OLD SHOREHAM INVESTMENT MANAGEMENT

When Does Behaviour Reduce To A Wealth -Path Constraint

Trigger Vectors, Filtration Enlargement And The Limits of Pathwise Representation

Working Paper V.0.1

Oladayo Oduwole
August, 2026

Behavioural Risk and Investment Engineering

Abstract

Suppose an investor changes strategy after a loss, drawdown, gain or period of underperformance. It is tempting to treat that behaviour as another constraint on the investment problem. Sometimes that is exactly right. Sometimes it is not.

I draw the boundary.

If the first trigger is a stopping time in the information available to the investor and the post-trigger continuation is fixed, the implementation rule defines an idempotent operator. Optimising expected utility after implementation is then the same problem as optimising ordinary expected utility over the image of that operator. The behaviour can be moved into the feasible set without changing the utility function.

The harder question is whether the feasible set can be described from wealth alone. A drawdown rule can depend on the whole history and still be determined by the wealth path. A benchmark-relative rule may depend only on the current state and still fail, because the same wealth path can trigger under one benchmark path and not under another. The problem is therefore not path dependence. It is missing state.

Fixed switching into a prescribed continuation leaves the reduction intact. Endogenous post-trigger choice and repeated intervention do not. They lead naturally to impulse control.

An order-invariance construction gives the practical implication. Two products can look the same under mean return, volatility, Sharpe ratio and other order-insensitive statistics yet produce different implemented outcomes once investor actions depend on the path.

Keywords: behavioural finance, path dependence, stopping times, portfolio constraints, implementation, filtration, impulse control
JEL: G11, G41, D81

1. The question

Take a simple investor rule: sell after a 20% drawdown.

There are two questions hiding inside it.

First, can the resulting investment problem be written as an ordinary expected-utility problem over a smaller feasible set?

Second, can that feasible set be identified from the investor's own wealth path?

The answers are different.

A fixed implementation rule can often be treated as a constraint. If the rule says what happens after the trigger, the implementation map has an image. Optimising over implemented strategies is then the same as optimising over that image.

But wealth may not contain enough information to describe the image. If the investor reacts to a benchmark, a peer portfolio or another external reference, the same wealth path can lead to different actions depending on what happened elsewhere.

That distinction matters. A running drawdown uses the entire history but is still wealth-measurable. A benchmark shortfall can use only today's two values and still require a larger state space. History is not the problem. Information is.

The paper makes four claims.

A pinned implementation rule can be reduced to conventional optimisation over the image of an idempotent implementation operator. A wealth-path representation exists only when the trigger is itself wealth-measurable. A fixed switch into another prescribed strategy preserves the reduction but not necessarily the welfare sign. Once the investor can choose the continuation or intervene repeatedly, the fixed-rule representation gives way to impulse control.

This sits beside the constrained-portfolio literature, not outside it. Grossman and Zhou (1993) solve optimal investment under a drawdown constraint. Basak and Shapiro (2001) solve optimal policies under Value-at-Risk constraints. I ask a prior question: **when does an observed or imposed implementation rule itself generate an equivalent constraint, and what state is needed to represent it?**

2. Setup

2.1 Market and wealth

Work on a filtered probability space

$$(\Xi, \mathcal{G}, \{\mathcal{G}_t\}_{t=0}^T, \mathbb{P}).$$

There is a riskless asset $B_t > 0$ and a collection of risky assets. Let $\mathcal{A}_{\mathcal{G}}$ be the admissible self-financing strategies adapted to $\mathcal{G}$. A strategy $\pi \in \mathcal{A}_{\mathcal{G}}$ generates strictly positive wealth

$$W^\pi = (W_0^\pi, \dots, W_T^\pi), \qquad W_0^\pi = 1.$$

Let

$$\Omega = (0, \infty)^{T+1}$$

be path space, with induced wealth-path law

$$\mu^\pi = \mathrm{Law}(W^\pi).$$

Investor welfare is evaluated with an increasing, strictly concave utility function U over terminal wealth.

I keep strategies and realised wealth paths separate. A wealth path need not reveal the strategy that generated it. That distinction becomes important later.

2.2 Wealth information and external information

For a realised wealth path $\omega \in \Omega$, define

$$\mathcal{F}_t^\omega = \sigma(\omega_0, \dots, \omega_t).$$

If an external reference process ν matters, the investor may instead act on

$$\mathcal{G}_t = \mathcal{F}_t^\omega \vee \mathcal{F}_t^\nu \vee \mathcal{H}_t,$$

where $\mathcal{H}_t$ contains any other observed state.

The difference between $\mathcal{F}^\omega$ and $\mathcal{G}$ is the central issue in the non-representation result.

2.3 Trigger statistics

Examples are enough to fix the idea. Loss from inception can be written

$$g_t^L = 1 - \omega_t.$$

Running drawdown is

$$g_t^D = 1 - \frac{\omega_t}{\max_{s \le t} \omega_s}.$$

Unrealised gain is

$$g_t^G = \omega_t - 1.$$

Shortfall against an external reference is

$$g_t^F = 1 - \frac{\omega_t}{\nu_t}.$$

The first three are determined by the wealth history. The fourth generally is not.

Again, the drawdown example is useful. It is highly path-dependent and still causes no representation problem because the path contains the required information.

2.4 First intervention

Let

$$\tau = (\tau^1, \dots, \tau^K) \in (0, \infty]^K$$

be the trigger levels and define

$$\rho(\tau, \omega) = \inf\{t : g_t^k \geq \tau^k \text{ for at least one } k\},$$

with $\rho = \infty$ if nothing fires.

For strategy π, write $\rho(\tau, \pi) = \rho(\tau, W^\pi)$.

The core results use the first intervention. Re-entry and recurrent decisions change the problem and are treated separately.

3. A fixed implementation rule is a feasible-set problem

The investor starts with strategy π. Once the trigger fires, a rule says what happens next.

Assumption A (pinned continuation). From the period after the first trigger observation onward, the investor follows a specified admissible continuation strategy π^R. For an absorbing exit, π^R is the riskless asset. There is no fresh optimisation at the intervention time.

Define the implementation operator

$$\mathcal{I}_\tau^R : \mathcal{A}_\mathcal{G} \to \mathcal{A}_\mathcal{G}$$

which follows π through ρ and π^R thereafter.

Its fixed-point set is

$$\mathcal{C}^R(\tau) = \{\pi \in \mathcal{A}_\mathcal{G} : \pi_t = \pi_t^R \text{ for all } t > \rho(\tau, \pi)\}.$$

Theorem 1 (implementation equivalence). Suppose ρ is a $\mathcal{G}$-stopping time, Assumption A holds and $\mathcal{I}_\tau^R(\mathcal{A}_\mathcal{G}) \subseteq \mathcal{A}_\mathcal{G}$. Then

$$\sup_{\pi \in \mathcal{A}_\mathcal{G}} \mathbb{E}\Big[U\Big(W_T^{\mathcal{I}_\tau^R(\pi)}\Big)\Big] = \sup_{\pi \in \mathcal{C}^R(\tau)} \mathbb{E}\big[U\big(W_T^\pi\big)\big].$$

Proof. For any admissible π, the implemented strategy $\mathcal{I}_\tau^R(\pi)$ belongs to $\mathcal{C}^R(\tau)$ by construction.

If $\pi \in \mathcal{C}^R(\tau)$ already follows the prescribed continuation, applying the operator again changes nothing. The strategy is unchanged up to the first trigger, so the first trigger time is unchanged. Hence

$$\mathcal{I}_\tau^R \circ \mathcal{I}_\tau^R = \mathcal{I}_\tau^R.$$

Thus

$$\mathcal{I}_\tau^R(\mathcal{A}_\mathcal{G}) = \mathcal{C}^R(\tau),$$

and optimisation over implemented strategies is optimisation over the image. □

The theorem is deliberately modest. It does not say the investor's rule is wise or foolish. It says that when the continuation is fixed, the rule can sit in the feasible set rather than in the utility function.

3.1 It is not the same as a level constraint

A first-hit rule is not generally equivalent to imposing

$$g_t \le \tau \qquad \forall t.$$

In discrete time the path can overshoot before action is possible. More importantly, a level constraint allows the original strategy again once the state returns inside the boundary. A first-hit rule changes the strategy after the first crossing.

The right feasible set is the image of the implementation operator, not a pointwise inequality pasted onto the original strategy.

4. When wealth alone is enough

Theorem 1 works with whatever information the investor actually uses. A separate question is whether wealth itself contains that information.

Proposition 1 (wealth-path representation). Suppose ρ is an $\mathscr{F}^\omega$-stopping time and the prescribed continuation is the bank account. Define

$$\mathfrak{C}^A(\tau) = \left\{\omega \in \Omega : \rho(\tau,\omega) = \infty \text{ or } \omega_s = \omega_\rho \frac{B_s}{B_\rho} \text{ for all } s > \rho(\tau,\omega)\right\}.$$

Then the absorbing implementation rule has a wealth-path image equal to the admissible implemented paths in $\mathfrak{C}^A(\tau)$.

This is a result about outcomes, not about recovering the portfolio strategy from the path. Different strategies can generate the same realised wealth path.

Loss-from-entry, running drawdown and own-position gain rules fit here because their firing times are determined by wealth history.

5. When wealth is not enough

Now take

$$g_t^F = 1 - \frac{\omega_t}{\nu_t},$$

where ν is a benchmark or another external reference.

The trigger is generally a stopping time under

$$\mathscr{F}^\omega \vee \mathscr{F}^\nu,$$

not under $\mathscr{F}^\omega$ alone.

Proposition 2 (no wealth-only representation under decision-relevant external state). Suppose there is a wealth path ω and two external-state paths ν and ν' such that the rule triggers on (ω,ν) and does not trigger on (ω,ν'). Then no set

$$C \subseteq \Omega$$

can represent the rule uniformly over reference paths.

Proof. If C represented the rule, the same wealth path ω would have to be classified differently depending on whether it was paired with ν or ν'. Membership in C depends only on ω. That is impossible. □

The qualification matters. External information is only a problem when it changes the decision while wealth is held fixed.

This is why I separate path dependence from state sufficiency. A drawdown rule can use the whole past and still live comfortably on wealth-path space. A benchmark rule can use today's values only and still require more state.

5.1 This happens in the data

An, Engelberg, Henriksson, Wang and Williams (2024) show that the disposition effect on an individual security changes with the state of the investor's wider portfolio. The security's own state is not enough.

Gelman, Reiter Gavish and Roussanov (2026) find that retail investors who fall behind a salient market index increase portfolio risk and buy securities with higher beta, idiosyncratic volatility and positive skewness. The response becomes stronger as the benchmark becomes more salient.

Both findings point to the same modelling lesson. If the decision uses portfolio or benchmark information, leave that information in the state space.

6. Fixed switching does not break the result

The pinned continuation need not be cash.

Suppose the investor switches after the trigger into a specified admissible strategy π^R. Theorem 1 still applies. The operator remains idempotent and the feasible-set reduction survives.

What no longer follows automatically is the sign of the value effect.

For an absorbing move to cash, one might assume

$$E[U(W_T^\pi) \mid \mathcal{G}_\rho] \geq U\left(W_\rho^\pi \frac{B_T}{B_\rho}\right)$$

on the trigger event.

For a risky switch, the corresponding condition is

$$E[U(W_T^\pi) \mid \mathcal{G}_\rho] \geq E[U(W_T^{\pi^R}) \mid \mathcal{G}_\rho].$$

There is no reason for either inequality to hold universally. A switch from a lagging asset may destroy value if the lag is temporary and improve value if the lag contains information.

The representation result does not need a welfare verdict.

7. Where the reduction stops

The fixed-rule result relies on a pinned continuation.

If the investor chooses the post-trigger action optimally at the intervention date, let the state be (t, w, z, i), where z contains external information and i is the current asset or regime. With proportional switching cost c, a standard impulse-control form is

$$\min\left(-\partial_t V - \mathcal{L}_i V, V - \mathcal{M}V\right) = 0,$$

with

$$\mathcal{M}V(t, w, z, i) = \sup_{j \neq i} V(t, w(1-c), z, j).$$

The continuation is now chosen, not pinned.

Repeated selling and re-entry creates the same problem. We have to specify when re-entry is possible, whether trigger statistics are rebased and whether later decisions depend on earlier interventions.

The classification is therefore simple:

Rule	Trigger information	Mathematical object
Fixed continuation	Wealth-measurable	Feasible-set reduction and wealth-path representation
Fixed continuation	Decision-relevant external state	Feasible-set reduction on enlarged filtration; no uniform wealth-only representation
Endogenous continuation or recurrent intervention	Wealth state	Impulse control
Endogenous continuation or recurrent intervention	External state	Impulse control on enlarged state

That table is the paper in practical form.

8. Conventional statistics can miss the implementation problem

A second result follows from the fact that the investor experiences returns in an order.

Definition 1 (order-invariant criterion). A statistic S of returns $R = (R_1, \dots, R_T)$ is order-invariant if

$$S(R_1, \dots, R_T) = S(R_{\pi(1)}, \dots, R_{\pi(T)})$$

for every permutation π.

Mean return, variance, Sharpe ratio and terminal wealth for a fixed multiset of multiplicative period returns are familiar examples.

A running trigger is not order-invariant.

Proposition 3 (order-invariant equality need not survive implementation). Let $\varepsilon_1, \dots, \varepsilon_T$ be exchangeable shocks and let $a_1, \dots, a_T$ be a deterministic exposure schedule. Define

$$R_t = r + a_t \varepsilon_t.$$

Let a second product use the same exposures in a different order. Then the two products agree in law under every order-invariant criterion. Their implemented terminal wealth laws can nevertheless differ for an investor whose action depends on the experienced path.

Proof. Exchangeability means that assigning the same deterministic exposures in a different order produces a return vector equal in law up to permutation. Every order-invariant statistic therefore has the same law under the two products.

If the implementation rule fires at different times on an event of positive probability and the prescribed continuation changes terminal wealth on that event, the implemented terminal laws differ. □

The claim is not that Sharpe ratio is defective. It simply answers a different question.

8.1 A numerical illustration

Take a ten-year monthly horizon, one risky asset and one riskless asset. One product uses exposure 1.5 in the first half and 0.6 in the second. The other reverses the schedule. The shocks are exchangeable.

The two products have essentially the same hold certainty equivalent and volatility. Under a 12% loss-from-inception rule followed by exit to the riskless asset, the front-loaded schedule triggers more often and earlier. In the stated simulation, the annualised implemented certainty equivalents differ by about 2.89 percentage points.

That number is not structural. Re-entry after 36 months reduces the gap to about 1.45 points. Re-entry after 12 months reduces it to about 0.48 points.

The result I need is only the existence result: products that look the same under order-insensitive statistics can behave differently once implementation depends on sequence.

9. Where this sits in portfolio theory

Drawdown constraints are not new.

Grossman and Zhou (1993) solve optimal investment when wealth must remain above a fixed fraction of its running maximum. Basak and Shapiro (2001) study optimal policies under Value-at-Risk constraints and show that the form of the risk rule changes optimal behaviour and equilibrium outcomes.

My question is different. I start with an implementation rule and ask whether it can be represented as a conventional constraint at all. If it can, I then ask whether wealth alone is sufficient to describe that constraint.

That matters when the implementation rule comes from observed investor behaviour rather than from a portfolio manager writing down an exogenous restriction. A drawdown-sensitive policy may fit on wealth-path space. A benchmark-sensitive policy may not. Calling both "risk tolerance" throws away the information needed to tell the difference.

10. Limits

The theory says nothing about which psychological mechanism generated the action.

A trigger can come from loss aversion, belief updating, regret, attention, a pre-committed rule or a genuine holding constraint. A hard threshold and a steep smooth hazard can also look very similar in observed data.

For empirical work, three things have to be kept explicit.

The information set has to be estimated rather than guessed. The continuation after intervention has to be specified. And a welfare claim needs an additional benchmark showing that the intervention was worse than the continuation it replaced.

Those are empirical and normative questions. They are separate from the representation result.

11. Conclusion

A behavioural investment rule does not automatically require a behavioural utility function.

If the first trigger is a stopping time and the post-trigger action is fixed, the rule defines an idempotent implementation operator. Expected-utility optimisation after implementation is then ordinary expected-utility optimisation over the operator's image.

The harder question is whether that image can be written as a set of wealth paths.

If wealth determines the trigger, yes. If external state can change the decision while wealth is held fixed, no. The missing object is not history. It is state.

Fixed switching leaves that logic intact. Endogenous continuation and repeated intervention do not. They lead to impulse control.

The result can be stated in one line:

The action determines whether behaviour becomes a feasible-set problem. The information used to trigger that action determines whether wealth alone is enough to represent it.

Appendix A. Formal path-space statement

Let Φ_τ^R map original wealth and relevant external state into implemented wealth under pinned continuation R.

If ρ is $\mathscr{F}^\omega$-measurable, there exists a measurable map

$$\phi_\tau^R : \Omega \to \Omega$$

such that

$$\Phi_\tau^R(\omega, z) = \phi_\tau^R(\omega)$$

for every admissible external state z.

If there are ω, z, z' such that

$$\Phi_\tau^R(\omega, z) \neq \Phi_\tau^R(\omega, z'),$$

no wealth-only map can represent the rule uniformly.

Appendix B. Permuted-exposure construction

Let monthly excess returns be

$$\varepsilon_t = \mu_e + \sigma_e Z_t,$$

with exchangeable Z_t. Wealth evolves as

$$W_t = W_{t-1}(1 + r_m + a_t \varepsilon_t).$$

Use 120 monthly periods, $r_m = 0.02/12$, annual volatility per unit exposure of 18%, Sharpe ratio 0.65 and common random numbers. Compare

$$\begin{aligned} a^F &= (1.5, \dots, 1.5, 0.6, \dots, 0.6), \\ a^B &= (0.6, \dots, 0.6, 1.5, \dots, 1.5), \end{aligned}$$

using 60 observations in each block.

The implementation rule moves to the riskless asset from the period after wealth first falls 12% below inception. With 100,000 paths, the schedules have essentially identical hold statistics under order-invariant criteria but materially different trigger incidence and implemented values. Under absorbing exit, the annualised implemented certainty-equivalent gap is approximately 2.89 percentage points. Re-entry assumptions reduce the magnitude as discussed in Section 8.1.

The construction is an existence example, not an estimate of investor welfare.

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.

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.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4), 373-413.

OLD SHOREHAM INVESTMENT MANAGEMENT

Traditional | Quantitative | Global

www.oldshoreham.com