\|g\|_\infty \leq 1$,

$$|(\Pi_t g)(x) - g(x)| \leq \mathbb{E}_x|S_t^G - 1| + 2\mathbb{P}_x(X_t \neq x).$$

Right continuity of X makes the second term vanish. Because X is finite, both convergences are uniform in x , so $\|\Pi_t - I\| \rightarrow 0$. This proves the continuity assertion used to define L_Π in (2.6).

A.2. Additive-functional semimartingale specification. Let N^X be the jump measure of the state process,

$$N^X((0, t] \times B) := \sum_{0 < u \leq t} \mathbf{1}_{\{X_{u-} \neq X_u, X_u \in B\}}, \quad t \geq 0, \quad B \subseteq X.$$

²⁴By truncation and monotone convergence, (2.1) extends to every nonnegative $\mathcal{F}_\infty^\circ$ -measurable Y , with both sides interpreted as $[0, +\infty]$ -valued random variables. We apply this extension to $Y = S_u^G g(X_u)$, which is nonnegative and $\mathcal{F}_\infty^\circ$ -measurable.

Thus $N^X((0, t] \times B)$ counts transitions up to time t whose destination lies in B . If $\ell_{ij} = (L_X)_{ij}$, its predictable compensator has, when $X_{t-} = x_i$, intensity ℓ_{ij} at $x_j \neq x_i$.

In addition to these state-transition shocks, allow shocks that affect discounting or growth without changing X at the same instant. Let W be an $(\mathcal{F}_t^\circ)$ -adapted d -dimensional process that is an $(\mathcal{F}_t)$ -Brownian motion under every $\mathbb{P}_x$, and let N be an integer-valued random measure on $\mathbb{R}_+ \times \mathbf{Z}$, where $(\mathbf{Z}, \mathcal{Z})$ is a measurable mark space. Under each $\mathbb{P}_x$, suppose that $N((0, t] \times B)$ is $\mathcal{F}_t^\circ$ -measurable and that N has predictable compensator $dt \nu(X_{t-}, dz)$ for a finite state-dependent kernel ν , and that N and N^X have no common jumps. Assume also that the driving shocks are compatible with the shifts: for $t, u \geq 0$ and $B \in \mathcal{Z}$,

$$W_u \circ \theta_t = W_{t+u} - W_t, \quad N((0, u] \times B) \circ \theta_t = N((t, t+u] \times B),$$

$\mathbb{P}_x$ -almost surely for every x .

An $(\mathcal{F}_t^\circ)$ -adapted real-valued process A is an *additive functional* relative to the shifts if $A_0 = 0$ and $A_{t+u} = A_t + A_u \circ \theta_t$ for $t, u \geq 0$, $\mathbb{P}$ -a.s. Then $\exp(A)$ is a strictly positive multiplicative functional. The following specification realizes the abstract processes in Assumption 2.2 in this way.

For $K \in \{S, G\}$, assume

$$\begin{aligned} \log K_t = & \int_0^t \mu_K(X_{u-}) du + \int_0^t \varsigma_K(X_{u-})^\top dW_u \\ & + \int_0^t \int_{\mathbf{X}} \xi_K(X_{u-}, y) N^X(du, dy) + \int_0^t \int_{\mathbf{Z}} \zeta_K(X_{u-}, z) N(du, dz). \end{aligned} \quad (\text{A.1})$$

Here $\mu_K : \mathbf{X} \rightarrow \mathbb{R}$ is the local drift of $\log K$, $\varsigma_K : \mathbf{X} \rightarrow \mathbb{R}^d$ is its continuous-shock exposure, $\xi_K : \mathbf{X} \times \mathbf{X} \rightarrow \mathbb{R}$ is the jump at a state transition, and $\zeta_K : \mathbf{X} \times \mathbf{Z} \rightarrow \mathbb{R}$ is the response to a non-transition jump; set $\xi_K(x, x) = 0$. Assume these functions are measurable and finite and that

$$\max_{x \in \mathbf{X}} \int_{\mathbf{Z}} \exp(\zeta_S(x, z) + \zeta_G(x, z)) \nu(x, dz) < \infty.$$

Since $\mathbf{X}$ and the jump intensities are finite, this exponential-moment condition ensures the finite-horizon moment required in Assumption 2.2. The shift identities above make the right-hand side of (A.1) additive, so its exponential is a strictly positive càdlàg multiplicative functional; the same conditions also give the continuity at zero in Assumption 2.3.

The four coefficient families isolate three sources of uncertainty. Brownian innovations generate continuous shocks without changing the state. A state transition may simultaneously change S and G through ξ_S and ξ_G . The random measure N generates discrete shocks that leave the current state unchanged through ζ_S and ζ_G . Because S and G share these driving processes, their continuous and jump exposures may be correlated and state dependent.

This construction builds on the multiplicative-functional framework of Hansen and Scheinkman (2009, Sections 3.1–3.5), while explicitly allowing discounting and cash-flow growth to respond to

shocks that do not change the finite-state economic state.²⁵ The distinction between regime transitions and shocks conditional on the regime is familiar from finite-state asset-pricing models (Bonomo and Garcia, 1996). Related continuous-time formulations with state-dependent diffusion and jump exposures include Duffie et al. (2000), Borovička et al. (2011), and Hansen and Scheinkman (2012); structural applications featuring rare jump risk include Gabaix (2012) and Wachter (2013).

For the growth-adjusted SDF $S^G = SG$, define the aggregate coefficients

$$\mu = \mu_S + \mu_G, \quad \varsigma = \varsigma_S + \varsigma_G, \quad \xi = \xi_S + \xi_G, \quad \zeta = \zeta_S + \zeta_G.$$

Then $\log S^G$ has the representation (A.1) with $(\mu_K, \varsigma_K, \xi_K, \zeta_K)$ replaced by $(\mu, \varsigma, \xi, \zeta)$.

A.3. Pricing generator and effective discounting. The next proposition maps the primitive shock exposures into the matrix generator defined abstractly in (2.6).

Proposition A.1 (Primitive pricing generator). *Under Assumptions 2.1–2.3 and the semimartingale specification above, let $\ell_{ij} = (L_X)_{ij}$. The pricing generator has entries*

$$(L_\Pi)_{ij} = \begin{cases} \ell_{ii} - \varrho(x_i), & i = j, \\ \ell_{ij} e^{\xi(x_i, x_j)}, & i \neq j, \end{cases} \quad (\text{A.2})$$

where $\|\cdot\|$ denotes the Euclidean norm on $\mathbb{R}^d$ and

$$\varrho(x) = -\mu(x) - \frac{1}{2} \|\varsigma(x)\|^2 - \int_{\mathbf{Z}} \left(e^{\zeta(x, z)} - 1 \right) \nu(x, dz).$$

In particular, L_Π and L_X have the same off-diagonal zero pattern.

Proof of Proposition A.1. Applying Itô's formula to $S_t^G = \exp(\log S_t^G)$ gives

$$\begin{aligned} \frac{dS_t^G}{S_{t-}^G} &= \left[\mu(X_{t-}) + \frac{1}{2} \|\varsigma(X_{t-})\|^2 \right] dt + \varsigma(X_{t-})^\top dW_t \\ &\quad + \int_{\mathbf{X}} \left[e^{\xi(X_{t-}, y)} - 1 \right] N^X(dt, dy) + \int_{\mathbf{Z}} \left[e^{\zeta(X_{t-}, z)} - 1 \right] N(dt, dz). \end{aligned}$$

For $g \in \mathbb{R}^X$, apply the product rule to $S_t^G g(X_t)$. At a transition from x_i to x_j , the product changes from $S_{t-}^G g(x_i)$ to $S_{t-}^G e^{\xi(x_i, x_j)} g(x_j)$. Compensation of N^X and N therefore gives predictable drift

$$S_{t-}^G \left\{ [\ell_{ii} - \varrho(x_i)] g(x_i) + \sum_{j \neq i} \ell_{ij} e^{\xi(x_i, x_j)} g(x_j) \right\}$$

when $X_{t-} = x_i$. Let $\mathcal{A}$ denote the matrix in (A.2). Under the stated integrability conditions, the compensated stochastic-integral terms are martingales, so the product formula gives

$$S_t^G g(X_t) = g(X_0) + \int_0^t S_{u-}^G (\mathcal{A}g)(X_{u-}) du + \mathcal{M}_t,$$

²⁵Hansen and Scheinkman (2009, Footnote 6) note that a larger filtration could be accommodated but assume, for simplicity, that the filtration is generated by the Markov state.

where $\mathcal{M}$ has zero expectation. Consequently, $\Pi_t g - g = \int_0^t \Pi_u(\mathcal{A}g) \, du$. Dividing by t and letting $t \downarrow 0$ yields $L_\Pi g = \mathcal{A}g$ for every g , which proves the claim. $\square$

Formula (A.2) separates transition and local pricing effects. For $i \neq j$, the physical transition rate ℓ_{ij} is multiplied by $e^{\xi(x_i, x_j)}$, the proportional response of the growth-adjusted SDF to that transition. Continuous shocks and non-transition jumps leave X unchanged and therefore enter only through the diagonal term $-\varrho(x_i)$.

To separate transition reweighting from local discounting, define $\ell_{ij}^\mathbb{Q} := \ell_{ij} e^{\xi(x_i, x_j)}$ for $i \neq j$, and let $L_X^\mathbb{Q}$ be the Markov generator with these off-diagonal entries. Define the state-specific *growth-adjusted effective discount rate*

$$\rho(x_i) := -(L_\Pi \mathbf{1})(x_i) = \varrho(x_i) - \sum_{j \neq i} \ell_{ij} \left[e^{\xi(x_i, x_j)} - 1 \right]. \quad (\text{A.3})$$

The pricing generator then has the Feynman–Kac decomposition

$$L_\Pi = L_X^\mathbb{Q} - \text{diag}(\rho), \quad L_\Pi \mathbf{1} = -\rho. \quad (\text{A.4})$$

Let $(\mathbb{Q}_x)_{x \in X}$ be the Markov family generated by $L_X^\mathbb{Q}$. The finite-state Feynman–Kac formula yields

$$(\Pi_t g)(x) = \mathbb{E}_x^\mathbb{Q} \left[\exp \left(- \int_0^t \rho(X_u) \, du \right) g(X_t) \right].$$

This representation is a matrix identity and does not require an equivalent change of measure on the original filtered space. Such a change of measure is available if $M_t^\mathbb{Q} := \exp(\int_0^t \rho(X_{u-}) \, du) S_t^G$ is a true martingale under every $\mathbb{P}_x$: then $\frac{d\mathbb{Q}_x}{d\mathbb{P}_x} \Big|_{\mathcal{F}_t} = M_t^\mathbb{Q}$ and X has generator $L_X^\mathbb{Q}$ under $\mathbb{Q}_x$. When G is nontrivial, $\mathbb{Q}$ is a growth-adjusted pricing law rather than, in general, the usual risk-neutral law, and ρ need not be the money-market short rate.

The row-sum identity in (A.4) gives ρ a direct price interpretation. From $z_t = e^{tL_\Pi} \mathbf{1}$,

$$z_t(x) = 1 - \rho(x)t + o(t), \quad \lim_{t \downarrow 0} y_t(x) = \rho(x).$$

Thus $\rho(x)$ is the short-horizon yield on the growth-adjusted unit payoff. By contrast, the long-run envelope yield is $-\lambda^* = -s(L_\Pi)$. The row-sum bounds for Metzler matrices imply

$$\min_{x \in X} \rho(x) \leq -\lambda^* \leq \max_{x \in X} \rho(x).$$

Combining (A.4) with Theorem 3.2, global valuation stability is equivalent to $s(L_X^\mathbb{Q} - \text{diag}(\rho)) < 0$. Hence positive effective discounting in every state is sufficient, but not necessary: a negative local rate can be offset by sufficiently rapid departure under the pricing-adjusted dynamics.

A.3.1. Recovery jumps and recovery intensity. In the one-way CRRA model of Section 6.2, the effective-discounting representation distinguishes the recovery jump multiplier from the physical recovery intensity. With the primitives of that model, (A.3) gives

$$\kappa_D^\Pi = \kappa \chi_{DN}^{\vartheta - \alpha}, \quad \rho_D = \varrho_D - \kappa(\chi_{DN}^{\vartheta - \alpha} - 1), \quad \lambda_D = -(\rho_D + \kappa_D^\Pi) = -\varrho_D - \kappa,$$

where κ_D^Π is the recovery intensity under the growth-adjusted pricing law. Changing the recovery multiplier $\chi_{DN}^{\vartheta-\alpha}$ alone changes this intensity and the effective discount rate by offsetting amounts. It therefore changes the nonzero off-diagonal entry of the pricing generator but not the disaster-class pricing rate. By contrast, increasing the physical intensity κ lowers λ_D one-for-one and also changes that off-diagonal entry, holding the other primitives fixed. Increasing pricing-adjusted recovery while holding ρ_D fixed is thus a different comparative static from changing the recovery jump in this primitive model. Appendix A.4 gives the local sensitivity to effective discounting for general multi-state blocks.

A.3.2. Separate consumption and cash-flow exposures. The power specification in Section 6.2 is convenient but not required. Retain the CRRA SDF in (6.9) and let consumption and cash-flow growth have separate log drifts μ_i^C, μ_i^G and Brownian exposure vectors $\sigma_i^C, \sigma_i^G \in \mathbb{R}^d$:

$$d \log C_t = \mu_i^C dt + (\sigma_i^C)^\top dW_t, \quad d \log G_t = \mu_i^G dt + (\sigma_i^G)^\top dW_t$$

between transitions in state i . Normalize $G_0 = 1$ and let $\chi_{ij} > 0$ and $\chi_{ij}^G > 0$ be the respective consumption and cash-flow multipliers at $i \rightarrow j$. With no non-transition jumps, Proposition A.1 gives

$$\varrho_i = \beta + \alpha \mu_i^C - \mu_i^G - \frac{1}{2} \|\sigma_i^G - \alpha \sigma_i^C\|^2, \quad m_{ij} = \chi_{ij}^{-\alpha} \chi_{ij}^G. \quad (\text{A.5})$$

In the one-way two-state economy, the class rates remain $\lambda_N = -\varrho_N$ and $\lambda_D = -\varrho_D - \kappa$. Their difference is

$$\begin{aligned} \lambda_D - \lambda_N &= (\mu_D^G - \mu_N^G) + \alpha(\mu_N^C - \mu_D^C) - \kappa \\ &\quad + \frac{1}{2}(\|\sigma_D^G\|^2 - \|\sigma_N^G\|^2) + \frac{1}{2}\alpha^2(\|\sigma_D^C\|^2 - \|\sigma_N^C\|^2) - \alpha[(\sigma_D^C)^\top \sigma_D^G - (\sigma_N^C)^\top \sigma_N^G]. \end{aligned}$$

The last term records the change in the instantaneous covariance of consumption and cash-flow growth. This extension separates cash-flow deterioration from consumption risk compensation. The main-text model sets $\mu_i^G = \vartheta \mu_i^C$, $\sigma_i^G = \vartheta \sigma_i^C$, and $\chi_{ij}^G = \chi_{ij}^\vartheta$.

A.4. Long-run sensitivity to effective discounting. The decomposition (A.4) also links the long-run pricing law to changes in effective discounting. Increasing ρ by $\delta \in \mathbb{R}^n$ while holding $L_X^\mathbb{Q}$ fixed replaces L_Π by $L_\Pi - \text{diag}(\delta)$.

Proposition A.2 (Effective-discount-rate sensitivity). *Suppose that $\mathcal{C}^* = \{\mathcal{C}^*\}$, and let φ and ψ be the right and left eigenvectors of L_Π in Theorem 4.1. For $\delta \in \mathbb{R}^n$ near zero, define*

$$L_\Pi(\delta) := L_\Pi - \text{diag}(\delta), \quad y^*(\delta) := -s(L_\Pi(\delta)).$$

Then y^ is differentiable at zero and, for every $h \in \mathbb{R}^n$,*

$$\left. \frac{d}{d\theta} y^*(\theta h) \right|_{\theta=0} = \sum_{i=1}^n \pi_i^{\text{LR}} h_i, \quad \pi^{\text{LR}} := \varphi \odot \psi. \quad (\text{A.6})$$

In particular, $\partial y^(0)/\partial \delta_i = \pi_i^{\text{LR}}$ and $\sum_{i=1}^n \pi_i^{\text{LR}} = 1$. More generally, if a real matrix-valued perturbation satisfies $L_\Pi(\theta) = L_\Pi + \theta H + o(\theta)$ in matrix norm, then*

$$\left. \frac{d}{d\theta} s(L_\Pi(\theta)) \right|_{\theta=0} = \psi^\top H \varphi. \quad (\text{A.7})$$

Proof. By Theorem 4.1, $\lambda^* = s(L_\Pi)$ is algebraically simple and strictly separated in real part from the remaining spectrum. Its continuation under a sufficiently small perturbation is therefore differentiable and remains the spectral bound. Differentiating the eigenvalue equation for $L_\Pi(\theta) = L_\Pi + \theta H + o(\theta)$ at zero and premultiplying by $\psi^\top$ gives (A.7), because $\psi^\top L_\Pi = \lambda^* \psi^\top$ and $\psi^\top \varphi = 1$. Taking $H = -\text{diag}(h)$ and changing sign yields (A.6); moreover, $\sum_i \pi_i^{\text{LR}} = \psi^\top \varphi = 1$. $\square$

Thus π_i^{LR} is the marginal pass-through from the effective discount rate in state x_i to the long-run envelope yield, as well as the stationary occupation weight of x_i under either Doob chain. Proposition D.2 implies that $\text{supp}(\pi^{\text{LR}}) = \mathcal{C}^*$, so sufficiently small effective-rate changes confined to classes outside $\mathcal{C}^*$ leave the global rate unchanged, although they may alter finite-maturity prices and asymptotic loadings. By contrast, (A.7) shows that changes in transition or risk-adjustment terms generally involve the two-sided weights $\psi_i \varphi_j$, rather than an expectation under π^{LR} .

The same calculation applies claim by claim. Fix a current class $\mathcal{C}$ and a nonnegative payoff g with nonempty corridor $\mathcal{K}(\mathcal{C}, g)$, and suppose that this corridor has a unique critical class $\mathcal{C}_c$. Let L_c be its diagonal block and let $\varphi_c, \psi_c \gg 0$ be corresponding right and left Perron vectors, normalized by $\psi_c^\top \varphi_c = 1$. If L_c is replaced by $L_c - \theta \text{diag}(h)$ for $h \in \mathbb{R}^{\mathcal{C}_c}$, while all other corridor blocks are held fixed, then $\mathcal{C}_c$ remains critical for all sufficiently small θ and

$$\left. \frac{d}{d\theta} [-\lambda(\mathcal{C}, g; \theta)] \right|_{\theta=0} = \sum_{x \in \mathcal{C}_c} \psi_c(x) \varphi_c(x) h(x),$$

where $\lambda(\mathcal{C}, g; \theta)$ is the resulting corridor rate. Sufficiently small effective-rate changes confined to classes outside $\mathcal{C}_c$ leave the corridor rate unchanged. If several corridor-critical classes are tied, the corridor rate is the maximum of their perturbed block rates; the corresponding long-run yield need not be differentiable when those rates have different first-order responses.

A.5. Rare-entry specialization. Because $e^{\xi(x_i, x_j)} > 0$, L_Π and L_X have the same off-diagonal zero pattern. Thus the primitive specification preserves every physical edge and, in particular, the class accessibility relation used in Section 4.

The fixed-recovery model in Section 7.3 follows directly from Proposition A.1. Partition the states into $\mathcal{N}$ and $\mathcal{D}$, hold the other coefficients in (A.1) fixed, and vary only the physical normal-to-disaster transition block F_ε . Set $D_\varepsilon = \text{diag}(F_\varepsilon \mathbf{1}_\mathcal{D})$ and $\Gamma_{ij} = e^{\xi(x_i, x_j)}$ for $x_i \in \mathcal{N}$ and $x_j \in \mathcal{D}$. If $L_\mathcal{N}$, $L_{\mathcal{D}\mathcal{N}}$, and $L_\mathcal{D}$ denote the pricing blocks at $\varepsilon = 0$, then

$$L_X(\varepsilon) = \begin{pmatrix} Q - D_\varepsilon & F_\varepsilon \\ H & U \end{pmatrix}, \quad L_\Pi(\varepsilon) = \begin{pmatrix} L_\mathcal{N} - D_\varepsilon & F_\varepsilon \odot \Gamma \\ L_{\mathcal{D}\mathcal{N}} & L_\mathcal{D} \end{pmatrix}.$$

Thus rare physical entry changes only the normal-state exit diagonals and the normal-to-disaster pricing block; disaster dynamics and recovery remain fixed. If $F_\varepsilon = \varepsilon F + O(\varepsilon^{1+\eta_F})$, the pricing entry block is $\varepsilon(F \odot \Gamma) + O(\varepsilon^{1+\eta_F})$ and has the same rarity exponent.

For the consumption-based application in Section 7.3.2, the physical entry intensity is $\varepsilon \ell$ and $\Delta = \lambda_\mathcal{D} - \lambda_\mathcal{N} > 0$. Physical flow balance gives $\pi_\varepsilon^P(\mathcal{N}) \varepsilon \ell = \pi_\varepsilon^P(\mathcal{D}) \kappa$, yielding the first identity in

(7.29). The pricing generator has the two distinct eigenvalues in (7.28), and its spectral projection at $\lambda_+(\varepsilon)$ is

$$\frac{L_\Pi(\varepsilon) - \lambda_-(\varepsilon)I}{\Delta(\varepsilon)}.$$

The disaster-state diagonal of this rank-one projection is the product of the corresponding normalized left and right Perron eigenvector entries, hence equals $\pi_\varepsilon^{\text{LR}}(D)$. Substituting the eigenvalues gives the second identity in (7.29). Expanding at $\varepsilon = 0$ gives

$$\pi_\varepsilon^P(D) = \frac{\ell}{\kappa}\varepsilon + O(\varepsilon^2), \quad \pi_\varepsilon^{\text{LR}}(D) = 1 - \frac{\ell\kappa(\chi_{ND}\chi_{DN})^{\vartheta-\alpha}}{\Delta^2}\varepsilon + O(\varepsilon^2).$$

Thus the constants in Corollary 7.8 are $c_P = \ell/\kappa$ and $c_{\text{LR}} = \ell\kappa(\chi_{ND}\chi_{DN})^{\vartheta-\alpha}/\Delta^2$; in particular, $\pi_\varepsilon^{\text{LR}}(N) = c_{\text{LR}}\varepsilon + O(\varepsilon^2)$.

The spectral decomposition expresses $z_t(N; \varepsilon)$ as a linear combination of $e^{\lambda_+(\varepsilon)t}$ and $e^{\lambda_-(\varepsilon)t}$. Its coefficients are fixed by

$$z_0(N; \varepsilon) = 1, \quad \partial_t z_t(N; \varepsilon)|_{t=0} = \lambda_N + \varepsilon\ell(\chi_{ND}^{\vartheta-\alpha} - 1),$$

which give (7.30). In particular,

$$\begin{aligned} \omega_\varepsilon &= \frac{\lambda_N + \varepsilon\ell(\chi_{ND}^{\vartheta-\alpha} - 1) - \lambda_-(\varepsilon)}{\Delta(\varepsilon)} \\ &= \frac{\lambda_+(\varepsilon) - \lambda_D + \varepsilon\ell\chi_{ND}^{\vartheta-\alpha}}{\Delta(\varepsilon)} > 0. \end{aligned}$$

With the primitives fixed,

$$\begin{aligned} \lambda_+(\varepsilon) &= \lambda_D + \frac{\ell\kappa(\chi_{ND}\chi_{DN})^{\vartheta-\alpha}}{\Delta}\varepsilon + O(\varepsilon^2), \\ \Delta(\varepsilon) &= \Delta + \ell \left(1 + \frac{2\kappa(\chi_{ND}\chi_{DN})^{\vartheta-\alpha}}{\Delta} \right) \varepsilon + O(\varepsilon^2). \end{aligned}$$

These expansions yield (7.31). Equating the two exponentials in (7.30), including their coefficients, gives the exact crossover in (7.32) when $0 < \omega_\varepsilon < 1/2$. Finally,

$$\log \frac{1 - \omega_\varepsilon}{\omega_\varepsilon} = \log(1/\varepsilon) - \log \omega + O(\varepsilon), \quad \Delta(\varepsilon) = \Delta + O(\varepsilon),$$

give its stated expansion, since $\varepsilon \log(1/\varepsilon) \rightarrow 0$.

These expansions hold with the primitives and $\Delta > 0$ fixed and are not uniform as $\Delta \downarrow 0$. Changing primitives affects both the loading and the eigenvalue gap, so a larger limiting gap need not imply an earlier finite-entry crossover. Small annual entry probabilities need not make the first-order loading or crossover approximations accurate. For example, at the 1.75% annual entry probability and $\kappa = 0.16$ in Figure 5, $\varepsilon\omega \approx 1.3473$, whereas $\omega_\varepsilon \approx 0.4241$; the coefficient-corrected crossover approximation is negative, whereas the exact crossover is 2.36 years. The figure therefore uses exact formulas throughout.

In the numerical benchmark, the reciprocal consumption multipliers satisfy $\chi_{ND}\chi_{DN} = 1$. The recovery jump reverses the discrete entry loss, not consumption-growth differences accumulated

during disaster. Jump amplitude cancels from the eigenvalues in (7.28) but still affects loadings and crossover maturities; nonreciprocal multipliers can also affect the eigenvalues.

APPENDIX B. NOTATION AND PRELIMINARIES

B.1. Notation. For scalar functions f and h , we write $f \sim h$, or say that f is *asymptotic to* h , if $f/h \rightarrow 1$ in the indicated limit, with h nonzero sufficiently near that limit. We use this notation both as maturity $t \rightarrow \infty$ and as the rare-entry parameter $\varepsilon \downarrow 0$.

The state space of the economic environment is the finite set $\mathbf{X} = \{x_1, \dots, x_n\}$. Let $\mathbb{R}^{\mathbf{X}}$ denote the vector space of real-valued functions on $\mathbf{X}$. For $g \in \mathbb{R}^{\mathbf{X}}$, write $g \geq 0$ if $g(x) \geq 0$ for every $x \in \mathbf{X}$, write $g \gg 0$ if $g(x) > 0$ for every $x \in \mathbf{X}$, and write $g > 0$ if $g \geq 0$ and $g \neq 0$. The positive cone is $\mathbb{R}_+^{\mathbf{X}} := \{g \in \mathbb{R}^{\mathbf{X}} : g \geq 0\}$. A linear operator $A : \mathbb{R}^{\mathbf{X}} \rightarrow \mathbb{R}^{\mathbf{X}}$ is *positive*, written $A \geq 0$, if $A\mathbb{R}_+^{\mathbf{X}} \subseteq \mathbb{R}_+^{\mathbf{X}}$. Equivalently, $A \geq 0$ if and only if $Ag \geq 0$ whenever $g \geq 0$.

For $g \in \mathbb{R}^{\mathbf{X}}$, the supremum norm is $\|g\|_{\infty} := \max_{x \in \mathbf{X}} |g(x)|$. For a linear operator A on $\mathbb{R}^{\mathbf{X}}$, $\|A\|_{\infty}$ denotes the operator norm induced by $\|\cdot\|_{\infty}$. Elsewhere, unsubscripted norms are fixed norms on the relevant finite-dimensional spaces, with matrix norms taken to be submultiplicative; on the shock space $\mathbb{R}^d$, $\|\cdot\|$ denotes the Euclidean norm.

For each $x \in \mathbf{X}$, let $e_x \in \mathbb{R}^{\mathbf{X}}$ denote the indicator function of state x : $e_x(x') := 1_{\{x\}}(x')$ for $x' \in \mathbf{X}$. Throughout the paper, we identify $\mathbb{R}^{\mathbf{X}}$ with $\mathbb{R}^n$ using the ordered basis $(e_{x_1}, \dots, e_{x_n})$, and use the same symbol for a function in $\mathbb{R}^{\mathbf{X}}$ and its coordinate column vector. Thus, for $g \in \mathbb{R}^{\mathbf{X}}$,

$$g_i := g(x_i), \quad g = \sum_{i=1}^n g_i e_{x_i}.$$

For $v \in \mathbb{R}^{\mathbf{X}}$, let $\text{diag}(v)$ denote the $n \times n$ diagonal matrix formed from the coordinates of v ; that is, the (i, j) -th entry of $\text{diag}(v)$ is v_i if $i = j$ and zero otherwise. For $u, v \in \mathbb{R}^{\mathbf{X}}$, their *Hadamard product*, denoted by $u \odot v$, is the componentwise product defined by $(u \odot v)_i := u_i v_i$ for $i = 1, \dots, n$. Moreover, for a subset $\mathcal{X} \subseteq \mathbf{X}$, let $v_{\mathcal{X}}$ denote the coordinate subvector indexed by $\mathcal{X}$. For subsets $\mathcal{X}, \mathcal{Y} \subseteq \mathbf{X}$, let $A(\mathcal{X}, \mathcal{Y})$ denote the submatrix of A with rows indexed by $\mathcal{X}$ and columns indexed by $\mathcal{Y}$; in particular, $A(\mathcal{X}, \mathcal{X})$ is the principal submatrix indexed by $\mathcal{X}$.

We similarly use the same symbol for a linear operator on $\mathbb{R}^{\mathbf{X}}$ and its matrix representation. For $A : \mathbb{R}^{\mathbf{X}} \rightarrow \mathbb{R}^{\mathbf{X}}$, the (i, j) -th matrix entry is

$$A_{ij} := (Ae_{x_j})(x_i), \quad i, j = 1, \dots, n.$$

Hence, for $g \in \mathbb{R}^{\mathbf{X}}$, $(Ag)(x_i) = (Ag)_i = \sum_{j=1}^n A_{ij} g_j$ for $i = 1, \dots, n$. Under this identification, A is positive if and only if its matrix is entrywise nonnegative, that is, $A \geq 0$ if and only if $A_{ij} \geq 0$ for all $i, j = 1, \dots, n$. Moreover, we write $A \gg 0$ whenever $A_{ij} > 0$ for all $i, j = 1, \dots, n$.

B.2. Spectral preliminaries. For a finite-dimensional linear operator A , we use $\sigma(A)$ to denote its spectrum. An eigenvalue in the spectrum is *semisimple* if its algebraic and geometric multiplicities coincide, or equivalently, if every Jordan block associated with it has size one. It is *algebraically*

simple if its algebraic multiplicity is one. Since geometric multiplicity cannot exceed algebraic multiplicity, every algebraically simple eigenvalue is semisimple, but the converse need not hold.

Furthermore, we use $s(A) := \max\{\Re \lambda : \lambda \in \sigma(A)\}$ to denote the spectral bound of A .

Lemma B.1 (Spectral bound of a Metzler matrix). *If A is Metzler, meaning that $A_{ij} \geq 0$ for $i \neq j$, then $s(A)$ is a real eigenvalue of A and is the unique spectral value with maximal real part. Equivalently, $s(A) \in \sigma(A) \cap \mathbb{R}$ and $\Re \mu < s(A)$ for every $\mu \in \sigma(A) \setminus \{s(A)\}$.*

Proof of Lemma B.1. See Theorem 7.2 of [Bátkai et al. \(2017, p. 82\)](#). □

For $z \notin \sigma(A)$, the *resolvent* of A is $R(z; A) := (zI - A)^{-1}$. The *Riesz spectral projection* of A associated with an eigenvalue $\lambda \in \sigma(A)$ is defined as

$$P_\lambda(A) := \frac{1}{2\pi i} \int_{\Gamma_\lambda} R(z; A) dz,$$

where Γ_λ is a positively oriented simple closed contour that encloses λ and no other point of $\sigma(A)$. $P_\lambda(A)$ is a projection satisfying $P_\lambda(A)^2 = P_\lambda(A)$. It is independent of the choice of Γ_λ and commutes with A . When the operator is clear from context, we suppress its argument and write P_λ . In general, the range and rank of $P_\lambda(A)$ are the generalized eigenspace associated with λ and the algebraic multiplicity of λ , respectively. If λ is algebraically simple and φ and ψ are corresponding right and left eigenvectors normalized by $\psi^\top \varphi = 1$, then $P_\lambda(A) = \varphi \psi^\top$.

B.3. Graph structure and Frobenius form of Metzler matrices. We collect several graph-theoretic concepts used in combinatorial matrix theory and finite-state Markov-chain analysis. Much of this subsection is adapted from [Rothblum \(1975\)](#) and [Schneider \(1986\)](#); see also [Beare and Toda \(2026, Section 3.1\)](#) for a review.

Let $A : \mathbb{R}^X \rightarrow \mathbb{R}^X$ be a linear operator whose matrix representation $A = (A_{ij})$ is Metzler, meaning that $A_{ij} \geq 0$ for $i \neq j$. The *directed graph* associated with A , denoted by $\mathcal{G}(A)$, records the off-diagonal zero pattern of A . It has vertex set $X = \{x_1, \dots, x_n\}$ and a directed edge from x_i to x_j , written $x_i \rightarrow x_j$, if and only if $i \neq j$ and $A_{ij} > 0$.

A *directed walk* of length $p \in \mathbb{N}_0$ is a sequence $(x_{i_0}, \dots, x_{i_p})$ of vertices of $\mathcal{G}(A)$ such that $x_{i_{r-1}} \rightarrow x_{i_r}$ for every $r = 1, \dots, p$. Directed walks with length zero are allowed. We say that x_i has *access to* x_j , or equivalently that x_j is *accessible from* x_i , written $x_i \rightsquigarrow x_j$, if there exists a directed walk from x_i to x_j . Thus every state has access to itself through the walk of length zero. A *directed path* is a directed walk with no repeated vertices. By deleting cycles from a walk, any directed walk from x_i to x_j can be reduced to a directed path with the same endpoints. Hence $x_i \rightsquigarrow x_j$ if and only if there is a directed path from x_i to x_j .

We say that two states $x_i, x_j \in X$ *communicate*, written $x_i \sim x_j$, if $x_i \rightsquigarrow x_j$ and $x_j \rightsquigarrow x_i$. Thus, communication is mutual accessibility. The relation $\sim$ is an equivalence relation and partitions X into disjoint equivalence classes $\mathcal{C}_1, \dots, \mathcal{C}_m \subseteq X$, for some $m \in \mathbb{N}$. These equivalence classes are called *classes*. Each class is a maximal subset of states whose elements communicate pairwise; equivalently, the classes are the *strongly connected components* of $\mathcal{G}(A)$.

A class $\mathcal{C} \subseteq \mathbf{X}$ is *closed*, also called *final* in combinatorial matrix theory, if no state in $\mathcal{C}$ has access to any state outside $\mathcal{C}$. Equivalently, $A_{ij} = 0$ whenever $x_i \in \mathcal{C}$ and $x_j \notin \mathcal{C}$. Thus, closed classes are precisely the *terminal strongly connected components* of $\mathcal{G}(A)$. When A is a Markov generator, this graph-theoretic classification coincides with the usual probabilistic one: a state is *recurrent* if and only if it belongs to a closed class, and *transient* if and only if it belongs to a nonclosed class.

If $\mathbf{X}$ consists of a single class, then that class is necessarily closed. The graph $\mathcal{G}(A)$ is then strongly connected, and the Metzler matrix A is called *irreducible*. In addition to the general spectral conclusion of Lemma B.1, Perron–Frobenius theory for irreducible Metzler matrices implies that $s(A)$ is an algebraically simple eigenvalue and there exist strictly positive left and right eigenvectors associated with $s(A)$.

The *reduced graph* of A , also called its *condensation graph* and denoted by $\mathcal{G}_r(A)$, is obtained by collapsing each class into a single vertex while retaining all directed connections between distinct classes. Its vertex set is therefore $\mathcal{C}(A) := \{\mathcal{C}_1, \dots, \mathcal{C}_m\}$, and it contains a directed edge $\mathcal{C}_k \rightarrow \mathcal{C}_\ell$ if and only if $k \neq \ell$ and $x_i \rightarrow x_j$ for some $x_i \in \mathcal{C}_k$ and $x_j \in \mathcal{C}_\ell$.

In the reduced graph, a *directed path* of length $p \in \mathbb{N}_0$ is a sequence $\gamma = (\mathcal{C}_{k_0}, \dots, \mathcal{C}_{k_p})$ of distinct vertices of $\mathcal{G}_r(A)$ such that $\mathcal{C}_{k_{r-1}} \rightarrow \mathcal{C}_{k_r}$ for every $r = 1, \dots, p$. We write $|\gamma| = p$ for the length of γ , that is, its number of directed edges. In particular, paths of length zero are allowed. We say that $\mathcal{C}_k$ has *access to* $\mathcal{C}_\ell$, written $\mathcal{C}_k \rightsquigarrow \mathcal{C}_\ell$, if there exists a directed path from $\mathcal{C}_k$ to $\mathcal{C}_\ell$. Allowing paths of length zero makes each class have access to itself. In particular, class-level accessibility agrees with state-level accessibility:

$$\mathcal{C}_k \rightsquigarrow \mathcal{C}_\ell \iff x_i \rightsquigarrow x_j \text{ for some } x_i \in \mathcal{C}_k, x_j \in \mathcal{C}_\ell.$$

Because all states within a class communicate, “some” may equivalently be replaced by “every.”

State-level accessibility is reflexive and transitive on $\mathbf{X}$ but need not be antisymmetric, because distinct states in the same class are mutually accessible. It therefore defines a preorder on $\mathbf{X}$. After quotienting by communication, class-level accessibility is antisymmetric and hence defines a partial order on $\mathcal{C}(A)$. Equivalently, the reduced graph $\mathcal{G}_r(A)$ is acyclic.

Two distinct classes $\mathcal{C}_k, \mathcal{C}_\ell \in \mathcal{C}(A)$ are said to be *comparable* if $\mathcal{C}_k \rightsquigarrow \mathcal{C}_\ell$ or $\mathcal{C}_\ell \rightsquigarrow \mathcal{C}_k$; otherwise they are *incomparable*. Following Rothblum (1975), a *chain of classes* is a nonempty collection of pairwise comparable classes; in particular, any singleton class forms a chain. Equivalently, its elements can be ordered as $\mathcal{C}_{k_0} \rightsquigarrow \mathcal{C}_{k_1} \rightsquigarrow \dots \rightsquigarrow \mathcal{C}_{k_q}$, for some $q \in \mathbb{N}_0$.

Since the class-level accessibility defines a partial order on $\mathcal{C}(A)$, we may label the classes so that

$$\mathcal{C}_k \rightsquigarrow \mathcal{C}_\ell \text{ for } k \neq \ell \implies k > \ell. \quad (\text{B.1})$$

After arranging the state coordinates class by class, with an arbitrary ordering within each class, A has the block lower-triangular Frobenius form

$$A = \begin{pmatrix} A_{11} & 0 & \cdots & 0 \\ A_{21} & A_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A_{m1} & A_{m2} & \cdots & A_{mm} \end{pmatrix}, \quad (\text{B.2})$$

where $A_{k\ell} = A(\mathcal{C}_k, \mathcal{C}_\ell)$ is the submatrix with rows indexed by $\mathcal{C}_k$ and columns indexed by $\mathcal{C}_\ell$. The ordering of the blocks in the Frobenius form is generally nonunique because some classes may be incomparable under accessibility.

Each diagonal block A_{kk} in (B.2) is Metzler and irreducible because its associated graph is the strongly connected subgraph induced by the class $\mathcal{C}_k$. For $k \neq \ell$, the off-diagonal block $A_{k\ell}$ is nonzero if and only if $\mathcal{C}_k \rightarrow \mathcal{C}_\ell$. Due to (B.1), every nonzero off-diagonal block lies strictly below the diagonal. Moreover, $\mathcal{C}_k$ is closed if and only if $A_{k\ell} = 0$ for every $\ell < k$.

B.4. Positivity and class-path expansions. A fundamental property of Metzler matrices is that they generate positive semigroups: if A is Metzler, then $e^{tA} \geq 0$ for every $t \geq 0$; see, for example, [Bátkai et al. \(2017, Theorem 7.1\)](#). The following lemma strengthens this result by characterizing the entrywise positivity pattern of e^{tA} in terms of accessibility in $\mathcal{G}(A)$. For any $g \in \mathbb{R}^X$, let $\text{supp}(g) := \{x \in X : g(x) \neq 0\} \subseteq X$ denote its support.

Lemma B.2. *Let $A : \mathbb{R}^X \rightarrow \mathbb{R}^X$ be a linear operator whose matrix representation is Metzler, and let $\mathcal{G}(A)$ be the associated directed graph. For $i, j = 1, \dots, n$, the following statements are equivalent:*

- (i) $x_i \rightsquigarrow x_j$ in $\mathcal{G}(A)$.
- (ii) $(e^{tA})_{ij} > 0$ for every $t > 0$.
- (iii) $(e^{tA})_{ij} > 0$ for some $t > 0$.

Moreover, for every $t > 0$, $g \in \mathbb{R}_+^X$, and $x \in X$, $(e^{tA}g)(x) > 0$ if and only if there exists $y \in \text{supp}(g)$ such that $x \rightsquigarrow y$ in $\mathcal{G}(A)$.

Proof of Lemma B.2. Choose a scalar $a \geq 0$ sufficiently large that $B = A + aI \geq 0$. Since A and I commute, $e^{tA} = e^{-at}e^{tB}$. Hence, for every $t > 0$, $(e^{tA})_{ij} > 0$ if and only if $(e^{tB})_{ij} > 0$. Moreover, $e^{tB} = \sum_{k=0}^{\infty} t^k B^k / k!$. Because $B \geq 0$, every term in this series is entrywise nonnegative. Hence for any $t > 0$, $(e^{tB})_{ij} > 0$ if and only if $(B^k)_{ij} > 0$ for some $k \in \mathbb{N}_0$.

We first prove (i) $\Rightarrow$ (ii). Suppose that $x_i \rightsquigarrow x_j$, and let $x_{r_0} \rightarrow x_{r_1} \rightarrow \cdots \rightarrow x_{r_k}$ be a directed walk from $x_i = x_{r_0}$ to $x_j = x_{r_k}$. If $k = 0$, then $(B^0)_{ij} = 1$. If $k \geq 1$, then, since B and A have the same off-diagonal entries, $B_{r_{\ell-1}r_\ell} = A_{r_{\ell-1}r_\ell} > 0$, for $\ell = 1, \dots, k$. The expansion

$$(B^k)_{ij} = \sum_{s_1, \dots, s_{k-1}=1}^n B_{is_1} B_{s_1 s_2} \cdots B_{s_{k-1} j}, \quad (\text{B.3})$$

therefore gives $(B^k)_{ij} \geq B_{r_0 r_1} B_{r_1 r_2} \cdots B_{r_{k-1} r_k} > 0$. In either case, $(B^k)_{ij} > 0$ for some $k \in \mathbb{N}_0$, and therefore $(e^{tA})_{ij} > 0$ for every $t > 0$.

The implication (ii) $\Rightarrow$ (iii) is immediate. It remains to prove (iii) $\Rightarrow$ (i). Suppose that $(e^{tA})_{ij} > 0$ for some $t > 0$. Then $(B^k)_{ij} > 0$ for some $k \in \mathbb{N}_0$. If $k = 0$, necessarily $i = j$, so the length-zero path gives $x_i \rightsquigarrow x_j$. If $k \geq 1$, the nonnegativity of the summands in (B.3) implies that there exist indices $s_1, \dots, s_{k-1}$ such that $B_{s_0 s_1} B_{s_1 s_2} \cdots B_{s_{k-1} s_k} > 0$, where $s_0 \equiv i$ and $s_k \equiv j$. Whenever $s_{\ell-1} \neq s_\ell$, $A_{s_{\ell-1} s_\ell} = B_{s_{\ell-1} s_\ell} > 0$, so $x_{s_{\ell-1}} \rightarrow x_{s_\ell}$ is an edge of $\mathcal{G}(A)$. Removing all consecutive repetitions $s_{\ell-1} = s_\ell$ therefore produces a directed walk from x_i to x_j in $\mathcal{G}(A)$, possibly of length zero. Hence $x_i \rightsquigarrow x_j$.

For the last statement, note that $(e^{tA}g)(x_i) = \sum_{j=1}^n (e^{tA})_{ij}g(x_j) = \sum_{x_j \in \text{supp}(g)} (e^{tA})_{ij}g(x_j)$ for some $i = 1, \dots, n$, $t > 0$ and $g \geq 0$. Because $e^{tA} \geq 0$ and $g(x_j) > 0$ for every $x_j \in \text{supp}(g)$, the sum is positive if and only if $(e^{tA})_{ij} > 0$ for some $x_j \in \text{supp}(g)$. The statement then follows from the equivalence between (i) and (ii). $\square$

When A is written in the Frobenius form (B.2), Lemma B.2 also shows that e^{tA} is block lower triangular for every $t \geq 0$. The claim is immediate at $t = 0$; if $t > 0$ and $k < \ell$, then $\mathcal{C}_k \not\rightsquigarrow \mathcal{C}_\ell$ by (B.1), and hence $(e^{tA})_{\mathcal{C}_k, \mathcal{C}_\ell} = 0$. Its block zero pattern need not, however, coincide with that of A : an indirect path from $\mathcal{C}_k$ to $\mathcal{C}_\ell$ makes the corresponding block of e^{tA} entrywise strictly positive for every $t > 0$, even when $A_{k\ell} = 0$. Lastly, if A is Metzler and irreducible, all states are mutually accessible, so Lemma B.2 implies that $e^{tA} \gg 0$ for every $t > 0$.

We next derive a path expansion for the resolvent of a Metzler matrix A in the Frobenius form (B.2). Because A is block triangular, $z \notin \sigma(A)$ also lies outside the spectrum of every diagonal block A_{kk} . Hence $R(z; A)$ and $R(z; A_{kk})$, $k = 1, \dots, m$, are all well defined. For $\mathcal{C}, \mathcal{C}' \in \mathcal{C}(A)$, let $\mathcal{D}(\mathcal{C}, \mathcal{C}')$ be the set of all directed paths in $\mathcal{G}_r(A)$ from $\mathcal{C}$ to $\mathcal{C}'$. For a directed path $\gamma = (\mathcal{C}_{k_0}, \dots, \mathcal{C}_{k_p})$, define its *resolvent contribution* by

$$\mathcal{R}_\gamma(z; A) := R(z; A_{k_0 k_0}) A_{k_0 k_1} R(z; A_{k_1 k_1}) \cdots A_{k_{p-1} k_p} R(z; A_{k_p k_p}). \quad (\text{B.4})$$

For a path of length zero, we have $\mathcal{R}_{(\mathcal{C}_k)}(z; A) = R(z; A_{kk})$.

Lemma B.3 (Resolvent path expansion). *Let $A : \mathbb{R}^X \rightarrow \mathbb{R}^X$ be a linear operator whose matrix representation is Metzler in its Frobenius form (B.2). Let $\mathcal{R}_\gamma(z; A)$ be the resolvent contribution from a path γ in $\mathcal{G}_r(A)$, defined in (B.4). For any $\mathcal{C}, \mathcal{C}' \in \mathcal{C}(A)$, the submatrix of the resolvent $R(z; A)$ with rows indexed by $\mathcal{C}$ and columns indexed by $\mathcal{C}'$ is given by*

$$[R(z; A)]_{\mathcal{C}, \mathcal{C}'} = \sum_{\gamma \in \mathcal{D}(\mathcal{C}, \mathcal{C}')} \mathcal{R}_\gamma(z; A), \quad z \in \mathbb{C} \setminus \sigma(A), \quad (\text{B.5})$$

with the convention that the sum is zero when $\mathcal{D}(\mathcal{C}, \mathcal{C}') = \emptyset$. Moreover, if $z > s(A)$ is real, every resolvent contribution $\mathcal{R}_\gamma(z; A)$ is entrywise strictly positive, and

$$[R(z; A)]_{\mathcal{C}, \mathcal{C}'} \begin{cases} \gg 0, & \text{if } \mathcal{C} \rightsquigarrow \mathcal{C}', \\ = 0, & \text{if } \mathcal{C} \not\rightsquigarrow \mathcal{C}'. \end{cases} \quad (\text{B.6})$$

Proof of Lemma B.3. Write $A = D + K$, where $D = \text{diag}(A_{11}, \dots, A_{mm})$ and $K = A - D$. Then

$$R(z; D) = \text{diag}(R(z; A_{11}), \dots, R(z; A_{mm})), \quad z \in \mathbb{C} \setminus \sigma(A). \quad (\text{B.7})$$

Since A is block triangular, $\sigma(A) = \bigcup_{k=1}^m \sigma(A_{kk})$, so $R(z; D)$ is well defined whenever $z \notin \sigma(A)$.

Since $R(z; D)$ is block diagonal and K is strictly lower block triangular—that is, block lower triangular with zero diagonal blocks—the product $KR(z; D)$ is also strictly lower block triangular. For a strictly lower block triangular matrix, each successive power moves its potentially nonzero blocks at least one block farther below the diagonal. Since there are only m diagonal blocks, $(KR(z; D))^m = 0$. Thus $KR(z; D)$ is nilpotent, with nilpotency index at most m .

Taking inverses in the identity $zI - A = (I - KR(z; D))(zI - D)$ and using the nilpotency of $KR(z; D)$ yields the finite resolvent expansion

$$R(z; A) = R(z; D)(I - KR(z; D))^{-1} = R(z; D) \sum_{p=0}^{m-1} (KR(z; D))^p, \quad z \in \mathbb{C} \setminus \sigma(A). \quad (\text{B.8})$$

Fix $z \in \mathbb{C} \setminus \sigma(A)$, and abbreviate $R_D := R(z; D)$. It follows from (B.8) and (B.7) that the (k, ℓ) block of $R(z; A)$ is

$$[R(z; A)]_{k\ell} = \sum_{p=0}^{m-1} \mathcal{R}_{k\ell}(p), \quad (\text{B.9})$$

where

$$\mathcal{R}_{k\ell}(p) := [R_D(KR_D)^p]_{k\ell} = R(z; A_{kk})[(KR_D)^p]_{k\ell}. \quad (\text{B.10})$$

We claim that, for $p = 0, \dots, m-1$,

$$\mathcal{R}_{k\ell}(p) = \sum_{\substack{\gamma \in \mathcal{D}(\mathcal{C}_k, \mathcal{C}_\ell) \\ |\gamma| = p}} \mathcal{R}_\gamma(z; A). \quad (\text{B.11})$$

Together with (B.9), identity (B.11) yields (B.5).

We prove (B.11) by induction on p , beginning with the case $p = 0$. Under the length-zero-path convention, the right-hand side of (B.11) gives $R(z; A_{kk})$ if $k = \ell$ and zero otherwise. This agrees with $\mathcal{R}_{k\ell}(0)$ as defined in (B.10).

Now suppose that (B.11) holds for some $0 \leq p \leq m-2$. It follows from (B.10) that

$$\mathcal{R}_{k\ell}(p+1) = [R_D(KR_D)^p(KR_D)]_{k\ell} = \sum_{s=1}^m [R_D(KR_D)^p]_{ks} [KR_D]_{s\ell} = \sum_{s=1}^m \mathcal{R}_{ks}(p) [KR_D]_{s\ell}.$$

Since R_D is block diagonal and K is strictly lower block triangular,

$$[KR_D]_{s\ell} = K_{s\ell}R(z; A_{\ell\ell}) = \begin{cases} A_{s\ell}R(z; A_{\ell\ell}), & s > \ell, \\ 0, & s \leq \ell. \end{cases}$$

Therefore, using the induction hypothesis,

$$\mathcal{R}_{k\ell}(p+1) = \sum_{s>\ell} \mathcal{R}_{ks}(p) A_{s\ell} R(z; A_{\ell\ell}) = \sum_{s>\ell} \sum_{\substack{\gamma \in \mathcal{D}(\mathcal{C}_k, \mathcal{C}_s) \\ |\gamma|=p}} \mathcal{R}_{\gamma}(z; A) A_{s\ell} R(z; A_{\ell\ell}).$$

Appending the edge $\mathcal{C}_s \rightarrow \mathcal{C}_\ell$ to a length- p path from $\mathcal{C}_k$ to $\mathcal{C}_s$ produces a length- $(p+1)$ path from $\mathcal{C}_k$ to $\mathcal{C}_\ell$. Conversely, every length- $(p+1)$ path from $\mathcal{C}_k$ to $\mathcal{C}_\ell$ has a unique final edge $\mathcal{C}_s \rightarrow \mathcal{C}_\ell$ and is obtained by appending this edge to its length- p initial segment. Consequently, the product in the summation is precisely $\mathcal{R}_{\gamma'}(z; A)$ with some $\gamma' \in \mathcal{D}(\mathcal{C}_k, \mathcal{C}_\ell)$ and $|\gamma'| = p+1$. Hence

$$\mathcal{R}_{k\ell}(p+1) = \sum_{\substack{\gamma \in \mathcal{D}(\mathcal{C}_k, \mathcal{C}_\ell) \\ |\gamma|=p+1}} \mathcal{R}_{\gamma}(z; A).$$

This completes the induction and proves (B.11).

Finally, let $z > s(A)$ be real. Since $z > s(A_{kk})$ for every $k = 1, \dots, m$, each diagonal-block resolvent admits the Laplace representation (e.g., [Bátkai et al., 2017](#), Proposition 9.33(a)) $R(z; A_{kk}) = \int_0^\infty e^{-zt} e^{tA_{kk}} dt$. Because A_{kk} is irreducible and Metzler, Lemma B.2 implies that $e^{tA_{kk}} \gg 0$ for every $t > 0$. Hence $R(z; A_{kk}) \gg 0$. Moreover, every coupling block along a directed path is nonzero and nonnegative. Alternating these coupling blocks with the strictly positive diagonal-block resolvents in (B.4) shows that every resolvent contribution $\mathcal{R}_{\gamma}(z; A)$ is entrywise strictly positive. The expansion (B.5), together with the empty-sum convention, now yields (B.6). $\square$

Equation (B.5) implies that the (k, ℓ) block of $R(z; A)$ vanishes whenever $\mathcal{D}(\mathcal{C}_k, \mathcal{C}_\ell) = \emptyset$. In particular, because the Frobenius form (B.2) is block lower triangular, there is no directed path from $\mathcal{C}_k$ to $\mathcal{C}_\ell$ when $k < \ell$. Hence $R(z; A)$ is block lower triangular for every $z \in \mathbb{C} \setminus \sigma(A)$.

By the Perron–Frobenius properties recorded in Appendix B.3, each irreducible block L_k has right and left Perron vectors $\varphi_k, \psi_k \gg 0$, which we normalize by $\psi_k^\top \varphi_k = 1$. Since $\lambda_k = s(L_k)$ is algebraically simple,

$$R(z; L_k) = \frac{\varphi_k \psi_k^\top}{z - \lambda_k} + H_k(z), \quad (\text{B.12})$$

where H_k is analytic near λ_k . Moreover, Lemma B.2 and the semigroup–resolvent formula imply that $R(\mu; L_k) \gg 0$ whenever $\mu > \lambda_k$; this is also the one-class case of Lemma B.3.

APPENDIX C. THE IRREDUCIBLE HANSEN–SCHEINKMAN BENCHMARK

Suppose that the pricing generator L_Π is irreducible. Perron–Frobenius theory then implies that $\lambda^* = s(L_\Pi)$ is algebraically simple and admits strictly positive right and left eigenvectors. Let $\varphi, \psi \gg 0$ be normalized by $\mathbf{1}^\top \varphi = 1$ and $\psi^\top \varphi = 1$, so that

$$L_\Pi \varphi = \lambda^* \varphi, \quad L_\Pi^\top \psi = \lambda^* \psi.$$

These normalizations uniquely determine both vectors. The Riesz projection at λ^* is $P_{\lambda^*} = \varphi \psi^\top$, and

$$\lim_{t \rightarrow \infty} e^{-\lambda^* t} \Pi_t = P_{\lambda^*}.$$

Indeed, Theorem 7.6(a) and Corollary 7.5 of [Bátkai et al. \(2017, p. 84\)](#), applied to the irreducible Metzler matrix $L_\Pi - \lambda^* I$, identify this limit with its spectral projection at zero. Consequently, for every $g \in \mathbb{R}^X$,

$$\lim_{t \rightarrow \infty} e^{-\lambda^* t} (\Pi_t g)(x) = \varphi(x) \psi^\top g, \quad x \in X.$$

The right eigenvector φ is the state-dependent long-run pricing factor, while $\psi^\top g$ is the payoff's loading on the dominant component. Since both eigenvectors are strictly positive, every nonzero $g \geq 0$ has a positive dominant component from every initial state and hence the common long-run yield $-\lambda^*$.

The right eigenvector also gives the finite-state Hansen–Scheinkman factorization ([Hansen and Scheinkman, 2009](#)). Since $\Pi_t \varphi = e^{\lambda^* t} \varphi$, the process

$$M_t^\varphi := e^{-\lambda^* t} S_t^G \frac{\varphi(X_t)}{\varphi(X_0)}, \quad t \geq 0,$$

is a strictly positive unit-mean martingale under every $\mathbb{P}_x$, and the growth-adjusted SDF factors as

$$S_t^G = e^{\lambda^* t} \frac{\varphi(X_0)}{\varphi(X_t)} M_t^\varphi. \quad (\text{C.1})$$

The factors in (C.1) separate the long-run exponential rate, a bounded state adjustment, and a martingale change of measure. Defining $d\mathbb{P}_x^\varphi / d\mathbb{P}_x|_{\mathcal{F}_t} = M_t^\varphi$ reweights state histories by their importance for long-run pricing.

Equivalently, the Doob h -transform of Π with $h = \varphi$ (see, e.g., [Doob, 1957](#)) is

$$P_t^\varphi := e^{-\lambda^* t} \text{diag}(\varphi)^{-1} \Pi_t \text{diag}(\varphi), \quad t \geq 0. \quad (\text{C.2})$$

It is a Markov semigroup and is the transition semigroup of X under $(\mathbb{P}_x^\varphi)_{x \in X}$. Its generator and unique stationary distribution are

$$L^\varphi = \text{diag}(\varphi)^{-1} (L_\Pi - \lambda^* I) \text{diag}(\varphi), \quad \pi^\varphi = \varphi \odot \psi.$$

Thus the transformed generator incorporates physical transition possibilities, discounting, cash-flow growth, and risk adjustment through the long-run pricing eigenfunction. The construction relies on $\varphi \gg 0$: under reducibility the right eigenvector can vanish, so neither the density nor the whole-space Doob transform is generally defined. Section 4.2 instead uses a transform on the exact support of the dominant eigenvector; Appendix D.3 develops this transform and the corresponding adjoint construction. In the notation of Appendix D.3, irreducibility gives $L_\Pi^\rightarrow = L^\varphi$ and

$$L_\Pi^\leftarrow = \text{diag}(\pi^\varphi)^{-1} (L_\Pi^\rightarrow)^\top \text{diag}(\pi^\varphi).$$

Thus $L_\Pi^\leftarrow$ is the stationary time reversal of the transformed pricing dynamics, not of the original physical dynamics generated by L_X . [Christensen \(2017, Section 2.1, p. 1506\)](#) calls the adjoint pricing operator *time reversed*. Here *graph reversal* refers only to edge reversal; under reducibility, the two transforms generally act on different supports and are not global time reversals.

APPENDIX D. PROOFS

D.1. Global valuation stability.

Proof of Proposition 3.1. Because Π_t is positive and the growth-adjusted SDF S^G is strictly positive, $z_t^{\text{env}} = \|\Pi_t \mathbf{1}\|_\infty = \|\Pi_t\|_\infty > 0$ for all $t \geq 0$. Finite-dimensional semigroup theory (see, e.g. [Bátkai et al., 2017](#), p. 48, Proposition 4.7) gives $\lim_{t \rightarrow \infty} t^{-1} \log \|\Pi_t\|_\infty = s(L_\Pi) = \lambda^*$. $\square$

Corollary D.1 (Ex ante pricing rate). *Suppose that the initial state has a full-support probability distribution μ on $\mathbf{X}$. Define the ex ante growth-adjusted zero-coupon price and yield by $z_t^\mu := \sum_{x \in \mathbf{X}} \mu(x) z_t(x)$ for $t \geq 0$ and $y_t^\mu := -t^{-1} \log z_t^\mu$ for $t > 0$. Under Assumptions 2.1–2.3,*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log z_t^\mu = \lambda^*, \quad \lim_{t \rightarrow \infty} y_t^\mu = -\lambda^*.$$

Proof of Corollary D.1. Since μ has full support, we have $\mu_\circ := \min_x \mu(x) > 0$. It follows that $z_t^\mu = \sum_{x \in \mathbf{X}} \mu(x) z_t(x) = \sum_{x \in \mathbf{X}} \mu(x) (\Pi_t \mathbf{1})(x) \leq \max_{x \in \mathbf{X}} (\Pi_t \mathbf{1})(x) = \|\Pi_t \mathbf{1}\|_\infty = z_t^{\text{env}}$. On the other hand, let $x_t^* \in \mathbf{X}$ denote a state such that $z_t^{\text{env}} = z_t(x_t^*)$. Then, $z_t^\mu = \sum_{x \in \mathbf{X}} \mu(x) z_t(x) \geq \mu(x_t^*) z_t(x_t^*) \geq \mu_\circ z_t^{\text{env}}$. Consequently, we have $0 < \mu_\circ z_t^{\text{env}} \leq z_t^\mu \leq z_t^{\text{env}}$, and hence $t^{-1} \log \mu_\circ + t^{-1} \log z_t^{\text{env}} \leq t^{-1} \log z_t^\mu \leq t^{-1} \log z_t^{\text{env}}$. The corollary then follows from the squeeze theorem and Proposition 3.1. $\square$

Proof of Theorem 3.2. Since Π_t is positive, its operator norm induced by the supremum vector norm satisfies

$$\|\Pi_t\|_\infty = \|\Pi_t \mathbf{1}\|_\infty = \max_{x \in \mathbf{X}} z_t(x) = z_t^{\text{env}}, \quad t \geq 0. \quad (\text{D.1})$$

Moreover, Proposition 4.7 of [Bátkai et al. \(2017\)](#) gives

$$\lambda^* = s(L_\Pi) = \inf_{t > 0} \frac{1}{t} \log \|\Pi_t\|_\infty. \quad (\text{D.2})$$

It follows from (D.1) and (D.2) that $\lambda^* < 0$ if and only if there exists $t > 0$ such that $\|\Pi_t\|_\infty < 1$, which is equivalent to $z_t^{\text{env}} < 1$. This proves the equivalence between (G2) and (G4). The equivalence between (G2) and (G3) follows immediately from Proposition 3.1.

We next show that (G4) implies both (G5) and (G1). Suppose that (G4) holds and hence there exists some $T > 0$ such that $q := z_T^{\text{env}} = \|\Pi_T\|_\infty < 1$. By continuity of the pricing semigroup, it follows that $C := \max_{0 \leq r \leq T} \|\Pi_r\|_\infty < \infty$. For any $t \geq 0$, write $t = kT + r$, where $k \in \mathbb{N}_0$ and $0 \leq r < T$. The semigroup property gives $\Pi_t = \Pi_T^k \Pi_r$, and therefore $\|\Pi_t\|_\infty \leq C q^k$. Since $q < 1$, we have $\|\Pi_t\|_\infty \rightarrow 0$ as $t \rightarrow \infty$ (equivalently, as $k \rightarrow \infty$). By (D.1), $z_t^{\text{env}} \rightarrow 0$, proving (G5). Next, partitioning $[0, \infty)$ into intervals of length T , together with $\|\Pi_t\|_\infty \leq C q^k$, gives

$$\int_0^\infty \|\Pi_t\|_\infty dt = \sum_{k=0}^\infty \int_{kT}^{(k+1)T} \|\Pi_t\|_\infty dt \leq CT \sum_{k=0}^\infty q^k = \frac{CT}{1-q} < \infty. \quad (\text{D.3})$$

Hence, for every $g \in \mathbb{R}_+^\mathbf{X}$ and $x \in \mathbf{X}$, $0 \leq \int_0^\infty (\Pi_t g)(x) dt \leq \|g\|_\infty \int_0^\infty \|\Pi_t\|_\infty dt < \infty$. Thus, global valuation stability holds, proving (G1).

Conversely, (G5) implies (G4): if $z_t^{\text{env}} \rightarrow 0$ as $t \rightarrow \infty$, then $z_T^{\text{env}} < 1$ for all sufficiently large T . It remains to show that (G1) implies (G4). Under global valuation stability, $\int_0^\infty z_t(x) dt < \infty$ for

every $x \in \mathbf{X}$. Using (D.1) and the finiteness of $\mathbf{X}$,

$$\int_0^\infty \|\Pi_t\|_\infty dt = \int_0^\infty \max_{x \in \mathbf{X}} z_t(x) dt \leq \sum_{x \in \mathbf{X}} \int_0^\infty z_t(x) dt < \infty.$$

Consequently, $\|\Pi_t\|_\infty < 1$ for some $t > 0$; otherwise the integral would be infinite. Equation (D.1) then gives $z_t^{\text{env}} < 1$, proving (G4). This establishes the equivalence of (G1)–(G5).

Under these conditions, (D.3) shows that $\int_0^\infty \|\Pi_t\|_\infty dt < \infty$. Hence, the operator-valued function $t \mapsto \Pi_t$ is Bochner integrable on $[0, \infty)$. Define $\tilde{V} : \mathbb{R}^{\mathbf{X}} \rightarrow \mathbb{R}^{\mathbf{X}}$ by $\tilde{V}g := \int_0^\infty \Pi_t g dt$, for $g \in \mathbb{R}^{\mathbf{X}}$. Norm integrability makes $\tilde{V}$ a finite linear operator, and positivity of Π_t makes $\tilde{V}$ positive. For $g \in \mathbb{R}_+^{\mathbf{X}}$, equation (2.9) gives $\tilde{V}g = Vg$. Thus $\tilde{V}$ is a positive linear extension of V , and the extension is unique because $\mathbb{R}^{\mathbf{X}} = \mathbb{R}_+^{\mathbf{X}} - \mathbb{R}_+^{\mathbf{X}}$. Relabeling $\tilde{V}$ as V gives $V = \int_0^\infty \Pi_t dt$. Since $\Pi_t = \exp(tL_\Pi)$, we have $\frac{d}{dt}\Pi_t = L_\Pi \Pi_t = \Pi_t L_\Pi$, for $t > 0$; see, for example [Bátkai et al. \(2017, Theorem 4.2\)](#). Therefore, for every $\tau > 0$, $L_\Pi \int_0^\tau \Pi_t dt = (\int_0^\tau \Pi_t dt) L_\Pi = \Pi_\tau - I$. Letting $\tau \rightarrow \infty$ and using $\Pi_\tau \rightarrow 0$ gives $L_\Pi V = V L_\Pi = -I$. Consequently, L_Π is invertible and $V = -L_\Pi^{-1}$. $\square$

D.2. Global spectral structure.

Proof of Theorem 4.1. Let $k_* \in \{1, \dots, m\}$ satisfy $\mathcal{C}_{k_*} = \mathcal{C}^*$ in the ordering compatible with the Frobenius form of L_Π in (4.1). We have

$$\lambda^* = s(L_{k_*}) > s(L_k) \quad \text{for all } k \in \{1, \dots, m\} \setminus \{k_*\}. \quad (\text{D.4})$$

Since L_{k_*} is irreducible Metzler, Perron–Frobenius theory implies that λ^* is an algebraically simple eigenvalue of L_{k_*} and the one-dimensional right and left eigenspaces admit strictly positive eigenvectors. Moreover, as the characteristic polynomial of L_Π is the product of those of its diagonal blocks and $\lambda^* \notin \sigma(L_k)$ for $k \neq k_*$, by (D.4), λ^* is algebraically simple for L_Π .

By (D.4), the resolvent identity $R(\lambda^*; L_k) = \int_0^\infty e^{-\lambda^* t} e^{tL_k} dt$ holds for $k \neq k_*$. Since L_k is irreducible Metzler, $e^{tL_k} \gg 0$ for all $t > 0$ by Lemma B.2. Consequently,

$$R(\lambda^*; L_k) \gg 0, \quad k \in \{1, \dots, m\} \setminus \{k_*\}. \quad (\text{D.5})$$

Let $\varphi = (\varphi_{c_1}^\top, \dots, \varphi_{c_m}^\top)^\top$, partitioned conformably with the Frobenius form of L_Π in (4.1), be a right eigenvector of L_Π associated with λ^* such that $(\lambda^* I - L_\Pi)\varphi = 0$. Then, it follows that

$$(\lambda^* I - L_k)\varphi_{c_k} = \sum_{\ell=1}^{k-1} L_{k\ell} \varphi_{c_\ell}, \quad k = 1, \dots, m, \quad (\text{D.6})$$

where the sum is zero when $k = 1$. First, we have

$$\varphi_{c_k} = 0, \quad k \in \{1, \dots, k_* - 1\}. \quad (\text{D.7})$$

If $k_* = 1$, (D.7) is vacuous. Otherwise, for $k = 1$, $\varphi_{c_1} = 0$ by the first equation in (D.6) and the invertibility of $\lambda^* I - L_1$. For $2 \leq k < k_*$, note that if $\varphi_{c_\ell} = 0$ for all $\ell = 1, \dots, k-1$, then the k -th equation in (D.6) and the invertibility of $\lambda^* I - L_k$ imply that $\varphi_{c_k} = R(\lambda^*; L_k) \sum_{\ell=1}^{k-1} L_{k\ell} \varphi_{c_\ell} = 0$. Hence (D.7) holds by induction.

For $k = k_*$, (D.6) and (D.7) imply that $(\lambda^* I - L_{k_*})\varphi_{C_{k_*}} = 0$. Since L_{k_*} is irreducible, we can choose $\varphi_{C_{k_*}} \gg 0$ to be a strictly positive right eigenvector of L_{k_*} associated with λ^* . Since $C_{k_*} \rightsquigarrow C_{k_*}$, the right-eigenvector support formula in (4.4) holds for $\mathcal{C} = C_{k_*}$.

Next, we show, again by induction, that

$$\varphi_{C_k} \begin{cases} \gg 0, & \text{if } C_k \rightsquigarrow C_{k_*}, \\ = 0, & \text{otherwise,} \end{cases} \quad k \in \{k_* + 1, \dots, m\}. \quad (\text{D.8})$$

If $k_* = m$, (D.8) is vacuous. Otherwise, for $k = k_* + 1$, the $(k_* + 1)$ -th equation in (D.6), the invertibility of $\lambda^* I - L_{k_*+1}$, and (D.7) jointly imply that $\varphi_{C_{k_*+1}} = R(\lambda^*; L_{k_*+1}) L_{k_*+1, k_*} \varphi_{C_{k_*}}$. By (D.5) and $\varphi_{C_{k_*}} \gg 0$, it follows that $\varphi_{C_{k_*+1}} \gg 0$ if $L_{k_*+1, k_*} \neq 0$, which holds exactly when $C_{k_*+1} \rightarrow C_{k_*}$, or equivalently $C_{k_*+1} \rightsquigarrow C_{k_*}$; otherwise, $L_{k_*+1, k_*} = 0$ implies $\varphi_{C_{k_*+1}} = 0$. Therefore, (D.8) holds for $k = k_* + 1$.

Now, suppose that (D.8) holds for $k \in \{k_* + 1, \dots, k'\}$ with $k_* + 1 \leq k' < m$. The $(k' + 1)$ -th equation in (D.6) and (D.5) give $\varphi_{C_{k'+1}} = R(\lambda^*; L_{k'+1}) \sum_{\ell=k_*}^{k'} L_{k'+1, \ell} \varphi_{C_\ell}$ with $R(\lambda^*; L_{k'+1}) \gg 0$. Then, $\varphi_{C_{k'+1}} \gg 0$ if there exists some $\ell \in \{k_*, \dots, k'\}$ such that $L_{k'+1, \ell} \neq 0$ and $\varphi_{C_\ell} \gg 0$. The latter conditions mean precisely $C_{k'+1} \rightarrow C_\ell$ and $C_\ell \rightsquigarrow C_{k_*}$ for some $\ell \in \{k_*, \dots, k'\}$, which is equivalent to $C_{k'+1} \rightsquigarrow C_{k_*}$. If C_{k_*} is not accessible from $C_{k'+1}$, then $L_{k'+1, \ell} \varphi_{C_\ell} = 0$ for every $\ell \in \{k_*, \dots, k'\}$, so $\varphi_{C_{k'+1}} = 0$. Consequently, (D.8) holds for $k = k' + 1$, and the full result holds by induction.

Since the ordering of distinct classes rules out $C_k \rightsquigarrow C_{k_*}$ for $k < k_*$, the right-eigenvector support formula in (4.4) holds due to (D.7) and (D.8).

Similarly, let $\psi = (\psi_{C_1}^\top, \dots, \psi_{C_m}^\top)^\top$ be a left eigenvector of L_Π associated with λ^* such that $(\lambda^* I - L_\Pi^\top)\psi = 0$. Its block equations are

$$(\lambda^* I - L_k^\top)\psi_{C_k} = \sum_{\ell=k+1}^m L_{\ell k}^\top \psi_{C_\ell}, \quad k = 1, \dots, m, \quad (\text{D.9})$$

where the sum is zero when $k = m$. Starting from $k = m$ and proceeding backward, (D.9) and the invertibility in (D.5) give $\psi_{C_k} = 0$ for every $k > k_*$. Choose $\psi_{C_{k_*}} \gg 0$ to be a left eigenvector of L_{k_*} associated with λ^* . For $k < k_*$, (D.9) becomes

$$\psi_{C_k} = R(\lambda^*; L_k^\top) \sum_{\ell=k+1}^{k_*} L_{\ell k}^\top \psi_{C_\ell}.$$

Since $R(\lambda^*; L_k^\top) \gg 0$, backward induction shows that $\psi_{C_k} \gg 0$ exactly when there exists $\ell \in \{k + 1, \dots, k_*\}$ such that $C_{k_*} \rightsquigarrow C_\ell$ and $C_\ell \rightarrow C_k$, or equivalently, when $C_{k_*} \rightsquigarrow C_k$; otherwise, $\psi_{C_k} = 0$. The class ordering also rules out $C_{k_*} \rightsquigarrow C_k$ for $k > k_*$. This proves the left-eigenvector support formula and hence (4.4).

The vectors constructed above are nonnegative and strictly positive on $\mathcal{C}^*$, so $\mathbf{1}^\top \varphi > 0$ and $\psi^\top \varphi \geq \psi_{C_*}^\top \varphi_{C_*} > 0$. We may therefore normalize them by $\mathbf{1}^\top \varphi = 1$ and $\psi^\top \varphi = 1$. Since λ^* is algebraically simple, its right and left eigenspaces are one-dimensional, and hence these normalizations make φ and ψ unique. Furthermore, the Riesz spectral projection is rank one, and is given by $P_{\lambda^*} = \varphi \psi^\top$.

Moreover,

$$e^{-\lambda^* t} \Pi_t = e^{-\lambda^* t} e^{tL_\Pi} [P_{\lambda^*} + (I - P_{\lambda^*})] = P_{\lambda^*} + e^{-\lambda^* t} e^{tL_\Pi} (I - P_{\lambda^*}), \quad t \geq 0.$$

By (3.1), the restriction of L_Π to the range of $I - P_{\lambda^*}$ has spectral bound strictly below λ^* . The second term therefore converges to zero as $t \rightarrow \infty$, proving (4.3). $\square$

Proof of Theorem 4.2. Lemma B.3 writes every block of $R(z; L_\Pi)$ as a finite sum of class-path products. At $z = \lambda^*$, the pole order of such a product is the number of classes from $\mathcal{C}^*$ on that path: each dominant diagonal-block resolvent has a simple Perron pole, whereas every lower-rate block resolvent is analytic and strictly positive at λ^* by Lemma B.3, applied to the corresponding diagonal block. The leading coefficients are nonnegative and nonzero, so positivity rules out cancellation. The largest resolvent pole order is therefore ν^* .

For a finite matrix, the pole order at an eigenvalue equals the size of its largest Jordan block. Set $N^* := (L_\Pi - \lambda^* I)P^*$. This matrix agrees with $L_\Pi - \lambda^* I$ on the range of P^* and is zero on $\ker P^*$, so it is nilpotent of index ν^* . Thus

$$e^{t(L_\Pi - \lambda^* I)} P^* = \sum_{j=0}^{\nu^*-1} \frac{t^j}{j!} (N^*)^j P^*.$$

By (3.1), the restriction of L_Π to $\ker P^*$ has spectral bound strictly below λ^* . Its normalized semigroup is therefore bounded by $Ce^{-\eta t}$ for suitable $C, \eta > 0$. Therefore, there exist constants $C, \eta > 0$ such that

$$e^{-\lambda^* t} \Pi_t = \sum_{j=0}^{\nu^*-1} \frac{t^j}{j!} (N^*)^j P^* + R_t, \quad \|R_t\| \leq Ce^{-\eta t}. \quad (\text{D.10})$$

Multiplication by $t^{1-\nu^*}$ gives (4.6). The limit is nonnegative because $\Pi_t \geq 0$, and it is nonzero because $(N^*)^{\nu^*-1} P^* \neq 0$. Finally, $N^* B^* = B^* N^* = 0$, so $L_\Pi B^* = B^* L_\Pi = \lambda^* B^*$. Semisimplicity is equivalent to $\nu^* = 1$, yielding the final claim. $\square$

D.3. Doob transforms and the long-run pricing law. This subsection gives the two-transform construction underlying the Markov interpretation in Section 4.2. Maintain Assumptions 2.1–2.3 and suppose that $\mathcal{C}^* = \{\mathcal{C}^*\}$, with eigenvectors normalized as in Theorem 4.1.

Restricting the dominant eigenvectors to their supports permits Doob h -transforms of the pricing semigroup and its adjoint (see, e.g., Doob, 1957). Let φ_+ and ψ_+ be the restrictions of φ and ψ to their respective supports, and let $L_\Pi^{\varphi_+}$ and $L_\Pi^{\psi_+}$ be the principal submatrices of L_Π indexed by $\text{supp}(\varphi)$ and $\text{supp}(\psi)$, respectively. Define

$$L_\Pi^{\rightarrow} := \text{diag}(\varphi_+)^{-1} (L_\Pi^{\varphi_+} - \lambda^* I) \text{diag}(\varphi_+), \quad (\text{D.11})$$

$$L_\Pi^{\leftarrow} := \text{diag}(\psi_+)^{-1} [(L_\Pi^{\psi_+})^\top - \lambda^* I] \text{diag}(\psi_+). \quad (\text{D.12})$$

Both transforms remove the common exponential rate λ^* ; the transpose in (D.12) reverses the retained edges. The following proposition describes the resulting Markov dynamics.




FIGURE 6. Original and transformed reduced graphs in the disaster-dominant economy. Panels (a)–(c) show $\mathcal{G}_r(L_\Pi)$, $\mathcal{G}_r(L_\Pi^\rightarrow)$, and $\mathcal{G}_r(L_\Pi^\leftarrow)$, respectively. Light shading identifies the globally pricing-dominant class $\mathcal{D}$. Double borders identify the unique closed class: $\mathcal{N}$ in panel (a) and $\mathcal{D}$ in panels (b) and (c). The graph in panel (b) has one vertex and no edges.

Proposition D.2 (Doob chains). *Under Assumptions 2.1–2.3, suppose that $\mathcal{C}^* = \{\mathcal{C}^*\}$ and let φ and ψ be the right and left eigenvectors of L_Π associated with λ^* , normalized as in Theorem 4.1. Then $L_\Pi^\rightarrow$ and $L_\Pi^\leftarrow$, defined in (D.11) and (D.12), are Markov generators on $\text{supp}(\varphi)$ and $\text{supp}(\psi)$, respectively; we call the associated Markov chains the right Doob chain and left Doob chain. Moreover,*

- (i) *the reduced graph $\mathcal{G}_r(L_\Pi^\rightarrow)$ is the subgraph of $\mathcal{G}_r(L_\Pi)$ induced by the classes that can reach $\mathcal{C}^*$, whereas $\mathcal{G}_r(L_\Pi^\leftarrow)$ is the edge reversal of the subgraph of $\mathcal{G}_r(L_\Pi)$ induced by the classes accessible from $\mathcal{C}^*$;*
- (ii) *$\mathcal{C}^*$ is the unique closed, and hence recurrent, class of both Doob chains;*
- (iii) *the unique stationary distributions of $L_\Pi^\rightarrow$ and $L_\Pi^\leftarrow$, extended by zero outside their respective state spaces, both equal $\pi^{\text{LR}} = \varphi \odot \psi$, which is supported on $\mathcal{C}^*$.*

Proof of Proposition D.2. The identities $L_\Pi^{\varphi+} \varphi_+ = \lambda^* \varphi_+$ and $(L_\Pi^{\psi+})^\top \psi_+ = \lambda^* \psi_+$, together with (D.11) and (D.12), imply $L_\Pi^\rightarrow \mathbf{1} = 0$ and $L_\Pi^\leftarrow \mathbf{1} = 0$. Both restricted matrices $L_\Pi^{\varphi+}$ and $L_\Pi^{\psi+}$ are Metzler; positive diagonal similarity preserves their off-diagonal zero patterns, while transposition reverses those patterns. Therefore, $L_\Pi^\rightarrow$ and $L_\Pi^\leftarrow$ are Markov generators on $\text{supp}(\varphi)$ and $\text{supp}(\psi)$ with the stated reduced graphs. Both reduced graphs are acyclic, being respectively an induced subgraph of $\mathcal{G}_r(L_\Pi)$ and the edge reversal of one. In each graph, $\mathcal{C}^*$ is closed and accessible from every class, and hence is the unique closed class.

Set $\mathbf{X}_\varphi := \text{supp}(\varphi)$ and $\mathbf{X}_\psi := \text{supp}(\psi)$. No pricing edge enters $\mathbf{X}_\varphi$ from outside: otherwise, the class containing its initial vertex would have access to $\mathcal{C}^*$ and would belong to $\mathbf{X}_\varphi$. Similarly, no pricing edge leaves $\mathbf{X}_\psi$, since the class containing its terminal vertex would then be accessible from $\mathcal{C}^*$ and would belong to $\mathbf{X}_\psi$. Consequently, taking the $\mathbf{X}_\varphi$ coordinates of $\psi^\top L_\Pi = \lambda^* \psi^\top$ and the $\mathbf{X}_\psi$ coordinates of $L_\Pi \varphi = \lambda^* \varphi$ gives $\psi_{\mathbf{X}_\varphi}^\top L_\Pi^{\varphi+} = \lambda^* \psi_{\mathbf{X}_\varphi}^\top$, and $L_\Pi^{\psi+} \varphi_{\mathbf{X}_\psi} = \lambda^* \varphi_{\mathbf{X}_\psi}$. Let $D_\varphi := \text{diag}(\varphi_+)$ and $D_\psi := \text{diag}(\psi_+)$. Using (D.11) and (D.12), we obtain

$$\begin{aligned}
 (\varphi_+ \odot \psi_{\mathbf{X}_\varphi})^\top L_\Pi^\rightarrow &= \psi_{\mathbf{X}_\varphi}^\top (L_\Pi^{\varphi+} - \lambda^* I) D_\varphi = 0, \\
 (\varphi_{\mathbf{X}_\psi} \odot \psi_+)^\top L_\Pi^\leftarrow &= \varphi_{\mathbf{X}_\psi}^\top [(L_\Pi^{\psi+})^\top - \lambda^* I] D_\psi = 0.
 \end{aligned}$$

Thus these vectors are stationary for $L_\Pi^\rightarrow$ and $L_\Pi^\leftarrow$, respectively. Their zero extensions to $\mathbf{X}$ both equal $\varphi \odot \psi$, because φ vanishes outside $\mathbf{X}_\varphi$ and ψ vanishes outside $\mathbf{X}_\psi$. Moreover, $\mathbf{1}^\top (\varphi \odot \psi) =$

$\psi^\top \varphi = 1$, and $\text{supp}(\varphi) \cap \text{supp}(\psi) = \mathcal{C}^*$: any other common class would communicate with $\mathcal{C}^*$. Because $\mathcal{C}^*$ is the unique closed class of each transformed graph, each generator has a unique stationary distribution. It is therefore $\pi^{\text{LR}} = \varphi \odot \psi$. $\square$

For $x \in \text{supp}(\varphi)$, the right Doob chain admits a change-of-measure interpretation. The process

$$M_t^\varphi := e^{-\lambda^* t} S_t^G \frac{\varphi(X_t)}{\varphi(x)}, \quad t \geq 0,$$

is a nonnegative unit-mean martingale and therefore defines a family of probability measures $(\mathbb{P}_x^\varphi)_{x \in \text{supp}(\varphi)}$ by $d\mathbb{P}_x^\varphi / d\mathbb{P}_x|_{\mathcal{F}_t} = M_t^\varphi$. Under $\mathbb{P}_x^\varphi$, the state process is Markov on $\text{supp}(\varphi)$ with generator $L_\Pi^\rightarrow$. The left Doob chain comes from the adjoint construction and does not generally correspond to another forward change of physical measure. When the economy is irreducible, the right transform is the whole-state Hansen–Scheinkman transform and the left is its stationary time reversal; Appendix C reviews this benchmark.

By Proposition D.2(ii)–(iii) and finite-state Markov-chain convergence, the distributions of both Doob chains converge from every initial state in their respective state spaces to π^{LR} . Thus a physically transient class can be the unique recurrent class under both transformed dynamics.

Figure 6 illustrates these results for the three-class graph $\mathcal{D} \rightarrow \mathcal{R} \rightarrow \mathcal{N}$ with $\mathcal{C}^* = \mathcal{D}$. Then $\text{supp}(\varphi) = \mathcal{D}$ and $\text{supp}(\psi) = \mathbf{X}$: the global rate is absent from prices starting in recovery or normal states, but prices starting in disaster can inherit it even when the payoff is delivered only after recovery. The right Doob chain is restricted to $\mathcal{D}$, while the left has the reversed reduced graph $\mathcal{N} \rightarrow \mathcal{R} \rightarrow \mathcal{D}$. Both have limiting distributions concentrated on $\mathcal{D}$, although the physical dynamics eventually enter the recurrent normal class $\mathcal{N}$. If the uniquely dominant class moves downstream along this chain, the right eigenvector’s support expands and the left eigenvector’s support contracts.

D.4. Pricing-corridor restriction and asymptotics.

Proof of Proposition 5.1. Suppose first that $\mathcal{K}(\mathcal{C}, g) = \emptyset$. Then no state in $\mathcal{C}$ has access to a state in $\text{supp}(g)$. Lemma B.2, applied to L_Π , gives $q_t(x, g) = (e^{tL_\Pi} g)(x) = 0$ for every $x \in \mathcal{C}$ and $t > 0$. The same conclusion holds at $t = 0$, because g must vanish on $\mathcal{C}$; otherwise $\mathcal{C} \in \mathcal{S}(g)$ and the fact that $\mathcal{C} \rightsquigarrow \mathcal{C}$ would place $\mathcal{C}$ in the corridor.

Now suppose that $\mathcal{K} = \mathcal{K}(\mathcal{C}, g) \neq \emptyset$. Choose $\mathcal{C}' \in \mathcal{K}$. By definition, there is a class $\mathcal{C}'' \in \mathcal{S}(g)$ such that $\mathcal{C} \rightsquigarrow \mathcal{C}' \rightsquigarrow \mathcal{C}''$. Hence $\mathcal{C} \rightsquigarrow \mathcal{C}''$, which implies both $\mathcal{C} \in \mathcal{K}$ and $\mathcal{C}'' \in \mathcal{S}(g) \cap \mathcal{K}$. In particular, $g_{\mathcal{K}} \neq 0$. The principal submatrix $L_{\mathcal{K}}$ retains exactly the pricing edges whose endpoints lie in $\mathbf{X}_{\mathcal{K}}$. Because $\mathcal{K}$ is a union of whole classes, these classes remain the classes of $L_{\mathcal{K}}$. Moreover, every class on a path between two corridor classes also lies in the corridor: it is accessible from $\mathcal{C}$ and has access to $\mathcal{S}(g)$. Thus the reduced graph of $L_{\mathcal{K}}$ is the subgraph of the reduced pricing graph induced by $\mathcal{K}$, with the same accessibility relation among its vertices.

To prove the semigroup restriction, let z be any sufficiently large real number and apply Lemma B.3 to L_Π and $L_{\mathcal{K}}$. A class in $\mathcal{S}(g)$ that is accessible from $\mathcal{C}$ belongs to $\mathcal{K}$, and every class on every

path from $\mathcal{C}$ to that payoff-support class also belongs to $\mathcal{K}$. Hence the path contributions to the two scalar resolvents coincide, giving

$$e_x^\top R(z; L_\Pi)g = e_x^\top R(z; L_{\mathcal{K}})g_{\mathcal{K}}.$$

For $z > \max\{s(L_\Pi), s(L_{\mathcal{K}})\}$, the two sides of the above equation are the Laplace transforms of $(e^{tL_\Pi}g)(x)$ and $(e^{tL_{\mathcal{K}}}g_{\mathcal{K}})(x)$, respectively. Uniqueness of the Laplace transform, together with continuity in t , proves (5.2) for every $t \geq 0$. Lastly, $\mathcal{K}(\mathcal{C}, g) \neq \emptyset$ means that there exists $y \in \text{supp}(g)$ such that $x \rightsquigarrow y$ for some, and hence every, $x \in \mathcal{C}$. Lemma B.2 then implies $q_t(x, g) = (e^{tL_\Pi}g)(x) > 0$ for every $t > 0$ and every $x \in \mathcal{C}$. $\square$

Lemma D.3 (Positive leading corridor coefficient). *Let $g \in \mathbb{R}_+^X$ be nonzero, fix $x \in \mathcal{C}$, and suppose that $\mathcal{K} = \mathcal{K}(\mathcal{C}, g)$ is nonempty. With λ , ν , $P_{\mathcal{K}}$, and $B_{\mathcal{K}}$ as in (5.3)–(5.4), set $N_{\mathcal{K}} := (L_{\mathcal{K}} - \lambda I)P_{\mathcal{K}}$. Then*

$$e_x^\top N_{\mathcal{K}}^{\nu-1} P_{\mathcal{K}} g_{\mathcal{K}} > 0, \quad (B_{\mathcal{K}} g_{\mathcal{K}})(x) > 0. \quad (\text{D.13})$$

Proof. Apply Lemma B.3 to the corridor generator $L_{\mathcal{K}}$. By Proposition 5.1, its reduced graph is induced by $\mathcal{K}$, and

$$e_x^\top R(z; L_{\mathcal{K}})g_{\mathcal{K}} = \sum_{\mathcal{C}_\ell \in \mathcal{S}(g) \cap \mathcal{K}} \sum_{\gamma: \mathcal{C} \rightsquigarrow \mathcal{C}_\ell} e_x^\top \mathcal{R}_\gamma(z; L_{\mathcal{K}})g_{\mathcal{C}_\ell}. \quad (\text{D.14})$$

Call the paths in this sum *relevant*. Every relevant path lies in $\mathcal{K}$, so it contains at most ν corridor-critical classes. Conversely, let $\mathcal{C}_{\ell_1} \rightsquigarrow \dots \rightsquigarrow \mathcal{C}_{\ell_\nu}$ be the ordered corridor-critical classes in a chain attaining length ν . Since its endpoints lie in the corridor, $\mathcal{C} \rightsquigarrow \mathcal{C}_{\ell_1}$ and $\mathcal{C}_{\ell_\nu} \rightsquigarrow \mathcal{C}_h$ for some $\mathcal{C}_h \in \mathcal{S}(g)$. Concatenating paths that realize these relations produces a directed walk through all ν critical classes. The reduced graph is acyclic, so this walk is a relevant directed path. Thus at least one relevant path contains exactly ν critical classes.

Fix such a path $\gamma = (\mathcal{C}_{k_0} \rightarrow \dots \rightarrow \mathcal{C}_{k_p})$. For each position h , set

$$G_{\gamma, h} := \begin{cases} \varphi_{k_h} \psi_{k_h}^\top, & \lambda_{k_h} = \lambda, \\ R(\lambda; L_{k_h}), & \lambda_{k_h} < \lambda. \end{cases}$$

Substituting the block Laurent expansion (B.12) into (B.4) shows that the coefficient of $(z - \lambda)^{-\nu}$ contributed by γ to (D.14) is

$$e_x^\top G_{\gamma, 0} L_{k_0 k_1} G_{\gamma, 1} \dots L_{k_{p-1} k_p} G_{\gamma, p} g_{\mathcal{C}_{k_p}}. \quad (\text{D.15})$$

This scalar is strictly positive. Indeed, every noncritical resolvent is strictly positive by Lemma B.3, applied to the corresponding diagonal block; every block Perron vector is strictly positive; every coupling block along the path is nonzero and nonnegative; and $g_{\mathcal{C}_{k_p}}$ is nonzero and nonnegative. The same argument applies to every relevant path containing ν critical classes. All remaining paths contribute a pole of lower order or a function analytic at λ . Hence the coefficient of $(z - \lambda)^{-\nu}$ in (D.14) is strictly positive.

Theorem 4.1 in the unique-critical case and Theorem 4.2 in the tied case identify ν as the nilpotency index of $N_{\mathcal{K}}$. On its dominant spectral subspace,

$$R(z; L_{\mathcal{K}})P_{\mathcal{K}} = \sum_{j=0}^{\nu-1} \frac{N_{\mathcal{K}}^j P_{\mathcal{K}}}{(z - \lambda)^{j+1}},$$

while $R(z; L_{\mathcal{K}})(I - P_{\mathcal{K}})$ is analytic at λ . Therefore, the strictly positive coefficient just identified equals $e_x^\top N_{\mathcal{K}}^{\nu-1} P_{\mathcal{K}} g_{\mathcal{K}}$. Dividing by $(\nu - 1)!$ proves (D.13). $\square$

Proof of Theorem 5.2. Write $\mathcal{K} = \mathcal{K}(\mathcal{C}, g)$, $\lambda = \lambda(\mathcal{C}, g)$, and $\nu = \nu(\mathcal{C}, g)$. The corridor-critical classes are exactly the globally dominant classes of $L_{\mathcal{K}}$, and their maximum chain length is ν . Hence Theorem 4.2, applied to $L_{\mathcal{K}}$, gives

$$\lim_{t \rightarrow \infty} t^{1-\nu} e^{-\lambda t} e^{tL_{\mathcal{K}}} = \frac{N_{\mathcal{K}}^{\nu-1} P_{\mathcal{K}}}{(\nu - 1)!} = B_{\mathcal{K}}.$$

Combining this limit with the exact restriction (5.2) yields $t^{1-\nu} e^{-\lambda t} q_t(x, g) \rightarrow (B_{\mathcal{K}} g_{\mathcal{K}})(x) =: c_x(g)$. Lemma D.3 shows that $c_x(g) > 0$. Thus $q_t(x, g) = c_x(g) t^{\nu-1} e^{\lambda t} (1 + o(1))$. Proposition 5.1(ii) ensures that $q_t(x, g) > 0$ for $t > 0$, so taking logarithms is valid. Dividing by $-t$ gives (5.6). $\square$

The leading coefficient also admits an eigenvector representation. Write $\mathcal{K} = \mathcal{K}(\mathcal{C}, g)$ and $\lambda = \lambda(\mathcal{C}, g)$. Theorem 4.2 gives $B_{\mathcal{K}} \neq 0$ and $L_{\mathcal{K}} B_{\mathcal{K}} = B_{\mathcal{K}} L_{\mathcal{K}} = \lambda B_{\mathcal{K}}$. Therefore, if $d_{\mathcal{K}} = \text{rank}(B_{\mathcal{K}})$, then $\text{ran}(B_{\mathcal{K}}) \subseteq \ker(L_{\mathcal{K}} - \lambda I)$ and $\text{ran}(B_{\mathcal{K}}^\top) \subseteq \ker(L_{\mathcal{K}}^\top - \lambda I)$. Hence a rank factorization may be chosen in terms of right and left eigenvectors $\varphi_{\mathcal{K},a}$ and $\psi_{\mathcal{K},a}$ of $L_{\mathcal{K}}$ at λ :

$$B_{\mathcal{K}} = \sum_{a=1}^{d_{\mathcal{K}}} \varphi_{\mathcal{K},a} \psi_{\mathcal{K},a}^\top. \quad (\text{D.16})$$

Although the individual factors need not be unique, their sum $B_{\mathcal{K}}$ is the canonical leading spectral coefficient. In particular,

$$c_x(g) = \sum_{a=1}^{d_{\mathcal{K}}} \varphi_{\mathcal{K},a}(x) \psi_{\mathcal{K},a}^\top g_{\mathcal{K}}.$$

Proof of Corollary 5.3. If the corridor is empty, Proposition 5.1 makes the strip price, and hence its integral, identically zero. Otherwise, Theorem 5.2 gives $q_t(x, g) \sim c_x(g) t^{\nu(\mathcal{C},g)-1} e^{\lambda(\mathcal{C},g)t}$ for $c_x(g) > 0$. The integral over any bounded maturity interval is finite by continuity of the pricing semigroup, while the displayed tail is integrable exactly when $\lambda(\mathcal{C}, g) < 0$. This proves the equivalence. $\square$

D.5. One-way disaster-recovery economy.

Proof of Proposition 6.1. We first establish (6.3). Conformably with the block representation of L_Π in (6.1), write

$$\Pi_t = e^{tL_\Pi} = \begin{pmatrix} A_t & B_t \\ H_t & C_t \end{pmatrix}, \quad t \geq 0.$$

The pricing semigroup satisfies the evolution equation

$$\frac{d}{dt}\Pi_t = L_\Pi \Pi_t, \quad \Pi_0 = I.$$

Together with the block representation of L_Π in (6.1), this gives

$$\frac{d}{dt}A_t = L_{\mathcal{N}}A_t, \quad \frac{d}{dt}B_t = L_{\mathcal{N}}B_t, \quad \frac{d}{dt}H_t = L_{\mathcal{D}}H_t + L_{\mathcal{D}\mathcal{N}}A_t, \quad \frac{d}{dt}C_t = L_{\mathcal{D}}C_t + L_{\mathcal{D}\mathcal{N}}B_t,$$

with initial conditions $A_0 = I$, $B_0 = 0$, $H_0 = 0$ and $C_0 = I$. It follows immediately that $A_t = e^{tL_{\mathcal{N}}}$, $B_t = 0$ and $C_t = e^{tL_{\mathcal{D}}}$. Applying the variation-of-constants formula (Duhamel's formula) to the equation for H_t yields $H_t = e^{tL_{\mathcal{D}}}H_0 + \int_0^t e^{(t-v)L_{\mathcal{D}}}L_{\mathcal{D}\mathcal{N}}A_v dv$. Substituting $H_0 = 0$ and $A_v = e^{vL_{\mathcal{N}}}$, then changing variables to $s = t - v$, gives

$$H_t = \int_0^t e^{sL_{\mathcal{D}}}L_{\mathcal{D}\mathcal{N}}e^{(t-s)L_{\mathcal{N}}} ds,$$

with the order of the matrix factors unchanged. This proves (6.3).

For part (i), fix $x \in \mathcal{N}$. If $g_{\mathcal{N}} = 0$, the block representation gives $q_t(x, g) = 0$ for every $t \geq 0$. If $g_{\mathcal{N}} \neq 0$, the pricing corridor is $\{\mathcal{N}\}$ and has rate $\lambda_{\mathcal{N}}$. Theorem 5.2 therefore gives $\lim_{t \rightarrow \infty} y_t(x, g) = -\lambda_{\mathcal{N}}$.

For part (ii), fix $x \in \mathcal{D}$. If $g_{\mathcal{N}} = 0$, then $g_{\mathcal{D}} \neq 0$ because g is nonzero, and the pricing corridor is $\{\mathcal{D}\}$. If $g_{\mathcal{N}} \neq 0$, the nonzero recovery block and irreducibility of the two diagonal blocks give the corridor $\{\mathcal{D}, \mathcal{N}\}$. The corresponding corridor rates are $\lambda_{\mathcal{D}}$ and $\max\{\lambda_{\mathcal{D}}, \lambda_{\mathcal{N}}\}$, respectively. The yield expansion in Theorem 5.2 gives the negatives of these rates as the limiting yields, including when $\lambda_{\mathcal{D}} = \lambda_{\mathcal{N}}$, since the $\log t/t$ correction vanishes as $t \rightarrow \infty$. $\square$

Proof of Corollary 6.2. Under (6.6), $\mathcal{D}$ is the unique globally dominant class. Only $\mathcal{D}$ can reach $\mathcal{D}$, whereas both $\mathcal{D}$ and $\mathcal{N}$ are accessible from $\mathcal{D}$. Theorem 4.1 therefore implies that $\lambda_{\mathcal{D}}$ is algebraically simple for L_Π , that the dominant right eigenvector is supported on $\mathcal{D}$, that the dominant left eigenvector has full support, and that the normalized semigroup converges to the associated rank-one Riesz projection.

With the state ordering in (6.1), the right eigenvector is $\Phi_{\mathcal{D}} = (0, \varphi_{\mathcal{D}}^\top)^\top$. Writing the left eigenvector as $(\xi^\top, \psi_{\mathcal{D}}^\top)^\top$, the normal-class block of $L_\Pi^\top(\xi^\top, \psi_{\mathcal{D}}^\top)^\top = \lambda_{\mathcal{D}}(\xi^\top, \psi_{\mathcal{D}}^\top)^\top$ gives $(\lambda_{\mathcal{D}}I - L_{\mathcal{N}}^\top)\xi = L_{\mathcal{D}\mathcal{N}}^\top\psi_{\mathcal{D}}$. Since $\lambda_{\mathcal{D}} > s(L_{\mathcal{N}})$, the matrix on the left is invertible and $\xi = R(\lambda_{\mathcal{D}}; L_{\mathcal{N}}^\top)L_{\mathcal{D}\mathcal{N}}^\top\psi_{\mathcal{D}}$. Thus the full-state eigenvectors are

$$\Phi_{\mathcal{D}} := \begin{pmatrix} 0 \\ \varphi_{\mathcal{D}} \end{pmatrix}, \quad \Psi_{\mathcal{D}} := \begin{pmatrix} R(\lambda_{\mathcal{D}}; L_{\mathcal{N}}^\top)L_{\mathcal{D}\mathcal{N}}^\top\psi_{\mathcal{D}} \\ \psi_{\mathcal{D}} \end{pmatrix} \gg 0. \quad (\text{D.17})$$

The block normalizations $\mathbf{1}^\top \varphi_{\mathcal{D}} = 1$ and $\psi_{\mathcal{D}}^\top \varphi_{\mathcal{D}} = 1$ give $\Psi_{\mathcal{D}}^\top \Phi_{\mathcal{D}} = 1$. Hence the Riesz projection is $P_{\mathcal{D}} := \Phi_{\mathcal{D}} \Psi_{\mathcal{D}}^\top$, and (4.3) gives

$$\lim_{t \rightarrow \infty} e^{-\lambda_{\mathcal{D}} t} \Pi_t = P_{\mathcal{D}}. \quad (\text{D.18})$$

Finally,

$$\Psi_{\mathcal{D}}^\top g = \psi_{\mathcal{D}}^\top g_{\mathcal{D}} + \psi_{\mathcal{D}}^\top L_{\mathcal{D}\mathcal{N}} R(\lambda_{\mathcal{D}}; L_{\mathcal{N}}) g_{\mathcal{N}},$$

where we used $R(\lambda_{\mathcal{D}}; L_{\mathcal{N}}^{\top})^{\top} = R(\lambda_{\mathcal{D}}; L_{\mathcal{N}})$. Since $\lambda_{\mathcal{D}} > s(L_{\mathcal{N}})$, the Laplace representation

$$R(\lambda_{\mathcal{D}}; L_{\mathcal{N}}) = \int_0^{\infty} e^{-\lambda_{\mathcal{D}} v} e^{v L_{\mathcal{N}}} dv$$

is convergent. Substituting this identity and applying $P_{\mathcal{D}}$ to g proves (6.7). Its coefficient is strictly positive for $x \in \mathcal{D}$, since $\varphi_{\mathcal{D}}(x) > 0$, $\Psi_{\mathcal{D}} \gg 0$, and $g \geq 0$ is nonzero.

For $x \in \mathcal{N}$, the block representation (6.3) gives $q_t(x, g) = (e^{t L_{\mathcal{N}}} g_{\mathcal{N}})(x)$. Since $L_{\mathcal{N}}$ is irreducible Metzler, $\lim_{t \rightarrow \infty} e^{-\lambda_{\mathcal{N}} t} e^{t L_{\mathcal{N}}} = \varphi_{\mathcal{N}} \psi_{\mathcal{N}}^{\top}$. The limit is positive when $g_{\mathcal{N}} \neq 0$, because $\varphi_{\mathcal{N}}, \psi_{\mathcal{N}} \gg 0$ and $g_{\mathcal{N}} \geq 0$. This proves part (i) of the corollary. $\square$

D.6. Rare-edge perturbations and crossover asymptotics. For each $\varepsilon > 0$ sufficiently small, let $\lambda^*(\varepsilon) := s(L_{\Pi}(\varepsilon))$ and let $\varphi(\varepsilon), \psi(\varepsilon) \gg 0$ be corresponding right and left eigenvectors of $L_{\Pi}(\varepsilon)$ associated with $\lambda^*(\varepsilon)$. For each limiting class $\mathcal{C} \in \mathcal{C}$, write $\varphi_{\mathcal{C}}(\varepsilon)$ and $\psi_{\mathcal{C}}(\varepsilon)$ for the restrictions of these vectors to $\mathcal{C}$. Normalize the vectors by $\mathbf{1}^{\top} \varphi_{\mathcal{C}^*}(\varepsilon) = 1$, and $\varphi(\varepsilon)^{\top} \psi(\varepsilon) = 1$.

In what follows, Lemma D.4 establishes the convergence of $\lambda^*(\varepsilon)$, $\varphi_{\mathcal{C}^*}(\varepsilon)$ and $\psi_{\mathcal{C}^*}(\varepsilon)$ as $\varepsilon \downarrow 0$, and Lemma D.5 establishes the rarity orders of the $\varphi_{\mathcal{C}}(\varepsilon)$ and $\psi_{\mathcal{C}}(\varepsilon)$ for $\mathcal{C} \neq \mathcal{C}^*$. Then, we use these two lemmas to prove Theorem 7.2 in Section 7.1.

Lemma D.4 (Eigenvalue and dominant-block eigenvector limits). *Assume the unique global pricing dominance condition (7.1). Let $\varphi_*, \psi_* \gg 0$ be the right and left Perron eigenvectors of L_{k^*} associated with λ^* , normalized by $\mathbf{1}^{\top} \varphi_* = 1$ and $\varphi_*^{\top} \psi_* = 1$. Then*

$$\lambda^*(\varepsilon) \rightarrow \lambda^*, \quad \varphi_{\mathcal{C}^*}(\varepsilon) \rightarrow \varphi_*, \quad \psi_{\mathcal{C}^*}(\varepsilon) \rightarrow \psi_*, \quad \text{as } \varepsilon \downarrow 0.$$

Proof of Lemma D.4. Since $L_{\Pi}(\varepsilon) \rightarrow L_{\Pi}$ and the spectral bound is continuous on finite-dimensional matrix spaces, $\lambda^*(\varepsilon) = s(L_{\Pi}(\varepsilon)) \rightarrow s(L_{\Pi}) = \lambda^*$ as $\varepsilon \downarrow 0$. The irreducibility of L_{k^*} implies that $\lambda^* = s(L_{k^*})$ is an algebraically simple eigenvalue of L_{k^*} . By (7.1), λ^* does not belong to the spectrum of any other diagonal block in the Frobenius form of L_{Π} . Therefore, λ^* is an algebraically simple eigenvalue of L_{Π} .

Let P_* be the Riesz spectral projection of L_{Π} associated with λ^* . Finite-dimensional perturbation theory implies that, for all sufficiently small $\varepsilon > 0$, the corresponding rank-one projection $P_*(\varepsilon)$ is well defined and $P_*(\varepsilon) \rightarrow P_*$; see Kato (1995, Chapter II, § 5.1, especially Theorem 5.1 and equation (5.2), pp. 107–108). Hence the one-dimensional eigenspaces $\text{ran } P_*(\varepsilon)$ converge to $\text{ran } P_*$. Let φ be the right eigenvector of L_{Π} spanning $\text{ran } P_*$ and normalized by $\mathbf{1}^{\top} \varphi_{\mathcal{C}^*} = 1$. The block eigenvector equation on $\mathcal{C}^*$ gives $L_{k^*} \varphi_{\mathcal{C}^*} = \lambda^* \varphi_{\mathcal{C}^*}$. By the uniqueness of the normalized Perron eigenvector of L_{k^*} , it follows that $\varphi_{\mathcal{C}^*} = \varphi_*$. Since the normalizing functional $v \mapsto \mathbf{1}^{\top} v_{\mathcal{C}^*}$ is nonzero on $\text{ran } P_*$, convergence of the one-dimensional right eigenspaces, together with the common normalization $\mathbf{1}^{\top} \varphi_{\mathcal{C}^*}(\varepsilon) = 1$, yields $\varphi(\varepsilon) \rightarrow \varphi$. Restricting this convergence to $\mathcal{C}^*$ gives $\varphi_{\mathcal{C}^*}(\varepsilon) \rightarrow \varphi_{\mathcal{C}^*} = \varphi_*$ as $\varepsilon \downarrow 0$.

It remains to identify the limit of the left eigenvectors under the normalization $\varphi(\varepsilon)^{\top} \psi(\varepsilon) = 1$. Since $\lambda^*(\varepsilon)$ is algebraically simple, $P_*(\varepsilon) = \varphi(\varepsilon) \psi(\varepsilon)^{\top}$. Fix any $x \in \mathcal{C}^*$. The convergence $\varphi_{\mathcal{C}^*}(\varepsilon) \rightarrow \varphi_*$

implies $\varphi_{c^*}(\varepsilon)(x) \rightarrow \varphi_*(x)$. Since $\varphi_*(x) > 0$, $\varphi_{c^*}(\varepsilon)(x)$ is bounded away from zero for all sufficiently small $\varepsilon > 0$. Hence, as $\varepsilon \downarrow 0$,

$$\psi(\varepsilon)^\top = \frac{e_x^\top P_*(\varepsilon)}{\varphi_{c^*}(\varepsilon)(x)} \rightarrow \frac{e_x^\top P_*}{\varphi_*(x)} =: \psi^\top.$$

Since λ^* is an algebraically simple eigenvalue of L_Π , the generalized left eigenspace coincides with the left eigenspace: $\text{ran}(P_*^\top) = \ker(L_\Pi^\top - \lambda^* I)$. Therefore, $\psi = P_*^\top e_x / \varphi_*(x) \in \ker(L_\Pi^\top - \lambda^* I)$, which implies that ψ is a left eigenvector of L_Π associated with λ^* . Moreover, $\varphi \in \text{ran}(P_*)$ implies $P_* \varphi = \varphi$. Since $x \in \mathcal{C}^*$ and $\varphi_{c^*} = \varphi_*$, it follows that $\psi^\top \varphi = e_x^\top P_* \varphi / \varphi_*(x) = e_x^\top \varphi / \varphi_*(x) = 1$. The block left-eigenvector equations imply that ψ_{c^*} is a strictly positive left Perron eigenvector of L_{k^*} . Moreover, the support characterization in Theorem 4.1 shows that φ and ψ have overlapping class support only on $\mathcal{C}^*$. Consequently, $1 = \varphi^\top \psi = \varphi_*^\top \psi_{c^*}$. The normalization and uniqueness of the left Perron vector of L_{k^*} thus give $\psi_{c^*} = \psi_*$. Restricting $\psi(\varepsilon) \rightarrow \psi$ to $\mathcal{C}^*$ completes the proof. $\square$

Lemma D.5 (Rarity orders on nondominant eigenvector blocks). *Under condition (7.1) and Assumption 7.1, for every $\mathcal{C} \in \mathcal{C} \setminus \{\mathcal{C}^*\}$, there exist $\bar{\varphi}_\mathcal{C}, \bar{\psi}_\mathcal{C} \gg 0$ such that*

$$\varphi_\mathcal{C}(\varepsilon) = \varepsilon^{d_{c\mathcal{C}^*}} (\bar{\varphi}_\mathcal{C} + o(1)), \quad \psi_\mathcal{C}(\varepsilon) = \varepsilon^{d_{c^*\mathcal{C}}} (\bar{\psi}_\mathcal{C} + o(1)), \quad \text{as } \varepsilon \downarrow 0, \quad (\text{D.19})$$

where $d_{c\mathcal{C}^*}, d_{c^*\mathcal{C}} \geq 0$ are the exponents defined in (7.2).

Proof of Lemma D.5. Let $\mathcal{C}^\circ := \mathcal{C} \setminus \{\mathcal{C}^*\}$ and $\mathbf{X}^\circ = \mathbf{X} \setminus \mathcal{C}^*$. After a fixed simultaneous permutation of block rows and columns, place the block indexed by $\mathcal{C}^*$ first and write

$$L_\Pi(\varepsilon) = \begin{pmatrix} L_{k^*k^*}(\varepsilon) & c(\varepsilon) \\ b(\varepsilon) & M(\varepsilon) \end{pmatrix}, \quad \varphi(\varepsilon) = \begin{pmatrix} \varphi_{c^*}(\varepsilon) \\ \varphi_{\mathbf{X}^\circ}(\varepsilon) \end{pmatrix}.$$

Here, $M(\varepsilon)$ is the principal submatrix indexed by $\mathbf{X}^\circ$ and the block column $b(\varepsilon)$ has entries $L_{kk^*}(\varepsilon)$ for $\mathcal{C}_k \in \mathcal{C}^\circ$. The lower block of the eigenvector equation is

$$(\lambda^*(\varepsilon)I - M(\varepsilon)) \varphi_{\mathbf{X}^\circ}(\varepsilon) = b(\varepsilon) \varphi_{c^*}(\varepsilon). \quad (\text{D.20})$$

Let $M_0(\varepsilon)$ retain the diagonal blocks of $M(\varepsilon)$ and those off-diagonal blocks whose edge-rarity exponent is zero, and set

$$E(\varepsilon) := M(\varepsilon) - M_0(\varepsilon).$$

Thus $E(\varepsilon)$ has zero diagonal blocks and retains exactly the positive finite-weight interclass blocks of $M(\varepsilon)$. It represents the rare edges among the classes in $\mathcal{C}^\circ$, whereas $M_0(\varepsilon)$ contains their within-class dynamics and non-rare interclass edges. In particular, $M_0(0) = M(0)$.

By (7.1), $s(M(0)) = \max_{\mathcal{C}_k \in \mathcal{C}^\circ} s(L_k) < \lambda^*$. Lemma D.4 gives $\lambda^*(\varepsilon) \rightarrow \lambda^*$, while $M_0(\varepsilon) \rightarrow M(0)$ and continuity of the spectral bound give $s(M_0(\varepsilon)) \rightarrow s(M(0))$. Therefore, $\lambda^*(\varepsilon) > s(M_0(\varepsilon))$ for all sufficiently small $\varepsilon > 0$. Hence the resolvent of $M_0(\varepsilon)$ at $\lambda^*(\varepsilon)$ admits the Laplace representation:

$$R_0(\varepsilon) := R(\lambda^*(\varepsilon); M_0(\varepsilon)) = \int_0^\infty e^{-\lambda^*(\varepsilon)t} e^{tM_0(\varepsilon)} dt \geq 0, \quad \varepsilon > 0 \text{ sufficiently small.}$$

Moreover, $R_0(\varepsilon) \rightarrow R(\lambda^*; M(0)) =: R_0(0)$ as $\varepsilon \downarrow 0$, so the family $\{R_0(\varepsilon)\}$ is uniformly bounded.

For all sufficiently small $\varepsilon \geq 0$, the positive within-class edges and the zero-weight interclass blocks of the limiting matrix persist. Hence the diagonal blocks of $M_0(\varepsilon)$ remain irreducible, and its class-level graph is exactly the zero-weight subgraph induced by $\mathcal{C}^\circ$. After a simultaneous permutation of these classes into Frobenius order, Lemma B.3 implies that the $(\mathcal{C}, \mathcal{C}')$ block of $R_0(\varepsilon)$ is strictly positive if $\mathcal{C} \rightsquigarrow_0 \mathcal{C}'$ and is zero otherwise.

Let $w_m = \min\{w_{h\ell} : h \neq \ell, 0 < w_{h\ell} < \infty\} > 0$. Such a positive minimum exists because $\mathcal{G}_w$ is strongly connected whereas its zero-weight subgraph is reducible. Since $\|E(\varepsilon)\| = O(\varepsilon^{w_m})$, the resolvent identity gives the uniformly convergent Neumann expansion

$$R(\lambda^*(\varepsilon); M(\varepsilon)) = R_0(\varepsilon) \sum_{n=0}^{\infty} (E(\varepsilon)R_0(\varepsilon))^n. \quad (\text{D.21})$$

Indeed, uniform boundedness of $R_0(\varepsilon)$ gives $\|E(\varepsilon)R_0(\varepsilon)\| = O(\varepsilon^{w_m}) < 1$ for all sufficiently small ε .

Combining (D.20) and (D.21) gives

$$\varphi_{x^\circ}(\varepsilon) = \sum_{n=0}^{\infty} R_0(\varepsilon) (E(\varepsilon)R_0(\varepsilon))^n b(\varepsilon) \varphi_{c^*}(\varepsilon). \quad (\text{D.22})$$

Because the series converges in norm, each block may be expanded term by term over the finitely many intermediate class indices. For $k \neq \ell$, write $w_{c_k c_\ell} := w_{k\ell}$. Fix $\mathcal{C} \in \mathcal{C}^\circ$. For $n \geq 0$, let $\mathcal{A}_{\mathcal{C}, n}$ consist of the tuples $\iota = (\mathcal{C}_0^+, \mathcal{C}_1^-, \mathcal{C}_1^+, \dots, \mathcal{C}_n^-, \mathcal{C}_n^+)$ in $(\mathcal{C}^\circ)^{2n+1}$ for which

$$\mathcal{C} \rightsquigarrow_0 \mathcal{C}_0^+, \quad \mathcal{C}_j^- \neq \mathcal{C}_{j-1}^+, \quad 0 < w_{\mathcal{C}_{j-1}^+ \mathcal{C}_j^-} < \infty, \quad \mathcal{C}_j^- \rightsquigarrow_0 \mathcal{C}_j^+, \quad j = 1, \dots, n, \quad w_{\mathcal{C}_n^+ \mathcal{C}^*} < \infty.$$

where $\rightsquigarrow_0$ denotes accessibility in the zero-weight subgraph induced by $\mathcal{C}^\circ$, including a path of length zero. Here $\mathcal{C}_j^-$ is the class reached by the j -th rare edge, and $\mathcal{C}_j^+$ is the class reached after the following zero-weight path segment. For $n \geq 1$, ι represents a walk of the form

$$\mathcal{C} \rightsquigarrow_0 \mathcal{C}_0^+ \xrightarrow{w>0} \mathcal{C}_1^- \rightsquigarrow_0 \mathcal{C}_1^+ \xrightarrow{w>0} \dots \xrightarrow{w>0} \mathcal{C}_n^- \rightsquigarrow_0 \mathcal{C}_n^+ \xrightarrow{w\geq 0} \mathcal{C}^*,$$

where the labeled arrows are single edges. For $n = 0$, the walk is $\mathcal{C} \rightsquigarrow_0 \mathcal{C}_0^+ \xrightarrow{w\geq 0} \mathcal{C}^*$. The block products in (D.22) are precisely

$$T_\iota(\varepsilon) := [R_0(\varepsilon)]_{\mathcal{C}\mathcal{C}_0^+} \prod_{j=1}^n \left([E(\varepsilon)]_{\mathcal{C}_{j-1}^+ \mathcal{C}_j^-} [R_0(\varepsilon)]_{\mathcal{C}_j^- \mathcal{C}_j^+} \right) b_{\mathcal{C}_n^+}(\varepsilon) \varphi_{\mathcal{C}^*}(\varepsilon), \quad (\text{D.23})$$

with the product ordered from $j = 1$ to $j = n$ and interpreted as the identity when $n = 0$. Thus

$$\varphi_{\mathcal{C}}(\varepsilon) = \sum_{n=0}^{\infty} \sum_{\iota \in \mathcal{A}_{\mathcal{C}, n}} T_\iota(\varepsilon). \quad (\text{D.24})$$

Associate with $\iota \in \mathcal{A}_{\mathcal{C}, n}$ the total edge weight

$$w(\iota) := \sum_{j=1}^n w_{\mathcal{C}_{j-1}^+ \mathcal{C}_j^-} + w_{\mathcal{C}_n^+ \mathcal{C}^*}. \quad (\text{D.25})$$

Assumption 7.1, Lemma D.4, and $R_0(\varepsilon) \rightarrow R_0(0)$ imply, for each fixed ι ,

$$T_\iota(\varepsilon) = \varepsilon^{w(\iota)} (H_\iota + o(1)), \quad H_\iota \gg 0. \quad (\text{D.26})$$

Reading the limiting product from right to left proves $H_\iota \gg 0$: a nonzero nonnegative edge matrix sends a strictly positive vector to a nonzero nonnegative vector, and the next strictly positive resolvent block restores strict positivity. Each ι also encodes at least one directed walk from $\mathcal{C}$ to $\mathcal{C}^*$ with weight $w(\iota)$; deleting cycles produces a path with no greater weight. Conversely, separate the final edge into $\mathcal{C}^*$ and decompose the preceding part of any directed path into maximal zero-weight segments and intervening positive-weight edges. This produces an element of some $\mathcal{A}_{\mathcal{C},n}$ with the same weight. Therefore,

$$\min_{\iota \in \bigcup_{n \geq 0} \mathcal{A}_{\mathcal{C},n}} w(\iota) = \min_{\gamma \in \mathcal{D}_w(\mathcal{C}, \mathcal{C}^*)} W(\gamma) = d_{\mathcal{C}\mathcal{C}^*}. \quad (\text{D.27})$$

Choose $N_{\mathcal{C}} \in \mathbb{N}_0$ such that $(N_{\mathcal{C}} + 1)w_{\text{m}} > d_{\mathcal{C}\mathcal{C}^*}$. Uniform boundedness of $R_0(\varepsilon)$, $b(\varepsilon)$, and $\varphi_{\mathcal{C}^*}(\varepsilon)$, together with $\|E(\varepsilon)\| = O(\varepsilon^{w_{\text{m}}})$, gives

$$\sum_{n=N_{\mathcal{C}}+1}^{\infty} \left\| [R_0(\varepsilon)(E(\varepsilon)R_0(\varepsilon))^n b(\varepsilon)\varphi_{\mathcal{C}^*}(\varepsilon)]_{\mathcal{C}} \right\| \leq C_0 \sum_{n=N_{\mathcal{C}}+1}^{\infty} (C_1 \varepsilon^{w_{\text{m}}})^n = O(\varepsilon^{(N_{\mathcal{C}}+1)w_{\text{m}}}) = o(\varepsilon^{d_{\mathcal{C}\mathcal{C}^*}}). \quad (\text{D.28})$$

Since $w(\iota) \geq nw_{\text{m}}$ for $\iota \in \mathcal{A}_{\mathcal{C},n}$, every minimizer in (D.27) occurs with $n \leq N_{\mathcal{C}}$. Define

$$\mathcal{A}_{\mathcal{C}}^{\min} := \left\{ \iota \in \bigcup_{n=0}^{N_{\mathcal{C}}} \mathcal{A}_{\mathcal{C},n} : w(\iota) = d_{\mathcal{C}\mathcal{C}^*} \right\}, \quad \bar{\varphi}_{\mathcal{C}} := \sum_{\iota \in \mathcal{A}_{\mathcal{C}}^{\min}} H_{\iota}.$$

This set is nonempty because the minimum over directed paths is attained, and (D.26) gives $\bar{\varphi}_{\mathcal{C}} \gg 0$. The part of (D.24) with $n \leq N_{\mathcal{C}}$ is finite: after division by $\varepsilon^{d_{\mathcal{C}\mathcal{C}^*}}$, its minimum-weight terms converge to $\bar{\varphi}_{\mathcal{C}}$ and all other terms vanish. Combining this fact with (D.28) yields $\varepsilon^{-d_{\mathcal{C}\mathcal{C}^*}} \varphi_{\mathcal{C}}(\varepsilon) \rightarrow \bar{\varphi}_{\mathcal{C}} \gg 0$, proving the first expansion in (D.19).

For the left eigenvector, apply the same path-expansion argument to $L_{\Pi}(\varepsilon)^{\top}$. Transposition preserves the limiting classes and their spectral bounds while reversing every edge without changing its weight. Thus the relevant minimum weight is $d_{\mathcal{C}^*\mathcal{C}}$ in the original graph. Lemma D.4 also gives $\psi_{\mathcal{C}^*}(\varepsilon) \rightarrow \psi_* \gg 0$, the only dominant-block convergence used above. Consequently, there exists $\bar{\psi}_{\mathcal{C}} \gg 0$ such that $\varepsilon^{-d_{\mathcal{C}^*\mathcal{C}}} \psi_{\mathcal{C}}(\varepsilon) \rightarrow \bar{\psi}_{\mathcal{C}}$, proving the second expansion in (D.19). $\square$

Proof of Theorem 7.2. Lemma D.4 gives $\lambda^*(\varepsilon) \rightarrow \lambda^*$ as $\varepsilon \downarrow 0$. For the dominant class, set $\bar{\varphi}_{\mathcal{C}^*} := \varphi_*$ and $\bar{\psi}_{\mathcal{C}^*} := \psi_*$, where φ_* and ψ_* are defined in Lemma D.4. Then, Lemma D.4, together with $d_{\mathcal{C}^*\mathcal{C}^*} = 0$, and Lemma D.5 imply that for each $\mathcal{C} \in \mathcal{C}$, there are vectors $\bar{\varphi}_{\mathcal{C}}, \bar{\psi}_{\mathcal{C}} \gg 0$ such that

$$\varphi_{\mathcal{C}}(\varepsilon) = \varepsilon^{d_{\mathcal{C}\mathcal{C}^*}} (\bar{\varphi}_{\mathcal{C}} + o(1)), \quad \psi_{\mathcal{C}}(\varepsilon) = \varepsilon^{d_{\mathcal{C}^*\mathcal{C}}} (\bar{\psi}_{\mathcal{C}} + o(1)), \quad \text{as } \varepsilon \downarrow 0. \quad (\text{D.29})$$

The normalization imposed at the beginning of this subsection gives $P_*(\varepsilon) = \varphi(\varepsilon)\psi(\varepsilon)^{\top}$ and $P_{*,\mathcal{C}\mathcal{C}'}(\varepsilon) = \varphi_{\mathcal{C}}(\varepsilon)\psi_{\mathcal{C}'}(\varepsilon)^{\top}$. Therefore, defining $\bar{P}_{\mathcal{C}\mathcal{C}'} := \bar{\varphi}_{\mathcal{C}}\bar{\psi}_{\mathcal{C}'}^{\top} \gg 0$, equation (D.29) yields

$$\varepsilon^{-(d_{\mathcal{C}\mathcal{C}^*} + d_{\mathcal{C}^*\mathcal{C}'})} P_{*,\mathcal{C}\mathcal{C}'}(\varepsilon) \rightarrow \bar{P}_{\mathcal{C}\mathcal{C}'}, \quad \text{as } \varepsilon \downarrow 0,$$

which proves (7.5). By the definition $\pi_{\varepsilon}^{\text{LR}}(x) = P_*(\varepsilon)(x, x)$, $\sum_{x \in \mathcal{C}} \pi_{\varepsilon}^{\text{LR}}(x) = \text{tr } P_{*,\mathcal{C}\mathcal{C}}(\varepsilon)$. Taking $\mathcal{C}' = \mathcal{C}$ in (7.5) and taking traces then gives (7.6).

Finally, if $\mathcal{C} \neq \mathcal{C}^*$, then $d_{\mathcal{C}\mathcal{C}^*} + d_{\mathcal{C}^*\mathcal{C}} > 0$. Otherwise, nonnegativity of the edge weights would give zero-weight paths in both directions between two distinct limiting classes, which yields a contradiction. Equation (7.6) therefore implies $\sum_{x \in \mathcal{C}} \pi_\varepsilon^{\text{LR}}(x) \rightarrow 0$ for every $\mathcal{C} \neq \mathcal{C}^*$. Since $P_*(\varepsilon)$ is a rank-one projection, $\sum_{x \in \mathcal{X}} \pi_\varepsilon^{\text{LR}}(x) = \text{tr } P_*(\varepsilon) = 1$. Hence $\sum_{x \in \mathcal{C}^*} \pi_\varepsilon^{\text{LR}}(x) = 1 - \sum_{x \in \mathcal{X} \setminus \mathcal{C}^*} \pi_\varepsilon^{\text{LR}}(x) \rightarrow 1$, as $\varepsilon \downarrow 0$. $\square$

We summarize the projection notation introduced in Sections 7.1 and 7.2 that will be used below. For every sufficiently small $\varepsilon > 0$, $P_*(\varepsilon)$ is the rank-one Riesz projection of $L_\Pi(\varepsilon)$ associated with its Perron eigenvalue $\lambda^*(\varepsilon)$, whereas $P_o(\varepsilon)$ is the Riesz projection of $L_\Pi(\varepsilon)$ associated with the spectral set $\Sigma_o(\varepsilon)$ defined in (7.8). At $\varepsilon = 0$, these projections satisfy $P_*(0) = P_*$ and $P_o(0) = P_o$, where P_* and P_o are Riesz projections of L_Π associated with its eigenvalue λ^* and spectral set Σ_o , respectively. Moreover, $s_o = s(L_\Pi|_{\ker P_*})$.

Lemma D.6. *Under Assumption 7.1 and the unique-dominant-class condition (7.1), there exist constants $\varepsilon_0, C, \delta > 0$ such that, for every $\varepsilon \in (0, \varepsilon_0]$ and $t \geq 0$, setting $\mathcal{E}_\varepsilon(t) := e^{tL_\Pi(\varepsilon)}(I - P_*(\varepsilon) - P_o(\varepsilon))$, we have*

$$e^{tL_\Pi(\varepsilon)} = e^{\lambda^*(\varepsilon)t} P_*(\varepsilon) + e^{tL_\Pi(\varepsilon)} P_o(\varepsilon) + \mathcal{E}_\varepsilon(t), \quad \|\mathcal{E}_\varepsilon(t)\| \leq C e^{(s_o - \delta)t}. \quad (\text{D.30})$$

Moreover, as $\varepsilon \downarrow 0$,

$$\|P_*(\varepsilon) - P_*\| + \|P_o(\varepsilon) - P_o\| = O(\|L_\Pi(\varepsilon) - L_\Pi\|), \quad \lambda^*(\varepsilon) = \lambda^* + O(\|L_\Pi(\varepsilon) - L_\Pi\|). \quad (\text{D.31})$$

Proof of Lemma D.6. Choose disjoint positively oriented contours Γ_* and Γ_o such that Γ_* encloses $\{\lambda^*\}$ and no other point of $\sigma(L_\Pi)$, while Γ_o satisfies $\Sigma_o \subset \text{int}(\Gamma_o)$ and $\overline{\text{int}(\Gamma_o)} \subset U_o$. Let $\Sigma_\dagger := \sigma(L_\Pi) \setminus (\{\lambda^*\} \cup \Sigma_o)$. If $\Sigma_\dagger \neq \emptyset$, finiteness of the spectrum permits a $\delta > 0$ and a positively oriented contour $\Gamma_\dagger$, disjoint from Γ_* and Γ_o , such that

$$\Sigma_\dagger \subset \text{int}(\Gamma_\dagger), \quad \max_{z \in \Gamma_\dagger} \Re z \leq s_o - \delta. \quad (\text{D.32})$$

If $\Sigma_\dagger = \emptyset$, fix any $\delta > 0$ and set $\Gamma_\dagger = \emptyset$.

Since $\lambda^*(\varepsilon) \rightarrow \lambda^*$, as shown in Theorem 7.2, and the spectrum is continuous under finite-dimensional matrix perturbations, there exists $\varepsilon_0 > 0$ such that, for every $\varepsilon \in [0, \varepsilon_0]$,

$$\sigma(L_\Pi(\varepsilon)) \cap \text{int}(\Gamma_*) = \{\lambda^*(\varepsilon)\}, \quad \sigma(L_\Pi(\varepsilon)) \cap \text{int}(\Gamma_o) = \Sigma_o(\varepsilon).$$

When $\Sigma_\dagger \neq \emptyset$, we may also choose ε_0 so that $\sigma(L_\Pi(\varepsilon)) \cap \text{int}(\Gamma_\dagger) = \sigma(L_\Pi(\varepsilon)) \setminus (\{\lambda^*(\varepsilon)\} \cup \Sigma_o(\varepsilon))$. Consequently, for every $\varepsilon \in [0, \varepsilon_0]$, the Riesz projections $P_*(\varepsilon)$ and $P_o(\varepsilon)$ have the contour integral representations

$$P_*(\varepsilon) = \frac{1}{2\pi i} \int_{\Gamma_*} R(z; L_\Pi(\varepsilon)) \, dz, \quad P_o(\varepsilon) = \frac{1}{2\pi i} \int_{\Gamma_o} R(z; L_\Pi(\varepsilon)) \, dz. \quad (\text{D.33})$$

When $\Sigma_\dagger \neq \emptyset$, define additionally $P_\dagger(\varepsilon) := (2\pi i)^{-1} \int_{\Gamma_\dagger} R(z; L_\Pi(\varepsilon)) \, dz$ for $\varepsilon \in [0, \varepsilon_0]$. In particular, for every $\varepsilon \in (0, \varepsilon_0]$, the spectral sets enclosed by Γ_* , Γ_o , and, when $\Sigma_\dagger \neq \emptyset$, by $\Gamma_\dagger$ form a partition of $\sigma(L_\Pi(\varepsilon))$. Therefore, $I - P_*(\varepsilon) - P_o(\varepsilon) = P_\dagger(\varepsilon)$ if $\Sigma_\dagger \neq \emptyset$ and equals zero if $\Sigma_\dagger = \emptyset$. Moreover, $L_\Pi(\varepsilon)P_*(\varepsilon) = \lambda^*(\varepsilon)P_*(\varepsilon)$. Multiplying the resulting decomposition of the identity by $e^{tL_\Pi(\varepsilon)}$ proves the identity in (D.30).

Next, we prove the bound in (D.30). If $\Sigma_{\dagger} = \emptyset$, then $I - P_*(\varepsilon) - P_o(\varepsilon) = 0$, and hence $\mathcal{E}_\varepsilon(t) = 0$, so the bound is immediate. Suppose now that $\Sigma_{\dagger} \neq \emptyset$ and hence $I - P_*(\varepsilon) - P_o(\varepsilon) = P_{\dagger}(\varepsilon)$. Let

$$A_\varepsilon := L_\Pi(\varepsilon)|_{\text{ran } P_{\dagger}(\varepsilon)}, \quad B_\varepsilon := L_\Pi(\varepsilon)|_{\ker P_{\dagger}(\varepsilon)}.$$

Since $P_{\dagger}(\varepsilon)$ is the Riesz projection associated with the spectrum enclosed by $\Gamma_{\dagger}$, under the invariant decomposition $\mathbb{C}^X = \text{ran } P_{\dagger}(\varepsilon) \oplus \ker P_{\dagger}(\varepsilon)$, the resolvent decomposes into

$$R(z; L_\Pi(\varepsilon)) = R(z; A_\varepsilon)P_{\dagger}(\varepsilon) + R(z; B_\varepsilon)(I - P_{\dagger}(\varepsilon)), \quad z \in \mathbb{C} \setminus \sigma(L_\Pi(\varepsilon)). \quad (\text{D.34})$$

Since the contour $\Gamma_{\dagger}$ encloses $\sigma(A_\varepsilon)$ but no point of $\sigma(B_\varepsilon)$, $R(z; B_\varepsilon)$ is holomorphic on an open neighborhood of $\overline{\text{int}(\Gamma_{\dagger})}$. It follows from Cauchy's theorem that

$$\int_{\Gamma_{\dagger}} e^{tz} R(z; B_\varepsilon)(I - P_{\dagger}(\varepsilon)) dz = 0. \quad (\text{D.35})$$

By (D.34), (D.35), and the Dunford–Taylor formula applied to A_ε with $f(z) = e^{tz}$ (see, e.g., Kato, 1995, Chapter I, equation (5.47), p. 44), we have

$$\frac{1}{2\pi i} \int_{\Gamma_{\dagger}} e^{tz} R(z; L_\Pi(\varepsilon)) dz = \frac{1}{2\pi i} \int_{\Gamma_{\dagger}} e^{tz} R(z; A_\varepsilon) dz P_{\dagger}(\varepsilon) = e^{tA_\varepsilon} P_{\dagger}(\varepsilon) = e^{tL_\Pi(\varepsilon)} P_{\dagger}(\varepsilon) = \mathcal{E}_\varepsilon(t). \quad (\text{D.36})$$

To analyze the resolvent $R(z; L_\Pi(\varepsilon))$ on the left side of (D.36), we set $\Delta_\varepsilon := L_\Pi(\varepsilon) - L_\Pi$, and note

$$zI - L_\Pi(\varepsilon) = (I - \Delta_\varepsilon R(z; L_\Pi))(zI - L_\Pi). \quad (\text{D.37})$$

Let $\Gamma = \Gamma_* \cup \Gamma_o \cup \Gamma_{\dagger}$. Because Γ is compact and does not intersect $\sigma(L_\Pi)$, the resolvent $R(z; L_\Pi)$ is continuous on Γ . Hence

$$M_\Gamma := \sup_{z \in \Gamma} \|R(z; L_\Pi)\| < \infty. \quad (\text{D.38})$$

Since $\|\Delta_\varepsilon\| \rightarrow 0$ as $\varepsilon \downarrow 0$, reducing ε_0 if necessary gives

$$\sup_{0 < \varepsilon \leq \varepsilon_0} \sup_{z \in \Gamma} \|\Delta_\varepsilon R(z; L_\Pi)\| \leq M_\Gamma \sup_{0 < \varepsilon \leq \varepsilon_0} \|\Delta_\varepsilon\| < \frac{1}{2}.$$

Consequently, for every $\varepsilon \in (0, \varepsilon_0]$ and $z \in \Gamma$, $I - \Delta_\varepsilon R(z; L_\Pi)$ is invertible, with $(I - \Delta_\varepsilon R(z; L_\Pi))^{-1} = \sum_{n=0}^{\infty} (\Delta_\varepsilon R(z; L_\Pi))^n$. Therefore, uniformly over $\varepsilon \in (0, \varepsilon_0]$ and $z \in \Gamma$,

$$\|(I - \Delta_\varepsilon R(z; L_\Pi))^{-1}\| \leq \sum_{n=0}^{\infty} \|\Delta_\varepsilon R(z; L_\Pi)\|^n = \frac{1}{1 - \|\Delta_\varepsilon R(z; L_\Pi)\|} \leq 2.$$

It then follows from (D.37) that $R(z; L_\Pi(\varepsilon)) = R(z; L_\Pi)(I - \Delta_\varepsilon R(z; L_\Pi))^{-1}$ for every $\varepsilon \in (0, \varepsilon_0]$ and $z \in \Gamma$, and

$$\sup_{0 < \varepsilon \leq \varepsilon_0} \sup_{z \in \Gamma} \|R(z; L_\Pi(\varepsilon))\| \leq \sup_{z \in \Gamma} \|R(z; L_\Pi)\| \sup_{0 < \varepsilon \leq \varepsilon_0} \sup_{z \in \Gamma} \|(I - \Delta_\varepsilon R(z; L_\Pi))^{-1}\| \leq 2M_\Gamma. \quad (\text{D.39})$$

By (D.32), (D.36) and (D.39), there exists a constant $C > 0$ such that, for every $t \geq 0$,

$$\sup_{0 < \varepsilon \leq \varepsilon_0} \|\mathcal{E}_\varepsilon(t)\| \leq \frac{|\Gamma_{\dagger}|}{2\pi} \sup_{0 < \varepsilon \leq \varepsilon_0} \sup_{z \in \Gamma_{\dagger}} \|R(z; L_\Pi(\varepsilon))\| e^{t \max_{z \in \Gamma_{\dagger}} \Re z} \leq C e^{(s_o - \delta)t},$$

where $|\Gamma_{\dagger}|$ denotes the total arclength of $\Gamma_{\dagger}$. This proves the bound in (D.30).

Finally, we prove the two estimates in (D.31). For $j \in \{*, \circ\}$, it follows from (D.33) that $P_j(\varepsilon) - P_j(0) = (2\pi i)^{-1} \int_{\Gamma_j} [R(z; L_\Pi(\varepsilon)) - R(z; L_\Pi)] dz$. By the resolvent identity $R(z; L_\Pi(\varepsilon)) - R(z; L_\Pi) = R(z; L_\Pi(\varepsilon)) \Delta_\varepsilon R(z; L_\Pi)$, (D.38), and (D.39), we have $\sup_{z \in \Gamma} \|R(z; L_\Pi(\varepsilon)) - R(z; L_\Pi)\| \leq 2M_\Gamma^2 \|\Delta_\varepsilon\|$ for every $\varepsilon \in (0, \varepsilon_0]$. Therefore,

$$\|P_j(\varepsilon) - P_j(0)\| \leq \frac{|\Gamma_j|}{2\pi} \sup_{z \in \Gamma} \|R(z; L_\Pi(\varepsilon)) - R(z; L_\Pi)\| \leq |\Gamma_j| M_\Gamma^2 \|\Delta_\varepsilon\| / \pi = O(\|\Delta_\varepsilon\|),$$

which proves the first estimate in (D.31). For the second estimate in (D.31), note that both $P_*(\varepsilon)$ and P_* have rank one. Moreover, it follows from $L_\Pi(\varepsilon)P_*(\varepsilon) = \lambda^*(\varepsilon)P_*(\varepsilon)$ and $L_\Pi P_* = \lambda^* P_*$ that

$$\lambda^*(\varepsilon) = \text{tr}(L_\Pi(\varepsilon)P_*(\varepsilon)), \quad \lambda^* = \text{tr}(L_\Pi P_*).$$

Consequently, $\lambda^*(\varepsilon) - \lambda^* = \text{tr}(L_\Pi(\varepsilon)P_*(\varepsilon) - L_\Pi P_*) = \text{tr}(\Delta_\varepsilon P_*(\varepsilon)) + \text{tr}(L_\Pi(P_*(\varepsilon) - P_*))$, and hence

$$|\lambda^*(\varepsilon) - \lambda^*| \leq |\text{tr}(\Delta_\varepsilon P_*(\varepsilon))| + |\text{tr}(L_\Pi(P_*(\varepsilon) - P_*))| = O(\|\Delta_\varepsilon\|).$$

This proves the second estimate in (D.31) and completes the proof. $\square$

Proof of Theorem 7.3. Apply the operator identity (D.30) in Lemma D.6 to g , evaluate at x , and set $R_{\varepsilon,t} := (\mathcal{E}_\varepsilon(t)g)(x)$. This gives (7.10); the bound on $\mathcal{E}_\varepsilon(t)$ gives its uniform remainder bound after enlarging the constant. The eigenvalue estimate in (7.11) follows from (D.31) and the definition of $r_L(\varepsilon)$, while (7.7) gives $a_*(x, g; \varepsilon) = a\varepsilon^d(1 + o(1))$ for some $a > 0$. This proves (7.11).

Fix $T < \infty$. Duhamel's identity and $\|L_\Pi(\varepsilon) - L_\Pi\| = r_L(\varepsilon)$ give $e^{tL_\Pi(\varepsilon)} - e^{tL_\Pi} = O(r_L(\varepsilon))$ uniformly over $0 \leq t \leq T$. Together with the projection estimate in (D.31), this proves the stated compact-maturity convergence of $b_t(x, g; \varepsilon)$.

To identify its limit, fix $\lambda \in \Sigma_\circ$ and set $N_\lambda := (L_\Pi - \lambda I)P_\lambda$. On the invariant subspace $\text{ran } P_\lambda$, the operator L_Π equals $\lambda I + N_\lambda$, where $N_\lambda^{\nu_\lambda} = 0$. Hence $e^{tL_\Pi}P_\lambda = e^{\lambda t} \sum_{j=0}^{\nu_\lambda-1} \frac{t^j}{j!} (L_\Pi - \lambda I)^j P_\lambda$. Summing over $\lambda \in \Sigma_\circ$, using $P_\circ = \sum_{\lambda \in \Sigma_\circ} P_\lambda$, and applying $e_x^\top(\cdot)g$ proves (7.12).

Suppose now that $\lambda_{\mathcal{K}} < s_\circ$. Exact corridor restriction gives $q_t(x, g; 0) = e_{x,\mathcal{K}}^\top e^{tL_{\mathcal{K}}} g_{\mathcal{K}}$ for every $t \geq 0$, where $e_{x,\mathcal{K}}$ and $g_{\mathcal{K}}$ denote the restrictions to $\mathbf{X}_{\mathcal{K}}$. Taking Laplace transforms for $\Re z$ sufficiently large yields

$$e_x^\top R(z; L_\Pi)g = e_{x,\mathcal{K}}^\top R(z; L_{\mathcal{K}})g_{\mathcal{K}}. \quad (\text{D.40})$$

Both sides are rational functions of z and agree on a right half-plane, so (D.40) extends as an identity wherever both sides are defined. For $\lambda \in \Sigma_\circ$, we have $\Re \lambda = s_\circ > \lambda_{\mathcal{K}} = s(L_{\mathcal{K}})$; hence $\lambda \notin \sigma(L_{\mathcal{K}})$, and the right-hand side is holomorphic near λ . The Laurent expansion of the left side of (D.40) is

$$e_x^\top R(z; L_\Pi)g = \sum_{j=0}^{\nu_\lambda-1} \frac{e_x^\top (L_\Pi - \lambda I)^j P_\lambda g}{(z - \lambda)^{j+1}} + h_\lambda(z),$$

where h_λ is holomorphic near λ . Holomorphy of the right-hand side of (D.40) therefore forces every coefficient in the principal part above to vanish. Equation (7.12) then gives $b_t(x, g; 0) \equiv 0$. $\square$

Proof of Corollary 7.4. Fix a sufficiently small $\varepsilon > 0$. By Assumption 7.1, $L_\Pi(\varepsilon)$ is irreducible. Theorem 4.1, specifically the rank-one Perron limit (4.3), applied to $L_\Pi(\varepsilon)$, therefore gives

$$e^{-\lambda^*(\varepsilon)t} q_t(x, g; \varepsilon) \rightarrow (P_*(\varepsilon)g)(x) = a_*(x, g; \varepsilon) > 0.$$

Here strict positivity follows because the right and left Perron vectors of the irreducible Metzler matrix $L_\Pi(\varepsilon)$ are strictly positive and $g \geq 0$ is nonzero. Thus $q_t(x, g; \varepsilon) = e^{\lambda^*(\varepsilon)t}(a_*(x, g; \varepsilon) + o(1))$, and taking logarithms gives

$$\lim_{t \rightarrow \infty} \left[-\frac{1}{t} \log q_t(x, g; \varepsilon) \right] = -\lambda^*(\varepsilon).$$

Taking $\varepsilon \downarrow 0$ on both sides proves the first limit in (7.13).

For the reverse order, fix $t > 0$ and set $\Delta_\varepsilon := L_\Pi(\varepsilon) - L_\Pi$. Duhamel's identity, also known as the variation-of-constants identity, gives

$$e^{tL_\Pi(\varepsilon)} - e^{tL_\Pi} = \int_0^t e^{(t-s)L_\Pi(\varepsilon)} \Delta_\varepsilon e^{sL_\Pi} ds.$$

Since $L_\Pi(\varepsilon) \rightarrow L_\Pi$, the two semigroup factors in the integrand are uniformly bounded for sufficiently small ε and $0 \leq s \leq t$. Hence, for some $C_t < \infty$ independent of ε , $\|e^{tL_\Pi(\varepsilon)} - e^{tL_\Pi}\| \leq C_t \|\Delta_\varepsilon\| \rightarrow 0$. Applying this operator convergence to g and evaluating at x gives $q_t(x, g; \varepsilon) \rightarrow q_t(x, g; 0)$. The corridor is nonempty by hypothesis, so Proposition 5.1(ii) gives $q_t(x, g; 0) > 0$. Continuity of the logarithm therefore yields

$$\lim_{\varepsilon \downarrow 0} \left[-\frac{1}{t} \log q_t(x, g; \varepsilon) \right] = -\frac{1}{t} \log q_t(x, g; 0).$$

Finally, the yield expansion (5.6) in Theorem 5.2 implies that the right-hand side converges to $-\lambda(\mathcal{C}, g) = -\lambda_{\mathcal{X}}$ as $t \rightarrow \infty$. This proves the second limit in (7.13). $\square$

Proof of Theorem 7.6. Under Assumption 7.5, $\Sigma_o = \{\lambda_o\}$ and $\lambda_o = \lambda_{\mathcal{X}}$ is algebraically simple. Hence (7.12) gives $b_t(x, g; 0) = e^{\lambda_o t} b_o(x, g; 0)$. Moreover, $d > 0$ makes the global coefficient $(P_*g)(x)$ vanish, while every remaining spectral term has real part strictly below λ_o . The full spectral expansion therefore gives

$$q_t(x, g; 0) = e^{\lambda_o t} b_o(x, g; 0) + o(e^{\lambda_o t}).$$

On the other hand, Theorem 5.2 gives $q_t(x, g; 0) = c_x(g) t^{\nu(\mathcal{C}, g) - 1} e^{\lambda_o t} (1 + o(1))$, with $c_x(g) > 0$. Comparing the two expansions yields $\nu(\mathcal{C}, g) = 1$ and $b_o(x, g; 0) = c_x(g) > 0$. Combining (7.10) and (7.15) proves (7.17).

For all small ε , both coefficients are positive and $\lambda^*(\varepsilon) > \lambda_o(\varepsilon)$, so equality of the modes has the unique solution (7.18). By (7.11) and (7.16),

$$\log \frac{b_o(x, g; \varepsilon)}{a_*(x, g; \varepsilon)} = d \log(1/\varepsilon) + \log \frac{b_o(x, g; 0)}{a} + o(1).$$

Since $\Delta(\varepsilon) \rightarrow \Delta > 0$, substitution gives (7.19) and $t_\varepsilon^c \sim (d/\Delta) \log(1/\varepsilon)$. In particular, $t_\varepsilon^c \rightarrow \infty$, so the crossover maturity is positive for all sufficiently small ε .

For the crossover-window limit, the defining equality at t_ε^c implies that, for every $s \in \mathbb{R}$,

$$\frac{a_*(x, g; \varepsilon) e^{\lambda^*(\varepsilon)(t_\varepsilon^c + s)}}{b_o(x, g; \varepsilon) e^{\lambda_o(\varepsilon)(t_\varepsilon^c + s)}} = e^{[\lambda^*(\varepsilon) - \lambda_o(\varepsilon)]s}.$$

Since $\lambda^*(\varepsilon) - \lambda_o(\varepsilon) \rightarrow \Delta$, the right-hand side converges to $e^{\Delta s}$ uniformly for s in compact sets. Moreover, $b_o(x, g; \varepsilon) \rightarrow b_o(x, g; 0) > 0$ and, for all sufficiently small ε , $|\lambda_o(\varepsilon) - \lambda_o| < \delta/2$. Hence, for each fixed $S < \infty$ and all $|s| \leq S$,

$$\frac{|R_{\varepsilon, t_\varepsilon^c + s}|}{b_o(x, g; \varepsilon) e^{\lambda_o(\varepsilon)(t_\varepsilon^c + s)}} \leq C_S e^{-(\delta/2)(t_\varepsilon^c - S)} \rightarrow 0.$$

Dividing (7.17) by the corridor term therefore shows that the full price divided by that term converges uniformly to $1 + e^{\Delta s}$. Dividing the global-to-corridor ratio by this limit proves (7.20).

For the refinement following the theorem, the eigenvalue bounds in (7.11) and (7.16) give $\Delta(\varepsilon) = \Delta + O(r_L(\varepsilon))$. Under (7.21),

$$\left(\frac{1}{\Delta(\varepsilon)} - \frac{1}{\Delta} \right) \log(1/\varepsilon) = O(r_L(\varepsilon) \log(1/\varepsilon)) = o(1).$$

Replacing the denominator in (7.19) by Δ therefore proves (7.22). $\square$

We conclude this subsection with the fixed-recovery specialization used to prove Corollaries 7.8 and 7.9. Retain the notation from Section 7.3. For every sufficiently small $\varepsilon > 0$, let $\varphi_\varepsilon, \psi_\varepsilon \gg 0$ be right and left Perron vectors of $L_\Pi(\varepsilon)$, normalized by $\mathbf{1}^\top \varphi_{\varepsilon, \mathcal{D}} = 1$ and $\psi_\varepsilon^\top \varphi_\varepsilon = 1$. Under this normalization, $\pi_\varepsilon^{\text{LR}} = \varphi_\varepsilon \odot \psi_\varepsilon$.

Lemma D.7 (Fixed-recovery first-order expansion). *Under Assumption 7.7 and $\lambda_{\mathcal{D}} > \lambda_{\mathcal{N}}$, let $\bar{\pi}_{\mathcal{N}}$ be the stationary distribution of Q . Then*

$$\pi_\varepsilon^P(\mathcal{D}) = \varepsilon c_P + o(\varepsilon), \quad c_P = -\bar{\pi}_{\mathcal{N}}^\top F U^{-1} \mathbf{1}_{\mathcal{D}} > 0. \quad (\text{D.41})$$

Moreover,

$$\lambda^*(\varepsilon) \rightarrow \lambda_{\mathcal{D}}, \quad \varphi_{\varepsilon, \mathcal{D}} \rightarrow \varphi_{\mathcal{D}}, \quad \varepsilon^{-1} \varphi_{\varepsilon, \mathcal{N}} \rightarrow R(\lambda_{\mathcal{D}}; L_{\mathcal{N}}) J \varphi_{\mathcal{D}} \gg 0, \quad (\text{D.42})$$

while $\psi_\varepsilon \rightarrow \Psi_{\mathcal{D}}$, with $\Psi_{\mathcal{D}}$ defined in (D.17).

Proof. Write the physical stationary distribution conformably as

$$(\pi_\varepsilon^P)^\top = ((\pi_{\varepsilon, \mathcal{N}}^P)^\top, (\pi_{\varepsilon, \mathcal{D}}^P)^\top).$$

The disaster block of $(\pi_\varepsilon^P)^\top L_X(\varepsilon) = 0$ is

$$(\pi_{\varepsilon, \mathcal{N}}^P)^\top F_\varepsilon + (\pi_{\varepsilon, \mathcal{D}}^P)^\top U = 0.$$

Since $s(U) < 0$, U is invertible and

$$(\pi_{\varepsilon, \mathcal{D}}^P)^\top = -(\pi_{\varepsilon, \mathcal{N}}^P)^\top F_\varepsilon U^{-1}.$$

Hence $\pi_\varepsilon^P(\mathcal{D}) = O(\varepsilon)$. The normal block of the stationary equation is $(\pi_{\varepsilon, \mathcal{N}}^P)^\top (Q - D_\varepsilon) + (\pi_{\varepsilon, \mathcal{D}}^P)^\top H = 0$. Since $D_\varepsilon = O(\varepsilon)$, every limit point of $\pi_{\varepsilon, \mathcal{N}}^P$ has unit mass and is stationary for Q . Irreducibility

therefore gives $\pi_{\varepsilon, \mathcal{N}}^P \rightarrow \bar{\pi}_{\mathcal{N}}$. Dividing by ε and using (7.24) yields

$$\varepsilon^{-1}(\pi_{\varepsilon, \mathcal{D}}^P)^\top \longrightarrow -\bar{\pi}_{\mathcal{N}}^\top F U^{-1}.$$

Because $-U^{-1} = \int_0^\infty e^{tU} dt \gg 0$, the limiting row vector is strictly positive and (D.41) follows.

Continuity of the spectral bound gives $\lambda^*(\varepsilon) \rightarrow \lambda_{\mathcal{D}}$, since $\lambda_{\mathcal{D}} > \lambda_{\mathcal{N}}$ is the unique spectral bound of $L_{\Pi}(0)$. The right Perron equation for (7.23) is

$$(\lambda^*(\varepsilon)I_{\mathcal{N}} - L_{\mathcal{N}} + D_{\varepsilon})\varphi_{\varepsilon, \mathcal{N}} = E_{\varepsilon}\varphi_{\varepsilon, \mathcal{D}}, \quad (\text{D.43})$$

$$L_{\mathcal{D}\mathcal{N}}\varphi_{\varepsilon, \mathcal{N}} + L_{\mathcal{D}}\varphi_{\varepsilon, \mathcal{D}} = \lambda^*(\varepsilon)\varphi_{\varepsilon, \mathcal{D}}. \quad (\text{D.44})$$

For small ε , the inverse in the first equation exists and converges to $R(\lambda_{\mathcal{D}}; L_{\mathcal{N}})$. Since $E_{\varepsilon} = O(\varepsilon)$, we obtain $\varphi_{\varepsilon, \mathcal{N}} = O(\varepsilon)$. By the second equation, every limit point of $\varphi_{\varepsilon, \mathcal{D}}$ is a right eigenvector of $L_{\mathcal{D}}$ at $\lambda_{\mathcal{D}}$. The normalization $\mathbf{1}^\top \varphi_{\varepsilon, \mathcal{D}} = 1$ and uniqueness of the normalized Perron vector therefore give $\varphi_{\varepsilon, \mathcal{D}} \rightarrow \varphi_{\mathcal{D}}$. Dividing (D.43) by ε and using $\varepsilon^{-1}E_{\varepsilon} \rightarrow J$ gives

$$\varepsilon^{-1}\varphi_{\varepsilon, \mathcal{N}} \longrightarrow R(\lambda_{\mathcal{D}}; L_{\mathcal{N}})J\varphi_{\mathcal{D}} \gg 0.$$

Since $\lambda_{\mathcal{D}}$ is a simple eigenvalue of $L_{\Pi}(0)$, continuity of its rank-one Riesz projection implies convergence of the corresponding left eigenspaces. We have $\varphi_{\varepsilon} \rightarrow \Phi_{\mathcal{D}}$, and the normalization $\psi_{\varepsilon}^\top \varphi_{\varepsilon} = 1$ therefore selects the limit $\Psi_{\mathcal{D}}$. Hence $\psi_{\varepsilon} \rightarrow \Psi_{\mathcal{D}}$, which proves (D.42). $\square$

Proof of Corollary 7.8. Equation (D.41) gives the physical-mass expansion with $c_P > 0$. Define, within this proof, $\bar{\varphi}_{\mathcal{N}} = R(\lambda_{\mathcal{D}}; L_{\mathcal{N}})J\varphi_{\mathcal{D}}$ and $c_{\text{LR}} = (\Psi_{\mathcal{D}})_{\mathcal{N}}^\top \bar{\varphi}_{\mathcal{N}}$. Lemma D.7 gives $\bar{\varphi}_{\mathcal{N}} \gg 0$ and $\psi_{\varepsilon, \mathcal{N}} \longrightarrow (\Psi_{\mathcal{D}})_{\mathcal{N}} \gg 0$, so $c_{\text{LR}} > 0$ and

$$\varepsilon^{-1}\pi_{\varepsilon}^{\text{LR}}(\mathcal{N}) = \varepsilon^{-1}\psi_{\varepsilon, \mathcal{N}}^\top \varphi_{\varepsilon, \mathcal{N}} \longrightarrow (\Psi_{\mathcal{D}})_{\mathcal{N}}^\top \bar{\varphi}_{\mathcal{N}} = c_{\text{LR}}.$$

Since $\pi_{\varepsilon}^{\text{LR}}$ has total mass one, its disaster mass is $1 - c_{\text{LR}}\varepsilon + o(\varepsilon)$, as claimed. $\square$

Proof of Corollary 7.9. Assumption 7.7 satisfies Assumption 7.1 with $\eta(\varepsilon) = \varepsilon^{\min\{\eta_F, 1\}}$: the normal diagonal changes by $O(\varepsilon)$, the rare entry block opens at order ε with relative remainder $O(\varepsilon^{\eta_F})$, and recovery is a zero-weight edge. Moreover, $r_L(\varepsilon) = O(\varepsilon)$ by (7.23)–(7.24).

For $x \in \mathcal{N}$ and $g = \mathbf{1}$, the rare entry edge gives $d_{\mathcal{N}\mathcal{D}} = 1$, and the payoff-side term is zero because $\mathcal{D} \in \mathcal{S}(\mathbf{1})$. Hence $d(\mathcal{N}, \mathbf{1}) = 1$. The limiting corridor is $\mathcal{K}(\mathcal{N}, \mathbf{1}) = \{\mathcal{N}\}$ and has rate $\lambda_{\mathcal{K}} = \lambda_{\mathcal{N}}$. Since $L_{\mathcal{N}}$ is irreducible, $\lambda_{\mathcal{N}}$ is algebraically simple and strictly dominates the real parts of its other eigenvalues. The spectrum of the limiting block-triangular generator is $\sigma(L_{\mathcal{N}}) \cup \sigma(L_{\mathcal{D}})$, so (7.25) makes $\lambda_{\mathcal{N}}$ the unique eigenvalue, other than $\lambda_{\mathcal{D}}$, with maximal real part. Thus Assumption 7.5 holds with $\lambda^* = \lambda_{\mathcal{D}}$ and $\lambda_{\mathcal{K}} = \lambda_{\circ} = \lambda_{\mathcal{N}}$.

Theorem 7.6 now applies. Equations (7.11) and (7.16), together with $r_L(\varepsilon) = O(\varepsilon)$, give the strictly positive coefficients a_x and b_x in the statement and

$$a_*(x, \mathbf{1}; \varepsilon) = a_x \varepsilon (1 + o(1)), \quad b_{\circ}(x, \mathbf{1}; \varepsilon) = b_x + O(\varepsilon), \quad \lambda^*(\varepsilon) - \lambda_{\circ}(\varepsilon) = \lambda_{\mathcal{D}} - \lambda_{\mathcal{N}} + O(\varepsilon).$$

Since $\varepsilon \log(1/\varepsilon) \rightarrow 0$, condition (7.21) holds. Substituting $d = 1$, $a = a_x$, $b_{\circ}(x, \mathbf{1}; 0) = b_x$, and $\Delta = \lambda_{\mathcal{D}} - \lambda_{\mathcal{N}}$ into (7.22) proves (7.26). $\square$

REFERENCES

- Ando, A. and Fisher, F. M. (1963). Near-decomposability, partition and aggregation, and the relevance of stability discussions. *International Economic Review*, 4(1):53–67.
- Applebaum, D. (2009). *Lévy Processes and Stochastic Calculus*. Cambridge University Press, second edition.
- Avrachenkov, K. E., Filar, J. A., and Howlett, P. G. (2013). *Analytic Perturbation Theory and Its Applications*. Society for Industrial and Applied Mathematics, Philadelphia.
- Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866.
- Barro, R. J. and Ursúa, J. F. (2008). Macroeconomic crises since 1870. *Brookings Papers on Economic Activity*, 39(1):255–350.
- Bátkai, A., Kramar Fijavž, M., and Rhandi, A. (2017). *Positive Operator Semigroups: From Finite to Infinite Dimensions*, volume 257 of *Operator Theory: Advances and Applications*. Birkhäuser.
- Beare, B. K. and Toda, A. A. (2026). Reducible Markov modulation, pole order, and tail behavior in random growth models. *Journal of Applied Probability*, pages 1–28.
- Bonomo, M. and Garcia, R. (1996). Consumption and equilibrium asset pricing: An empirical assessment. *Journal of Empirical Finance*, 3(3):239–265.
- Borovička, J., Hansen, L. P., Hendricks, M., and Scheinkman, J. A. (2011). Risk-price dynamics. *Journal of Financial Econometrics*, 9(1):3–65.
- Borovička, J., Hansen, L. P., and Scheinkman, J. A. (2016). Misspecified recovery. *The Journal of Finance*, 71(6):2493–2544.
- Borovička, J. and Stachurski, J. (2021). Stability of equilibrium asset pricing models: A necessary and sufficient condition. *Journal of Economic Theory*, 193:105227.
- Christensen, T. M. (2017). Nonparametric stochastic discount factor decomposition. *Econometrica*, 85(5):1501–1536.
- Courtois, P.-J. (1975). Error analysis in nearly-completely decomposable stochastic systems. *Econometrica*, 43(4):691–709.
- Doob, J. L. (1957). Conditional Brownian motion and the boundary limits of harmonic functions. *Bulletin de la Société Mathématique de France*, 85:431–458.
- Duffie, D. and Garman, M. B. (1986). Intertemporal arbitrage and the Markov valuation of securities. Working paper.
- Duffie, D., Pan, J., and Singleton, K. J. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376.
- Friedland, S. and Schneider, H. (1980). The growth of powers of a nonnegative matrix. *SIAM Journal on Algebraic and Discrete Methods*, 1(2):185–200.
- Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics*, 127(2):645–700.
- Garman, M. B. (1985). Towards a semigroup pricing theory. *The Journal of Finance*, 40(3):847–861.
- Giglio, S., Kelly, B. T., and Kozak, S. (2024). Equity term structures without dividend strips data. *The Journal of Finance*, 79(6):4143–4196.
- Gourio, F. (2012). Disaster risk and business cycles. *American Economic Review*, 102(6):2734–2766.

- Hansen, L. P. (2012). Dynamic valuation decomposition within stochastic economies. *Econometrica*, 80(3):911–967.
- Hansen, L. P. and Scheinkman, J. A. (2009). Long-term risk: An operator approach. *Econometrica*, 77(1):177–234.
- Hansen, L. P. and Scheinkman, J. A. (2012). Pricing growth-rate risk. *Finance and Stochastics*, 16(1):1–15.
- Hasler, M. and Marfè, R. (2016). Disaster recovery and the term structure of dividend strips. *Journal of Financial Economics*, 122(1):116–134.
- Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.
- Kato, T. (1995). *Perturbation Theory for Linear Operators*. Springer, Berlin, second edition.
- Krishnamurthy, A. and Muir, T. (2025). How credit cycles across a financial crisis. *Journal of Finance*, 80(3):1339–1378.
- Martin, I. (2012). On the valuation of long-dated assets. *Journal of Political Economy*, 120(2):346–358.
- Meyer, C. D. (1989). Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems. *SIAM Review*, 31(2):240–272.
- Muir, T. (2017). Financial crises and risk premia. *The Quarterly Journal of Economics*, 132(2):765–809.
- Nakamura, E., Steinsson, J., Barro, R. J., and Ursúa, J. F. (2013). Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics*, 5(3):35–74.
- Norris, J. R. (1997). *Markov Chains*. Cambridge University Press.
- Qin, L. and Linetsky, V. (2016). Positive eigenfunctions of Markovian pricing operators: Hansen–Scheinkman factorization, Ross recovery, and long-term pricing. *Operations Research*, 64(1):99–117.
- Qin, L. and Linetsky, V. (2017). Long-term risk: A martingale approach. *Econometrica*, 85(1):299–312.
- Ross, S. A. (2015). The recovery theorem. *The Journal of Finance*, 70(2):615–648.
- Rothblum, U. G. (1975). Algebraic eigenspaces of nonnegative matrices. *Linear Algebra and its Applications*, 12(3):281–292.
- Schneider, H. (1986). The influence of the marked reduced graph of a nonnegative matrix on the Jordan form and on related properties: A survey. *Linear Algebra and its Applications*, 84:161–189.
- Schweitzer, P. J. (1968). Perturbation theory and finite Markov chains. *Journal of Applied Probability*, 5(2):401–413.
- Simon, H. A. and Ando, A. (1961). Aggregation of variables in dynamic systems. *Econometrica*, 29(2):111–138.
- van Binsbergen, J. H., Brandt, M. W., and Koijen, R. S. J. (2012). On the timing and pricing of dividends. *American Economic Review*, 102(4):1596–1618.
- van Binsbergen, J. H., Hueskes, W., Koijen, R. S. J., and Vrugt, E. B. (2013). Equity yields. *Journal of Financial Economics*, 110(3):503–519.

- Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, 68(3):987–1035.
- Wu, D. (2025). A disaster explanation of equity term structures. *Review of Finance*, 29(5):1437–1465.
- Yin, G. G. and Zhang, Q. (2012). *Continuous-Time Markov Chains and Applications: A Two-Time-Scale Approach*, volume 37 of *Stochastic Modelling and Applied Probability*. Springer.
- Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61(1):57–84.