Fredkin's Paradox in Portfolio Management

Portfolio Management Research Paper

Michael Walmer | Walmer Family Fund

August 27, 2026 · *Draft for review — not approved for publication*

> **Research question.** When does additional analysis of a portfolio choice improve the portfolio, and when does it merely consume attention while leaving the economic outcome substantially unchanged?

Abstract

This paper applies **Fredkin's Paradox**—the observation that choices can become hardest when the alternatives appear most alike—to routine portfolio implementation. The apparent paradox is not that close choices are unimportant in every respect. Rather, it is that the **effort devoted to resolving them can exceed the plausible economic value of resolving them**. In a portfolio setting, this can show up as repeated comparison of nearly overlapping index funds, parameter debates over adjacent rebalancing bands, fee comparisons that omit trading costs, or prolonged selection of a temporary tax-loss-harvesting substitute.

The recommended response is not careless simplification. It is a disciplined sequence: first identify any **material difference** in exposure, implementation, tax treatment, liquidity, governance, or account constraints; then use a predefined scorecard and stop rule; finally, resolve a genuine residual tie with a documented default or, where appropriate, a neutral split. This preserves analytical capacity for decisions where the dispersion of plausible outcomes is genuinely large. The analysis uses VOO and VTI as an illustrative—not prescriptive—case: their daily adjusted returns were highly correlated over the ten years ending August 27, 2026, but the funds remain structurally different and should not be treated as universally interchangeable.[3] [4] [10]

Core conclusion	Practical interpretation
Difficulty is not a proxy for importance.	The hardest operational choice may have the smallest expected effect on the portfolio.
High correlation is not equivalence.	Similar funds still require a basic review of benchmark, exposure, costs, liquidity, taxes, and vehicle structure.
A stopping rule is a risk control.	It limits delay, cash drag, analysis fatigue, and repeated fitting of policy to historical noise.
A tie-breaker is acceptable only after a material-difference screen.	A 50/50 split can be a governance device; it is not automatic diversification or a substitute for tax and risk analysis.

The paradox: decision effort becomes misallocated

For purposes of this paper, **Fredkin's Paradox** labels a common choice condition: two alternatives are judged so close on the criteria that matter that the decision-maker finds it unusually difficult to select

one. The paradox becomes financially relevant when the increasing time, attention, and emotional load devoted to the decision are **out of proportion to the expected difference in portfolio outcomes**.

The behavioral mechanism is credible even though most decision-deferral research is not portfolio-specific. In a vignette study, Berens and Funke distinguish deferring a decision from refusing the currently available options, and report that decision deferral was associated with selection difficulty as well as situational and individual factors.[1] Dhar's foundational work on a no-choice option also treats the relative attractiveness of alternatives as relevant to whether a choice is deferred.[2] Portfolio work adds a concrete channel that consumer-choice studies do not need to model: while an intended allocation is left unresolved, the portfolio can remain off target or partially in cash.

The key analytical mistake is to ask only, **"Which option is better?"** The better question is two-layered:

Layer	Question	Appropriate output
Object-level decision	Which fund, threshold, or substitute better serves the stated portfolio objective?	A documented selection, or a documented tie.
Meta-decision	Is another hour of research likely to improve the selection enough to justify its cost?	Continue, delegate, apply a default, or stop.

A sound process treats the second question as independently important. Research time has value, delay can change the asset allocation actually held, and unbounded comparison can invite retrospective rationalization. The aim is not to reduce diligence; it is to **match diligence to decision materiality**.

From psychological paradox to portfolio drag

The opportunity cost of attention

Every investment team has a finite research budget. Spending disproportionate time deciding between nearly identical core vehicles can displace work on questions with wider outcome dispersion: assessing a concentrated exposure, evaluating an active mandate's real sources of return, testing a balance-sheet risk, or examining an allocation that is materially off policy. The opportunity cost is therefore not simply the analyst's salary or an abstract value of time. It is the foregone chance to examine a decision where an error would alter expected return, downside risk, liquidity, or governance more substantially.

This does **not** imply that broad-market fund selection is trivial. FINRA notes that an exchange-traded product's prospectus and related documents contain its investment objectives, risks, fees, and expenses, and emphasizes that funds tracking the same index may do so in different ways.[5] The implication is narrower: once those potentially material dimensions have been checked and no meaningful distinction remains, endlessly ranking the residual differences has a low expected payoff.

Cash drag is a conditional implementation cost

Cash drag is often stated too categorically. Holding cash while a decision is pending does not guarantee a loss: the subsequent return of the intended risky allocation is uncertain. The correct statement is conditional. If a delay leaves the portfolio below its intended risk exposure, the **realized effect of delay** over that interval is the return of the intended implementation minus the return earned on the displaced

capital, net of costs. The result can be positive or negative, but it is a real exposure decision rather than a neutral pause.

A practical consequence follows. If the only unresolved question is which of two already-approved, operationally suitable core funds to use, the default should favor **timely execution of the policy**, subject to normal trading controls. The decision should not remain open merely because the alternatives have become psychologically harder to separate.

Fine-tuning can turn into data snooping

The label **HARKing** is sometimes applied to this problem, but it is not the most precise term. HARKing concerns presenting a hypothesis developed after observing results as though it were specified beforehand. The closer portfolio-research danger is **data snooping** or **backtest overfitting**: repeatedly revisiting the same history to select the most attractive-looking parameter, fund, or rule. White defines data snooping as reusing a data set for inference or model selection and explains that the apparently satisfactory result can stem from the method rather than genuine merit.[8]

This matters for rebalancing rules. A historical test may show that a 4.5% drift threshold slightly outperformed 5.0% in one sample. Unless the process controlled for the number of specifications considered, used an information cutoff, and was validated out of sample, the apparent advantage may be indistinguishable from selection noise. Bailey and coauthors developed a formal framework for assessing the probability of backtest overfitting in investment simulations.[9] The governance response is to set a small, economically motivated family of candidate rules in advance and favor robustness, simplicity, and executability over the historical winner by a marginal amount.

Hidden drag	Mechanism	What a disciplined process tests
Attention displacement	Near-tie analysis crowds out work on higher-dispersion decisions.	Does more review change a decision with meaningful economic range?
Implementation delay	A policy allocation remains unresolved or in a temporary holding.	What exposure is actually held during the delay, and is it deliberate?
Parameter overfitting	Adjacent rules are optimized against the same historical path.	Was the candidate set limited, specified ex ante, and tested beyond the selection sample?
False precision	Basis-point comparisons omit larger one-time or operational frictions.	Are trading, tax, governance, and tracking effects included alongside the stated fee?
Decision fatigue	Repeated comparison lowers quality or delays later choices.	Has the decision met its pre-set stopping condition?

Similar is not identical: an ETF implementation example

VOO and VTI are often placed at the center of the "index-splitting" debate because both provide broad U.S. equity exposure at a stated 0.03% expense ratio as displayed on their official profiles on August 27, 2026.[3]

[4] They are, however, not the same mandate. VOO seeks to track the S&P 500 Index using full replication and is dominated by large U.S. companies.[3] VTI seeks to track the Morningstar US Total Market Index, which Vanguard describes as representing approximately 100% of the investable U.S. market and including large-, mid-, small-, and micro-cap stocks; it uses sampling.[4]

The figure below is deliberately designed to convey both sides of the issue. The paths exhibit extremely high historical daily co-movement, but they do not sit on top of each other. Thus, the figure supports an **implementation-efficiency** argument—not a claim that either vehicle is always the right answer.

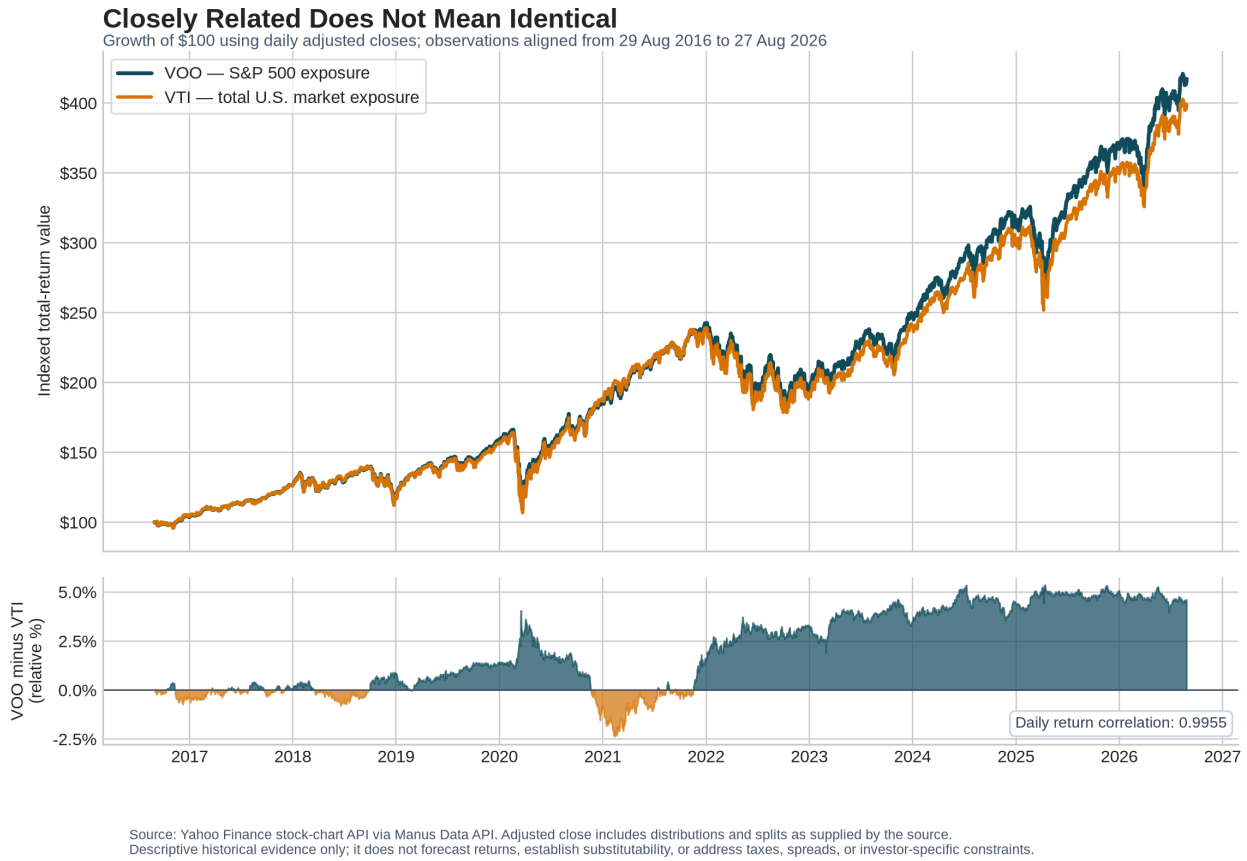


Figure 1 note. The calculation uses 2,512 aligned daily simple-return observations derived from adjusted closes supplied by the Yahoo Finance stock-chart API. The daily-return correlation was **0.9955**. Over the same sample, the calculated annualized return was **15.37%** for VOO and **14.85%** for VTI, while annualized daily-return volatility was **18.03%** and **18.32%**, respectively. These are descriptive historical calculations, not forecasts. The full calculation methodology and data summary are provided in the accompanying data appendix.[10]

Illustrative comparison, as of August 27, 2026	**VOO**	**VTI**	**Why the difference may matter**
Stated expense ratio	0.03%	0.03%	No stated fee difference in this snapshot.[3] [4]

Illustrative comparison, as of August 27, 2026	VOO	VTI	Why the difference may matter
Benchmark scope	S&P 500 / large-cap U.S. equity	Investable U.S. market / includes smaller companies	The portfolio's intended small- and micro-cap exposure may differ.[3] [4]
Index implementation	Full replication	Sampling	Implementation and tracking methods are not identical.[3] [4]
Daily-return correlation in this study	-	0.9955 with VOO	High co-movement is historical, period-specific, and not a guarantee.[10]
One-time trading effects	Depends on order, market conditions, and account	Depends on order, market conditions, and account	ETF shares are subject to bid–ask spreads, which need not be captured by the expense ratio.[5]

The appropriate conclusion is neither "debate forever" nor "ignore the difference." It is: **make the material comparison once, record the conclusion, and use a stopping rule.** A portfolio that deliberately targets the broadest investable domestic-equity exposure has a reason to prefer VTI; a policy designed around a large-cap benchmark has a reason to prefer VOO. If neither distinction affects the policy objective or the account's implementation constraints, residual comparison should be brief.

When the paradox does—and does not—apply

The word "similar" should be earned, not assumed. Correlation is a useful descriptive statistic but is insufficient as a substitution test. It can be high because two funds share large constituents while still masking meaningful differences in small-cap exposure, index construction, securities lending, tracking, tax characteristics, bid–ask spreads, or legal structure. FINRA explicitly cautions that exchange-traded products vary in objectives, costs, risks, and implementation, and that bid–ask spreads are part of the trading experience.[5]

Do not fast-track the choice when…	Why the question is materially different
The alternatives change the portfolio's strategic exposure, factor tilt, duration, credit quality, geography, currency, or concentration.	The risk budget itself may change.
The account is taxable, constrained, or subject to a transaction that could have tax consequences.	Vehicle choice can affect realized tax treatment, reporting, or compliance.

Do not fast-track the choice when…	Why the question is materially different
Liquidity, capacity, spreads, premiums/discounts, derivatives use, securities lending, or trading venue differ materially.	Implementation shortfall can dominate a small difference in stated fees.
A fund's structure, domicile, counterparty exposure, leverage, or regulatory status changes.	The legal and operational risk is not a basis-point preference.
The selection is hard to reverse, has large size, or creates governance precedent.	The cost of a mistaken decision is elevated.
The alternatives are merely *claimed* to be substitutes without an exposure or mandate review.	The premise of the paradox has not yet been established.

Tax-loss harvesting is a particularly important exception. The SEC's investor guidance states that a wash sale may arise when a security is sold at a loss and substantially identical securities are purchased within 30 days before or after the sale; it directs readers to IRS Publication 550.[7] That timing rule explains why a temporary substitute selection can merit real care. It does **not** establish that any two named ETFs are or are not "substantially identical." The status of a particular transaction is fact-specific and should be handled through appropriate tax review rather than a generic tie-breaker.

A decision architecture for bypassing the paradox

The solution is best understood as **procedural alpha**: improving the reliability, speed, and auditability of decisions without pretending that process alone creates excess market return. The following architecture is intentionally conservative. It makes rapid resolution available only after a minimum evidence threshold has been met.

Step 1: Write the decision in one sentence

State the portfolio objective, the account type, the size and timing, and the alternatives. For example: "Select one approved U.S. core-equity vehicle for a long-term taxable account, with a target exposure defined by the investment-policy statement." A vague question such as "Which is better?" makes comparison unbounded.

Step 2: Apply a material-difference gate

Screen the alternatives against the factors in Section 4. This gate is pass/fail, not a weighted score. If any factor is material or uncertain, the choice graduates to ordinary due diligence. If no factor is material, classify the decision as a **near-equivalent implementation choice**.

Step 3: Use a standard scorecard, then stop

For a near-equivalent choice, use a compact scorecard that is completed once and saved with the decision record. A standard template prevents the analyst from inventing criteria after seeing a preferred answer.

Criterion	Evidence to review	Decision rule
Mandate fit	Benchmark, exposure scope, replication method, holdings method	Must match the policy objective.
Cost and trading	Expense ratio, brokerage cost, estimated spread, available order controls	Must meet the portfolio's predefined implementation standards.
Liquidity and operations	Creation/redemption ecosystem, trading conditions, custodian support, settlement	Must be operationally adequate for the trade size and process.
Tax and account treatment	Account restrictions, distribution characteristics, planned loss harvesting	Escalate if the consequence is uncertain or material.
Governance	Approved-list status, recordkeeping, monitoring, replacement process	Must be documentable and repeatable.

A useful **satisficing** rule is: choose the first approved instrument that clears all mandatory criteria, unless another candidate has a pre-specified and economically material advantage. The rule is not "choose carelessly." It is "do not manufacture a ranking after all required conditions are satisfied."

Step 4: Use a decision-value budget

Further research is warranted only if its **expected value of information** exceeds its cost. A simple governance approximation is:

[\text{Continue researching only if } p \times \Delta U \times A > C_r + C_d]

where (p) is the probability that additional work uncovers a decision-changing fact, (\Delta U) is the expected improvement in portfolio utility if that fact is found, (A) is the amount of capital or decision horizon affected, (C_r) is the research cost, and (C_d) is the cost of delay. The expression is not a precise personal-finance calculator. It is a forcing device: before opening the next spreadsheet tab, articulate the plausible mechanism by which the next hour could change the portfolio meaningfully.

A practical time budget may be set by decision class rather than by false numerical precision. For example, a pre-approved core-vehicle choice that has passed all gates may receive one documented review session. A new strategic allocation, a tax-sensitive transition, or a concentrated investment receives a larger budget and independent challenge. The threshold should be fixed in the investment-governance process—not reverse engineered from a favorable backtest.

Step 5: Resolve the genuine tie with a pre-committed rule

If two candidates pass all mandatory filters and neither has a material advantage, the decision record should invoke one of the following rules. The choice of rule should itself be specified in advance.

Tie-resolution rule	Appropriate use	Important caveat
Existing-position default	Avoiding unnecessary turnover or operational complexity.	Do not use it to ignore a material exposure or tax problem.
Approved-list first pass	Several equivalent options meet a published standard.	The approval process must be current and credible.
Lowest all-in implementation cost	Costs can be measured consistently for the actual trade.	"All-in" includes spread and execution, not only the stated expense ratio.
50/50 split	A true economic tie where two vehicles may be held without worsening operations or taxes.	It is a governance compromise, not new diversification when holdings overlap heavily.
Delegated or randomized tie-break	The alternatives are genuinely exchangeable and execution is the only remaining question.	The equivalence test must precede, not follow, delegation or randomization.

The 50/50 rule deserves restraint. It can prevent a false sense that a single basis point must contain decisive information. But splitting VOO and VTI, for example, does not create two independent sources of U.S. equity risk; it creates a combined exposure with substantial overlap. A split is therefore suitable only when it maintains the intended policy exposure, preserves operational simplicity, and does not create avoidable tax or monitoring complexity.

Step 6: Record the decision and schedule the next review

A short decision memo should identify the alternatives, material-difference check, data date, selected rule, rationale, and next review trigger. Reopening a settled near-equivalent choice should require new information—such as a benchmark change, fee change, tax event, liquidity deterioration, or policy revision—not merely renewed discomfort.

Applying the framework to common manifestations

Index splitting

The first task is to determine whether two broad-market funds truly serve the same policy role. In the VOO–VTI illustration, the answer depends on whether the investment policy seeks large-cap U.S. equity or the investable domestic market inclusive of smaller-cap exposure.[3] [4] When the policy answers that question, the fund selection becomes implementation. The decision should then use the standard scorecard and stop rule, rather than repeated performance-chart comparisons.

Marginal rebalancing bands

Rebalancing is intended to restore the portfolio to its chosen risk profile; it is not a market-timing technique.[6] A 4.5% versus 5.0% drift debate is therefore material only to the degree that the difference changes the portfolio's risk, turnover, tax burden, monitoring load, and implementation costs. Vanguard describes threshold rebalancing as triggering when allocation change exceeds a specified threshold and emphasizes the need to balance method, costs, and adherence to the underlying plan.[6]

A robust policy states the target allocation, the permitted drift definition, the monitoring cadence, the tax-aware execution hierarchy, and the exception process. If adjacent threshold candidates all satisfy those governance requirements and historical ranking is unstable across sensible samples, choosing the simpler rule may be more defensible than selecting the narrow historical winner.

Provider and fee selection

A one-basis-point annual expense-ratio difference equals **$1 per $10,000 per year before compounding**. That arithmetic is real, but it is not the whole comparison. A single trade executed across a wider spread, a less efficient order, or a tax-sensitive event can outweigh a very small stated fee difference. FINRA notes both that prospectus disclosures are central to fund comparison and that ETF shares are subject to bid–ask spreads.[5] An all-in cost review should therefore be brief but complete.

Tax-loss-harvesting substitutes

A temporary substitute can preserve broad market exposure during the wash-sale window, but the decision is not merely a question of price-chart similarity. The wash-sale rule turns on a fact-specific "substantially identical" standard within a 30-day-before-or-after window.[7] The correct process is to pre-approve substitute families with tax review, document the trade path, and treat a compliance uncertainty as a reason to escalate—not as a prompt to keep comparing nearly correlated charts.

Governance tools that make the process durable

The paradox is most damaging when it is rediscovered on every trade. The remedy is institutional memory: an approved universe, a short scorecard, a list of genuine escalation triggers, and a decision log. These tools reduce avoidable research but preserve the evidence necessary to explain why a choice was considered equivalent at the time it was made.

Tool	Minimum contents	Benefit
Near-equivalent decision log	Objective, options, data date, passed filters, tie-breaker, reviewer, next trigger	Prevents repeated relitigation and preserves auditability.
Approved substitute map	Core holding, permitted substitute families, account constraints, tax-review status	Speeds response to tax-loss or operational events without improvisation.
Rebalancing policy card	Target, drift calculation, cadence, tax-aware execution sequence, exceptions	Converts a recurring emotional decision into a rule-governed process.
Research budget matrix	Decision class, maximum initial review time, escalation threshold, required reviewer	Moves effort toward choices with larger plausible consequences.
Robustness protocol	Candidate rules specified before testing, holdout period, implementation costs, reasonableness review	Limits selection on historical noise.

A disciplined workflow also separates **reversible** from **irreversible** decisions. A small, liquid, approved core-fund purchase is usually more reversible than a private investment, a large in-kind transition, or a tax event. Reversibility does not eliminate cost, but it lowers the value of exhaustive pre-decision optimization and strengthens the case for a clear default.

Limitations and boundaries

This paper does not claim that Fredkin's Paradox is a formally estimated factor in portfolio returns, nor that every close investment choice has negligible impact. It is a decision-governance lens. The framework should never override material differences in risk exposure, taxes, legal structure, liquidity, concentration, or client/investment-policy constraints.

The VOO–VTI analysis is also intentionally limited. It uses daily adjusted-close data for a fixed ten-year sample and summarizes historical co-movement. It does not test expected utility, after-tax return, individual trade execution, tracking error versus a common benchmark, fund flows, securities lending, or future correlation regimes. The realized performance gap visible in the chart is precisely why high correlation must not be translated into "identical."

Finally, a time budget should not become a pretext for avoiding difficult work. It is most useful when the alternatives have already been shown to be near-equivalent on the dimensions that govern the specific portfolio decision. A material new fact is always a valid reason to reopen the analysis.

Conclusion

The practical lesson of Fredkin's Paradox is not to care less about portfolio construction. It is to care **at the right resolution**. Core vehicles, rebalancing rules, and temporary substitutes should receive enough diligence to establish mandate fit, implementation soundness, tax awareness, and governance quality. Beyond that point, the incremental return on further deliberation can be lower than the cost of delay and the opportunity cost of diverted attention.

A portfolio process is strongest when it can say, with evidence, both **why two options are not identical** and **why their residual differences are not large enough to justify another round of analysis**. In that narrow but common setting, a predefined default, satisficing filter, or carefully permitted split is not intellectual surrender. It is a disciplined decision to preserve scarce analytical capital for choices where the portfolio outcome is genuinely at stake.

> **Compliance note.** This is research and analysis only, not personalized financial, investment, tax, or legal advice. It does not recommend any security, fund, allocation, tax treatment, or trading action.

References

1. Berens, S., & Funke, J. (2020). *A vignette study of option refusal and decision deferral as two forms of decision avoidance.* PLOS ONE.
2. Dhar, R. (1997). *Consumer Preference for a No-Choice Option.* Journal of Consumer Research.
3. Vanguard S&P 500 ETF (VOO), official fund profile.
4. Vanguard Morningstar Total Stock Market ETF (VTI), official fund profile.
5. Financial Industry Regulatory Authority. *Exchange-Traded Funds and Products.*
6. Vanguard. *Rebalancing your portfolio.*
7. U.S. Securities and Exchange Commission, Investor.gov. *Wash Sales.*

8. White, H. (2000). *A Reality Check for Data Snooping*. Econometrica, 68(5), 1097–1126.
9. Bailey, D. H., Borwein, J. M., López de Prado, M., & Zhu, Q. J. (2017). *The Probability of Backtest Overfitting*. Journal of Computational Finance.
10. Data appendix: reproducible VOO–VTI adjusted-return analysis, supplied with this report.