Liquidity Portfolios

Swaminathan Balasubramaniam* Philip H. Dybvig† Michael Tehranchi‡

Preliminary Draft: 2026-04-17

Abstract

Cash management by treasuries motivates us to study optimal investments in assets with various degrees of liquidity to meet cash flow shocks over time. Typically, it is optimal to buy many assets with differing degrees of liquidity, with more liquid assets having lower hold-to-maturity returns. Relatively illiquid assets are liquidated infrequently, in states demanding large cash flows, while liquid assets are liquidated more often. The option to liquidate at various times implies a convex liquidity premium, interpretable as an option value, even with a single representative investor. We develop a liquidity beta formula that prices liquidity risk: assets whose liquidation values are low precisely when liquidity needs are high earn higher expected returns. In a multi-period framework, we show that the usual pecking order of selling the most liquid asset first can fail because preserving the liquid asset as a buffer against future shocks can be more valuable, and our calibration shows that this reversal arises at empirically plausible levels of liquidity shock persistence. The firm's liability structure also matters: even for a given average liquidity need, changes in the dispersion of funding shocks affect continuation values by changing the frequency with which assets are liquidated.

Keywords: Liquidity risk, Asset allocation, Liquidity beta, Corporate treasury.

JEL Classification: G11, G12, G32.

*NEOMA Business School, France; s.balasubramaniam@neoma-bs.fr

†Olin Business School, Washington University in St. Louis; dybvig@wustl.edu

‡Department of Pure Mathematics and Mathematical Statistics, University of Cambridge; m.tehranchi@statslab.cam.ac.uk

1 Introduction

Firms face liquidity shocks in many forms. They often finance assets with outside funding that may be withdrawn or called at short notice. In the ordinary course of business, they may also need to make unforeseen purchases or meet emergency cash needs. Corporate treasuries therefore play a central role in ensuring that firms have sufficient cash to meet these obligations while allocating resources across assets with different degrees of liquidity. We study a fully stochastic cash-management problem in the spirit of Baumol (1952), in which cash demands, asset values, and asset liquidation values are all stochastic. The value of each asset to the buyer can be written as the buyer's private valuation of an American put option for liquidation over time. The convexity of value, usually understood as a clientele effect, can occur even with a single representative agent, due to convexity of the option valuation. We further show that the standard pecking order of selling the most liquid asset first fails dynamically whenever liquidity shocks are sufficiently persistent.

Our basic model has three dates. At date zero an agent assembles a portfolio of risky assets. At an interim date a random liquidity shock arrives, representing a margin call, a redemption demand, or a debt payment that must be financed from asset sales. At the terminal date, retained units pay their continuation values. What distinguishes assets is the ratio of the interim liquidation value to the continuation value, which we call the liquidation ratio. The agent sells first the assets that deliver the most cash per unit of foregone continuation value. Because the shock is random, the identity of the cutoff asset varies across states, and it is this state dependence that drives the pricing results.

Our first main result establishes that the optimal liquidation policy has a threshold structure. In each state, all assets with liquidation ratios above a cutoff are sold in full, the marginal asset is partially sold to exactly meet the requirement, and all assets below the cutoff are retained. The cutoff is determined endogenously by the portfolio and the shock realization. A technical subtlety arises when the shock distribution has mass points: the set of Lagrange multipliers consistent with optimality can then be an interval rather than a single number. We resolve this by selecting the *minimal supporting multiplier*. This delivers a canonical single-valued cutoff representation and supports the pricing results that follow.

Our second main result derives asset prices from the date-zero portfolio choice. Each asset provides two services. The first is its terminal continuation value. The second is a liquidity service: in states where the shock is large enough to draw the asset into the liquidation set, it delivers cash at the margin. The liquidity premium is the expected marginal-utility-weighted excess of this cash service over the continuation value. As the shock grows larger, the agent is forced to sell assets with progressively worse liquidation ratios, so the marginal cost of liquidity rises at an increasing rate. This generates a convex relationship between asset prices and liquidity with a single representative investor, the same qualitative finding as Amihud and Mendelson (1986), but without requiring heterogeneous investors.[1]

An asset derives value from its liquidity only in the states where it is actually liquidated. If it is not liquidated, its value simply equals the discounted future liquidation value. This gives rise to a dual source of value. First, an asset commands a higher price if it tends to have a large liquidation value precisely when the liquidity shock is severe. Second, it commands a further premium if it supplies cash precisely when other assets in the market fail to do so. Like a put option, a liquid asset is most valuable when the aggregate shock is high and alternative sources of liquidity are scarce. The downside is limited by the asset's terminal value, which acts as a floor. Assets are therefore ranked for liquidation purposes in descending order of their liquidation-to-continuation ratio. We formalize this intuition through a liquidity beta: the expected marginal-utility-weighted cash service an asset provides when it enters the liquidation set. The cross-sectional ordering of prices is determined by this beta. A high liquidation ratio commands a large premium only when the asset enters the liquidation set during genuine scarcity, not in states of abundant liquidity. This connects to the dual-role characterization of Geromichalos et al. (2023) and to the market-to-funding liquidity interaction of Brunnermeier and Pedersen (2009).

We also apply the pricing framework to liability design. The distribution of the liquidity shock is itself a choice variable: a firm's funding structure, liability maturity, and access to credit

[1] Amihud and Mendelson (1986) find a concave relationship between *returns* and illiquidity. Since prices are the inverse of returns, their result is consistent with a convex relationship between prices and liquidity

lines all shape the distribution of interim funding needs. Under risk neutrality, firm value rises when the liability structure places more mass on large shocks, because this activates the liquidation service of high-ratio assets more often. Under risk aversion the precautionary motive reverses the incentive: marginal utility is highest in the worst states, so smoother funding structures become optimal. This tension accounts for the finding of Berlin and Mester (1999) that core retail deposits stabilize bank funding, and for the mechanism of Huang and Ratnovski (2011) whereby short-term wholesale funding creates fragility by concentrating mass on the worst shock realizations.

The framework extends to a multiperiod economy with trading at every date. At each date and state, the discounted shadow value of an asset equals the larger of its discounted immediate liquidation payoff and its expected continuation value. On states where the liquidity multiplier is strictly positive, this yields the familiar representation of the shadow price as the larger of its immediate liquidation value and its risk-adjusted continuation value. The asset is liquidated only when selling immediately dominates preserving it as a buffer against future shocks. As a result, the heuristic of selling the most liquid asset first is not generally optimal. An asset that is very liquid today may be more valuable to retain as a future option. We illustrate this with a simple counterexample and connect it to the patterns of Chernenko and Sunderam (2016), who document that mutual funds trade across the portfolio rather than exhausting their most liquid positions, and of Andonov et al. (2025), who show that pension funds sell equities while retaining cash as a buffer.

We sharpen this dynamic result in Proposition 5, which characterizes the pecking-order reversal in closed form. In a two-period economy with a perfectly liquid buffer asset and a secondary asset whose liquidation quality deteriorates under stress, selling the secondary asset first is strictly optimal whenever shock persistence is sufficiently high relative to a threshold that depends on the extent of stress deterioration. The threshold is lower when shocks are more persistent and when secondary-market liquidity deteriorates more sharply under stress, because both forces raise the option value of preserving the buffer for a potential second shock. The proposition thus identifies shock persistence and stress-state deterioration in liquidation quality as the two primitives that jointly govern whether the static pecking order survives in a dynamic setting.

The shadow prices in our multiperiod model have a natural connection to arbitrage theory with transaction costs. Jouini and Kallal (1995) show that a bid-ask spread is arbitrage-free if and only if some process lying between the bid and ask prices is a martingale under an equivalent measure. In our model, the shadow price of each asset is exactly such a process: it lies between the liquidation price and the purchase price, and it follows a martingale whenever the agent holds positive quantities of that asset. The shadow price therefore has a clean interpretation as a no-arbitrage price within the liquidity-constrained economy.

Section 4 studies the implication of our theoretical results quantitatively through a calibrated three-asset model, mapping the framework to empirical evidence on mutual-fund cash buffers (Chernenko and Sunderam, 2016), corporate bond illiquidity (Bao et al., 2011), redemption-shock magnitudes (Ma et al., 2022), and the persistence of institutional funding pressure (Andonov et al., 2025). Three findings emerge. First, the curvature of the price-liquidity schedule is sensitive to the liquid-buffer share: as the buffer thins, the premium concentrates disproportionately at the top of the quality hierarchy, driven entirely by where the shock distribution sits relative to the endogenous band boundaries rather than by any change in the average liquidity need. Second, the closed-form reversal threshold from Proposition 5 lies well below the range of shock-persistence estimates in the data, so the static pecking order fails at the baseline calibration and indeed for every empirically plausible persistence value. Third, an asset with a stable liquidation ratio can earn a higher liquidity premium than one with high liquidation value in normal-states. This shows that liquidity measures derived from normal-periods alone may be poor guides to calculate realized liquidity premia when stress-state effect. differs sharply across assets.

Related literature. Our paper contributes to four strands of research. First, it relates to corporate liquidity management, where cash and liquid assets serve as inventory to meet stochastic transaction demand (Baumol, 1952; Tobin, 1956; Miller and Orr, 1966). A comprehensive review of this field is provided by Bolton et al. (2024). In modern financial contexts, this demand for liquidity is driven by interim obligations such as redemption demands in mutual funds (Chernenko and Sunderam, 2016; Goldstein et al., 2017) or the need to roll over short-term debt (He and Xiong, 2012).We extend this inventory approach by considering a range of assets with varying liquidation values and evaluating the marginal value of liquidity through shadow prices, regardless of whether the agent currently holds those assets.

Second, we contribute to the literature on liquidity management under funding constraints and dynamic portfolio selection with transaction costs. Holmström and Tirole (2001) analyze a firm's willingness to pay a premium for liquidity within an agency framework; their result emerges as a special case of our portfolio-based approach. The effective transaction cost in our setting is captured endogenously by the gap between the shadow price of an asset and its liquidation value. Jansen and Werker (2022) empirically define the shadow cost of illiquidity as the reduction in expected returns an investor accepts to convert an illiquid asset into a liquid one, finding that these costs are large for short-horizon investors and vary across asset classes. Our paper provides the theoretical foundation for these findings by showing that the shadow price is a state-contingent process. While much of the existing literature derives closed-form solutions by imposing specific functional forms on utility and asset payoffs (Constantinides, 1986; Vayanos, 1998), we value assets based on their state-contingent role in meeting future liquidity needs without such restrictions. Ang et al. (2014) characterize the welfare cost borne by an investor who cannot trade continuously in an illiquid asset; our multiperiod result characterizes the complementary option value preserved by retaining liquid assets. Proposition 5 makes this option value precise by identifying the persistence threshold at which the insurance value of the buffer exceeds its current-period liquidation advantage, and our calibration confirms this threshold is breached at empirically plausible persistence levels.

Third, we connect to the corporate finance literature on liability structure and bank fragility. Firms can mitigate liquidity risk by matching illiquid assets with stable liabilities, such as core deposits or long-term relationship lending, which are less sensitive to aggregate shocks (Berlin and Mester, 1999; Hanson et al., 2015; Doerr, 2024). Our liability design result shows that shock distributions concentrated away from the liquidation thresholds maximize expected continuation value under risk aversion, formalizing this patient-investor view. Conversely, Huang and Ratnovski

(2011) identify wholesale funding as a source of fragility. In our model this corresponds to a shock distribution with heavy upper-tail mass that creates the dark side of debt effects: the firm is repeatedly pushed toward the worst liquidation thresholds.

Fourth, the paper connects to the literature on liquidity risk and asset pricing and to the optimal liquidation literature. Pástor and Stambaugh (2003) demonstrate that stock returns vary cross-sectionally with fluctuations in aggregate market liquidity. Our framework highlights the value of liquid assets in states of tight aggregate liquidity, where demand is high and other liquid resources are scarce, providing a structural interpretation of their finding. While Acharya and Pedersen (2005) propose a linear factor model for liquidity risk, our model generates a strictly convex relationship between liquidation value and asset prices. The convexity arises through the optimal stopping rule rather than clientele effects, as in Amihud and Mendelson (1986). On the liquidation side, traditional views suggest investors should exhaust their most liquid assets first (Ma et al., 2022). Our multiperiod theorem shows this pecking order fails once future liquidity risk is accounted for, and Proposition 5 delivers a closed-form persistence threshold at which the reversal occurs; the calibration of Section 4 shows this threshold is comfortably exceeded at empirically observed shock-persistence levels. This connects to the dual role of asset safety and liquidity documented by Geromichalos et al. (2023). Pegoraro et al. (2025) study liquidity transformation in interval funds; our framework provides a tractable structural model for the portfolio behavior they document.

The remainder of the paper is organized as follows. Section 2 presents the three-date model, derives the optimal liquidation policy, the asset pricing formula, the comparative statics, the liquidity beta decomposition, and the liability design application. Section 3 extends the framework to a general multiperiod economy and proves the shadow-price theorem. Section 4 provides a quantitative calibration of the theoretical model. Section 5 concludes. The Appendix contains all proofs.

2 Model

We begin with a simple three-date economy that isolates the core mechanism. Section 3 extends the framework to a general multi-period trading economy, of which the present section is a special case.

2.1 Environment and timing

There are three dates, $t = 0$, $t = 1$, and $t = 2$. At $t = 0$ an agent assembles a portfolio $b = (b_1, \ldots, b_n) \in \mathbb{R}^n_+$, where b_i denotes the number of units held in asset i. At $t = 1$ the agent observes a random liquidity shock $y \geq 0$ and chooses a vector of sales $s = (s_1, \ldots, s_n)$. Each asset i can be sold at $t = 1$ at the interim liquidation value $L_i > 0$, raising $L_i s_i$ units of liquidity, and yields a terminal payoff per retained unit equal to V_i at $t = 2$.

The agent's objective is to choose a portfolio at $t = 0$ and a liquidation policy at $t = 1$ to maximize expected utility of terminal wealth. At $t = 0$ the agent buys assets at market prices. At $t = 1$ the agent observes the liquidity shock y and sells a subset of the portfolio to meet it, retaining the rest until $t = 2$. The interim liquidation value $L_i > 0$ is the price at which asset i can be sold in the secondary market at $t = 1$; it may be random and heterogeneous across assets, and is known at $t = 1$.

To isolate the role of liquidity shocks, we work conditional on the date-0 information set. The state at $t = 1$ consists of the realization of y together with the realized cross section of asset values; throughout, we treat asset values as given within the $t = 1$ state and focus on how variation in y moves the cutoff. The random variable y represents an interim obligation such as a margin call, a redemption demand, or a scheduled debt payment that must be met from asset sales at $t = 1$.

Assumption 1 (Strict feasibility)**.** *The liquidity shock satisfies $y \geq 0$ almost surely and $y < \sum_{i=1}^{n} b_i L_i$ almost surely.*

Strict feasibility rules out the boundary case in which the liquidity requirement equals the total liquidation capacity of the portfolio. As shown in the appendix, the boundary $y = \sum_i b_i L_i$ is the only case that generates an unbounded minimizer set in the dual problem, so strict feasibility is exactly what is needed to guarantee that the shadow price of liquidity is finite.

2.2 Optimal liquidation

The liquidation problem at $t = 1$ is to choose sales $s \in [0, b]$ to maximize retained terminal value subject to meeting the liquidity requirement:

$$\max_{s \in [0,b]} \sum_{i=1}^{n} (b_i - s_i) V_i \quad \text{subject to} \quad \sum_{i=1}^{n} L_i s_i \geq y. \tag{1}$$

Here V_i denotes the terminal payoff per unit of asset i retained to $t = 2$, which is known at $t = 1$. [2]

Assumption 2 (Distinct liquidation ratios)**.** *The liquidation-to-continuation ratios satisfy $L_i/V_i \neq L_j/V_j$ for all $i \neq j$ almost surely. Assets are indexed (state by state) so that*

$$\frac{L_1}{V_1} > \frac{L_2}{V_2} > \cdots > \frac{L_n}{V_n},$$

that is, in decreasing order of liquidation attractiveness.

The assumption of distinct ratios is without essential loss of generality. It ensures that the optimal liquidation threshold is unique for every realization of y, which simplifies the statement of results in this section. It is not imposed in the general model of Section 3.

The objective preserves as much terminal value as possible while meeting the liquidity need y through asset sales. Because the V_i are heterogeneous, each unit of asset i that is sold sacrifices V_i units of continuation value while raising L_i units of liquidity. The efficient tradeoff is to sell assets in decreasing order of L_i/V_i: those with the highest liquidity per unit of foregone continuation value go first. Since problem (1) is linear in s, the optimum is a corner solution with this threshold structure.

[2] Because this is a three-date model there is no discounting between $t = 1$ and $t = 2$, so V_i coincides with the $t = 2$ payoff; the multi-period extension in Section 3 replaces V_i with a discounted continuation value.

For $x \geq 0$, define the dual objective

$$\phi(x; b, y) := \sum_{i=1}^{n} b_i \max\{V_i,\ xL_i\} - xy. \tag{2}$$

The function ϕ is convex and piecewise linear in x, with breakpoints at $x_k := V_k/L_k$ for $k = 1, \ldots, n$. Equivalently, define the liquidation ratio $\ell_i := L_i/V_i$, so that $x_k = 1/\ell_k$. The variable x is the shadow price of liquidity: it measures the marginal cost, in continuation-value units, of relaxing the liquidity requirement by one unit. Define the argmin correspondence

$$X^*(b, y) := \arg\min_{x \geq 0}\ \phi(x; b, y).$$

The subgradient optimality condition for ϕ yields a threshold characterization: the breakpoint $x_k = V_k/L_k$ belongs to $X^*(b, y)$ if and only if

$$\sum_{i=1}^{k-1} b_i L_i \ \leq\ y\ \leq\ \sum_{i=1}^{k} b_i L_i. \tag{3}$$

When y is continuously distributed, both inequalities in (3) are strict for almost every realization and $X^*(b, y)$ is generically a singleton. When y is discrete, both inequalities may bind simultaneously, making $X^*(b, y)$ a nondegenerate interval. In that case the shadow price of liquidity is set valued even though the optimal allocation is pinned down. To obtain a single-valued shadow price and a single-valued pricing kernel, we select the infimum of the argmin set.

Definition 1 (Minimal supporting multiplier and cutoff ratio)**.** *The minimal supporting multiplier is*

$$\bar{x}(b, y) := \inf\ X^*(b, y),$$

and the induced liquidation cutoff ratio is

$$l^*(b, y)\ :=\ \frac{1}{\bar{x}(b, y)}.$$

At the marginal asset k, $l^*(b, y) = \ell_k = L_k/V_k$ is the liquidation-to-continuation ratio of the threshold asset. The identity $\max\{V_i, \bar{x}(b, y)L_i\} = V_i \max\{1, \ell_i/l^*(b, y)\}$ *holds for every asset i, so the cutoff ratio $l^*(b, y)$ is the natural object for expressing liquidation decisions and asset prices.*

The selector $\bar{x}(b, y)$ is the lower-endpoint supporting multiplier associated with the optimal liquidation rule. In knife-edge states where $X^*(b, y)$ is a nondegenerate interval, this convention picks the smallest multiplier consistent with optimality and thereby preserves the cutoff interpretation in terms of the marginal asset. Strict feasibility (Assumption 1) ensures $\bar{x}(b, y) < \infty$ for every realization of y, so $l^*(b, y)$ is well defined and strictly positive.

Theorem 1 (Optimal liquidation)**.** *Under Assumptions 2 and 1, the optimal liquidation policy has a threshold structure. There exists a unique index $k = k(b, y) \in \{1, \ldots, n\}$ satisfying*

$$\sum_{i=1}^{k-1} b_i L_i \; < \; y \; \leq \; \sum_{i=1}^{k} b_i L_i, \tag{4}$$

with the convention that an empty sum equals zero. The cutoff ratio is $l^*(b, y) = \ell_k = L_k/V_k$, and the optimal sales are

$$s_i^*(b, y) = \begin{cases} b_i & \text{if } i < k, \\ \alpha(b, y) & \text{if } i = k, \\ 0 & \text{if } i > k, \end{cases} \tag{5}$$

where $\alpha(b, y) = \left(y - \sum_{i=1}^{k-1} b_i L_i\right)/L_k \in [0, b_k]$ is the fraction of the marginal asset sold to exactly meet the requirement.

The proof is in Appendix A.1. It establishes the primal-dual equivalence, characterizes $X^*(b, y)$ via the subgradient condition, shows that $\bar{x}(b, y) = x_k = V_k/L_k$ is the minimal supporting minimizer under strict feasibility, and derives the liquidation rule from the dual solution.

The economic logic of Theorem 1 is transparent. The linearity of problem (1) in s forces the optimum to a corner: sell first the assets with the highest ratio L_i/V_i and stop once the requirement is met. Selling one unit of asset i raises L_i units of liquidity at the cost of V_i units of continuation value, so the tradeoff is governed by L_i/V_i, not by L_i alone. What is not immediate is the endogeneity of the cutoff ratio $l^*(b, y)$: the threshold depends jointly on the portfolio b, the realization of y, and the entire cross section of (L_i, V_i) pairs. The selection $l^*(b, y) = 1/\bar{x}(b, y)$ pins this threshold as a single-valued function of (b, y). We use this minimal-supporting selection throughout the paper as the canonical cutoff representation, including in the pricing discussion below.

The novelty of this framework relative to standard knapsack arguments lies in three features: the cutoff $l^*(b, y)$ is state dependent through the random shock y; the minimal supporting selection yields a single-valued cutoff even when y is discrete; and this selection generates a well-defined, single-valued liquidity service term in asset prices at $t = 0$.

When $y = 0$, the optimal choice is $s^*(b, 0) = 0$, and the cutoff ratio is immaterial; our selection rule continues to deliver a well-defined $l^*(b, 0)$.

2.3 Asset pricing

We roll back to $t = 0$ and embed the liquidation problem in a portfolio choice problem. The agent chooses $b \in \mathbb{R}^n_+$ to maximize expected utility of terminal wealth subject to a budget constraint:

$$\max_{b \in \mathbb{R}^n_+} \mathbb{E}\big[U(\bar{g}(b, y))\big] \quad \text{subject to} \quad \sum_{i=1}^{n} b_i P_i \leq w, \tag{6}$$

where $w > 0$ is initial wealth, P_i is the $t = 0$ price of asset i, and

$$\bar{g}(b, y) := \sum_{i=1}^{n} b_i \max\{V_i,\, \bar{x}(b, y) L_i\} - y\,\bar{x}(b, y)$$

is the terminal value of the retained portfolio under optimal liquidation. Since $l^*(b, y) = 1/\bar{x}(b, y)$, the last term may equivalently be written as $-y/l^*(b, y)$ in the cutoff-ratio representation. The utility function U is strictly increasing and strictly concave.

Let $\lambda > 0$ be the Lagrange multiplier on the budget constraint at an interior optimum b^*, and define the stochastic discount factor

$$M(y) := \frac{1}{\lambda} U'(\bar{g}(b^*, y)).$$

Theorem 2 (Asset pricing with liquidity service)**.** *Suppose b^* is an interior optimum of* (6) *with multiplier $\lambda > 0$. Then there exist a strictly positive random variable $M(y)$ and a nonnegative random variable $X(y)$ such that, for each asset i with $b_i^* > 0$,*

$$P_i = \mathbb{E}[M(y) \max\{V_i, L_i X(y)\}] . \tag{7}$$

Moreover, $X(y)$ may be chosen to take values in the endpoint set of minimizers of $\phi(\cdot; b^*, y)$. If $X^*(b^*, y)$ is a singleton almost surely, then $X(y) = \bar{x}(b^*, y) = 1/l^*(b^*, y)$ almost surely, and (7) *reduces to*

$$P_i = \mathbb{E}\left[M(y) \max\left\{V_i, \frac{L_i}{l^*(b^*, y)}\right\}\right] . \tag{8}$$

The proof is in Appendix A.2. In singleton-cutoff states the pricing kernel is uniquely pinned down by the optimal cutoff. When $X^*(b^*, y)$ is a nondegenerate interval, the first-order condition admits multiple valid subgradient selections; the theorem therefore delivers the pricing rule in the form (7), with $X(y)$ supported on the endpoint set of minimizers.

Remark 1 (Canonical minimal-supporting representation)**.** *Throughout the paper we use the minimal supporting selection $\bar{x}(b^*, y) = \inf X^*(b^*, y)$, equivalently $l^*(b^*, y) = 1/\bar{x}(b^*, y)$, as the canonical single-valued cutoff representation. In the generic singleton case this coincides with the unique pricing kernel. In interval-minimizer states it provides the lower-endpoint single-valued representation used in the examples, comparative statics, and calibration.*

To simplify notation in what follows, we write $l^*(y) := l^*(b^*, y)$ and use the minimal-supporting selection as the canonical single-valued pricing representation. Under that representation, the pricing formula (8) reflects a liquidity service payoff. Asset i delivers not simply V_i units of continuation value at $t = 2$ but also, in states where it is drawn into the liquidation set, liquidity at $t = 1$. The kernel $\max\{V_i, L_i/l^*(y)\}$ is the total service value per unit. In states where $L_i/V_i < l^*(y)$, the asset is below the cutoff and is not sold; it contributes only its continuation value V_i. In states where $L_i/V_i \geq l^*(y)$, the asset is at or above the cutoff; each unit provides $L_i/l^*(y)$ units of liquidity, which exceeds the foregone continuation value V_i because $L_i/l^*(y) = V_i(L_i/V_i)/l^*(y) \geq V_i$ (equivalently, the positive-part term $(L_i/l^*(y) - V_i)^+$ is strictly positive in these states). The endogenous cutoff $l^*(y)$ also implies that the liquidity premium is a convex function of the shock y. As y increases, the cutoff $l^*(y)$ falls and the agent is forced to liquidate assets with progressively lower L_i/V_i ratios; the marginal cost of meeting an additional unit of liquidity need therefore rises as y rises. While Amihud and Mendelson (1986) derive a concave relationship between returns and illiquidity in an equilibrium with heterogeneous investor holding horizons, our model generates a convex premium within a single representative-agent framework. The convexity arises from the optimal stopping rule: the marginal asset sacrificed at cutoff $l^*(y)$ has a progressively worse liquidity-to-continuation ratio as the shock intensifies, so the per-unit cost of liquidity is increasing in y.

2.4 Comparative statics

In a one-period interim shock environment, the cutoff ratio $l^*(b, y)$ responds to changes in the liquidity environment in intuitive ways. We record two results; proofs are in Appendix A.3.

Proposition 1 (Monotonicity in the liquidity requirement)**.** *Fix* (b, L, V). *If* $y_1 < y_2$, *then* $l^*(b, y_2) \leq l^*(b, y_1)$; *equivalently,* $\bar{x}(b, y_2) \geq \bar{x}(b, y_1)$.

A larger shock forces the agent deeper into the liquidation set, drawing in assets with lower L_i/V_i ratios. The cutoff l^* falls and the shadow price $\bar{x}$ rises. An immediate consequence is that the total service value $\max\{V_i, L_i/l^*(y)\}$ is weakly increasing in y for every asset i: assets provide more liquidity service in states of high aggregate demand.

Proposition 2 (Improving liquidation values raises the cutoff)**.** *Fix* (b, y, V) *and an initial profile* L *with induced cutoff* $l^*(b, y) = L_k/V_k$. *Consider an alternative profile* $\widetilde{L}$ *with* $\widetilde{L}_i \geq L_i$ *for all* $i \leq k$, *strict inequality for at least one such i, and with the relative ordering of ratios $\widetilde{L}_i/V_i$ preserved. Then*

$$l^*(b, y; \widetilde{L}) \;\geq\; l^*(b, y; L);$$

equivalently, $\bar{x}(b, y; \widetilde{L}) \leq \bar{x}(b, y; L)$.

When the assets with high liquidation ratios become more liquid, the same requirement can be met by selling a set whose marginal member has a higher ratio; the cutoff rises and the shadow price falls. This clarifies how improvements in secondary market quality reduce the cost of interim financing through the threshold structure.

Example 1: Convex pricing from dispersion in the liquidity shock

This example shows how greater dispersion in the liquidity shock y, holding the mean fixed, translates into stronger convexity of prices in liquidation value. The mechanism is the endogenous threshold $l^*(y)$: a wider support for y generates a wider range of realized cutoffs, exposing high-liquidation-value assets to more states in which they provide liquidity service.

Consider $n = 5$ assets with liquidation values

$$(L_1, L_2, L_3, L_4, L_5) = (50,\ 100,\ 200,\ 300,\ 400),$$

with the agent holding one unit of each asset, $b_i \equiv 1$. The terminal payoff is normalized to $V = 1000$ per unit. Under the risk-neutral benchmark $M(y) \equiv 1$, the pricing formula (8) reduces to

$$P_i = \frac{V}{\lambda}\,\mathbb{E}\left[\max\left\{1, \frac{L_i}{L^*(y)}\right\}\right].$$

Define the normalized price index $p(L) := \mathbb{E}[\max\{1, L/L^*(y)\}]$. In this example the realized cutoffs are singleton-valued, so the canonical minimal-supporting representation coincides with the unique pricing kernel. For each realization of y, the map $L \mapsto \max\{1, L/L^*(y)\}$ is convex and piecewise linear in L, so $p(L)$ is convex as an expectation.

We compare three distributions of y, each with mean $\mathbb{E}[y] = 800$, chosen so that the induced thresholds take values in $\{400, 300, 200, 100, 50\}$.

In Case A, y takes values 600 and 1000 with equal probability. The corresponding thresholds are $L^*(600) = 300$ and $L^*(1000) = 100$. Computing the expectation yields

$$p(50) = 1, \quad p(100) = 1, \quad p(200) = \tfrac{3}{2}, \quad p(300) = 2, \quad p(400) = \tfrac{8}{3}.$$

In Case B, y takes values 500, 850, and 1050 with equal probability, giving mean 800. The thresholds are $L^*(500) = 300$, $L^*(850) = 200$, and $L^*(1050) = 50$. The price index becomes

$$p(50) = 1, \quad p(100) = \tfrac{4}{3}, \quad p(200) = 2, \quad p(300) = \tfrac{17}{6}, \quad p(400) = \tfrac{34}{9}.$$

In Case C, y takes values 400, 700, 850, 1000, and 1050 with equal probability, again yielding mean 800. Each asset becomes the marginal asset in exactly one state. The price index is

$$p(50) = 1, \quad p(100) = \tfrac{6}{5}, \quad p(200) = \tfrac{9}{5}, \quad p(300) = \tfrac{5}{2}, \quad p(400) = \tfrac{49}{15}.$$

Across the three cases, the slope of $p(L)$ steepens substantially as dispersion increases, even though the mean liquidity requirement is unchanged at 800. This is because greater dispersion in y generates a wider range of realized cutoffs, exposing high-L assets to more states in which they provide liquidity service.

To pin down the level of prices in a one-agent competitive equilibrium with unit supplies, market clearing requires $b_i^* = 1$ for all i. Substituting into the pricing formula and solving for the scale factor gives

$$P_i = w\,\frac{p(L_i)}{\sum_{j=1}^{n} p(L_j)}.$$

In Case A ($\sum_i p(L_i) = 49/6$),

$$(P_{50}, P_{100}, P_{200}, P_{300}, P_{400}) = w\left(\tfrac{6}{49}, \tfrac{6}{49}, \tfrac{9}{49}, \tfrac{12}{49}, \tfrac{16}{49}\right).$$

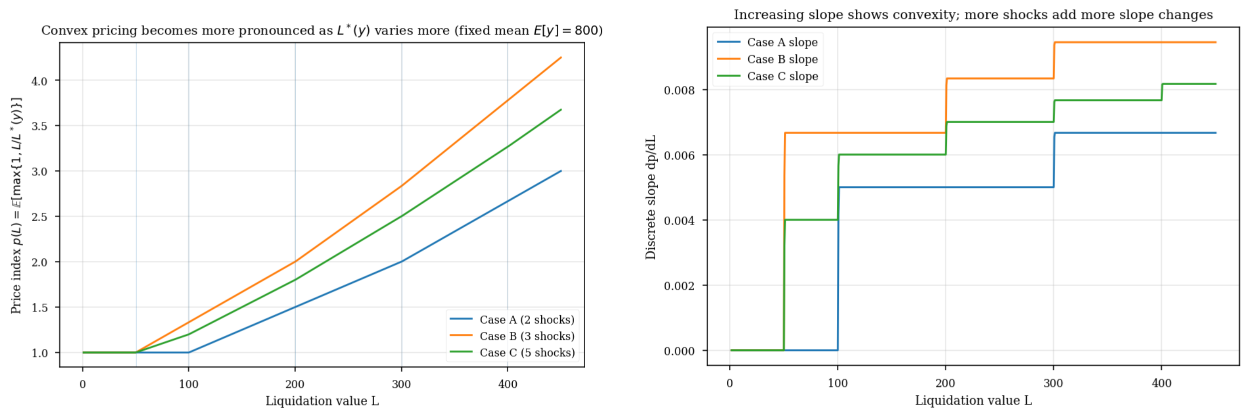

(a) **Example 1a.** The figure plots the price index $p(L) = \mathbb{E}[\max\{1, L/L^*(y)\}]$ from Example 1 as a function of the liquidation value L. The horizontal axis reports L, and the vertical axis reports the corresponding value of $p(L)$. Each curve corresponds to a different discrete distribution of the liquidity requirement y with the same mean.

(b) **Example 1b.** The figure plots a discrete slope approximation to $\mathrm{d}p(L)/\mathrm{d}L$ across the same grid of liquidation values L as in panel (a). The horizontal axis reports L, and the vertical axis reports the corresponding discrete slope of the price index.

Figure 1: Price index and slope in Example 1.

In Case B ($\sum_i p(L_i) = 197/18$),

$$(P_{50}, P_{100}, P_{200}, P_{300}, P_{400}) = w\left(\frac{18}{197}, \frac{24}{197}, \frac{36}{197}, \frac{51}{197}, \frac{68}{197}\right).$$

In Case C ($\sum_i p(L_i) = 293/30$),

$$(P_{50}, P_{100}, P_{200}, P_{300}, P_{400}) = w\left(\frac{30}{293}, \frac{36}{293}, \frac{54}{293}, \frac{75}{293}, \frac{98}{293}\right).$$

Example 2: Multiple optimal portfolios

This example illustrates two features of the date-$t = 0$ portfolio choice that are not apparent from the three-date model's structure alone: the optimality of holding many assets under random liquidity needs, and the possibility that multiple portfolios attain the same objective value at a fixed price vector.

Consider $n = 9$ bonds, each paying terminal value $V = 500$ at $t = 2$. Table 1 reports the liquidation value and purchase price for each bond. The agent has initial wealth $w = 1209$ at $t = 0$, and the liquidity requirement y takes values 400, 700, 1000, 1025, and 1050 with equal probability.

Table 1: Bond characteristics in Example 2.

Bond i	Liquidation value L_i	Purchase price P_i
1	30	100
2	80	150
3	120	160
4	160	230
5	240	260
6	280	305
7	300	315
8	320	360
9	360	374

Solving the date $t = 0$ program under the given prices yields the following two optimal portfolios, among others. The first is a laddered allocation that holds one unit each of bonds at indices $\{1, 3, 5, 7, 9\}$, corresponding to liquidation values 30, 120, 240, 300, and 360, with zero holdings elsewhere:

$$b_1 = b_3 = b_5 = b_7 = b_9 = 1, \qquad b_i = 0 \text{ for } i \notin \{1, 3, 5, 7, 9\}.$$

The second is a sparse fractional allocation:

$$b_1 = 0.963611, \quad b_3 = 0.208333, \quad b_5 = 1.354167, \quad b_9 = 1.944444, \quad b_i = 0 \text{ otherwise.}$$

Both portfolios attain the same value of the date $t = 0$ objective.

The laddered structure reflects a diversification motive that arises specifically from random liquidity needs. Each realization of y induces a different threshold $L^*(b, y)$, so a bond with an intermediate liquidation value earns a liquidity premium precisely by being the marginal asset in intermediate shock states. Spreading holdings across the liquidation value spectrum ensures that the marginal asset at each realization of y is always a bond already in the portfolio, so no liquidity service value is foregone.

The multiplicity of optima arises because, with discrete shocks, $\bar{g}(b, y)$ is piecewise linear in b for each y, and the date-$t = 0$ objective can be flat across portfolios that induce the same liquidation thresholds and expected losses. The price vector therefore supports an entire set of optimal portfolios rather than a unique one. Importantly, the liquidation policy at $t = 1$ remains well defined state by state for every portfolio in this set: the cutoff $L^*(b, y)$ is pinned down by b and the realization of y through Theorem 1, and the pricing representation of Theorem 2 continues to hold for any portfolio in the optimal set, with the minimal-supporting selection providing the canonical cutoff representation used here.

2.5 Application: Liquidity Beta

Under the canonical minimal-supporting representation of Theorem 2, the pricing formula (8) implies a decomposition that separates each asset's return into a pure continuation component and a liquidity service component. Writing $a^+ := \max\{a, 0\}$,

$$P_i = \mathbb{E}[M(y)\,V_i] + \mathbb{E}\left[M(y)\left(\frac{L_i}{l^*(y)} - V_i\right)^+\right]. \tag{9}$$

The first term is the present value of the continuation payoff. The second term is the *liquidity premium*: the excess service the asset provides in states where it enters the liquidation set. Here $\left(\frac{L_i}{l^*(y)} - V_i\right)^+$ is positive exactly in the states where asset i enters the liquidation set, equivalently when $L_i/V_i \geq l^*(y)$. We define asset i's *liquidity beta* as the function of its liquidation value L_i:

$$\beta^{\text{liq}}(L_i) \;:=\; \mathbb{E}\left[M(y)\left(\frac{L_i}{l^*(y)} - V_i\right)^+\right], \tag{10}$$

so that $P_i = \mathbb{E}[M(y)V_i] + \beta^{\text{liq}}(L_i)$. Writing the beta as $\beta^{\text{liq}}(L_i)$ makes explicit that the cross-sectional dispersion in prices is driven by dispersion in liquidation values: assets with higher L_i enter the liquidation set more often and in more states, generating a larger premium. The beta is increasing in L_i (holding V_i fixed) because a higher liquidation value both widens the set of states in which asset i is drawn above the cutoff and raises the payoff $L_i/l^*(y)$ conditional on being there.

Proposition 3 (Liquidity beta ordering)**.** *Fix b^* and two assets i and j. If $\beta^{\text{liq}}(L_i) \geq \beta^{\text{liq}}(L_j)$, then $P_i \geq P_j$.*

What distinguishes this ordering from a simple ranking by L_i or V_i is the role of the endogenous cutoff $l^*(y)$. A high ratio L_i/V_i commands a large premium only if the asset is drawn into the liquidation set in states where liquidity is scarce — states where $l^*(y)$ is low and $\bar{x}(y)$ is high. An asset whose high ratio is relevant only in states of abundant liquidity provides little incremental service. The endogeneity of the cutoff generates a state-dependent liquidity beta $\beta^{\text{liq}}(L_i)$ even when the cross section of (L_i, V_i) pairs is deterministic.

The decomposition (9) allows us to distinguish sharply between two properties of an asset that are often conflated. Following the dual role of assets in Geromichalos et al. (2023) and the market-to-funding liquidity interaction of Brunnermeier and Pedersen (2009), we map them to two distinct primitives:

- **Liquidity:** The static ratio L_i/V_i. High liquidity means an asset can be converted to cash with minimal sacrifice of continuation value; it governs whether the asset enters the liquidation set for a given $l^*(y)$.

- **Safety (Liquidity Beta):** The dynamic sensitivity $\text{Cov}(L_i/V_i,\ \bar{x}(y))$. An asset is "safe" if its liquidation-to-continuation ratio remains robust when the shadow price of liquidity $\bar{x}(y)$ is highest, i.e., when the shock y is most severe.

Assets with liquidation ratios negatively correlated with $\bar{x}(y)$ — those whose L_i/V_i ratios collapse precisely when $l^*(y)$ is most demanding — earn lower returns because they fail to provide safety when it is most valuable.

Our endogenous liquidity beta contrasts with the exogenous-friction approach of Acharya and Pedersen (2005), where liquidity risk is priced via assumed transaction costs. Here the premium arises from the agent's optimal liquidation strategy: the market price of liquidity risk is the shadow price $\bar{x}(b, y)$ itself, which is determined in equilibrium by the portfolio and the shock distribution. The framework also formalises the "liquidity as insurance" intuition of Holmström and Tirole

(2001): the liquidity service in (9) pays off precisely when internal resources are insufficient to cover y, and the portfolio acts as a ladder of liquidity options whose exercise sequence is governed by the state-contingent cutoff $l^*(y)$. This yields a clear empirical implication for the cross-section of returns: during liquidity crunches, the widening of spreads is driven by the interaction between an asset's static liquidation ratio and its dynamic safety profile.

2.6 Application: Liability Design and Firm Value

The distribution of y is partly a choice variable: a firm's funding structure, the maturity of its liabilities, and the availability of committed credit lines all affect the distribution of interim liquidity needs. We abstract from institutional details and treat the distribution f over y as the firm's instrument, holding the portfolio b^* fixed. Let $\mathbb{E}_f[\,\cdot\,]$ denote expectations under f. Under the canonical minimal-supporting representation, the market value of the firm is

$$\mathcal{V}(f) := \sum_{i=1}^{n} b_i^* P_i(f) = \mathbb{E}_f[M(y)\,H(y)]\,,$$

where

$$H(y) := \sum_{i=1}^{n} b_i^* \max\left\{V_i,\ \frac{L_i}{l^*(y)}\right\}$$

aggregates the total service values across the portfolio. By Proposition 1, $l^*(y)$ is weakly decreasing in y, so $H(y)$ is weakly increasing: a larger shock draws more assets into the liquidation set, expanding the aggregate service provided.

Proposition 4 (Liability structure and firm value: risk-neutral benchmark)**.** *Suppose U is linear, so $M(y)$ is constant. Then $\mathcal{V}(f) \propto \mathbb{E}_f[H(y)]$. If f first-order stochastically dominates g, then $\mathcal{V}(f) \geq \mathcal{V}(g)$.*

The result follows from the monotonicity of H: a distribution placing more mass on larger values of y yields higher expected service and thus higher firm value. As y rises, more assets are drawn into the liquidation set, each providing liquidity service $L_i/l^*(y)$ that exceeds its continuation value V_i.

Holding the mean liquidity need fixed, a funding structure with more upper-tail mass, such as short-term wholesale funding with rollover risk, generates higher $\mathcal{V}$ in this benchmark because it more frequently activates the liquidation service of high-ratio assets. A retail deposit base, by contrast, produces a distribution of y concentrated near its mean; it activates the liquidation service less often and produces lower firm value. The direction of the effect is reversed under risk aversion, where marginal utility is high precisely in states of large y, creating a precautionary motive for smoother liability structures.

The design of the liability structure is therefore a primary determinant of the firm's exposure to the liquidation boundary $l^*(y)$. In the language of Proposition 4, shifting f toward distributions with lower mean and less upper-tail mass reduces the probability of crossing the thresholds $\{S_k\}$ at which $H(y)$ jumps, insulating continuation values V_i from forced realisation. Berlin and Mester (1999) show that "core" liabilities such as retail deposits provide exactly this stabilising role, enabling the funding of illiquid, high-L/V assets without exposing the portfolio to forced liquidation. In our framework, such a liability structure lowers the sensitivity of f to aggregate shocks, reduces the frequency with which $l^*(y)$ is breached, and thereby raises expected $\mathcal{V}$ under risk aversion. Geographical diversification of the funding base reinforces this effect by reducing the co-movement of deposit outflows with aggregate liquidity shocks (Doerr, 2024), which in our model corresponds to reducing the variance of y and shifting f toward the precautionary-optimal distribution.

Conversely, relying on concentrated wholesale funding creates what Huang and Ratnovski (2011) call the "dark side" of short-term debt: the liability structure itself becomes a source of high-beta liquidity risk. In our framework, this corresponds to a distribution f with heavy upper-tail mass, which pushes the firm repeatedly toward the liquidation boundary and forces the sale of assets with progressively lower L_i/V_i ratios — exactly the poor-safety assets that earn low $\beta^{\text{liq}}(L_i)$ premia. Hanson et al. (2015) argue that banks function as patient investors by matching illiquid, high-continuation-value assets with "quiet" liabilities that are unlikely to withdraw during stress; in the model, this corresponds to choosing f concentrated away from the thresholds $\{S_k\}$, preserving the portfolio's option to hold assets to maturity and protecting the long-run continuation values V_i.

Remark 2. *Under general U, the effect of a mean-preserving spread in y on $\mathcal{V}(f) = \mathbb{E}_f[M(y)H(y)]$ depends on the curvature of $y \mapsto M(y)H(y)$. Because $H(y)$ is a step function in y with jumps at the threshold levels $\{S_k := \sum_{i=1}^{k} b_i^* L_i\}$, it is neither convex nor concave in any smooth sense. A mean-preserving spread result therefore requires either a smoothed approximation to H or an explicit curvature condition on $M(y)H(y)$ stated region by region. The risk-neutral first-order stochastic dominance result is the clean, unconditional statement available in this model.*

The three-date framework isolates the core mechanism — optimal threshold liquidation, the endogenous shadow price of liquidity, and the resulting cross-sectional pricing implications — in the simplest possible setting. Section 3 embeds this structure in a general multi-period trading economy, where continuation values are replaced by discounted future payoffs and the cutoff ratio $l^*(y)$ evolves endogenously alongside the agent's portfolio.

3 A Multi-period Model with Trading

The three-date framework captures the essential mechanism, but the shadow-pricing result extends naturally to a multi-period economy in which the agent can trade at every date.

3.1 Environment and trading

We consider a finite-horizon economy with dates $t = 0, 1, \ldots, T$ on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t=0}^{T}, \mathbb{P})$. There are n tradable assets indexed by $i \in \{1, \ldots, n\}$. The agent enters at $t = 0$ with an initial portfolio $\theta_{-1} = (\theta_{1,-1}, \ldots, \theta_{n,-1})$, where $\theta_{i,-1} \geq 0$ for all i.

At each date $t \leq T - 1$, the agent observes the $\mathcal{F}_t$-measurable liquidity shock $y_t \in \mathbb{R}$ and then trades. Trading at date t consists of purchases $b_{i,t} \geq 0$ and sales $\ell_{i,t} \geq 0$, both $\mathcal{F}_t$-measurable, with resulting next-period holdings $\theta_{i,t} \geq 0$ satisfying the accounting identity

$$\theta_{i,t} = \theta_{i,t-1} + b_{i,t} - \ell_{i,t}, \qquad i = 1, \ldots, n, \quad t = 0, \ldots, T - 1. \tag{11}$$

Asset i can be purchased at date t at the adapted purchase price $P_{i,t} \geq 0$ and sold at the adapted liquidation price $L_{i,t} \geq 0$. The pair $(P_{i,t}, L_{i,t})$ is $\mathcal{F}_t$-measurable for each i and t. These prices may be stochastic and heterogeneous across assets; all trading decisions are adapted. At the terminal date, $L_{i,T}$ is the payoff per unit of asset i, and terminal wealth is

$$w_T = \sum_{i=1}^{n} \theta_{i,T-1} L_{i,T}. \tag{12}$$

The liquidity shock must be covered each period by net trading proceeds:

$$y_t \;\leq\; \sum_{i=1}^{n} \big(\ell_{i,t} L_{i,t} - b_{i,t} P_{i,t}\big), \qquad t = 0, \ldots, T-1. \tag{13}$$

The agent chooses an admissible policy $\pi := \{(b_t, \ell_t, \theta_t)\}_{t=0}^{T-1}$ to maximize expected utility of terminal wealth,

$$\max_{\pi} \; \mathbb{E}\big[u(w_T)\big],$$

where $u : \mathbb{R}_+ \to \mathbb{R}$ is strictly increasing, strictly concave, and continuously differentiable. We assume an optimal policy π^* exists.

Assumption 3 (Integrability)**.** *Under admissible policies, $w_T \in L^2(\mathbb{P})$ and $\mathbb{E}[u(w_T)]$ is well defined and finite.*

Assumption 4 (No immediate arbitrage)**.** *For all i and $t \leq T-1$, $L_{i,t} \leq P_{i,t}$ almost surely.*

The no-immediate-arbitrage assumption rules out the possibility of purchasing and immediately reselling any asset at a profit. It ensures the shadow price of each asset lies between its liquidation and purchase price.

Assumption 5 (Dynamic strict feasibility)**.** *For each $t \leq T-1$, almost surely at each date-t state, if $y_t > 0$ then*

$$y_t \;<\; \sum_{i=1}^{n} \theta_{i,t-1} \, L_{i,t}.$$

Dynamic strict feasibility says that when the shock is positive, the agent can meet it without exhausting the entire portfolio. The condition is evaluated on the realized $(L_{i,t}, \theta_{i,t-1}, y_t)$ within each date-t state. As in the three-date model, this rules out the boundary case in which the liquidity constraint forces complete liquidation, which is the only case that would generate an unbounded multiplier on (13).

When y_t has atoms, the set of Lagrange multipliers consistent with optimality on (13) may be an interval rather than a singleton, exactly as in the three-date framework. To obtain single-valued shadow prices, we select the *minimal supporting multiplier* at each date and state. This is the dynamic analogue of the selection $\bar{x}(b, y) = \inf X^*(b, y)$ used in the three-date model (Section 2.2). The selection is defined precisely in Appendix A.7; Assumption 5 guarantees it is finite almost surely.

3.2 Shadow prices and the pricing theorem

The main result of this section characterizes the optimal trading policy through a multiplier process (ξ_t) and the associated discounted shadow prices $(\rho_{i,t})$. On states where $\xi_t > 0$, the undiscounted shadow prices $\phi_{i,t} := \rho_{i,t}/\xi_t$ are additionally well defined. Under the minimal-supporting convention, $\rho_{i,t}$ is the primary globally defined dynamic object; $\phi_{i,t}$ is used only in statements restricted to $\{\xi_t > 0\}$.

Theorem 3 (State-price density and shadow prices)**.** *Suppose $\pi^* = \{(b_t^*, \ell_t^*, \theta_t^*)\}_{t=0}^{T-1}$ is an optimal policy with terminal wealth w_T^*. Under Assumptions 3–5, the liquidity constraint* (13) *binds at every date:*

$$y_t = \sum_{i=1}^{n} \big(\ell_{i,t}^* L_{i,t} - b_{i,t}^* P_{i,t}\big), \qquad t = 0, \ldots, T-1.$$

Moreover, there exist an adapted nonnegative process $(\xi_t)_{t=0}^{T}$ and, for each asset i, an adapted nonnegative process $(\rho_{i,t})_{t=0}^{T}$, the discounted shadow prices*, with*

$$\xi_T = u'(w_T^*), \qquad \rho_{i,T} = \xi_T L_{i,T},$$

with $\xi_T > 0$ almost surely, and for each $t \le T-1$, $\xi_t > 0$ almost surely on the event $\{y_t > 0\}$. The discounted shadow prices satisfy, for each i and $t = 0, \dots, T-1$, the global recursion

$$\rho_{i,t} = \max\{\xi_t L_{i,t},\ \mathbb{E}_t[\rho_{i,t+1}]\}, \tag{14}$$

the global bounds

$$\xi_t L_{i,t} \ \le\ \rho_{i,t} \ \le\ \xi_t P_{i,t}, \tag{15}$$

and the following complementary-slackness conditions:

$$\begin{aligned} \ell^*_{i,t} > 0 \ &\Rightarrow\ \rho_{i,t} = \xi_t L_{i,t}, \\ b^*_{i,t} > 0 \ &\Rightarrow\ \rho_{i,t} = \xi_t P_{i,t}, \\ \theta^*_{i,t} > 0 \ &\Rightarrow\ \rho_{i,t} = \mathbb{E}_t[\rho_{i,t+1}]. \end{aligned} \tag{16}$$

The proof is in Appendix A.7. Under the minimal-supporting convention, ξ_t may equal zero on some states with $y_t = 0$, so $\rho_{i,t}$ is the primary global object. On states where $\xi_t > 0$, define the *undiscounted shadow price*

$$\phi_{i,t} := \rho_{i,t}/\xi_t.$$

Dividing (15) and (16) by ξ_t gives, on $\{\xi_t > 0\}$, the undiscounted bounds $L_{i,t} \le \phi_{i,t} \le P_{i,t}$, the recursion

$$\phi_{i,t} = \max\left\{L_{i,t},\ \frac{1}{\xi_t}\,\mathbb{E}_t\left[\xi_{t+1}\,\phi_{i,t+1}\right]\right\}, \tag{17}$$

and the undiscounted complementary-slackness conditions

$$\begin{aligned} \ell^*_{i,t} > 0 \ &\Rightarrow\ \phi_{i,t} = L_{i,t}, \\ b^*_{i,t} > 0 \ &\Rightarrow\ \phi_{i,t} = P_{i,t}, \\ \theta^*_{i,t} > 0 \ &\Rightarrow\ \phi_{i,t} = \frac{1}{\xi_t}\,\mathbb{E}_t\left[\xi_{t+1}\,\phi_{i,t+1}\right]. \end{aligned}$$

Liquidation policy in the multiperiod model. In the three-date model, the liquidation problem at the interim date reduces to a static linear program, and the optimal liquidation rule has a sharp "sell-first" ordering: assets are sold in decreasing order of their liquidation ratio $\ell_i := L_i/V_i$, with a single cutoff ℓ^* that pins down the marginal asset. In the multiperiod setting, this static ordering need not carry over. While the "pecking order" theory of liquidation suggests that investors should exhaust their most liquid assets first to minimise transaction costs (Ma et al., 2022), our multiperiod extension shows this is not a general property of optimal policy. When future shadow prices of liquidity (ξ_{t+1}) are accounted for, an agent may optimally preserve her most liquid assets to guard against future liquidity-beta risk, instead selling assets with higher *current* L/V ratios. The reason is that an asset's liquidity value is no longer summarised by its current ratio alone: retaining a liquid asset preserves the option to meet future shocks, and this option value depends on the joint outlook for shocks and secondary-market conditions.

The multiperiod liquidation policy is instead characterised by the shadow-price system. Selling one unit of asset i at date t delivers discounted liquidity value $\xi_t L_{i,t}$, while holding preserves expected discounted continuation value $\mathbb{E}_t[\rho_{i,t+1}]$. Liquidation is optimal when the former dominates the latter:

$$\rho_{i,t} = \max\left\{\xi_t L_{i,t},\ \mathbb{E}_t[\rho_{i,t+1}]\right\}, \qquad \ell^*_{i,t} > 0 \ \Rightarrow\ \rho_{i,t} = \xi_t L_{i,t}.$$

On states where $\xi_t > 0$, dividing by ξ_t yields the equivalent undiscounted comparison:

$$\phi_{i,t} = \max\left\{L_{i,t},\ \frac{1}{\xi_t}\mathbb{E}_t[\xi_{t+1}\phi_{i,t+1}]\right\}, \qquad \ell^*_{i,t} > 0 \ \Rightarrow\ \phi_{i,t} = L_{i,t}.$$

Thus, even an asset that is "most liquid" today — or even *fully liquid* in the sense that $L_{i,t}$ equals its terminal payoff for all t — need not be sold first, because its value as a future-shock buffer can exceed its immediate liquidation value. Example 3.3 provides a simple counterexample. This mechanism provides a theoretical micro-foundation for the "horizontal selling" behaviour (trading across the portfolio rather than exhausting the most liquid positions) documented empirically by

Chernenko and Sunderam (2016), and for the complex liquidation patterns in Andonov et al. (2025), where pension funds meet funding shocks by selling equities while sparing cash to maintain a buffer against future needs. It also connects to the broader evidence on liquidity management in less-than-daily-liquid vehicles; see Pegoraro et al. (2025) on liquidity transformation and portfolio behaviour in interval funds.

3.3 Precautionary Pecking-Order Reversal

To sharpen the intuition from Theorem 3, we specialize to a two-interim-period economy with dates $t = 1, 2$ and terminal date $t = 3$. There are two assets with identical terminal continuation value, normalized to one. The *buffer* asset 1 is perfectly liquid in all states, $\ell_{1,t} \equiv 1$. The *secondary* asset 2 has liquidation-to-continuation ratio $\ell_{2,1} = \ell_2^N < 1$ at $t = 1$ — its normal-state ratio — and $\ell_{2,2} \in \{\ell_2^N, \ell_2^S\}$ at $t = 2$, with $\ell_2^S < \ell_2^N$, so secondary-market conditions tighten in the stress realization. The shock $y_1 = \bar{y} > 0$ arrives at $t = 1$ almost surely, and $y_2 = \bar{y}$ arrives at $t = 2$ with conditional probability $\rho_y \in (0, 1)$ — the persistence of the funding-need process — and $y_2 = 0$ otherwise. The portfolio is assumed deep enough to meet $\bar{y}$ using either asset at $t = 1$, but asset 2 alone cannot cover shocks at both dates, so asset 1 is the unique effective buffer against a second shock.[3]

When shocks are persistent, preserving the buffer at $t = 1$ sacrifices its current liquidity advantage but avoids the sharply higher cost of meeting a second shock with a deteriorated asset 2. The following proposition identifies the break-even persistence at which this insurance motive overturns the static pecking order.

Proposition 5 (Pecking-Order Reversal)**.** *Under the risk-neutral benchmark, let Policy M sell the buffer at $t = 1$ and rely on asset 2 at $t = 2$, and let Policy D sell asset 2 at $t = 1$ and preserve the buffer for $t = 2$. Policy D yields strictly higher expected terminal wealth than Policy M if and only*

[3] This is the minimal feasibility configuration underlying the two-policy comparison; it is satisfied by the calibration in Section 4.

if

$$\rho_y \ > \ \rho_y^* \ := \ \frac{1/\ell_2^N - 1}{1/\ell_2^S - 1}. \tag{18}$$

The proof is in Appendix A.8.

Whenever $\rho_y > \rho_y^*$, the static pecking order "sell the most liquid asset first" fails dynamically: it is optimal to sell the *less* liquid asset at $t = 1$ in order to preserve the buffer as insurance against the potential second shock. The threshold ρ_y^* is decreasing in the stress ratio ℓ_2^S/ℓ_2^N, so the reversal region is large when shocks are persistent *and* when secondary-market liquidity deteriorates sharply under stress.

Proposition 5 is a concrete implication of Theorem 3: liquidation of asset i is optimal at date t only when its immediate liquidation payoff dominates the risk-adjusted continuation value. Preserving the buffer at $t = 1$ raises that continuation value because it retains the option to meet the potential $t = 2$ shock without paying the stress-state excess cost $C^S := 1/\ell_2^S - 1$. The threshold ρ_y^* is the break-even persistence at which this option value exactly equals the buffer's current-period liquidation advantage.

Example 3: Sell-most-liquid-first can fail

We provide a minimal numerical counterexample showing that the static heuristic "sell the most liquid asset first" is not optimal in the multiperiod model, even when the most liquid asset is *fully* liquid in the sense that its liquidation value equals its terminal payoff at every date.

Setup. There are two interim dates $t = 1, 2$ and a terminal date $t = 3$. The agent holds

$$b_A = 1, \qquad b_B = 2$$

units of assets A and B, each paying $V_{A,3} = V_{B,3} = 10$ at $t = 3$. Asset A is fully liquid at all dates ($L_{A,1} = L_{A,2} = 10$); asset B is less liquid, and its secondary-market value deteriorates further at

$t = 2$:

$$L_{B,1} = 9, \qquad L_{B,2} = 5.$$

A liquidity need $y_1 = 9$ arrives at $t = 1$ with certainty. At $t = 2$, $y_2 = 9$ with probability p and $y_2 = 0$ otherwise.

Policy M (myopic): sell the most liquid asset at $t = 1$. Selling one unit of A at $t = 1$ raises $10 \geq 9$ and meets y_1. The agent enters $t = 2$ holding $(0, 2)$. If $y_2 = 9$, both units of B must be sold for $2 \times 5 = 10 \geq 9$, leaving terminal wealth 0; if $y_2 = 0$, terminal wealth is 20. Hence

$$\mathbb{E}[W_3 \mid \mathrm{M}] = 20(1 - p).$$

Policy D (dynamic): sell B at $t = 1$, preserve A as a buffer. Selling one unit of B at $t = 1$ raises $9 = y_1$. The agent enters $t = 2$ holding $(1, 1)$. If $y_2 = 9$, sell A for $10 \geq 9$ and retain B, giving terminal wealth 10; if $y_2 = 0$, terminal wealth is 20. Hence

$$\mathbb{E}[W_3 \mid \mathrm{D}] = 20 - 10p.$$

Comparison. For any $p > 0$, $20 - 10p > 20(1 - p)$, so Policy D strictly dominates Policy M. The failure of the myopic rule is not an artefact of A being partially liquid: even though $L_{A,t} = V_{A,3}$ at every date, selling A first forces the agent to rely on the deteriorated asset B at $t = 2$, destroying $10p$ units of expected terminal wealth. The recursion (14) of Theorem 3 correctly accounts for this cost: the discounted shadow price $\rho_{A,1}$ exceeds $\xi_1 L_{A,1}$ by the option value of having A available at $t = 2$, so the complementary-slackness condition $\ell^*_{A,1} > 0 \Rightarrow \rho_{A,1} = \xi_1 L_{A,1}$ is violated under Policy M whenever $p > 0$.

3.4 Shadow Prices and Asset Values

The discounted shadow price $\rho_{i,t}$ is the globally defined process characterizing the dynamic value of asset i at date t. On states where $\xi_t > 0$, the undiscounted shadow price $\phi_{i,t} = \rho_{i,t}/\xi_t$ is additionally well defined. The discounted shadow price satisfies the recursion

$$\rho_{i,t} = \max\{\xi_t L_{i,t},\ \mathbb{E}_t[\rho_{i,t+1}]\};$$

its statewise undiscounted counterpart on $\{\xi_t > 0\}$ is the undiscounted recursion (17).

The sell-versus-hold comparison is sharpest in discounted units: selling asset i at date t delivers $\xi_t L_{i,t}$, while holding delivers $\mathbb{E}_t[\rho_{i,t+1}]$. Liquidation is therefore optimal when the former dominates the latter, as encoded in (14). On $\{\xi_t > 0\}$, dividing by ξ_t recovers the undiscounted comparison in terms of $\phi_{i,t}$.

The recursion (14) encodes this comparison in every state and at every date, and within each date-t state the realized cross section $(L_{i,t}, P_{i,t})$ is treated as given while y_t drives the cutoff.

This generalises the shadow-price approach used in life-cycle models of illiquidity (**?**). By constructing a single-valued supporting multiplier process and the associated discounted and undiscounted shadow-price system explicitly from primitives, we provide a structural micro-foundation for the shadow costs estimated empirically in that literature. The "liquidation as optimal stopping" interpretation also connects to the timing problem in Ang et al. (2014): where they characterise the option value *lost* by an investor who cannot trade continuously, our framework characterises the complementary option value *preserved* by retaining a liquid asset. The convex liquidity premium in the three-date model, which increases as the shock forces the agent to liquidate assets with progressively worse L_i/V_i ratios, is the static counterpart of this dynamic option price.

Definition 2 (Shadow price and state-price density)**.** *The* discounted shadow price *of asset i at date t is $\rho_{i,t}$, the globally defined process asserted by Theorem 3. The* undiscounted shadow price *$\phi_{i,t} := \rho_{i,t}/\xi_t$ is defined on states where $\xi_t > 0$. The* state-price density *at date t is ξ_t. At the terminal date, $\xi_T = u'(w_T^*)$ is the marginal utility of terminal wealth; for $t \le T-1$, ξ_t is the*

Lagrange multiplier on the liquidity constraint (13)*. Under the minimal-supporting convention it is nonnegative in general and strictly positive on states where* $y_t > 0$*, where it equals the marginal utility cost of tightening that constraint by one unit.*

On states where $\xi_t > 0$, the complementary-slackness conditions also imply useful no-trade implications. If $\phi_{i,t} > L_{i,t}$, the agent sells none of asset i at date t. If $\phi_{i,t} < P_{i,t}$, the agent buys none of asset i at date t. Thus strict inequality identifies dates at which the corresponding trading margin is inactive.

By recursion (14),

$$\mathbb{E}_t[\rho_{i,t+1}] \leq \rho_{i,t}$$

for all i and $t \leq T-1$, so $(\rho_{i,t})_{t=0}^{T}$ is a supermartingale. Define the first active trading date

$$\hat{\tau}_i := \min\{t \geq 0 : \ell^*_{i,t} > 0 \text{ or } b^*_{i,t} > 0\}.$$

Then, on histories for which the asset is still carried and no active trade has yet occurred, that is, for $t < \hat{\tau}_i$, the discounted holding condition in Theorem 3 implies $\rho_{i,t} = \mathbb{E}_t[\rho_{i,t+1}]$, so the discounted shadow price is a martingale up to first active trade.

For the liquidation interpretation, define instead

$$\tau_i^{\text{sell}} := \min\{t \geq 0 : \ell^*_{i,t} > 0\},$$

the first date at which the agent begins to liquidate asset i. On states where $\xi_t > 0$, the undiscounted shadow price $\phi_{i,t}$ is analogous to the date-t value of an American claim with exercise payoff $L_{i,t}$, with sale times τ_i^{sell} playing the role of exercise times. The liquidity friction generates a bid-ask spread $[L_{i,t}, P_{i,t}]$ within which this shadow value is determined by the agent's optimal trading policy.

This connection relates the framework to Jouini and Kallal (1995), who show that the bid-ask spread is arbitrage-free if and only if some process lying between the bid and ask prices is a martingale under an equivalent measure. In our setting, the discounted shadow price $(\rho_{i,t})$ is a martingale up to the first active trading date and the Snell-envelope dominance relation $\rho_{i,t} \geq \mathbb{E}_t[\rho_{i,t+1}]$ globally. On states where $\xi_t > 0$, the corresponding undiscounted shadow price $\phi_{i,t}$ lies in $[L_{i,t}, P_{i,t}]$. The state-price density (ξ_t) provides the risk adjustment that maps between the two representations.

4 Quantitative Calibration

The model delivers three qualitative results: the convex price-liquidity schedule, the reversal of the static liquidation pecking order under persistent shocks, and the divergence between static liquidity and dynamic safety rankings. This section explores the real-world implications of these results quantitatively using empirical evidence on liquid-buffer management, interim liquidity needs, and stress-state deterioration in liquidation values.

Three findings emerge from the calibration. First, convex pricing is quantitatively sensitive to the liquid-buffer share: the computation shows that the curvature of the price-liquidity schedule changes sign as the buffer share moves from eight to twelve percent of the portfolio, concentrating sharply at the high-quality end of the cross section when the buffer is thin. What changes is not the average liquidity need but the amount of internal liquidity available before the endogenous cutoff is forced into lower-ranked assets. Second, under the baseline calibration, the empirically plausible static pecking order — liquidate the most liquid asset first — is not dynamically optimal: the closed-form reversal threshold $\rho_y^* = 0.320$ lies well below the calibrated persistence $\rho_y = 0.50$, so the dynamic optimum preserves the buffer. Third, assets that rank first by normal-state liquidation quality can earn lower liquidity premia than assets with stable ratios once the stochastic discount factor loads heavily on stress states; the inversion threshold is $M_S^* = 8.8$ in the two-state exercise.

4.1 Empirical Inputs and Mapping into Model Primitives

We calibrate a three-asset baseline. Asset 1 is a cash-like or short-term Treasury-like buffer; Asset 2 is a medium-liquidity bond-like instrument; and Asset 3 is an illiquid alternative sleeve. Continuation values are normalized to $V_1 = V_2 = V_3 = 1$, so that liquidation values L_i directly encode liquidation quality and the liquidation ratio $\ell_i = L_i/V_i = L_i$ for each asset. We allow liquidation values to differ across aggregate states, writing L_i^N for the normal-state value and L_i^S for the stress-state value; the liquidation ratios inherit this notation as ℓ_i^N and ℓ_i^S.

Portfolio shares. We use the mutual-fund cash-buffer evidence in Chernenko and Sunderam (2016) to motivate a baseline liquid-buffer share $b_1 = 0.08$. That paper documents liquid-asset shares averaging around eight percent of net asset value across open-end mutual funds, with substantial variation across funds facing different redemption exposures. Morris et al. (2017) show that managers hoard additional liquidity when redemption uncertainty rises, suggesting the effective buffer share can be elevated during stress; this motivates varying b_1 over $[0.04, 0.12]$ in the sensitivity analysis for Result 1. The remaining portfolio share is allocated across the bond-like and illiquid sleeves in proportions intended to reflect a realistic liquid-core/illiquid-satellite structure: $b_2 = 0.52$ and $b_3 = 0.40$. These values are transparent normalizations and are not directly identified from the cited data.

Liquidation ratios. Asset 1 is the buffer by design: it can be converted at par with negligible sacrifice of continuation value, so $\ell_1^N = \ell_1^S = 1.00$. For Asset 2, Bao et al. (2011) document that investment-grade corporate bond illiquidity, measured by the γ statistic, rises materially in stress periods relative to normal conditions, suggesting a significant deterioration in effective liquidation quality. We use the magnitude of that stress deterioration to discipline the spread between ℓ_2^N and ℓ_2^S, and set $\ell_2^N = 0.98$ and $\ell_2^S = 0.94$ as our baseline calibration. These values are not uniquely identified by the data; they are chosen so that the order of magnitude of the normal-to-stress deterioration is consistent with the evidence in Bao et al. (2011). Asset 3 represents an illiquid alternative sleeve. We set $\ell_3^N = 0.93$ and $\ell_3^S = 0.91$, reflecting an already-thin normal-state margin with a smaller additional stress deterioration. These values are transparent normalizations chosen to preserve the ordering $\ell_1 > \ell_2 > \ell_3$ in both states.

Liquidity-need process. The interim liquidity requirement y_t follows a two-point Markov process:

$$y_t \in \{0, \bar{y}\}, \qquad \Pr(y_t = \bar{y}) = q, \qquad \Pr(y_{t+1} = \bar{y} \mid y_t = \bar{y}) = \rho_y.$$

The evidence in Andonov et al. (2025) suggests substantial persistence in institutional cash-flow pressure: pension funds facing negative flows in one quarter are significantly more likely to face further outflows in subsequent quarters, with autocorrelation patterns broadly consistent with $\rho_y \in [0.40, 0.65]$. We set $\rho_y = 0.50$ as our baseline, situated conservatively within that range, and vary it over $[0.20, 0.80]$ in the pecking-order analysis. The unconditional shock probability is $q = 0.10$. The large-redemption episodes documented in Ma et al. (2022) motivate a shock size on the order of ten to fifteen percent of liquidation capacity; we set $\bar{y} = 0.12$ in the baseline. The risk-neutral benchmark $M(y) \equiv 1$ is used for the main figures; a robustness check uses the exponential specification $M(y) \propto \exp(2y)$, normalized so that $\mathbb{E}[M(y)] = 1$ under the shock distribution.

Liquidation pecking-order benchmark. Ma et al. (2022) establish that mutual funds liquidate their most liquid assets first in the cross section: holdings in liquid instruments are drawn down before less liquid positions. In our model this corresponds to the static pecking order $\ell_1^N > \ell_2^N > \ell_3^N$, which holds under the normal-state calibration. The calibration for Result 2 quantifies when the dynamic optimum departs from this empirically observed static benchmark.

Table 2 records the calibrated inputs together with their empirical discipline and status. All subsequent price schedules, premia, and rankings are model outputs under these inputs.

Table 2: Calibrated Inputs

Parameter	Interpretation	Baseline	Source / discipline	Status
b_1	Liquid-buffer share	0.08	Cash-buffer evidence (Chernenko and Sunderam, 2016)	Baseline calibrated value
b_2	Bond-sleeve share	0.52	Liquid-core/satellite; normalization	Normalization
b_3	Illiquid-sleeve share	0.40	Liquid-core/satellite; normalization	Normalization
ℓ_1^N	Buffer: normal liq. ratio	1.00	Par conversion	Normalization
ℓ_1^S	Buffer: stress liq. ratio	1.00	Stable by construction	Normalization
ℓ_2^N	Bond: normal liq. ratio	0.98	Magnitude: bond illiquidity (Bao et al., 2011)	Baseline calibrated value
ℓ_2^S	Bond: stress liq. ratio	0.94	Magnitude: stress deterioration (Bao et al., 2011)	Baseline calibrated value
ℓ_3^N	Illiquid: normal liq. ratio	0.93	Order-of-magnitude; illiquid sleeve	Normalization
ℓ_3^S	Illiquid: stress liq. ratio	0.91	Mild stress deterioration	Normalization
$\bar{y}$	Shock size (frac. of capacity)	0.12	Large-redemption scale (Ma et al., 2022)	Baseline calibrated value
q	Unconditional shock prob.	0.10	Redemption-event frequency	Baseline calibrated value
ρ_y	Shock persistence	0.50	Clustering range (Andonov et al., 2025)	Baseline; varied in $[0.20, 0.80]$
$M(y)$	Stochastic discount factor	$\equiv 1$	Risk-neutral benchmark	Baseline benchmark
		$\propto e^{2y}$	SDF increasing in shock	Robustness specification

4.2 Convex Pricing and the Role of Liquid Buffers

The pricing formula of Theorem 2 implies that each asset's price decomposes as

$$P_i \;=\; \mathbb{E}[M(y)\,V_i] + \beta^{\text{liq}}(L_i),$$

where the liquidity premium $\beta^{\text{liq}}(L_i)$ is increasing in L_i. Under the risk-neutral benchmark and with $V_i = 1$, the normalized price index

$$p(\ell_i) \;:=\; \mathbb{E}\left[\max\left\{1,\; \frac{\ell_i}{l^*(y)}\right\}\right]$$

summarizes the cross-sectional shape of prices as a function of liquidation ratio. Because $p(\ell_i) \geq 1$ for all ℓ_i, every asset weakly commands a premium over its continuation value. For the buffer asset with $\ell_1 = 1.00$, the premium is strictly positive in any state where the shock is large enough to drive the cutoff into the buffer band — that is, whenever $l^*(y) < 1$. Because the buffer asset has the highest liquidation ratio, it earns a liquidity premium strictly above continuation value in those states even though its liquidation ratio is at par. A large value of Δp means that the valuation gain from moving from a moderately liquid asset to the buffer is much greater than the gain from moving from the illiquid sleeve to the bond asset: the price-liquidity schedule is steep near the top of the quality hierarchy.

Shock distribution for the convex pricing exercise. To isolate the role of the liquid buffer, we fix a six-point shock distribution

$$y \;\in\; \{0,\; 0.03,\; 0.07,\; 0.11,\; 0.35,\; 0.75\}$$

with probabilities $(0.80, 0.04, 0.04, 0.04, 0.04, 0.04)$ respectively. The non-zero shock values are chosen so that they straddle the Asset-1 liquidation band boundary $S_1(b_1) = b_1$ for different values of b_1: $y = 0.03$ crosses S_1 for $b_1 = 0.04$; $y = 0.07$ for $b_1 = 0.06$; $y = 0.11$ for $b_1 = 0.08$; and $y = 0.35$ falls in the Asset-2 band for $b_1 = 0.12$. The design ensures that the curvature comparison across buffer shares reflects genuine variation in the internal liquidity structure, not a trivial boundary effect. The unconditional shock probability is 0.20; the mean shock is $\mathbb{E}[y] = 0.040$ of portfolio liquidation capacity. The expectation in $p(\ell_i)$ is a probability-weighted sum across the six shock values, with the cutoff $l^*(y)$ determined analytically from Theorem 1 for each (b_1, y) pair.

Result 1 (convex pricing). *Holding the shock distribution fixed, the curvature of the price-liquidity schedule changes sign as the buffer share moves from thin to thick. A thin buffer concentrates the liquidity premium at the top of the quality hierarchy; a thick buffer spreads it toward the lower end. What changes is not the average liquidity need but the amount of internal liquidity available before the endogenous cutoff is forced into lower-ranked assets.*

The mechanism is transparent. When b_1 is large (0.12), moderate shocks $y \leq S_1 = 0.12$ stay within the buffer band, so the cutoff remains at $\ell_1 = 1.00$ and the liquidity premium accrues mainly to Asset 1. Only the large shock ($y = 0.75$, probability 0.04) reaches the Asset-3 band. When b_1 shrinks to 0.04, the same shock distribution pushes the cutoff into the Asset-2 band more frequently ($S_1 = 0.04$), so the bond asset earns a premium across a wider set of states while the buffer continues to earn a premium in all states where it is not marginal. The distribution of the premium shifts down the quality ladder as the buffer thins — not because the shocks change, but because the balance-sheet structure does.

Define the curvature statistic

$$\Delta p \ := \ \left[p(\ell_2^N) - p(\ell_3^N)\right] - \left[p(1) - p(\ell_2^N)\right].$$

When $\Delta p < 0$, the increment from the bond asset to the buffer exceeds the increment from the illiquid sleeve to the bond asset: the premium schedule is steep at the top of the quality hierarchy. When $\Delta p > 0$, the schedule is steeper at the lower end. Table 3 reports Δp and the model-implied liquidity premia, computed as model outputs from the pricing formula of Theorem 2.

Figure 2 illustrates the finding. The left panel plots $p(\ell)$ against ℓ for each buffer share; the right panel plots Δp as a function of b_1. The computed values in Table 3 show that Δp transitions from -0.00116 at $b_1 = 0.04$ to $+0.00047$ at $b_1 = 0.12$: as the buffer thins, the premium concentrates disproportionately at the buffer asset and the convex schedule steepens at the top of the hierarchy. This is a consequence not of any change in the shock distribution but entirely of where the balance-sheet capacity sits relative to the endogenous band boundaries.

Table 3: Curvature of the price-liquidity schedule by liquid-buffer share. All entries are model outputs computed analytically from Theorem 1 and the pricing formula of Theorem 2, using the six-point shock distribution of Section 4.2 and the calibrated inputs of Table 2. The final column reports which asset bands are reached across the support of y for each buffer share.

b_1	β^{liq}: Asset 2	β^{liq}: Asset 3	$p(\ell_2^N)/p(\ell_3^N)$	Δp	Bands reached
0.04	0.00215	0.00000	1.00215	−0.00116	1, 2, 3
0.06	0.00215	0.00000	1.00215	−0.00116	1, 2, 3
0.08	0.00215	0.00000	1.00215	−0.00034	1, 2, 3
0.12	0.00215	0.00000	1.00215	+0.00047	1, 2, 3

Note. $\Delta p < 0$ *indicates that the price increment from Asset 2 to the buffer exceeds the increment from Asset 3 to Asset 2: the premium is steep at the top of the quality hierarchy.* $\Delta p > 0$ *(*$b_1 = 0.12$*) indicates the reverse. The shift in sign as* b_1 *falls from* 0.12 *to* 0.04 *reflects redistribution of the liquidity premium toward the buffer as internal liquidity thins. Asset 2's premium is* 0.00215 *across all buffer shares because the shock* $y = 0.75$ *(probability* 0.04*) falls in the Asset-2 band for every* b_1*; Asset 3's premium is zero because strict feasibility prevents full portfolio exhaustion.*

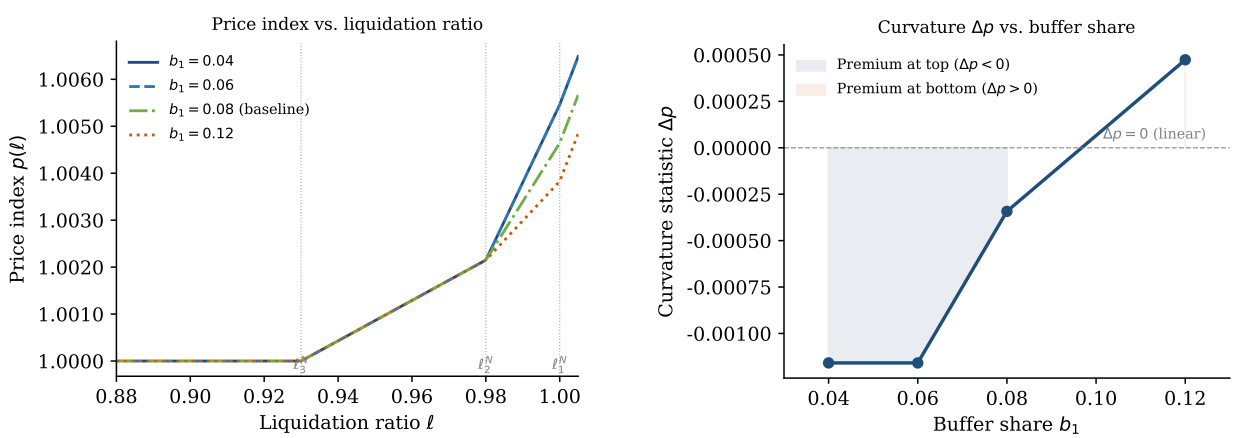

Figure 2: Convex pricing and the role of the liquid buffer. The left panel plots the normalized price index $p(\ell) = \mathbb{E}[\max\{1, \ell/l^*(y)\}]$ against the liquidation ratio ℓ for four values of the buffer share b_1. The right panel plots the curvature statistic Δp (defined in Section 4.2) against b_1. The figure isolates the effect of internal liquidity management by varying the liquid-buffer share while holding the shock distribution fixed; what changes is the amount of internal liquidity available before the endogenous cutoff is forced into lower-ranked assets, not the average liquidity need. Both panels use the calibrated inputs of Table 2 under the risk-neutral benchmark $M(y) \equiv 1$.

4.3 Dynamic Liquidation: When the Pecking Order Reverses

The multiperiod theorem (Theorem 3) implies that the static pecking order — sell the most liquid asset first — is optimal only when the option value of retaining the buffer for future shocks is sufficiently small. This option value rises with the persistence of the liquidity need ρ_y and with the stress-state deterioration ℓ_2^S/ℓ_2^N: when both are large, the buffer becomes too valuable to deploy immediately and the agent optimally sells Asset 2 first, even though $\ell_1^N > \ell_2^N$ in the current period.

The formal characterization of the reversal is given by Proposition 5 in Section 3.3. That result specializes Theorem 3 to a two-period economy with a buffer asset (perfectly liquid in all states, $\ell_1 \equiv 1$) and a secondary asset whose liquidation ratio deteriorates from ℓ_2^N in normal times to $\ell_2^S < \ell_2^N$ under stress. Under risk neutrality, selling Asset 2 first (Policy D) strictly dominates the static pecking order (Policy M) if and only if shock persistence exceeds the threshold

$$\rho_y \;>\; \rho_y^* \;:=\; \frac{1/\ell_2^N - 1}{1/\ell_2^S - 1}, \tag{19}$$

as established in equation (18). Intuitively, ρ_y^* is the break-even persistence at which the option value of retaining the buffer exactly equals its current-period liquidation advantage: the numerator measures the immediate cost of selling Asset 2 rather than the buffer, and the denominator measures the additional future cost of relying on a stress-deteriorated Asset 2 at $t = 2$.

Result 2 (pecking-order reversal). *Under the baseline calibration, the static pecking order is not dynamically optimal. Proposition 5 gives the closed-form threshold* $\rho_y^* = 0.320$ *at the calibrated liquidation ratios* $\ell_2^N = 0.98$ *and* $\ell_2^S = 0.94$*. Since the calibrated persistence* $\rho_y = 0.50$ *strictly exceeds this threshold, the dynamic optimum preserves the buffer and sells Asset 2 first.*

Figure 3 maps the boundary (19) in the $(\rho_y, \ell_2^S/\ell_2^N)$ plane, holding all other parameters at their baseline values. The figure partitions the plane into two regions: the lower-left region where the static order is optimal and the upper-right region where Policy D dominates. The baseline calibration sits inside the reversal region at $(\rho_y, \ell_2^S/\ell_2^N) = (0.50, 0.959)$, with $\rho_y^* = 0.320 < 0.50$.

As b_1 shrinks, the boundary shifts inward because a thinner buffer provides more concentrated option value per unit of continuation value.

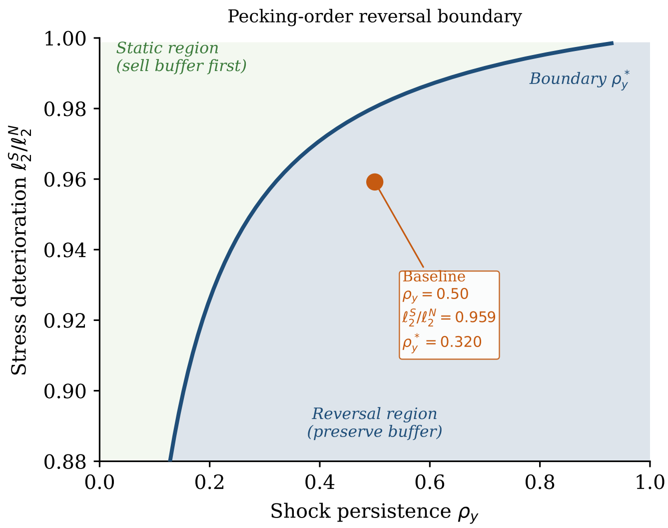

Figure 3: Pecking-order reversal boundary (Proposition 5). The locus $\rho_y^* = (1/\ell_2^N - 1)/(1/\ell_2^S - 1)$ partitions the $(\rho_y, \ell_2^S/\ell_2^N)$ plane into the static region (lower-left, sell buffer first) and the reversal region (upper-right shaded, preserve buffer). The baseline calibration $(\rho_y, \ell_2^S/\ell_2^N) = (0.50, 0.959)$ lies inside the reversal region ($\rho_y^* = 0.320 < 0.50$). Risk-neutral benchmark $M(y) \equiv 1$.

4.4 Safety versus Liquidity in the Cross Section

The distinction between liquidity and safety is state contingent: an asset may rank highly by normal-state liquidation quality and yet provide little liquidity service in the states where the shadow price of liquidity is highest. The liquidity beta decomposition of equation (9) makes this precise. The premium $\beta^{\text{liq}}(L_i)$ depends on the full state-contingent profile (ℓ_i^N, ℓ_i^S), weighted by the SDF $M(y)$ and the indicator that the asset lies above the cutoff in each state. Under risk aversion, $M(y)$ is elevated in stress states; an asset whose liquidation ratio deteriorates precisely in those states earns less premium than its normal-state ratio alone would suggest.

To make this contrast quantitatively sharp, we compare two bond-like assets within the two-state (N, S) framework:

- **Bond-B** (the baseline Asset 2): $\ell^N = 0.98$, $\ell^S = 0.94$. Higher normal-state ratio, but the ratio deteriorates by 4.1% in stress.

- **Bond-A** (a stable alternative): $\ell^N = \ell^S = 0.96$. Lower normal-state ratio, but invariant to aggregate conditions.

Under a large shock $y = 0.75$, which falls in the Asset-3 liquidation band for both states, the model-implied cutoffs are $l^{*N} = 0.93$ and $l^{*S} = 0.91$. The per-unit liquidity services, computed from Definition 1 and Theorem 1, are:

$$\text{Normal state:} \quad \text{Bond-B: } 0.98/0.93 - 1 = 0.054; \quad \text{Bond-A: } 0.96/0.93 - 1 = 0.032.$$
$$\text{Stress state:} \quad \text{Bond-B: } 0.94/0.91 - 1 = 0.033; \quad \text{Bond-A: } 0.96/0.91 - 1 = 0.055.$$

Bond-B provides more liquidity service in the normal state but less in the stress state. Its ratio falls from 0.98 to 0.94 while the cutoff l^{*S} also falls to 0.91, leaving Bond-B with a weaker relative position precisely when $M(y)$ is highest. Bond-A's stable ratio means its service rises in stress relative to the normal state, so its premium receives a proportionally larger contribution from the states that carry the highest SDF weight. The ranking inversion is therefore strongest under the risk-averse SDF robustness specification and absent under risk neutrality.

Result 3 (safety versus liquidity). *Under risk neutrality, Bond-B earns a higher liquidity premium than Bond-A, reflecting its higher normal-state ratio. Under sufficient risk aversion, this ranking inverts: Bond-A's stable ratio earns a higher premium once the SDF loads heavily on the stress states in which Bond-B's service deteriorates. The inversion threshold is $M_S^* = 8.8$, derived analytically from the model.*

Table 4 reports both assets' liquidity premia under the risk-neutral benchmark ($M_S = 1$) and under moderate risk aversion ($M_S = 5$, with $\pi_N = 0.90$ and $\pi_S = 0.10$). All entries are model outputs computed from Theorem 2 and equation (10).

The table makes three points precise. First, under risk neutrality the ordering tracks the normal-state ratio: Bond-B $>$ Bond-A because $0.98 > 0.96$, and the premium gap is 0.0172.

Table 4: Static liquidity versus dynamic safety: model-implied liquidity premia. Bond-B has the higher normal-state ratio but deteriorates in stress; Bond-A has a lower but stable ratio. All entries are model outputs. "Enters liq. set" indicates whether $\ell_i^S \geq l^{*S}$ for shock $y = 0.75$; it is a model output, not a calibrated input.

Asset	ℓ^N	ℓ^S	Detn	Liq. set	β^{liq} (RN)	β^{liq} ($M_S = 5$)	Rank
Buffer (Asset 1)	1.00	1.00	0.0%	Both	0.0776	0.1172	1st
Bond-B (detn.)	0.98	0.94	4.1%	Both	0.0517	0.0649	2nd
Bond-A (stable)	0.96	0.96	0.0%	Both	0.0345	0.0565	3rd
Ranking by ℓ^N:				Buffer > Bond-B > Bond-A			
β^{liq} *ranking (RN):*				Buffer > Bond-B > Bond-A			
β^{liq} *ranking ($M_S = 5$):*				Buffer > Bond-B > Bond-A (gap −51%)			
β^{liq} *ranking ($M_S > 8.8$):*				**Buffer** > **Bond-A** > **Bond-B** *(inversion)*			

Note. Model outputs from the two-state pricing formula (Theorem 2) with $\pi_N = 0.90$, $\pi_S = 0.10$, $y = 0.75$, $l^{*N} = 0.93$, $l^{*S} = 0.91$, $b = (0.08, 0.52, 0.40)$. The inversion threshold $M_S^* = 8.8$ equates $\beta^{\text{liq}}(\text{Bond-A}) = \beta^{\text{liq}}(\text{Bond-B})$ and is obtained analytically. The ranking inversion is strongest under the risk-averse robustness specification.

Second, under moderate risk aversion ($M_S = 5$) the gap compresses to 0.0084 — a 51% reduction — because Bond-A's stable ratio generates a proportionally larger stress-state contribution to its premium. Third, once M_S exceeds $M_S^* = 8.8$, the gap reverses: Bond-A's premium overtakes Bond-B's, and the normal-state liquidity ranking is fully inverted. This suggests that normal-period liquidity measures alone — the bid-ask spread, the γ statistic of Bao et al. (2011), or turnover ratios — may be poor guides to realized liquidity premia if stress-state deterioration differs sharply across assets.

The calibration delivers three findings that sharpen the qualitative theory. Taken together, they show that balance-sheet liquidity management, shock persistence, and stress-state liquidation quality are first-order determinants of both the level and curvature of liquidity premia in the model.

First, the liquid-buffer share governs the distribution of the convex premium across the quality hierarchy. At the baseline $b_1 = 0.08$, the curvature statistic $\Delta p = -0.00034$: the price increment from the bond asset to the buffer exceeds the increment from the illiquid sleeve to the bond asset, so the premium is concentrated at the top. As b_1 falls to 0.04, Δp deepens to -0.00116; at $b_1 = 0.12$ it reverses to $+0.00047$. The sign change is the main finding: it reflects a redistribution of the premium across the quality hierarchy as internal liquidity shifts, driven entirely by where the shock distribution sits relative to the endogenous band boundaries, not by any change in the average liquidity need.

Second, the static pecking order fails at the calibrated persistence level. As established in Proposition 5, the closed-form reversal threshold is $\rho_y^* = 0.320$ at the baseline liquidation ratios. Since the calibrated baseline $\rho_y = 0.50$ strictly exceeds this threshold, the dynamic optimum preserves the buffer. This result is not sensitive to the exact value of ρ_y within the range $[0.40, 0.65]$ suggested by the evidence in Andonov et al. (2025): every value in that range exceeds $\rho_y^* = 0.320$.

Third, stress-state safety matters separately from normal-state liquidity, and the divergence has a precise quantitative threshold. Under risk neutrality the deteriorating bond (Bond-B) earns 0.0517 versus 0.0345 for the stable bond (Bond-A). Under $M_S = 5$, the gap compresses from 0.0172 to 0.0084, a 51% reduction. At $M_S^* = 8.8$ the ranking inverts: Bond-A overtakes Bond-B because its stable ratio provides more service in the stress states that receive the highest utility weight. The ranking inversion is a consequence of the SDF loading on stress states rather than of the normal-state liquidation hierarchy.

5 Conclusion

This paper studies cash management when a firm holds assets with different degrees of liquidity and faces stochastic interim funding needs. In a baseline three-date environment, the optimal liquidation policy has a threshold structure: conditional on a liquidity shock, assets are liquidated in descending order of their liquidation ratio, with the cutoff determined endogenously by the size of the shock and the composition of the portfolio. This threshold gives rise to a shadow price of liquidity, which is the central pricing object in the model and summarizes the marginal value of relaxing the funding constraint.

Our main pricing result is that an asset is valued not only for its continuation payoff, but also for the liquidity service it provides in the states in which it is drawn into the liquidation set.

This gives liquidity an option-like value and generates a convex relationship between prices and liquidity even with a single representative investor: as shocks grow larger, the firm is forced to liquidate assets with progressively worse liquidation ratios, so the marginal cost of liquidity rises endogenously. We also derive a liquidity beta formula that prices liquidity risk through the state-contingent cash service an asset provides when funding needs are high. This distinction shows that static liquidity and dynamic safety need not coincide: assets that appear liquid in normal times may command lower premia than assets with more stable liquidation values in stress states. The same framework yields implications for liability design. Under risk aversion, stable funding structures reduce exposure to the worst liquidation thresholds and thereby protect continuation values.

We then extend the analysis to a multiperiod setting in which assets can be traded repeatedly over time. In that environment, the static heuristic of selling the most liquid asset first is not generally optimal, because a liquid asset may be more valuable as a buffer against future shocks than as an immediate source of cash. We characterize this force through a shadow-price recursion in which each asset's value equals the larger of its current liquidation value and its risk-adjusted continuation value. This delivers a sharp pecking-order reversal result: when shock persistence is sufficiently high relative to the deterioration in liquidation quality under stress, it is strictly optimal to sell a less liquid asset first in order to preserve the more liquid asset as a buffer. Our calibration shows that this reversal arises at empirically plausible parameter values. More broadly, our paper provides a unified model of liquidity management, liability structure, and asset pricing, and delivers new empirical predictions for liquidation behavior and liquidity premia.

References

Acharya, V. V. and Pedersen, L. H. (2005). Asset pricing with liquidity risk. *Journal of financial Economics*, 77(2):375–410.

Amihud, Y. and Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of financial Economics*, 17(2):223–249.

Andonov, A., Jansen, K. A., and Rauh, J. D. (2025). When cash flows turn negative: Liquidity-driven selling by pension funds. *Available at SSRN 5498920*.

Ang, A., Papanikolaou, D., and Westerfield, M. M. (2014). Portfolio choice with illiquid assets. *Management Science*, 60(11):2737–2761.

Bao, J., Pan, J., and Wang, J. (2011). The illiquidity of corporate bonds. *Journal of Finance*, 66(3):911–946.

Baumol, W. J. (1952). The transactions demand for cash: An inventory theoretic approach. *The Quarterly Journal of Economics*, pages 545–556.

Berlin, M. and Mester, L. J. (1999). Deposits and relationship lending. *The Review of Financial Studies*, 12(3):579–607.

Bolton, P., Chen, H., and Wang, N. (2024). The marginal value of cash: Corporate savings, investment, and financing. *Annual Review of Financial Economics*, 16.

Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *The review of financial studies*, 22(6):2201–2238.

Chernenko, S. and Sunderam, A. (2016). Liquidity transformation in asset management: Evidence from the cash holdings of mutual funds. Technical report, National Bureau of Economic Research.

Constantinides, G. M. (1986). Capital market equilibrium with transaction costs. *Journal of political Economy*, 94(4):842–862.

Doerr, S. (2024). Bank geographic diversification and funding stability. *Available at SSRN 4788627*.

Geromichalos, A., Herrenbrueck, L., and Lee, S. (2023). Asset safety versus asset liquidity. *Journal of Political Economy*, 131(5):1172–1212.

Goldstein, I., Jiang, H., and Ng, D. T. (2017). Investor flows and fragility in corporate bond funds. *Journal of Financial Economics*, 126(3):592–613.

Hanson, S. G., Shleifer, A., Stein, J. C., and Vishny, R. W. (2015). Banks as patient fixed-income investors. *Journal of Financial Economics*, 117(3):449–469.

He, Z. and Xiong, W. (2012). Dynamic debt runs. *The Review of Financial Studies*, 25(6):1799–1843.

Holmström, B. and Tirole, J. (2001). Lapm: A liquidity-based asset pricing model. *the Journal of Finance*, 56(5):1837–1867.

Huang, R. and Ratnovski, L. (2011). The dark side of bank wholesale funding. *Journal of Financial Intermediation*, 20(2):248–263.

Jansen, K. A. E. and Werker, B. J. M. (2022). The shadow costs of illiquidity. *Journal of Financial and Quantitative Analysis*, 57(7):2693–2723.

Jouini, E. and Kallal, H. (1995). Martingales and arbitrage in securities markets with transaction costs. *Journal of Economic Theory*, 66(1):178–197.

Ma, Y., Xiao, K., and Zeng, Y. (2022). Mutual fund liquidity transformation and reverse flight to liquidity. *The Review of Financial Studies*, 35(10):4674–4711.

Miller, M. H. and Orr, D. (1966). A model of the demand for money by firms. *The Quarterly journal of economics*, 80(3):413–435.

Morris, S., Shim, I., and Shin, H. S. (2017). Redemption risk and cash hoarding by asset managers. *Journal of Monetary Economics*, 89:71–87.

Pástor, L. and Stambaugh, R. F. (2003). Liquidity risk and expected stock returns. *Journal of Political economy*, 111(3):642–685.

Pegoraro, S., Shive, S., and Zambrana, R. (2025). Democratizing illiquid assets: Liquidity transformation and performance in interval funds. *Available at SSRN 5365061*.

Tobin, J. (1956). The interest-elasticity of transactions demand for cash. *The review of Economics and Statistics*, 38(3):241–247.

Vayanos, D. (1998). Transaction costs and asset prices: A dynamic equilibrium model. *The Review of Financial Studies*, 11(1):1–58.

A Appendix: Proofs

Throughout this appendix we maintain Assumptions 2 (distinct liquidation ratios, $L_1/V_1 > \cdots > L_n/V_n$) and 1 (strict feasibility, $y < \sum_i b_i L_i$ a.s.) without repeated reference. We index assets in decreasing order of L_i/V_i and write $x_k := V_k/L_k$ for the breakpoints of ϕ, with the conventions $x_0 := +\infty$ and $x_{n+1} := 0$. All sums with an empty index set equal zero.

A.1 Proof of Theorem 1

We proceed in four steps.

Step 1: Dual representation. The primal problem is

$$\max_{s\in[0,b]} \sum_i (b_i - s_i)V_i \quad \text{subject to} \quad \sum_i L_i s_i \geq y.$$

Rewriting as a minimization: minimize $\sum_i s_i V_i$ subject to the liquidity constraint, since $\sum_i (b_i - s_i)V_i = \sum_i b_i V_i - \sum_i s_i V_i$ differs only by the constant $\sum_i b_i V_i$. Introduce a multiplier $x \geq 0$ on the liquidity constraint. For fixed x, the Lagrangian relaxation decomposes across assets:

$$\min_{s\in[0,b]} \sum_i s_i(V_i - xL_i) + xy.$$

Each term $s_i(V_i - xL_i)$ is minimized by setting $s_i = 0$ when $xL_i \leq V_i$ and $s_i = b_i$ when $xL_i > V_i$: sell the asset fully if the liquidity value xL_i exceeds the continuation cost V_i. The dual function is

$$d(x) = xy - \sum_{i:\, xL_i > V_i} b_i(xL_i - V_i),$$

and the dual problem is $\max_{x\geq 0} d(x)$. Writing $d(x) = \sum_i b_i V_i - \phi(x; b, y)$ where

$$\phi(x; b, y) := \sum_i b_i \max\{V_i, xL_i\} - xy,$$

strong duality for this linear program gives

$$\text{primal optimal retained value} = \sum_i b_i V_i - \min_{x \geq 0} \phi(x; b, y).$$

Step 2: Subgradient characterization of $X^*(b, y)$. The function $\phi(\cdot; b, y)$ is convex and piecewise linear. Since assets are indexed in decreasing order of L_i/V_i, the breakpoints satisfy $x_1 = V_1/L_1 < x_2 = V_2/L_2 < \cdots < x_n = V_n/L_n$. On the interval (x_k, x_{k+1}), the set $\{i : xL_i > V_i\} = \{1, \ldots, k\}$, so

$$\phi'(x; b, y) = \sum_{i=1}^{k} b_i L_i - y, \quad x \in (x_k, x_{k+1}).$$

A value x^* minimizes ϕ over $\mathbb{R}_+$ if and only if $0 \in \partial_x \phi(x^*)$. At a non-breakpoint the subdifferential is a singleton, and optimality requires the slope to equal zero. At the breakpoint $x_k = V_k/L_k$, the left and right derivatives are

$$\phi'(x_k^-) = \sum_{i=1}^{k-1} b_i L_i - y \qquad \text{and} \qquad \phi'(x_k^+) = \sum_{i=1}^{k} b_i L_i - y,$$

so $0 \in \partial_x \phi(x_k)$ if and only if

$$\sum_{i=1}^{k-1} b_i L_i \;\leq\; y \;\leq\; \sum_{i=1}^{k} b_i L_i. \tag{20}$$

Step 3: Finiteness of $\bar{x}(b, y)$ under strict feasibility. Recall $x_k = V_k/L_k$ with $x_1 < x_2 < \cdots < x_n$. For $x > x_n = V_n/L_n$, every asset satisfies $xL_i > V_i$, so $\phi'(x) = \sum_i b_i L_i - y$, which is strictly positive by strict feasibility. Hence ϕ is strictly increasing on (x_n, ∞) and no minimizer lies there, giving $X^*(b, y) \subseteq [0, x_n]$ and $\bar{x}(b, y) \leq x_n < \infty$. The boundary case $y = \sum_i b_i L_i$ is the only situation that makes this slope equal to zero for all $x \geq x_n$, which would generate an unbounded minimizer set; strict feasibility rules out that case precisely. For $y > 0$, the right derivative of ϕ at $x = 0$ equals $-y < 0$, so the minimizer is strictly positive and $l^*(b, y) = 1/\bar{x}(b, y)$ is finite. For $y = 0$, $x = 0$ is optimal, consistent with zero liquidation required.

Step 4: Unique cutoff, minimal selection, and the liquidation rule. By Step 2, $x_k \in X^*(b, y)$ if and only if (20) holds. Because the cumulative capacities $\sum_{i=1}^{k} b_i L_i$ are strictly increasing in k (since $b_i L_i > 0$), the regions defined by (20) for different k can overlap only at boundary points; with $y \in (0, \sum_i b_i L_i)$, exactly one index k satisfies (4), confirming uniqueness. The minimizer set $X^*(b, y)$ is contained in $[x_k, x_{k+1}]$ and the breakpoint x_k belongs to it, so $\bar{x}(b, y) = \inf X^*(b, y) = x_k = V_k/L_k$. Hence $l^*(b, y) = 1/\bar{x}(b, y) = L_k/V_k$.

At the dual solution $x^* = V_k/L_k$, the condition $x^* L_i \gtrless V_i$ reduces to $L_i/V_i \gtrless L_k/V_k$: for $i < k$ (higher ratio), $x^* L_i > V_i$ and $s_i^* = b_i$; for $i > k$ (lower ratio), $x^* L_i < V_i$ and $s_i^* = 0$; for $i = k$, $x^* L_k = V_k$ and s_k is free in $[0, b_k]$, pinned by the binding constraint $\sum_i L_i s_i = y$:

$$s_k^* = \alpha(b, y) := \frac{y - \sum_{i=1}^{k-1} b_i L_i}{L_k} \in [0, b_k],$$

where the bounds follow from the two inequalities in (4).

Expressing the liquidation rule in terms of the index k: $s_i^* = b_i$ if $i < k$, $s_i^* = 0$ if $i > k$, and $s_k^* = \alpha(b, y)$ if $i = k$. This is equation (5). The knife-edge case $y = \sum_{i=1}^{k-1} b_i L_i$ gives $\alpha = 0$; the cutoff $l^* = L_k/V_k$ is still pinned by the subgradient condition and the allocation is well defined. □

Lemma 1 (Convex-combination selection)**.** *Let ξ and η be integrable $\mathbb{R}^n$-valued random vectors such that*

$$\mathbb{E}\big[\min\{\delta \cdot \xi, \delta \cdot \eta\}\big] \le 0 \qquad \textit{for all } \delta \in \mathbb{R}^n.$$

Then there exists a random variable $\theta \in [0, 1]$ such that

$$\mathbb{E}\big[\theta\xi + (1 - \theta)\eta\big] = 0.$$

Proof of Lemma 1. Let

$$\zeta := \eta - \xi, \qquad c := -\mathbb{E}[\eta].$$

We seek a random variable $\theta \in [0,1]$ such that

$$\mathbb{E}[\theta\zeta] = c,$$

because this is equivalent to

$$\mathbb{E}[\theta\xi + (1-\theta)\eta] = 0.$$

Consider the set

$$\mathcal{A} := \{\mathbb{E}[\theta\zeta] : \theta \text{ is a random variable taking values in } [0,1]\} \subset \mathbb{R}^n.$$

Since ζ is integrable, the set $\mathcal{A}$ is bounded. Choose a sequence $(\theta_m)_m$ of $[0,1]$-valued random variables such that

$$\|\mathbb{E}[\theta_m\zeta] - c\| \longrightarrow \inf_{\theta\in[0,1]} \|\mathbb{E}[\theta\zeta] - c\|.$$

By Komlós' theorem, there exists a subsequence, not relabeled, whose Cesàro means

$$\bar{\theta}_m := \frac{1}{m}\sum_{j=1}^{m}\theta_j$$

converge almost surely to a random variable θ^* taking values in $[0,1]$. Because $|\bar{\theta}_m\zeta| \leq |\zeta|$ and ζ is integrable, dominated convergence yields

$$\mathbb{E}[\bar{\theta}_m\zeta] \to \mathbb{E}[\theta^*\zeta].$$

Hence $\mathbb{E}[\theta^*\zeta]$ attains the minimum distance from c over $\mathcal{A}$.

Suppose, for contradiction, that $\mathbb{E}[\theta^*\zeta] \neq c$. Write

$$\mathbb{E}[\theta^*\zeta] = c + \varepsilon \qquad \text{with } \varepsilon \neq 0.$$

For any $[0,1]$-valued random variable θ and any $t \in [0,1]$, the convex combination $t\theta + (1-t)\theta^*$

is again $[0, 1]$-valued, so by optimality of θ^* the function

$$f(t) := \|t\mathbb{E}[\theta\zeta] + (1-t)\mathbb{E}[\theta^*\zeta] - c\|^2$$

is minimized at $t = 0$. Its right derivative at zero must therefore be nonnegative:

$$0 \le f'(0+) = 2\,\varepsilon \cdot \big(\mathbb{E}[\theta\zeta] - c - \varepsilon\big).$$

Using $c = -\mathbb{E}[\eta]$ and $\zeta = \eta - \xi$, this becomes

$$\mathbb{E}\big[\theta\,(\varepsilon\cdot\eta - \varepsilon\cdot\xi)\big] \ge \|\varepsilon\|^2 + \varepsilon\cdot\mathbb{E}[\eta].$$

Rearranging,

$$\mathbb{E}\big[\theta\,\varepsilon\cdot\xi + (1-\theta)\,\varepsilon\cdot\eta\big] \ge \|\varepsilon\|^2 > 0.$$

Now choose

$$\theta := \mathbf{1}_{\{\varepsilon\cdot\xi\le\varepsilon\cdot\eta\}}.$$

Then

$$\theta\,\varepsilon\cdot\xi + (1-\theta)\,\varepsilon\cdot\eta = \min\{\varepsilon\cdot\xi, \varepsilon\cdot\eta\},$$

so the hypothesis of the lemma gives

$$\mathbb{E}\big[\theta\,\varepsilon\cdot\xi + (1-\theta)\,\varepsilon\cdot\eta\big] = \mathbb{E}\big[\min\{\varepsilon\cdot\xi, \varepsilon\cdot\eta\}\big] \le 0,$$

a contradiction. Therefore $\mathbb{E}[\theta^*\zeta] = c$, and hence

$$\mathbb{E}[\theta^*\xi + (1-\theta^*)\eta] = 0.$$

Setting $\theta := \theta^*$ completes the proof. □

A.2 Proof of Theorem 2

We prove the pricing representation in four steps.

Step 1: Continuity of the minimized dual value. For each (b, y) define

$$g(b, y) := \min_{x \geq 0} \phi(x; b, y).$$

By the primal-dual equivalence in Appendix A.1, this minimized dual value coincides with the terminal value of the retained portfolio under optimal liquidation, so

$$\bar{g}(b, y) = g(b, y).$$

For each realization of (L, V, y), the function $\phi(\cdot; b, y)$ is convex and piecewise linear in x, with breakpoints $x_k = V_k/L_k$, $k = 1, \ldots, n$. Hence the minimum is attained on the finite set $\{0, x_1, \ldots, x_n\}$. Therefore

$$g(b, y) = \min\big\{\phi(0; b, y), \phi(x_1; b, y), \ldots, \phi(x_n; b, y)\big\}$$

for each realization. Since each map $b \mapsto \phi(0; b, y)$ and each map $b \mapsto \phi(x_k; b, y)$ is affine, hence continuous, it follows that $b \mapsto g(b, y)$ is continuous as the minimum of finitely many continuous functions. Thus $b \mapsto g(b, y) = \bar{g}(b, y)$ is continuous for each realization of (L, V, y). In the theorem, we take b^* to be a given interior optimum of the portfolio problem, with associated multiplier $\lambda > 0$.

Step 2: Endpoint payoff vectors. Let

$$\underline{x}(y) := \inf X^*(b^*, y) = \bar{x}(b^*, y), \qquad \overline{x}(y) := \sup X^*(b^*, y),$$

and define

$$q_i(y) := \max\{V_i, \underline{x}(y)L_i\}, \qquad r_i(y) := \max\{V_i, \overline{x}(y)L_i\}, \qquad i = 1, \dots, n.$$

Because $g(\cdot, y)$ is the minimum of finitely many affine functions, its directional derivative at b^* in direction δ is

$$Dg(b^*; \delta)(y) = \min\{\delta \cdot q(y),\ \delta \cdot r(y)\}.$$

Step 3: First-order condition on the budget hyperplane. Let

$$B := \left\{ b \in \mathbb{R}^n_+ : \sum_{i=1}^{n} b_i P_i \le w \right\}.$$

Fix $\varepsilon \in \mathbb{R}^n$ and define

$$\delta := \varepsilon - \frac{P \cdot \varepsilon}{\|P\|^2} P,$$

so that $\delta \cdot P = 0$. Since b^* is interior relative to the budget hyperplane, $b^* + t\delta \in B$ for all sufficiently small $|t|$. Optimality of b^* therefore implies the directional first-order inequality

$$\mathbb{E}[U'(\bar{g}(b^*, y)) \min\{\delta \cdot q(y),\ \delta \cdot r(y)\}] \le 0$$

for every such direction δ.

Define the projected random vectors

$$\xi(y) := U'(\bar{g}(b^*, y)) \left(q(y) - \frac{P \cdot q(y)}{\|P\|^2} P \right), \qquad \eta(y) := U'(\bar{g}(b^*, y)) \left(r(y) - \frac{P \cdot r(y)}{\|P\|^2} P \right).$$

Since $\delta = \varepsilon - (P \cdot \varepsilon / \|P\|^2) P$, the inequality above is equivalent to

$$\mathbb{E}\big[\min\{\varepsilon \cdot \xi(y),\ \varepsilon \cdot \eta(y)\}\big] \le 0 \qquad \text{for all } \varepsilon \in \mathbb{R}^n.$$

By Lemma 1, there exists a random variable $\theta(y) \in [0, 1]$ such that

$$\mathbb{E}\big[\theta(y)\xi(y) + (1 - \theta(y))\eta(y)\big] = 0.$$

Expanding this identity gives

$$\mathbb{E}\big[U'(\bar{g}(b^*, y))\big(\theta(y)q(y) + (1 - \theta(y))r(y)\big)\big] = \frac{P \cdot \mathbb{E}\big[U'(\bar{g}(b^*, y))\big(\theta(y)q(y) + (1 - \theta(y))r(y)\big)\big]}{\|P\|^2} P.$$

Hence the expectation vector on the right-hand side is proportional to P. Writing the proportionality constant as $1/\lambda$, with $\lambda > 0$, yields

$$P = \frac{1}{\lambda}\mathbb{E}\big[U'(\bar{g}(b^*, y))\big(\theta(y)q(y) + (1 - \theta(y))r(y)\big)\big].$$

Therefore, for each asset i with $b_i^* > 0$,

$$P_i = \frac{1}{\lambda}\mathbb{E}\Big[U'(\bar{g}(b^*, y))\Big(\theta(y)\max\{V_i, \underline{x}(y)L_i\} + (1 - \theta(y))\max\{V_i, \overline{x}(y)L_i\}\Big)\Big]. \tag{21}$$

Step 4: Random-X representation. Define

$$M(y) := \frac{1}{\lambda}U'(\bar{g}(b^*, y)).$$

Then (21) becomes

$$P_i = \mathbb{E}\Big[M(y)\Big(\theta(y)\max\{V_i, \underline{x}(y)L_i\} + (1 - \theta(y))\max\{V_i, \overline{x}(y)L_i\}\Big)\Big].$$

If desired, enlarge the probability space to support an independent $\upsilon \sim \mathrm{Unif}[0, 1]$, and define

$$X(y, \upsilon) = \underline{x}(y)\mathbf{1}_{\{\upsilon \le \theta(y)\}} + \overline{x}(y)\mathbf{1}_{\{\upsilon > \theta(y)\}}.$$

Then $X(y,\upsilon)$ takes values in the endpoint set of minimizers of $\phi(\cdot;b^*,y)$ and

$$P_i = \mathbb{E}[M(y)\max\{V_i, L_i X(y,\upsilon)\}],$$

which is (7).

If $X^*(b^*,y)$ is a singleton almost surely, then

$$\underline{x}(y) = \bar{x}(y) = \bar{x}(b^*,y) = \frac{1}{l^*(b^*,y)} \quad \text{a.s.,}$$

and the representation reduces to

$$P_i = \mathbb{E}\left[M(y)\max\left\{V_i, \frac{L_i}{l^*(b^*,y)}\right\}\right],$$

which is (8). □

A.3 Proofs of Propositions 1 and 2

Proof of Proposition 1. For any $y > 0$, let $k(y)$ be the unique index satisfying (4), so $l^*(b,y) = L_{k(y)}/V_{k(y)}$ by Theorem 1. In the ratio ordering, k increases as the cutoff falls deeper into the ranking (toward lower L_i/V_i). Let $y_1 < y_2$. Suppose for contradiction that $k(y_2) < k(y_1)$ (the cutoff moved to a higher-ratio asset). Then

$$\sum_{i=1}^{k(y_2)} b_i L_i \le \sum_{i=1}^{k(y_1)-1} b_i L_i < y_1 < y_2,$$

where the first inequality uses $k(y_2) < k(y_1)$, and the strict middle inequality uses the left-strict part of (4) for $k(y_1)$. This contradicts $y_2 \le \sum_{i=1}^{k(y_2)} b_i L_i$. Hence $k(y_2) \ge k(y_1)$, meaning the cutoff moves weakly deeper in the ranking, so $L_{k(y_2)}/V_{k(y_2)} \le L_{k(y_1)}/V_{k(y_1)}$ since ratios are strictly decreasing in k. Therefore $l^*(b,y_2) \le l^*(b,y_1)$ and, taking reciprocals, $\bar{x}(b,y_2) \ge \bar{x}(b,y_1)$. □

Proof of Proposition 2. Fix (b, y, V) and initial profile L with cutoff $l^*(b, y) = L_k/V_k$. By Theorem 1:

$$\sum_{i=1}^{k-1} b_i L_i < y \le \sum_{i=1}^{k} b_i L_i.$$

Under the alternative profile $\widetilde{L}$ with $\widetilde{L}_i \ge L_i$ for all $i \le k$ (strict for at least one) and the relative ordering of ratios $\widetilde{L}_i/V_i$ preserved:

$$\sum_{i=1}^{k} b_i \widetilde{L}_i \ \ge \ \sum_{i=1}^{k} b_i L_i \ \ge \ y.$$

Hence the requirement can be met using only assets $i \le k$ under profile $\widetilde{L}$, so the new cutoff index $\widetilde{k}$ satisfies $\widetilde{k} \le k$. Since ratios are strictly decreasing in index, $\widetilde{L}_{\widetilde{k}}/V_{\widetilde{k}} \ge L_k/V_k$. Therefore $l^*(b, y; \widetilde{L}) \ge l^*(b, y; L)$ and, taking reciprocals, $\bar{x}(b, y; \widetilde{L}) \le \bar{x}(b, y; L)$. □

A.4 Liquidity option decomposition

Under the canonical minimal-supporting representation, for any asset i and any realization of y, the identity

$$\max\left\{V_i, \frac{L_i}{l^*(y)}\right\} = V_i + \left(\frac{L_i}{l^*(y)} - V_i\right) \mathbf{1}_{\{L_i/V_i \ge l^*(y)\}}$$

holds because the second term equals $L_i/l^*(y) - V_i > 0$ when $L_i/V_i > l^*(y)$, equals zero when $L_i/V_i = l^*(y)$, and the indicator is zero when $L_i/V_i < l^*(y)$. Taking expectations against $M(y)$ and applying Theorem 2 yields equation (9).

A.5 Proof of Proposition 3

Appendix A.4 gives $P_i = \mathbb{E}[M(y)V_i] + \beta^{\text{liq}}(L_i)$ for any asset i, where $\beta^{\text{liq}}(\cdot)$ is defined in (10). The stated condition $\beta^{\text{liq}}(L_i) \ge \beta^{\text{liq}}(L_j)$ together with the common first term gives $P_i \ge P_j$. □

A.6 Proof of Proposition 4

Suppose U is linear, so $M(y)$ is a positive constant m and $\mathcal{V}(f) = m \cdot \mathbb{E}_f[H(y)]$. We show $H(y)$ is weakly increasing in y.

By Proposition 1, $l^*(y)$ is weakly decreasing in y: a larger shock weakly lowers the cutoff ratio $l^*(y)$, which raises $L_i/l^*(y)$ and thereby increases each service term $\max\{V_i, L_i/l^*(y)\}$. Since $H(y) = \sum_i b_i^* \max\{V_i, L_i/l^*(y)\}$ is a sum of such terms with positive weights, $H(y)$ is weakly increasing in y. Moreover, $H(y)$ is a step function in y because $l^*(y)$ only changes when y crosses a cumulative liquidation capacity threshold; monotonicity therefore holds without requiring differentiability.

If f first-order stochastically dominates g, then $\mathbb{E}_f[H(y)] \geq \mathbb{E}_g[H(y)]$ by the standard criterion. Hence $\mathcal{V}(f) \geq \mathcal{V}(g)$. □

A.7 Proof of Theorem 3

We maintain Assumptions 3–5 and assume an optimal policy $\pi^* = \{(b_t^*, \ell_t^*, \theta_t^*)\}_{t=0}^{T-1}$ exists with terminal wealth $w_T^* = \sum_i \theta_{i,T-1}^* L_{i,T}$.

Step 1: Tightness of the liquidity constraints. Fix $t \leq T-1$. Because u is strictly increasing, leaving slack at the liquidity constraint at date t is never optimal. Formally, if $y_t < \sum_i(\ell_{i,t}^* L_{i,t} - b_{i,t}^* P_{i,t})$ on a positive-probability set, then one can reduce sales (or increase purchases) by a sufficiently small $\varepsilon > 0$ on that set while preserving feasibility of all later constraints and strictly increasing w_T^*, contradicting optimality. Hence

$$y_t = \sum_{i=1}^{n} \big(\ell_{i,t}^* L_{i,t} - b_{i,t}^* P_{i,t}\big), \qquad t = 0, \ldots, T-1.$$

Step 2: Lagrangian and KKT conditions. Introduce adapted nonnegative multipliers $\xi_t \geq 0$ on (13) for $t = 0, \ldots, T-1$, adapted multipliers $\eta_{i,t} \geq 0$ on the accounting identities (11), and nonnegativity multipliers $\alpha_{i,t}, \beta_{i,t}, \gamma_{i,t} \geq 0$ on $\theta_{i,t} \geq 0$, $b_{i,t} \geq 0$, $\ell_{i,t} \geq 0$ respectively. The Lagrangian is

$$\mathcal{L} = \mathbb{E}\Bigg[u(w_T) + \sum_{t=0}^{T-1} \xi_t \Bigg(\sum_{i=1}^{n} (\ell_{i,t} L_{i,t} - b_{i,t} P_{i,t}) - y_t \Bigg) + \sum_{t=0}^{T-1} \sum_{i=1}^{n} \eta_{i,t} (\theta_{i,t-1} + b_{i,t} - \ell_{i,t} - \theta_{i,t}) + \sum_{t=0}^{T-1} \sum_{i=1}^{n} (\alpha_{i,t}\theta_{i,t} + \beta_{i,t} b_{i,t} + \gamma_{i,t} \ell_{i,t}) \Bigg].$$

Setting $\xi_T := u'(w_T^*)$ and collecting terms, the stationarity conditions with respect to $\ell_{i,t}$ and $b_{i,t}$ give

$$\xi_t L_{i,t} - \eta_{i,t} + \gamma_{i,t} = 0, \tag{22}$$

$$\eta_{i,t} - \xi_t P_{i,t} + \beta_{i,t} = 0. \tag{23}$$

Since $\beta_{i,t}, \gamma_{i,t} \geq 0$, these imply

$$\xi_t L_{i,t} \; \leq \; \eta_{i,t} \; \leq \; \xi_t P_{i,t}.$$

Stationarity with respect to $\theta_{i,t}$ for $t \leq T - 2$ gives

$$\mathbb{E}_t[\eta_{i,t+1}] - \eta_{i,t} + \alpha_{i,t} = 0,$$

and for $t = T - 1$,

$$\mathbb{E}_{T-1}[\xi_T L_{i,T}] - \eta_{i,T-1} + \alpha_{i,T-1} = 0.$$

Since $\alpha_{i,t} \geq 0$, both conditions give $\eta_{i,t} \geq \mathbb{E}_t[\eta_{i,t+1}]$ (where we set $\eta_{i,T} := \xi_T L_{i,T}$), so $(\eta_{i,t})_{t=0}^T$ is a supermartingale. Complementary slackness yields:

$$\ell_{i,t}^* > 0 \;\Rightarrow\; \gamma_{i,t} = 0 \;\Rightarrow\; \eta_{i,t} = \xi_t L_{i,t},$$
$$b_{i,t}^* > 0 \;\Rightarrow\; \beta_{i,t} = 0 \;\Rightarrow\; \eta_{i,t} = \xi_t P_{i,t},$$
$$\theta_{i,t}^* > 0 \;\Rightarrow\; \alpha_{i,t} = 0 \;\Rightarrow\; \eta_{i,t} = \mathbb{E}_t[\eta_{i,t+1}].$$

Step 3: Snell envelope construction. Given $(\xi_t)_{t=0}^T$, define $(\eta_{i,t}^*)_{t=0}^T$ as the Snell envelope of $(\xi_t L_{i,t})_{t=0}^T$:

$$\eta_{i,T}^* := \xi_T L_{i,T}, \qquad \eta_{i,t}^* := \max\{\xi_t L_{i,t},\ \mathbb{E}_t[\eta_{i,t+1}^*]\}, \quad t \leq T-1.$$

We show by backward induction on t that $\eta_{i,t} \geq \eta_{i,t}^*$ for all t. The base case $t = T$ holds by definition. Suppose $\eta_{i,t} \geq \eta_{i,t}^*$ for some $t \in \{1, \ldots, T\}$. Then dual feasibility implies

$$\eta_{i,t-1} \geq \max\{\xi_{t-1} L_{i,t-1},\ \mathbb{E}_{t-1}[\eta_{i,t}]\} \geq \max\{\xi_{t-1} L_{i,t-1},\ \mathbb{E}_{t-1}[\eta_{i,t}^*]\} = \eta_{i,t-1}^*.$$

The process $(\eta_{i,t}^*)$ also satisfies the dual feasibility bound $\xi_t L_{i,t} \leq \eta_{i,t}^* \leq \eta_{i,t} \leq \xi_t P_{i,t}$, so it is dual feasible. By the optimality of π^*, we can therefore replace $\eta_{i,t}$ with $\eta_{i,t}^*$ in the KKT system. Henceforth write $\eta_{i,t}$ for this Snell-envelope selection.

Step 4: Nonnegativity of (ξ_t), statewise positivity on $\{y_t > 0\}$, and the shadow price definition. Since $u' > 0$, we have $\xi_T = u'(w_T^*) > 0$ almost surely. For $t \leq T-1$, the multiplier on (13) is nonnegative by construction, so $\xi_t \geq 0$ almost surely. Under the minimal-supporting convention, zero may belong to the admissible multiplier set in states with $y_t = 0$, so strict positivity need not hold globally.

We therefore work first with the discounted shadow value

$$\rho_{i,t} := \eta_{i,t}, \qquad i = 1, \ldots, n, \quad t = 0, \ldots, T.$$

By construction,

$$\xi_t L_{i,t} \le \rho_{i,t} \le \xi_t P_{i,t}.$$

Moreover, the Snell-envelope recursion in Step 3 yields

$$\rho_{i,t} = \max\big\{\xi_t L_{i,t},\ \mathbb{E}_t[\rho_{i,t+1}]\big\}, \qquad t \le T-1,$$

which is globally well defined.

On the event $\{\xi_t > 0\}$, define the shadow price

$$\phi_{i,t} := \frac{\rho_{i,t}}{\xi_t} = \frac{\eta_{i,t}}{\xi_t}.$$

Dividing the recursion above by ξ_t gives

$$\phi_{i,t} = \max\left\{L_{i,t},\ \frac{1}{\xi_t}\mathbb{E}_t[\xi_{t+1}\phi_{i,t+1}]\right\}$$

on $\{\xi_t > 0\}$, which is equation (17). Likewise, the bound $\xi_t L_{i,t} \le \rho_{i,t} \le \xi_t P_{i,t}$ yields

$$L_{i,t} \le \phi_{i,t} \le P_{i,t}$$

on $\{\xi_t > 0\}$.

Finally, we show that $\xi_t > 0$ almost surely on $\{y_t > 0\}$. Fix $t \le T-1$. By Step 1, on $\{y_t > 0\}$ the liquidity constraint binds with a strictly positive funding need. If $\xi_t = 0$ on a positive-probability subset of $\{y_t > 0\}$, then the binding liquidity constraint would carry zero marginal value there, so a small relaxation of the requirement on that subset would have no first-order value, contradicting optimality of the multiplier system for a strictly increasing objective. Hence $\xi_t > 0$ almost surely on $\{y_t > 0\}$.

The complementary-slackness conditions now translate on $\{\xi_t > 0\}$:

$$\ell^*_{i,t} > 0 \Rightarrow \phi_{i,t} = L_{i,t}, \qquad b^*_{i,t} > 0 \Rightarrow \phi_{i,t} = P_{i,t}, \qquad \theta^*_{i,t} > 0 \Rightarrow \phi_{i,t} = \frac{1}{\xi_t}\mathbb{E}_t[\xi_{t+1}\phi_{i,t+1}].$$

Step 5: Minimal supporting selection under discrete shocks. When y_t has atoms, the set of multipliers ξ_t consistent with KKT optimality at a given state ω may be a nondegenerate interval:

$$\Xi_t(\omega) := \left\{\xi_t(\omega) : \text{KKT conditions hold at } (\omega, t)\right\}.$$

Define the *minimal supporting multiplier*

$$\underline{\xi}_t(\omega) := \operatorname{ess\,inf}\ \Xi_t(\omega).$$

Assumption 5 rules out the boundary case $y_t = \sum_i \theta_{i,t-1} L_{i,t}$ state by state, which is the only configuration that generates an unbounded Ξ_t in the linear liquidation subproblem at date t (by the same mechanism as in the three-date linear program, Appendix A.1). Hence $\underline{\xi}_t < \infty$ almost surely and the minimal supporting multiplier is well defined. It is nonnegative by construction, and by Step 4 it is strictly positive on $\{y_t > 0\}$.

Using $\xi = \underline{\xi}$ throughout the construction above yields a single-valued multiplier process (ξ_t) and a single-valued discounted shadow-value process $(\rho_{i,t})$. On the event $\{\xi_t > 0\}$ this in turn yields single-valued shadow prices $(\phi_{i,t})$ via $\phi_{i,t} = \rho_{i,t}/\xi_t$. This completes the proof. □

A.8 Proof of Proposition 5 (Precautionary Pecking-Order Reversal Boundary)

Setup and notation. Fix the two-interim-period environment of the proposition. Normalize all terminal continuation values to one. Write

$$C^N \;:=\; \frac{1}{\ell_2^N} - 1 \;>\; 0, \qquad C^S \;:=\; \frac{1}{\ell_2^S} - 1 \;>\; 0, \tag{24}$$

for the *excess continuation-value cost* (relative to the buffer) of raising one unit of cash using asset 2 in the normal state and the stress state respectively. Since $\ell_2^S < \ell_2^N < 1$, both quantities are strictly positive and $C^S > C^N > 0$.

Because the terminal continuation value of each unit is one, *expected terminal wealth equals the expected number of units retained*. Both policies meet the $t = 1$ requirement $\bar{y}$ with equality; the difference in terminal wealth therefore arises solely from the expected number of continuation-value units sacrificed across both dates.

Cost accounting for Policy M (myopic pecking order). At $t = 1$, asset 1 (the buffer) is sold. Since $\ell_{1,1} = 1$, each unit of cash raised sacrifices exactly one unit of continuation value. There is no excess cost at $t = 1$.

At $t = 2$, the shock $y_2 = \bar{y}$ arrives with probability ρ_y. The buffer has been exhausted, so the agent must use asset 2 in the stress state $\ell_{2,2} = \ell_2^S$, incurring excess cost C^S per unit of cash raised.

The expected excess cost under Policy M, per unit of $\bar{y}$, is

$$\mathrm{EC}^M \;=\; 0 + \rho_y\, C^S. \tag{25}$$

Cost accounting for Policy D (dynamic buffer preservation). At $t = 1$, asset 2 is sold. Since $\ell_{2,1} = \ell_2^N$, each unit of cash raised sacrifices $1/\ell_2^N$ units of continuation value, an excess cost of C^N per unit of cash raised.

At $t = 2$, if the shock $y_2 = \bar{y}$ arrives (probability ρ_y), the preserved buffer covers it: $\ell_{1,2} = 1$ gives zero excess cost. If $y_2 = 0$ (probability $1 - \rho_y$), no further liquidation is needed.

The expected excess cost under Policy D, per unit of $\bar{y}$, is

$$\mathrm{EC}^D \;=\; C^N + \rho_y \cdot 0 \;=\; C^N. \tag{26}$$

Comparison and threshold. Policy D yields strictly higher expected terminal wealth than Policy M if and only if $\mathrm{EC}^D < \mathrm{EC}^M$, that is,

$$C^N \;<\; \rho_y \, C^S.$$

Rearranging and substituting the definitions (24),

$$\rho_y \;>\; \frac{C^N}{C^S} \;=\; \frac{1/\ell_2^N - 1}{1/\ell_2^S - 1} \;=:\; \rho_y^*.$$

This is (18). Because $C^S > C^N > 0$ and $\rho_y \in (0, 1)$, the threshold satisfies $\rho_y^* \in (0, 1)$, so the reversal region $(\rho_y^*, 1)$ is non-empty.

Uniqueness and boundary cases. At $\rho_y = \rho_y^*$ the two policies yield equal expected terminal wealth; for $\rho_y < \rho_y^*$ Policy M dominates. The threshold is strictly decreasing in the stress ratio ℓ_2^S/ℓ_2^N: a larger deterioration of secondary-market liquidity under stress raises C^S relative to C^N, lowering ρ_y^* and expanding the reversal region. As $\ell_2^S \to 0$, $C^S \to \infty$ and $\rho_y^* \to 0$: the buffer is almost always worth preserving regardless of persistence. As $\ell_2^S \to \ell_2^N$, $C^S \to C^N$ and $\rho_y^* \to 1$: the reversal region vanishes because the stress-state deterioration provides no additional incentive to preserve the buffer.

Link to the Snell recursion (Theorem 3). Under the risk-neutral benchmark, ξ_t is constant and the Snell recursion (17) reduces to

$$\phi_{i,t} = \max\left\{L_{i,t},\ \mathbb{E}_t\left[\phi_{i,t+1}\right]\right\}.$$

For asset 1 (the buffer) at $t = 1$ under Policy D:

$$\phi_{1,1} = \max\left\{L_{1,1},\ \mathbb{E}_1[\phi_{1,2}]\right\} = \max\{1,\ 1\} = 1.$$

The holding condition $\theta^*_{1,1} > 0$ is consistent since the buffer is retained and delivers $\phi_{1,2} = 1$ in both states at $t = 2$.

For asset 2 at $t = 1$ under Policy D, the agent sells ($\ell^*_{2,1} > 0$) and so the complementary-slackness condition requires $\phi_{2,1} = L_{2,1}$. The $t = 2$ shadow price satisfies $\phi_{2,2} = \max\{L_{2,2}, \mathbb{E}_2[\phi_{2,3}]\} = \max\{\ell_2^S, 1\} = 1$ (since $\ell_2^S < 1$ and the terminal value is one). Thus the continuation value of asset 2 at $t = 1$ is $\mathbb{E}_1[\phi_{2,2}] = 1$, while its liquidation value is $\ell_2^N < 1$; the agent sells, consistent with $\phi_{2,1} = \ell_2^N$.

Policy D is optimal (Theorem 3 complementary-slackness) precisely when the option value of the buffer – the expected savings from avoiding the stress-state cost C^S – exceeds its current-period liquidation advantage. This option value equals $\rho_y C^S$, and the liquidation advantage of the buffer at $t = 1$ relative to asset 2 equals C^N. Equating these yields $\rho^*_y = C^N/C^S$, confirming that (18) is the break-even persistence implied by the general Snell optimality conditions. □