OLD SHOREHAM INVESTMENT MANAGEMENT

Behaviourally Realisable Return

Forecasting the Investment an Investor Can Actually Hold

Working Paper V.0.1

Oladayo Oduwole

September, 2026

Behavioural Risk and Investment Engineering

Behaviourally Realisable Return

Forecasting the Investment an Investor Can Actually Hold

Oladayo Oduwole

Old Shoreham Investment Management

Disclosure. The author is a director of Old Shoreham Investment Management (OSIM), which designs and manufactures structured investment products, including protected notes.

Working paper, Version 1.1, 3 September 2026

Abstract

Investment products are forecast as return paths, yet investors are matched to them through static risk categories. This paper defines Behaviourally Realisable Return as the annualised certainty equivalent of terminal wealth once a forecast implementation policy, the investor's predicted pattern of intervention and continuation, has acted on a product's forecast paths. Welfare remains conventional expected utility and behaviour enters only through implementation, so any validated investor representation can supply the policy. Implementation separates into incidence, timing and consequence, and product families must be compared at their respective optima under common behavioural uncertainty. In a three-period complete-market example a floored note suppresses a loss-triggered exit that ordinary de-risking does not, and the optimal floor rises with risk aversion, from about 85 percent under log utility to par at a relative risk aversion of five, where the note is dominated. The behavioural advantage supports 33 to 59 basis points of incremental upfront margin.

Keywords: behavioural finance, suitability, implementation risk, structured products, certainty equivalent

Practical Applications

Advisers match clients to investments using risk questionnaires, then hand them a product whose path they may not be able to hold. This article gives a way to score that gap in return units before the recommendation is made, so two otherwise suitable products can be compared on what the client is expected to keep rather than on what the product is expected to earn. It also sets a commercial boundary: in the calibration here the behavioural advantage of a protected note supports only 33 to 59 basis points of extra upfront margin, well below typical structured-product pricing, so the design pays only for clients who want high exposure and carry a tight loss boundary.

1. Introduction

Investment analysis is forward-looking. Portfolios are assessed using expected returns, volatility, correlations, scenarios and distributions of future wealth. Suitability often treats the investor more statically. A client answers a questionnaire, receives a score and is matched to an investment whose path remains dynamic.

The mismatch matters because the investor can alter the strategy after recommendation. Selling after a fall, reducing risk, failing to rebalance, switching investments or re-entering later changes subsequent wealth. The terminal distribution actually experienced can therefore differ from the distribution used to justify the original investment.

The ex-post distinction between investment returns and investor returns is established. Dichev (2007), Friesen and Sapp (2007), Morningstar (2026) and the recent critique by Fulkerson et al. (2026) approach that gap from different directions. Behavioural research documents heterogeneous selling, reference dependence, panic selling and dynamic inconsistency. The question here is narrower and ex ante: before an investment is selected, can a forecast of likely implementation behaviour change which otherwise-suitable investment is expected to deliver the greatest value to this investor?

Behaviourally Realisable Return is designed for that comparison. A product generates forecast paths. Those paths create experienced states. An investor-response model maps those states into probabilities of holding, de-risking, selling, switching or adding risk. A continuation rule determines what follows. BRR is the annualised certainty equivalent of the resulting terminal-wealth distribution.

The welfare interpretation is deliberately conventional. Following the planner-doer language of Thaler and Shefrin (1981), BRR evaluates the forecast actions of the later doer from the ex-ante planner's terminal-wealth criterion. It does not reward an action merely because it is behavioural. If $V_j(i)$ is expected utility when investment i is implemented as intended and $J_j(i)$ is expected utility under the forecast implementation policy, define the implementation effect

$$L_j(i) = V_j(i) - J_j(i).$$

When the intended policy is already optimal over the admissible strategy set and intervention contains no new information, $L_j(i) \geq 0$. Outside that case the sign can reverse; the worked example contains such a continuation.

The paper's main contribution is operational rather than a new preference theory. Implementation is separated into three objects: *incidence*, whether intervention occurs; *timing*, when it occurs; and *consequence*, what the intervention does to subsequent wealth. A protective payoff and a lower-risk constant mix can therefore solve different parts of the same implementation problem.

This distinction leads to two practical tests. First, product families must be compared symmetrically: best family against best family under the same behavioural uncertainty. Second, any behavioural advantage must exceed the incremental manufacturing and distribution cost required to obtain it. The resulting breakeven margin converts an implementation benefit into a commercial boundary.

The paper uses a three-period binomial model because every state can be audited. The market is complete, so the protected note is a replicable dynamic strategy rather than a free source of value. The honest comparison is between constant-mix exposure and an option-replicating nonlinear payoff that a retail investor may be unable or unwilling to reproduce and maintain. The tree also makes clear what is not general. Annual monitoring creates atoms in experienced losses and therefore literal cliffs and corner solutions. A monthly lognormal robustness exercise removes those artefacts while preserving the underlying implementation wedge.

Utility matters as well. Under log utility the illustrative unrestricted note family can beat the constant-mix family. Under CRRA with $\gamma > 1$ in the same tree, the ranking reverses because the investor's conventional optimum moves towards or below the first behavioural boundary, leaving less implementation drag for protection to remove. The product-design claim is therefore conditional, not universal.

The paper proceeds from the definition of BRR to experienced-state variables, investor-response information, continuation, and the incidence-timing-consequence decomposition. It then develops ranking, uncertainty, family and margin tests, positions them against the stop-loss, portfolio-insurance, planner-doer and behavioural-portfolio literatures, and works through the exact tree. A robustness section varies utility and monitoring frequency before the paper turns to practical implementation and empirical burden.

2. Behaviourally Realisable Return

BRR is a forward-looking return measure for the investor-product system rather than for the product alone. Start with the terminal-wealth distribution implied by the investment's intended rule. Apply the investor's forecast implementation policy to the intervening path, including any sale, de-risking, switch, re-entry or other continuation. Evaluate the resulting terminal wealth using the same conventional expected-utility function that would be used to value the intended strategy, convert that value to a certainty equivalent, and annualise it. The full notation and derivation are reported in Online Appendix A.

The distinction is between intended and implemented wealth. Hold certainty equivalent asks what the product is worth if the investor follows the stated rule. BRR asks what the same product is worth after the investor's predicted actions have altered the path. Their difference is implementation drag. The drag need not be positive in every conceivable strategy set, but when the intended policy is already optimal over the admissible set and intervention contains no new information, an unplanned deviation cannot improve ex-ante welfare.

Because BRR is expressed in annualised certainty-equivalent return units, it can be compared directly across otherwise suitable products. No behavioural utility is required: behaviour changes the implementation policy, not the welfare criterion. This separation allows any investor representation with validated out-of-sample predictive power to supply the implementation policy without changing the valuation framework.

3. Experienced path variables

Investors experience sequences rather than terminal distributions.

A portfolio that loses 15% in a month need not produce the same behaviour as one that loses 15% over three years. An investment that eventually recovers after a long period underwater can be harder to hold than another investment with the same terminal return. A positive absolute return can still provoke intervention when a benchmark has risen much faster.

Following the notation in Oduwole (2026a), relevant path coordinates include loss from inception,

$$g_t^L = 1 - W_t, \tag{1}$$

$$g_t^D = 1 - \frac{W_t}{\max_{s \le t} W_s}, \tag{2}$$

$$g_t^G = W_t - 1, \tag{3}$$

$$g_t^F = 1 - \frac{W_t}{R_t}. \tag{4}$$

Other relevant quantities can include time underwater, time since visible progress, irregularity of outcomes and the frequency with which the investment is observed.

These quantities are derived from the forecast paths. They are not a second independent market dataset.

Some coordinates require information beyond the product's own wealth path. Reference shortfall, for example, requires the reference asset or benchmark to be simulated alongside the product.

This matters because the investor can respond to what the investment failed to capture as well as to what it lost.

4. FREBO and implementation

FREBO is a five-dimensional description of investment behaviour intended to capture differences that a single risk score is not designed to represent.

The five dimensions are Fortress, Rhythm, Engagement, Blueprint and Opportunism.

Their scoring methodology is proprietary. Their roles in the implementation framework can still be stated directly.

Table 1. FREBO dimensions and implementation channels

FREBO dimension	Meaning	Implementation channel	Typical action
Fortress	Protection, preservation and response to loss exposure	Loss g^L, drawdown g^D	De-risk, liquidate, move toward safety
Rhythm	Preference for regular outcomes and visible progress	Time underwater, time since progress, outcome irregularity	Switch, reduce exposure, seek cadence
Engagement	Monitoring and emotional interaction with money in motion	Observation frequency and attention intensity	Increases opportunities for intervention across channels
Blueprint	Attachment to plans, goals and defined journeys	Persistence with intended strategy or deviation from plan	Hold, rebalance to plan, resist discretionary change
Opportunism	Appetite to seize upside or asymmetric opportunity	Relative shortfall g^F, salient upside elsewhere	Add risk, switch, re-risk

The mapping is deliberately not forced into five identical hard thresholds.

Fortress

Fortress concerns capital preservation and response to loss. The natural path coordinates are loss from inception and drawdown from a peak.

A high-Fortress investor may be more likely to reduce risk when one of these coordinates becomes sufficiently uncomfortable.

Rhythm

Rhythm concerns when outcomes arrive.

Long periods underwater, irregular distributions or extended stretches without visible progress can cause intervention independently of the final expected return.

Engagement

Engagement concerns monitoring intensity and emotional interaction with investments while outcomes remain unresolved.

It differs from a conventional threshold coordinate because it can alter the observation process itself. An investor who checks an investment every day has more opportunities to encounter and act on a boundary than an otherwise identical investor who reviews it quarterly.

Engagement therefore changes incidence across multiple channels.

Blueprint

Blueprint concerns adherence to an intended plan, goal or predefined journey.

It does not map cleanly onto one of the four trigger coordinates above, and there is no reason to invent a fifth merely to make the framework symmetrical.

Its natural role is in persistence. A stronger Blueprint coordinate may increase the probability that an investor follows the intended policy, rebalances back towards a plan or treats a temporary path state as insufficient reason to alter course.

A future empirical model may support a distinct deviation-from-plan statistic. This paper does not assume one.

Opportunism

Opportunism concerns attraction to salient upside and willingness to re-risk when another opportunity appears more compelling.

The natural coordinate is relative shortfall. Participation-limited investments can expose an investor to this channel even while protecting the loss channel.

This is important for structured products. Improving fit on Fortress does not imply improvement on Opportunism.

The worked example below activates only the Fortress channel. It is therefore a one-coordinate demonstration of BRR rather than a complete FREBO implementation.

5. Actions and continuation

An implementation policy maps experienced states and investor characteristics into actions.

A useful action set is

$$\mathcal{A} = \{\text{hold}, \text{sell}, \text{de-risk}, \text{switch}, \text{add risk}\}.$$

These categories can be made more granular. A sale can be partial. De-risking can involve cash, bonds or a lower-risk portfolio. Switching requires an identified destination. Adding risk can involve additional capital or a change in exposure.

The continuation after intervention is economically essential.

Selling an equity fund and remaining in cash produces one terminal distribution. Selling and returning later produces another. Switching directly into the benchmark produces another.

The action also determines whether a protective feature remains relevant after intervention.

An investor who liquidates a protected note and reinvests the proceeds risklessly can preserve the economic value of its discounted floor under the conditions developed below.

An investor who sells the note and switches into equities no longer owns the contractual floor.

The distinction links directly to FREBO. A Fortress-driven intervention into safety and an Opportunism-driven switch into a risky alternative can be triggered by the same product but produce completely different consequences.

6. Incidence, timing and consequence

Let τ_ij denote the first intervention time and A_ij the corresponding action.

For action a at date t, let

$$W^1_{ij,T}(t,a)$$

denote terminal wealth after the continuation rule has been applied.

The utility consequence of the intervention on the same underlying market path is

$$\Delta U^a_{ij,t} = U_j(W^0_{iT}) - U_j\left(W^1_{ij,T}(t,a)\right). \tag{5}$$

In expected-utility units, total implementation effect is

$$L_{ij} = V_{ij} - J_{ij}. \tag{6}$$

For a first-intervention model,

$$\boxed{L_{ij} = \sum_{t=1}^{T} \sum_{a \in \mathcal{A}} p^a_{ij,t} s^a_{ij,t}.} \tag{7}$$

where

$$p^a_{ij,t} = \mathbb{P}(\tau_{ij} = t, A_{ij} = a), \tag{8}$$

$$s^a_{ij,t} = \mathbb{E}\left[\Delta U^a_{ij,t} \mid \tau_{ij} = t, A_{ij} = a\right]. \tag{9}$$

This separates three economic margins.

Incidence describes how likely intervention is.

Timing describes when it occurs.

Consequence describes what the action does to terminal wealth when it occurs.

A product can therefore have a high intervention rate and modest implementation drag where the interventions are cheap.

A second product can have fewer interventions but greater drag if those interventions occur at damaging points in the path.

This decomposition becomes particularly useful in comparing conventional de-risking with payoff protection.

7. Risk-budgeted BRR and the path-constraint limit

BRR is applied only after conventional suitability and financial-risk constraints have defined the admissible set. It is therefore a ranking device within that set, not a substitute for risk budgeting and not a new behavioural efficient frontier. Using post-intervention volatility as the primary risk budget would be misleading because a damaging sale into cash can mechanically lower subsequent volatility while reducing terminal wealth.

A known, stable and finely observed behavioural boundary is an important limiting case. Once the trigger and continuation rule are pinned down, the implementation problem can collapse to an ordinary wealth-path constraint, so BRR adds little optimisation structure. Its incremental content is greatest when the boundary is uncertain, monitoring is imperfect, continuation differs across investors or product families alter intervention incidence and consequence differently. Discrete monitoring can also create visible cliffs when a positive-probability state crosses a hard boundary; with continuous state distributions the same mechanism is normally smooth or kinked rather than literally discontinuous. The formal statements are in Online Appendix B.

8. Product ranking and the benchmark test

Suppose product A has higher value under continuous holding:

$$V_{Aj} > V_{Bj}.$$

Implementation reverses the ranking where

$$J_{Bj} > J_{Aj}.$$

Since $J\_ij = V\_ij - L\_ij$, the reversal occurs precisely when

$$\boxed{L_{Aj} - L_{Bj} > V_{Aj} - V_{Bj}.} \tag{10}$$

The identity itself is elementary.

Its importance is that the behavioural advantage and the conventional disadvantage are measured in comparable units.

8.1 Threshold uncertainty

The relevant behavioural boundary is rarely known exactly before product selection. BRR therefore treats the investor's threshold or response parameters as a distribution and averages expected utility over that uncertainty before converting the result to a certainty equivalent. This ordering matters under nonlinear utility and avoids averaging already-transformed return measures. The point estimate is recovered as the special case in which the distribution collapses to one value.

8.2 Best family against best family

A fair comparison optimises both product families under the same behavioural uncertainty. A single protected note should not be compared with the best member of a conventional allocation family, and a single constant-mix portfolio should not be compared with an optimised note family. Feasible-set restrictions also matter: a family that permits sub-par floors and participation above 100 percent is economically different from one restricted to full capital protection. The worked example therefore compares each family at its own optimum under the same threshold distribution.

8.3 Breakeven margin

Behavioural fit has commercial value only up to the incremental cost it can support. The breakeven margin is the additional upfront manufacturing or distribution margin at which the best protected structure and the best conventional allocation have equal BRR. It converts an implementation advantage into a price boundary. A payoff may improve implementation and still be a poor product if the cost of manufacturing and distributing it exceeds that boundary. Formal definitions of threshold uncertainty, symmetric family optimisation and the breakeven margin are in Online Appendix C.

Sharpe ratio and BRR answer different questions. Sharpe evaluates excess return relative to volatility under a specified investment rule; BRR evaluates certainty-equivalent return after the investor's forecast implementation policy has acted on the path. A technically efficient portfolio can therefore have a lower BRR for a particular investor, while a product with lower conventional efficiency can have higher implemented value. BRR does not displace the Sharpe ratio. It adds the investor-specific question of how much of the investment's forecast value is expected to survive implementation.

9. Empirical and theoretical antecedents

BRR sits between portfolio insurance, stop-loss rules, self-control, behavioural portfolio choice and investor-return measurement. The distinctions matter because none of those literatures makes investor-specific forecast implementation value the product-selection criterion used here.

9.1 Stop-loss rules and dynamic strategies

The loss-trigger rule in the worked example is a stop-loss policy. Kaminski and Lo (2014) show that such rules can add or destroy value depending on the return process; under random-walk prices they reduce expected return, while momentum can create conditions in which they add value. BRR does not claim that stopping after a loss is optimal. It treats the stop as a forecast implementation response and asks how much value remains after that response.

Leland (1980), Brennan and Solanki (1981), and Benninga and Blume (1985) establish that portfolio insurance is costly and can be dominated by ordinary risk reduction under standard preferences. Dybvig (1988) shows that path-dependent dynamic strategies can be cost-inefficient relative to distributionally equivalent alternatives. That observation provides a natural interpretation of implementation drag: when the investor's path-dependent trading is uninformed and the intended strategy is optimal, BRR cannot exceed the conventional certainty equivalent merely because the path-dependent strategy is packaged as behaviour.

The binomial market used below is complete. The floored note is therefore replicable by dynamic trading in the risky and riskless assets. Its distinctive economic role in the example is not creation of an unattainable payoff. It is packaging a nonlinear exposure that a retail investor may be unable or unwilling to replicate and maintain without intervention.

9.2 Planner, doer and welfare

Thaler and Shefrin (1981) model self-control as a conflict between a farsighted planner and a myopic doer. BRR adopts a narrow version of that separation. Welfare is measured from the ex-ante planner's stated terminal-wealth utility, while the later doer's behaviour is forecast rather than treated as welfare-improving by definition. Oduwole (2026a) develops the commitment and filtration issues more formally.

This choice is important. If behavioural actions were evaluated using a preference system constructed to rationalise those same actions, superior BRR could become tautological. The present framework instead asks whether the forecast doer behaviour raises or lowers the planner's conventional terminal-wealth value.

9.3 Path constraints and behavioural portfolio theory

Grossman and Zhou (1993) study optimal portfolio choice subject to drawdown constraints, and Basak and Shapiro (2001) show how formal risk constraints alter portfolio decisions. When an exit-stable behavioural boundary is known and finely observed, Oduwole (2026a) shows that the BRR problem can reduce to this familiar path-constraint class. Threshold uncertainty and continuation heterogeneity are therefore central to the distinct content of the present paper.

Shefrin and Statman (2000) develop Behavioral Portfolio Theory and explicitly define a behavioural efficient frontier. Das et al. (2010) integrate mental accounting with Markowitz optimisation. BRR uses neither aspiration-based utility nor mental-account risk definitions. It retains conventional welfare and conditions realised value on forecast implementation, which is why Section 7 uses a risk-budgeted value function rather than claiming a new behavioural efficient frontier.

9.4 Behavioural action evidence

Kahneman and Tversky (1979) establish reference-dependent decision-making. Shefrin and Statman (1985), Odean (1998), Kaustia (2010), and Ben-David and Hirshleifer (2012) document systematic selling responses to gains and losses. Dhar and Zhu (2006) document investor heterogeneity; An et al. (2024) show that selling of an individual position depends on the wider portfolio; Elkind et al. (2022) show that panic selling is predictable from investor and account characteristics; Heimer et al. (2025) provide evidence that intended and later risk-taking can diverge; and Gelman et al. (2026) document re-risking associated with falling behind salient benchmarks.

These findings support an investor-response model. They do not validate a universal hard threshold, and they do not validate FREBO. The empirical burden remains prospective prediction of action and continuation.

9.5 Investor-return gaps and the 2026 dispute

Dichev (2007) and Friesen and Sapp (2007) distinguish investment performance from the returns investors actually earn through their cash-flow timing. Morningstar (2026) reports that the average dollar invested in US mutual funds and ETFs earned 8.7% annually over the decade to 31 December 2025, against a 9.9% aggregate total return. The same report shows wide heterogeneity across product types and particularly large gaps in highly tradeable, volatile categories.

Fulkerson, Jordan, Riley and Yan (2026) challenge the interpretation of the usual aggregate return-gap statistic, arguing that retrospective internal-rate-of-return calculations mix genuine timing with a hindsight component. That dispute reinforces rather than resolves the problem addressed here. BRR does not infer behaviour from an aggregate dollar-weighted residual. It forecasts specified actions conditional on paths and maps those actions directly into subsequent wealth.

9.6 Security design and assignment

Calvet, Célérier, Sodini and Vallée (2023) find that the introduction of capital-guaranteed products in Sweden was associated with greater household risk-taking and higher expected financial portfolio returns, particularly among households with low ex-ante risk appetite. The result shows that payoff geometry can alter participation and portfolio choice; it does not establish that protection universally raises welfare.

Oduwole (2026b) develops the assignment problem. Investor information has economic value only if it changes the preferred product, allocation or payoff after costs. The best-family and breakeven-margin tests below are the corresponding objects in the present calibration.

10. Product design and continuation

Behavioural product design can work through incidence or consequence.

A lower-risk conventional allocation principally reduces incidence. Smaller path movements make a given loss or drawdown boundary less likely to be reached.

A contractual floor can reduce consequence.

Consider a product with maturity floor F. Suppose the issuer is default-free, liquidation occurs at arbitrage-free value, and the investor places liquidation proceeds in the riskless asset until maturity.

At date t, the product value must be at least the value of the floor component,

$$\frac{F}{(1+r)^{T-t}},$$

because any residual optional payoff is non-negative.

Liquidation at that value followed by riskless reinvestment therefore produces at least

$$F$$

at maturity.

The result holds across liquidation thresholds in that continuation class.

It does not follow that the floor prevents intervention.

An investor can still sell.

The economic result is that intervention need not destroy the protected terminal wealth.

10.1 The continuation restriction

This result belongs to a particular implementation channel.

A Fortress-driven investor who sells the note and moves to cash retains the floor economics under the stated assumptions.

An Opportunism-driven investor who sells the note and switches into the risky asset does not.

The contractual guarantee remains attached to the note, not to the investor after the note has been sold.

The worked model illustrates the distinction. At a 10% loss threshold, an investor liquidating the 90% note at the DD node and reinvesting risklessly retains a BRR of 5.72%. Replacing the continuation with a switch into the risky underlying gives approximately

$$6.15\%.$$

The outcome is better in this particular tree, but the guarantee has disappeared. Future wealth is again exposed to the underlying asset.

The sign can therefore vary, while the structural point does not.

A product feature can protect one FREBO channel and leave another untouched or even increase its exposure.

10.2 The design problem

Protection and de-risking should therefore be regarded as alternative implementation technologies.

Linear de-risking changes the return process on every path.

A floor uses option budget to reshape selected states.

Their relative value depends on market economics, investor behaviour and continuation.

Calvet et al. (2023) provide empirical evidence that nonlinear protection can alter household willingness to take risk. The portfolio-insurance literature provides the counterweight that protection is not costless and can be dominated by ordinary allocation.

BRR supplies a common unit in which the two can be compared.

11. Worked example

Each period in this section represents one year. The model is a complete binomial market, so every contingent payoff used below is dynamically replicable. The comparison is therefore between implementation technologies, not between a replicable note and an otherwise unavailable source of return. The risky asset begins at $S_0 = 1$ and moves each year by $u = 1.35$ or $d = 0.85$, with real probability one half for either move. The riskless gross return is $R_f = 1.02$. Utility is logarithmic, $U(W) = \log W$, and the horizon is three years.

11.1 Hold value of the risky asset

After three periods, the possible terminal values are:

Table 2. Terminal wealth distribution of the risky asset

Up moves	Terminal wealth	Probability
0	$0.85^3 = 0.614125$	1/8
1	$1.35(0.85)^2 = 0.975375$	3/8
2	$(1.35)^2(0.85) = 1.549125$	3/8
3	$(1.35)^3 = 2.460375$	1/8

Expected log wealth is

$$\begin{aligned}\mathbb{E}[\log W_T] &= \frac{1}{8}\log(0.614125) + \frac{3}{8}\log(0.975375)\\ &\quad + \frac{3}{8}\log(1.549125) + \frac{1}{8}\log(2.460375)\\ &= 0.206378.\end{aligned}$$

The terminal certainty equivalent is $\widehat{C}_U = e^{0.206378} = 1.229218$. Annualising gives

$$C_U = 1.229218^{1/3} - 1 = 0.071214 = 7.12\%.$$

11.2 A 90% floored note

Consider the payoff

$$W_N(T) = 0.90 + \pi \max(S_T - 1, 0). \tag{11}$$

The present value of the floor is $0.90/1.02^3 = 0.848090$. This leaves an option budget of 0.151910. The risk-neutral up probability is

$$q = \frac{1.02 - 0.85}{1.35 - 0.85} = 0.34.$$

The at-the-money call is worth

$$C_0 = \frac{q^3(1.460375) + 3q^2(1-q)(0.549125)}{1.02^3} = 0.172527.$$

Participation is therefore

$$\pi = \frac{0.151910}{0.172527} = 0.880501 = 88.05\%.$$

The note's terminal payoffs are

$$0.90, \quad 0.90, \quad 1.383505, \quad 2.185861$$

for zero, one, two and three up moves respectively.

Expected log wealth is

$$\begin{aligned}\mathbb{E}[\log W_{N,T}] &= \frac{1}{2}\log(0.90) + \frac{3}{8}\log(1.383505) + \frac{1}{8}\log(2.185861) \\ &= 0.166804.\end{aligned}$$

The terminal certainty equivalent is 1.181522, giving $C_N = 5.72\%$.

The risky asset is therefore approximately 140 basis points better under continuous holding.

11.3 A 12% loss threshold

Consider an investor whose Fortress channel causes liquidation after a loss of 12% or more from inception.

The threshold wealth level is 0.88. After one down move the risky asset is worth 0.85, so the investor exits and reinvests risklessly.

Terminal wealth on every down-first path becomes

$$0.85(1.02)^2 = 0.884340.$$

The resulting expected implemented log wealth is 0.157388 and the terminal certainty equivalent is 1.170450, so $\mathrm{BRR}_U = 5.39\%$.

There is an observation-grid effect inside this calculation.

The investor is described as having a 12% loss threshold, but the market is observed only once each year. The first down move is 15%, so liquidation occurs three percentage points beyond the nominal threshold.

The realised trigger is therefore partly a property of the investor and partly a property of the monitoring grid.

This is where Engagement can enter even in a one-coordinate Fortress example. More frequent observation changes which crossings are seen and how much overshoot occurs.

11.4 Mark-to-market value of the note

After one down move, the note is worth

$$V_{N,1}^{D} = 0.918775.$$

Its mark-to-market loss is therefore $1 - 0.918775 = 0.081225$, or approximately 8.12%.

After two down moves,

$$V_{N,2}^{DD} = \frac{0.90}{1.02} = 0.882353.$$

The mark-to-market loss there is approximately 11.76%.

At a 12% threshold, neither node triggers, so $\mathrm{BRR}_N = 5.72\%$. The note therefore beats the full-risk position for this investor, $5.72\% > 5.39\%$.

11.5 The note also has a cliff

The note is not behaviourally smooth.

Its first economically costly boundary occurs at the first-period down node:

$$\theta_N = 1 - 0.918775 = 8.1225\%.$$

For thresholds above approximately 8.12%, the first down move does not trigger. BRR remains approximately 5.72%.

For thresholds below approximately 8.12%, the investor exits the note after a first-period down move. The investor then gives up the remaining upside option on paths that subsequently recover.

BRR falls to approximately 4.89%, a discrete step from 5.72% to 4.89%.

The fall is approximately 83 basis points.

A separate event occurs at the DD node for thresholds below 11.76%. The investor may technically liquidate there, but the upside option is already worthless. Selling at fair value and placing the proceeds in the riskless asset gives

$$0.882353(1.02) = 0.90.$$

That is exactly what holding the note to maturity would have produced.

Incidence has changed.

Consequence has not.

This difference is central to the robustness of the note.

11.6 The linear allocation

Now consider an annually rebalanced allocation with risky weight x.

Following an up move,

$$R_U(x) = 1.02 + 0.33x.$$

Following a down move,

$$R_D(x) = 1.02 - 0.17x.$$

The 12% intervention boundary creates two important discontinuities.

Two consecutive down moves first cross the boundary when

$$[R_D(x)]^2 = 0.88.$$

This occurs at

$$x = \frac{1.02 - \sqrt{0.88}}{0.17} \approx 0.4819.$$

A single down move reaches the boundary when

$$1.02 - 0.17x = 0.88,$$

which gives

$$x = \frac{0.14}{0.17} = 0.823529.$$

The 82.35% solution is therefore a corner generated by the intervention rule.

Immediately below that boundary, at $x = 0.8235$, the one-period down factor is 0.880005, fractionally above the 0.88 trigger. The resulting BRR is approximately 6.23%, while hold value is approximately 6.62%.

At the known 12% threshold:

Table 3. Hold value and BRR at a known 12% loss threshold

Product	**Hold value**	**BRR**
Unprotected risky asset	7.12%	5.39%
90% floored note	5.72%	5.72%
82.35% risky / 17.65% riskless	6.62%	**6.23%**

11.7 BRR across the risky weight

The shape of BRR makes the corner visible.

The large second jump is approximately 117 basis points.

This is larger than the note's roughly 83-basis-point first-down cliff.

The more important difference is location.

The note's costly cliff sits near an 8.12% threshold. The linear portfolio's large cliff at a 12% threshold sits in the middle of the behavioural range used below.

11.8 Piecewise BRR under annual monitoring

The steps in the preceding table can be derived exactly from the annual tree. The closed-form piecewise expressions are useful for auditing the two intervention boundaries but are not needed for the economic argument, so they are reported in Online Appendix D. The discontinuities arise from discrete monitoring and positive probability mass at a small number of experienced-loss states; they are features of this pedagogical tree rather than a general prediction of BRR.

Table 4. BRR as a function of risky-asset weight at a 12% threshold

Risky weight	BRR
45.00%	5.00%
48.18%	5.17%
48.19%	4.90%
60.00%	5.42%
70.00%	5.81%
80.00%	6.15%
82.35%	**6.23%**
82.36%	5.06%
90.00%	5.21%
100.00%	5.39%

11.9 One basis point

Keep the risky allocation fixed at 82.35%.

With a 12.00% threshold, the down-state value 0.880005 sits just above the trigger, producing 6.23% BRR. Move the threshold to 11.99%, and the boundary becomes 0.8801. Since $0.880005 < 0.8801$, the investor exits after one down move and BRR falls to 5.06%.

The note remains at 5.72%.

This does not establish that linear de-risking is intrinsically fragile. It establishes that a linear allocation chosen using an exactly known behavioural boundary can be fragile to error in that boundary.

11.10 Re-optimisation under a known threshold

Re-optimising the linear portfolio for each known threshold gives:

Table 5. BRR-maximising risky weight under a known loss threshold

Known threshold	BRR-maximising risky weight	Best BRR
11.00%	76.47%	6.04%
11.25%	77.94%	6.09%
11.50%	79.41%	6.13%
11.75%	80.88%	6.18%
12.00%	82.35%	6.23%

The weights are determined by the intervention boundary:

$$x^*(\theta) = \frac{0.02 + \theta}{0.17} \tag{12}$$

over this region.

For a known threshold, some conventional allocation beats the fixed 90% note throughout the range.

That is the pointwise result in the annual tree. Section 12 shows that it reverses under monthly monitoring at a known 12% threshold, so the ranking is not treated as a general finding.

11.11 Threshold uncertainty and symmetric optimisation

Now suppose one investment must be selected before the threshold is known.

The linear family is optimised over x.

The note family must also be optimised over floor F .

The relevant comparison is therefore

$$\sup_F \mathrm{BRR}_N(F;\kappa) \qquad \text{against} \qquad \sup_x \mathrm{BRR}_M(x;\kappa).$$

The annual tree admits only a small number of experienced loss states, so only the locations of the cut points matter for the ranking. The following uniform bands are deliberately illustrative sensitivity bands around economically active cut points; they are not estimates of a population threshold distribution. A field implementation should estimate κ from observed investor behaviour. In the FREBO interpretation, Fortress would shift the posterior mass assigned to tighter or looser loss boundaries rather than mechanically determine one exact threshold.

Three illustrative uniform threshold distributions give:

Table 6. Best-family comparison under threshold uncertainty

Threshold distribution	Best linear BRR	Fixed 90% note	Best note BRR	Best floor
$U[8\%, 16\%]$	5.64%	5.70%	**5.88%**	84.82%
$U[10\%, 14\%]$	5.83%	5.72%	**6.08%**	87.0%
$U[10\%, 20\%]$	6.04%	5.72%	**6.19%**	82.6%

Under log utility and annual monitoring, the optimised nonlinear family wins in all three distributions tested. Section 12 shows that this ranking reverses under CRRA with $\gamma > 1$ in the same tree.

That result changes the interpretation of the example.

The earlier comparison between the best linear allocation and one fixed 90% note was asymmetric.

Best-of-family against best-of-family produces a different answer.

11.12 What the optimal floors actually are

The winning floors in the previous table are all below the initial investment amount.

That matters.

At a floor of approximately 84.82%, participation is around

$$116.3\%.$$

The optimum is not knife-edge: floors from roughly 83.9% to 85.7% remain within half a basis point of the maximum BRR.

At a floor of 87.0%, participation is approximately

$$104.4\%.$$

At a floor of 82.6%, participation is approximately

$$128.5\%.$$

The winning structures are therefore partial-protection, geared-participation products.

They are not conventional principal-protected notes.

The option budget explains the result. At a 2% riskless rate, a 100% maturity floor consumes almost all of the issue price.

11.13 Requiring protection at par

Now impose

$$F \geq 1.$$

Within the feasible zero-coupon range, the optimum occurs at

$$F = 1.$$

The residual option budget supports only about 33.43% upside participation. The resulting BRR is approximately 3.84%

under each of the three threshold distributions, since the note does not approach any of those behavioural boundaries.

Table 7. Effect of imposing a par-floor constraint

Threshold distribution	**Best linear**	**Best unrestricted note**	**Best note with $F \geq 1$**
$U[8\%, 16\%]$	5.64%	**5.88%**	3.84%
$U[10\%, 14\%]$	5.83%	**6.08%**	3.84%
$U[10\%, 20\%]$	6.04%	**6.19%**	3.84%

The unconstrained result therefore should not be read as evidence that capital-guaranteed products dominate conventional allocation.

It says something narrower and more interesting.

In this low-rate tree, the best nonlinear solution trades some downside floor for more than full participation in upside.

A higher-volatility market can change both the option budget and the likelihood that a sub-par structure reaches its own intervention boundary. The companion calibration addresses that different geometry.

11.14 Breakeven manufacturing margin

The nonlinear advantage is economically small.

Let m be an upfront margin deducted from the note's option budget.

For the fixed 90% note under $U[8\%, 16\%]$, the zero-margin BRR of 5.70% exceeds the best linear allocation of 5.64% by only six basis points a year.

The fixed note can absorb only about 12.6 bp of upfront margin before losing that advantage.

This is roughly four basis points a year over the three-year horizon.

The fixed 90% note already loses to the best linear portfolio in the other two distributions, so it has no positive margin capacity there.

Optimising the floor increases the available surplus.

This is the commercial boundary.

Table 8. Breakeven manufacturing margin under threshold uncertainty

Threshold distribution	Best linear	Best note	Breakeven margin	Annual equivalent
$U[8\%, 16\%]$	5.64%	5.88%	**51 bp**	17 bp
$U[10\%, 14\%]$	5.83%	6.08%	**48 bp**	16 bp
$U[10\%, 20\%]$	6.04%	6.19%	**33 bp**	11 bp

A 3% upfront margin is between roughly six and nine times the maximum upfront behavioural surplus in this calibration.

Relative to the fixed 90% structure's six-basis-point annual advantage in the only uncertainty band where it wins, the difference is larger still.

At $F = 0.90$, participation under margin m is

$$\pi(m) = \frac{1 - 0.90/1.02^3 - m}{0.172527}. \tag{13}$$

With $m = 3\%$, participation falls to approximately 70.66%.

That margin does not merely reduce the behavioural advantage.

It overwhelms it.

11.15 Relation to the companion calibration

The margin numbers above are much smaller than those obtained in the higher-volatility calibration used in the companion product-design work, but there is no contradiction. The amount a protective feature can economically support is determined by the implementation drag available to be removed and by the conventional return surrendered to obtain the protection.

Both are functions of the underlying return geometry.

A quiet tree with 35% up moves, 15% down moves and a 2% riskless rate produces a relatively small behavioural surplus.

A market with larger drawdowns, a larger risk premium and materially different option economics can generate a much larger implementation wedge and therefore a larger breakeven margin.

The margin is not a universal property of BRR or of structured products.

It is an output of the market and behavioural calibration.

For that reason, the breakeven-margin surface

$$m^* = m^*(\sigma, \mu, r, \kappa, F, \text{continuation})$$

is a more meaningful object than one quoted spread.

11.16 Secondary-market haircut

The floor result assumes fair-value execution.

At the DD node, the 90% note is worth

$$0.882353.$$

A 1% secondary-market haircut reduces executable value to

$$0.873529.$$

Riskless reinvestment gives approximately 0.891 at maturity.
The note's BRR falls by approximately nine basis points a year.
Liquidity therefore belongs inside the product comparison.
A contractual floor held to maturity is not a promise of frictionless secondary-market value.

11.17 Continuation changes the meaning of protection

At a 10% threshold, the DD node breaches the investor's loss boundary.
Under a Fortress-style continuation, the investor sells the note and reinvests risklessly.
The 90% terminal floor is preserved economically and BRR remains

$$5.72\%.$$

Now replace the continuation with an Opportunism-style switch into the risky asset.
The investor no longer owns the floor.
In this particular tree the resulting BRR rises to approximately

$$\boxed{6.15\%.}$$

The higher number is not a protection result.
It is the outcome of renewed risky exposure.
This is why the FREBO-to-action mapping matters. A product can be robust to one implementation channel and exposed to another.

11.18 What the example establishes

The tree establishes mechanisms, not universal product rankings. It shows exactly how incidence, timing and consequence enter implemented value; how a nonlinear payoff can alter the consequence of an exit; why product families must be optimised symmetrically under threshold uncertainty; and how little incremental margin the resulting advantage can support.

Several numerical features are deliberately local to annual monitoring. The 82.35% optimum, the one-basis-point fragility and the literal cliffs occur because the annual tree places positive probability on a few discrete experienced-loss states. They should be read as transparent demonstrations of state entry, not as predictions that a finely monitored market will contain the same corners.

The log-utility best-family result is also conditional. A more risk-averse investor chooses less risky exposure even before implementation is considered, leaving less behavioural drag for a protected payoff to remove. The robustness section quantifies that reversal.

The durable product-design condition is therefore comparative: protection is useful only when the implementation loss it removes exceeds both the conventional value surrendered to manufacture it and its incremental cost. This condition is most plausible for an investor who wants substantial upside exposure while carrying a relatively tight loss boundary.

12. Robustness: utility and monitoring frequency

The exact tree is useful because every number is auditable, but two of its most visible results depend on annual monitoring and its product-family ranking depends on log utility. This section separates those artefacts from the mechanism that survives.

12.1 CRRA utility

Replace log utility with

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma}, \qquad \gamma \neq 1,$$

with the log case recovered as $\gamma \to 1$. Re-optimising both families in the annual tree changes the ranking materially.

Table 9. Utility sensitivity of the best-family comparison

Utility	Best linear BRR	Best note BRR	Family ranking
$\gamma = 1$, $U[8, 16]$	5.64%	5.88%	Note
$\gamma = 1$, $U[10, 14]$	5.83%	6.08%	Note
$\gamma = 1$, $U[10, 20]$	6.04%	6.19%	Note
$\gamma = 2$	4.41–4.52%	4.21%	Linear
$\gamma = 3$	3.75–3.80%	3.42%	Linear
$\gamma = 5$	3.08%	2.83%	Linear

For $\gamma = 2$, the best linear values across $U[8, 16]$, $U[10, 14]$ and $U[10, 20]$ are 4.41%, 4.44% and 4.52%, respectively; the best note is approximately 4.21% in each band. For $\gamma = 3$ the best linear values lie between 3.75% and 3.80%, against 3.42% for the note. Reported figures are rounded to the basis point.

The floor choice moves in the same direction as risk aversion.

Table 10. Optimal note floor by risk aversion under $\theta \sim U[8\%, 16\%]$: annual-tree and monthly-monitoring calibrations

CRRA coefficient	Annual tree	Monthly lognormal
$\gamma = 1$	84.82%	about 86%
$\gamma = 2$	91.2%	about 92%
$\gamma = 3$	96.6%	about 96%
$\gamma = 5$	about 101%	100.0%

The pattern is more useful than the log-utility win on its own. As risk aversion rises, the best floor rises towards par. At $\gamma = 2$ the optimum is close to the fixed 90% structure used earlier in the paper. By $\gamma = 5$ the annual-tree optimiser is fractionally above par, at about 101%, while the monthly optimum is at par; in both cases the note is already dominated by the constant-mix alternative. Protection therefore becomes more expensive at the same time that there is less implementation drag left to remove.

The mechanism is straightforward. The conventional risky-weight optimum is approximately 71% at $\gamma = 2$, 47% at $\gamma = 3$ and 28% at $\gamma = 5$. As risk aversion rises, the investor's own optimum moves towards or below the first annual-monitoring intervention boundary. There is then less implementation drag for a floor to remove, while the 2% riskless rate still makes protection expensive in option-budget terms.

The log-utility result should therefore be read as identifying a target configuration rather than a generic investor. Nonlinear protection is most promising when desired financial exposure is high relative to tolerated loss. In FREBO language, that is the economic configuration associated with high Opportunism and high Fortress, subject to empirical validation of the mapping.

12.2 Monthly monitoring

A second robustness exercise replaces the annual binomial tree with monthly lognormal returns matched to the risky asset's annual log moments, approximately 6.88% mean and 23.1% volatility. The horizon remains three years, the riskless rate remains 2%, intervention moves wealth to cash, $\theta \sim U[8\%, 16\%]$, and the log-utility experiment uses 200,000 simulated paths.

Table 11. Annual tree and monthly-monitoring comparison under log utility

Object	Annual tree	Monthly lognormal
Risky asset hold CE	7.12%	7.10%
Risky asset BRR	5.39%	5.11%
Best constant-mix BRR, $U[8, 16]$	5.64%	5.13%
Best risky weight	about 83.5%	about 85–90%
Best note BRR	5.88%	5.42%
Best note floor	84.82%	about 86%
Best note participation	116.3%	about 103%
Exit probability, best linear	discrete-state dependent	about 47%
Exit probability, best note	discrete-state dependent	about 39%
Breakeven incremental upfront margin	51 bp	about 59 bp

The monthly experiment preserves the economically important results while removing the annual-grid geometry. The implementation wedge remains material: the risky asset falls from about 7.10% hold certainty equivalent to 5.11% BRR. Under log utility the optimised note still beats the best constant mix, and the breakeven incremental margin remains of the same order as in the tree.

What disappears is equally important. At a known 12% threshold the note produces about 5.44% against about 5.15% for the best linear allocation, reversing the annual-tree pointwise ranking. The 82.35% corner disappears, the constant-mix optimum becomes smooth and interior, and the one-basis-point fragility disappears with it. Literal cliffs and corners are therefore discrete-monitoring phenomena, not headline claims of BRR.

At $\gamma = 2$ the monthly product comparison is approximately a tie and is sensitive to the precise band and specification; by $\gamma \geq 3$ the linear family wins. Taken together with Table 9, this supports a conditional product-design statement: geared partial protection pays only where the investor wants enough risky exposure for implementation drag to be economically large.

12.3 What survives

Four conclusions survive both changes in specification. First, hold value and implemented value can differ materially. Second, incidence and consequence are distinct design channels. Third, uncertainty requires best-family rather than fixed-product comparisons. Fourth, the behavioural surplus is small enough that an incremental breakeven-margin test is essential. The exact location of annual-tree cliffs, the 82.35% corner and an unconditional note-family advantage do not survive and are not treated as general results.

13. Practical use

A BRR system requires three components.

The first is a forecast model for each candidate investment.

The second is an investor-response model informed by FREBO and any other information shown empirically to predict subsequent action.

The third is a continuation model describing what happens after intervention.

The path variables used by the response model are generated from simulated investment and reference paths.

Where behavioural parameters are uncertain, the model should integrate over their posterior distribution rather than substitute a single point estimate and behave as though it were known.

For investor j, let

$$\mathcal{S}_j$$

be the investment set already permitted by conventional suitability requirements such as capacity for loss, liquidity, horizon, concentration, knowledge, experience and regulatory eligibility.

The BRR choice is

$$i_j^* \in \arg\max_{i \in \mathcal{S}_j} \mathrm{BRR}_i(\kappa_j). \tag{14}$$

For a product family, the same logic applies over design parameters.

FREBO does not have to identify a latent threshold exactly.

Its commercial and investment value lies in whether it changes the posterior distribution enough to alter the preferred investment, allocation or payoff design.

That is the assignment question developed in Oduwole (2026b).

14. Model risk and empirical burden

BRR is a forecast and should be treated as one.

14.1 Product-model risk

Expected returns, volatility, interest rates, correlations and option assumptions can all be wrong.

An accurate behavioural model applied to poor market forecasts still produces a poor BRR forecast.

14.2 Investor-response risk

A FREBO profile does not turn behaviour into certainty.

The relevant empirical test is whether it improves prediction of subsequent implementation behaviour and, more importantly, whether that improvement changes expected implemented value or product assignment.

14.3 Monitoring intensity

Monitoring is part of the behavioural process rather than an incidental implementation detail.

Section 11.3 showed that a nominal 12% rule exits only after a realised 15% loss under annual observation. Section 12 shows that this overshoot geometry changes under monthly monitoring.

A finer observation grid can change intervention frequency, timing and overshoot.

Engagement therefore has a natural empirical role even where loss thresholds themselves remain unchanged.

14.4 Continuation risk

The same trigger can produce different BRR values depending on the subsequent action.

Moving to cash, switching to a benchmark, partially de-risking and later re-entering are economically different processes.

A behavioural model that predicts only the trigger but not the continuation is incomplete.

14.5 Product transport

Behaviour estimated in an ordinary fund need not transport unchanged to a protected note.

The product can change both the experienced path and the meaning attributed to that path.

A 10% decline in an ordinary equity fund may not be experienced like the same mark-to-market decline in a product with a contractual floor.

Until direct evidence exists, novel products should therefore be evaluated across plausible response specifications.

14.6 Execution and credit

The protection result assumes a solvent issuer and fair-value liquidation.

Dealer spreads, illiquidity and issuer default can all break the floor result.

These are ordinary financial risks and should be represented explicitly.

15. Conclusion

BRR asks a narrow forward-looking question: among investments already acceptable on conventional suitability grounds, which is expected to deliver the greatest terminal-wealth value after this investor's likely implementation behaviour is taken into account?

The framework does not require behavioural utility. The ex-ante planner evaluates terminal wealth using an ordinary utility function; the later doer's actions are forecast. The implementation effect is then the difference between intended-policy and implemented-policy expected utility. Its economically useful decomposition is incidence, timing and consequence. A lower-risk allocation can reduce the probability of reaching a behavioural boundary, while a nonlinear payoff can alter the wealth consequence if intervention occurs. Continuation determines whether either benefit survives after the initial action.

The worked tree makes those channels auditable but also exposes the limits of the example. Annual monitoring creates discrete experienced-loss states, which generate the 82.35% corner, the one-basis-point fragility and literal BRR cliffs. Monthly monitoring removes those features. They are illustrations of discrete state entry, not general theorems. Likewise, the log-utility result that an unrestricted floored-note family beats the constant-mix family is not preference-free. Under CRRA with $\gamma > 1$ in the same tree, the linear family wins because the investor's conventional optimum is already sufficiently conservative that little implementation drag remains to be removed. The design response is monotone: the optimal floor rises from about 85% under log utility to par or above at $\gamma = 5$, where the note is already dominated.

The more durable product-design condition is

$$\boxed{\text{implementation value removed} > \text{conventional value surrendered} + \text{incremental product cost.}}$$

This condition identifies the economically interesting investor more sharply than a generic label such as risk-averse. Protection is most likely to pay for an investor who wants substantial risky exposure but is unlikely to tolerate the path required to hold it. In the FREBO vocabulary, high Fortress combined with high Opportunism is the natural candidate configuration, but that mapping remains an empirical hypothesis rather than a maintained assumption.

Threshold uncertainty and continuation are therefore central. With a known, exit-stable trigger under fine monitoring, the problem can collapse to a conventional wealth-path constraint, as developed in Oduwole (2026a). BRR adds more when the behavioural boundary is uncertain, monitoring is imperfect, continuation differs across investors, or payoff families change incidence and consequence in different ways.

The correct commercial benchmark is best family against best family after incremental costs. In the annual log-utility calibration, the unrestricted nonlinear family supports only 33–51 basis points of incremental upfront margin across the three threshold bands; the monthly lognormal exercise produces roughly 59 basis points. These magnitudes are not universal. Their usefulness is that they make the behavioural claim falsifiable in money terms.

The empirical burden is now clear. Investor-side information such as FREBO must predict subsequent intervention and continuation out of sample, and the prediction must be strong enough to change product assignment or implemented value after costs. If it does not, BRR remains an interesting accounting framework. If it does, behavioural measurement becomes an input to portfolio and payoff design rather than an additional questionnaire score.

References

An, L., Engelberg, J., Henriksson, M., Wang, B. and Williams, J. (2024). The portfolio-driven disposition effect. *Journal of Finance, 79(5), 3459–3495.*

Basak, S. and Shapiro, A. (2001). Value-at-Risk-based risk management: optimal policies and asset prices. *Review of Financial Studies, 14(2), 371–405.*

Ben-David, I. and Hirshleifer, D. (2012). Are investors really reluctant to realize their losses? Trading responses to past returns and the disposition effect. *Review of Financial Studies, 25(8), 2485–2532.*

Benninga, S. and Blume, M. (1985). On the optimality of portfolio insurance. *Journal of Finance, 40(5), 1341–1352.*

Brennan, M. J. and Solanki, R. (1981). Optimal portfolio insurance. *Journal of Financial and Quantitative Analysis, 16(3), 279–300.*

Calvet, L. E., Célérier, C., Sodini, P. and Vallée, B. (2023). Can security design foster household risk-taking? *Journal of Finance, 78(4), 1917–1966.*

Das, S., Markowitz, H., Scheid, J. and Statman, M. (2010). Portfolio optimization with mental accounts. *Journal of Financial and Quantitative Analysis, 45(2), 311–334.*

Dhar, R. and Zhu, N. (2006). Up close and personal: investor sophistication and the disposition effect. *Management Science, 52(5), 726–740.*

Dichev, I. D. (2007). What are stock investors' actual historical returns? Evidence from dollar-weighted returns. *American Economic Review, 97(1), 386–401.*

Dybvig, P. H. (1988). Inefficient dynamic portfolio strategies or how to throw away a million dollars in the stock market. *Review of Financial Studies, 1(1), 67–88.*

Elkind, D., Kaminski, K., Lo, A. W., Siah, K. W. and Wong, C. H. (2022). When do investors freak out? Machine learning predictions of panic selling. *Journal of Financial Data Science, 4(1), 11–39.*

Friesen, G. C. and Sapp, T. R. A. (2007). Mutual fund flows and investor returns: an empirical examination of fund investor timing ability. *Journal of Banking & Finance, 31(9), 2796–2816.*

Fulkerson, J. A., Jordan, B. D., Riley, T. B. and Yan, Q. (2026). Bad timing does not cost investors 15% of their funds' returns: an examination of Morningstar's "Mind the Gap" study. *Financial Analysts Journal, 82(3), 34–42.*

Gelman, M., Reiter Gavish, L. and Roussanov, N. L. (2026). FOMO Economics: External reference-dependence and risk-taking in household portfolios. *Working paper, SSRN 6307978.*

Grossman, S. J. and Zhou, Z. (1993). Optimal investment strategies for controlling drawdowns. *Mathematical Finance, 3(3), 241–276.*

Heimer, R., Iliewa, Z., Imas, A. and Weber, M. (2025). Dynamic inconsistency in risky choice: evidence from the lab and field. *American Economic Review, 115(1), 330–363.*

Kahneman, D. and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica, 47(2), 263–291.*

Kaminski, K. M. and Lo, A. W. (2014). When do stop-loss rules stop losses? *Journal of Financial Markets, 18, 234–254.*

Kaustia, M. (2010). Prospect theory and the disposition effect. *Journal of Financial and Quantitative Analysis, 45(3), 791–812.*

Leland, H. E. (1980). Who should buy portfolio insurance? *Journal of Finance, 35(2), 581–594.*

Morningstar (2026). Mind the Gap 2026. *Morningstar Portfolio and Planning Research, 6 August 2026.*

Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance, 53(5), 1775–1798.*

Oduwole, O. (2026a). When Does Behaviour Reduce to a Wealth-Path Constraint? Trigger Vectors, Filtration Enlargement and the Limits of Pathwise Representation. *Working paper.*

Oduwole, O. (2026b). Does Behavioural Product Design Pay for Itself? Assignment Value, Manufacturing Spread and the Boundary for Behavioural Instruments. *Working paper.*

Shefrin, H. and Statman, M. (1985). The disposition to sell winners too early and ride losers too long: theory and evidence. *Journal of Finance, 40(3), 777–790.*

Shefrin, H. and Statman, M. (2000). Behavioral portfolio theory. *Journal of Financial and Quantitative Analysis, 35(2), 127–151.*

Thaler, R. H. and Shefrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy, 89(2), 392–406.*

OLD SHOREHAM INVESTMENT MANAGEMENT

Traditional | Quantitative | Global

www.oldshoreham.com