RESEARCH MEMO

Direct Indexing as a Lifecycle Decision

Why the exit determines whether the entry was worth making

Steven Ede, CFA

September 2026

Modeled scenario. All figures hypothetical. No proprietary manager data is used.

1 · SUMMARY OF CONCLUSIONS

Most direct indexing analysis begins with a funded sleeve and ends when the harvesting stops. Both ends of that treatment are wrong. The sleeve has to be funded from somewhere, and in a fully invested taxable portfolio the cost of raising that cash can exceed several years of harvesting benefit. At the other end, the strategy manufactures the problem it was hired to solve: by year ten the account holds a large, deeply appreciated, low-basis position that is harder to unwind than whatever was sold to fund it.

> Evaluate direct indexing as a three-part lifecycle, funding cost, harvest yield, and exit path, and treat the exit as the binding constraint. A client without a credible non-recognition exit is buying tax deferral, not tax savings, and should size the allocation accordingly.

Findings

1 · Harvesting defers tax; it does not erase it. Every dollar of loss harvested reappears as a dollar of additional embedded gain at exit. In the sample sleeve, $9,375,000 of cumulative harvested losses produces exactly $9,375,000 of additional terminal gain versus a passively held index fund. This is an identity, not an estimate, and it holds regardless of the growth path. (§5)

2 · Consequently, a liquidation exit is worth only the time value of the deferral. Net present value of the strategy on a full-liquidation exit is $571,000 on a $25,000,000 sleeve, roughly 2.3% of funding, spread over ten years. Real, but far below the headline tax alpha figures the category markets on. (§6.1)

3 · The charitable exit is worth 3.4× the liquidation exit. Contributing the terminal sleeve to a donor advised fund raises net present value to $1,941,000, because the deferred gain is never recognized. The entire $1,370,000 difference is attributable to the exit path alone, not to better harvesting, a better manager, or a better market. (§6.2)

4 · The strategy's value scales almost linearly with charitable intent. Between 0% and 100% of the sleeve donated, net present value moves from $571,000 to $1,941,000 in near-perfect proportion. Charitable intent is therefore not a footnote to the recommendation. It is the largest single input to it. (§6.3)

5 · Funding matters as much as harvesting. Selling appreciated legacy positions to raise the $25,000,000 would have cost $2,261,000 in immediate tax and permanently reduced the invested base. A right of substitution with a grantor trust funded the sleeve at zero recognized gain. That saving alone exceeds the entire net present value of the liquidation-exit case. (§3)

6 · The growth rate does not affect the liquidation result. Because the extra exit gain equals cumulative losses harvested regardless of the path taken, the liquidation case is a pure function of the discount rate. This is counterintuitive and worth stating plainly: a bullish or bearish market view is not a reason to adopt or reject the strategy. (§7)

Three standing rules

1. Never fund a direct indexing sleeve by realizing gain while a non-recognition funding route exists.
2. Do not adopt direct indexing without an identified exit. Absent one, model it as deferral and size it against the client's actual liquidity horizon.
3. Substitute the lots you least want to sell, not the lots with the most harvesting potential. Legacy basis is irrelevant to a cash-funded sleeve.

What would change these conclusions

- **A step-up at death.** If the sleeve is held to death, the deferred gain is eliminated by basis step-up and the economics resemble the charitable case without requiring charitable intent. The analysis here assumes a liquidation or gift inside the modeled horizon.
- **Ongoing contributions.** The harvest schedule assumes a single funding event and no subsequent additions. Regular contributions reset harvestable basis and materially extend the sleeve's productive life.
- **The rate at which losses are used.** The base case values harvested losses at the same rate applied to the exit gain. Where losses offset income taxed at a higher rate than the exit, the benefit rises substantially. This is a client-specific fact, addressed in §7.
- **Tracking error tolerance.** The model prices tax outcomes only. It assigns no cost to the tracking error a harvested portfolio accumulates against its benchmark.

◆

2 · THE SAMPLE PORTFOLIO AND THE FUNDING PROBLEM

The modeled client holds a $25,000,000 taxable allocation across five exchange traded funds, a majority in broad large-capitalization index exposure, with the remainder spread across sector funds. The portfolio is diversified but is not a clean index: the sector positions represent legacy allocation decisions the client had been steadily moving away from. Embedded gain across the account is roughly 38% of market value, or approximately $9,500,000. There is no meaningful cash.

That last sentence is the whole problem, and it is the one most direct indexing literature skips. The client wants a direct indexing sleeve because he wants harvestable losses. A direct indexing sleeve requires cash. Every dollar of cash available to him is currently inside an appreciated position.

Raising $25,000,000 by selling would recognize approximately $9,500,000 of long-term gain. At a combined 23.8% federal rate, that is $2,261,000 of tax due at funding, leaving $22,739,000 to invest. The client would begin the strategy having paid, on day one, more than the entire ten-year net present value of the strategy under a liquidation exit.

> For any fully invested taxable portfolio, the funding decision dominates the manager selection decision.

3 · ENTRY: THE RIGHT OF SUBSTITUTION

3.1 The mechanism

The client's broader balance sheet included an irrevocable grantor trust that had recently received a large cash distribution not yet deployed. That fact made a different route available.

Under IRC §675(4)(C), a trust instrument may grant a power of administration, exercisable in a non-fiduciary capacity, to reacquire the trust corpus by substituting other property of an equivalent value. Holding that power is itself what causes the trust to be treated as a grantor trust for income tax purposes.

The consequence is the useful part. Because the grantor is treated as owner of the trust corpus, an exchange between the grantor and the trust is a transaction with oneself. No gain is recognized. The client transferred appreciated securities to the trust; the trust delivered $25,000,000 of cash.

The result solved two problems with one transaction. The sleeve was funded with the full $25,000,000 and zero recognized gain, and the trust's idle cash was put to work in a diversified portfolio it was happy to hold.

Measured against the sell-to-fund alternative, the substitution preserved $2,261,000 of tax and roughly 9% of the invested base.

3.2 Which lots to substitute

The lots transferred were the lowest-basis, oldest legacy positions. The reasoning was primarily portfolio construction rather than tax: these were holdings the client had been trying to reduce for years and had never sold, precisely because the embedded gain made selling painful. The substitution moved them out of the personal portfolio without a taxable event and into an entity whose objectives they suited.

One argument that does not apply. It might be suggested that low-basis lots should be retained on the theory that embedded gain is what harvesting works against. That reasoning fails. Harvesting occurs in the new, cash-funded sleeve, which begins at a basis equal to its market value. The legacy positions' basis has no bearing on the sleeve's harvesting capacity. Low basis in an old position does not create harvesting potential; it creates a future tax liability.

One argument that does apply, and cuts the other way. Assets held in an irrevocable grantor trust excluded from the grantor's estate receive no basis adjustment at death. The IRS confirmed this in Revenue Ruling 2023-2. Low-basis assets held personally until death receive a step-up that eliminates the embedded gain entirely; the same assets held in the trust carry their basis forward and the gain is eventually taxed.

This inverts the conventional use of a substitution power. Standard basis-management planning holds the lowest-basis assets on the grantor's personal balance sheet and the highest-basis assets in the trust, and practitioners commonly exercise the swap power late in life to pull appreciated assets back into the estate for exactly this reason. Moving low-basis lots into the trust runs against that convention.

So there is a genuine tradeoff, and it should be stated rather than assumed away.

- Substituting the low-basis lots resolves the concentration the client wanted resolved, and forfeits step-up on that value.
- Substituting the high-basis lots preserves step-up, but leaves the concentrated legacy positions in the personal portfolio, where reducing them would require the taxable sale the substitution was designed to avoid.

Two considerations weigh on which is right, and neither is universal.

Estate size relative to the exemption. A client well below the exclusion amount gains little from removing appreciation from the estate and should weight step-up heavily. A client well above it gains from the freeze, and the forfeited step-up is the price of that benefit.

Reversibility. The substitution power is bidirectional and may be exercised repeatedly. Nothing about this transaction is permanent: the client can later swap cash or high-basis assets into the trust to reacquire the low-basis positions, restoring them to the estate in time for a step-up. The decision made at funding is therefore an opening position rather than a final one, which materially lowers the cost of getting it wrong, provided the power is actively managed rather than forgotten.

The narrow claim that survives. Against the alternative of selling legacy positions to fund the sleeve, substitution is strictly better: it achieves the same portfolio outcome with no gain recognized. Against the alternative of substituting different lots, the answer depends on the client's estate profile and on whether the swap power will be revisited. The memo's framework does not resolve that question, and it should not be presented as though it does.

3.3 Conditions the treatment depends on

Eligibility was determined by tax counsel, not by the investment team. The general requirements are worth stating because they are conditions, not formalities:

- Equivalent value is required; equivalent basis is not. This is precisely what makes lot selection a live design choice rather than a mechanical one.
- Under Revenue Ruling 2008-22, estate inclusion under §§2036 and 2038 is avoided only where the trustee has a fiduciary obligation, under the trust agreement or local law, to ensure the substituted and acquired properties are of genuinely equivalent value, and where the power cannot be exercised to shift benefits among beneficiaries. Valuation discipline is the condition the entire treatment rests on.
- The power should be non-assignable, for creditor protection purposes.
- The power should not be held by a person with a substantial adverse interest, to avoid gift tax exposure.

Any team contemplating this route should treat the valuation of the substituted securities as the control point. With exchange traded funds priced daily in a liquid market, that burden is light. With illiquid or closely held assets, it is the whole engagement.

◆

4 · THE HARVEST SCHEDULE

Harvest assumptions in this category are frequently sourced from manager marketing material, which is both proprietary and optimistic. This model instead uses the shape reported in published academic work.

Israelov and Lu (2022) find losses harvested of approximately 13% of portfolio value in year one, falling to high single digits in year two, and into low single digits thereafter. The schedule modeled here follows that shape:

Yields are applied to the initial funding amount, so the schedule sums to a stated percentage of capital committed rather than drifting with market value.

YEAR	LOSS YIELD	LOSSES HARVESTED	TAX BENEFIT
1	13.0%	$3,250,000	$773,500

2	8.0%	$2,000,000	$476,000
3	5.0%	$1,250,000	$297,500
4	3.5%	$875,000	$208,250
5	2.5%	$625,000	$148,750
6	1.8%	$450,000	$107,100
7	1.3%	$325,000	$77,350
8	1.0%	$250,000	$59,500
9	0.8%	$200,000	$47,600
10	0.6%	$150,000	$35,700
Total	**37.5%**	**$9,375,000**	**$2,231,250**

Two features of this curve deserve emphasis.

The benefit is severely front-loaded. The first three years account for roughly 69% of the lifetime harvest. Year ten contributes $35,700, an immaterial figure on a $25,000,000 account, and less than the management fee on the sleeve.

The decay is structural, not a market call. Harvesting requires positions trading below their cost basis. As the portfolio appreciates, fewer positions qualify, and each harvest lowers basis further, making the next one less likely. Equity markets rise in roughly two years of every three. A sleeve receiving no new contributions therefore exhausts its own opportunity set by construction.

This is the source of the widely cited seven-to-ten year productive life. It is not a forecast; it is arithmetic.

5 · THE CENTRAL IDENTITY: HARVESTING DEFERS, IT DOES NOT ERASE

This section contains the memo's central claim, and it is the one most likely to be resisted.

Harvesting a loss requires selling a position below basis and repurchasing equivalent exposure. The economic position is unchanged. What changes is the tax basis, which is now lower by exactly the amount of the loss realized.

So each dollar harvested has two effects: it delivers a tax benefit today, and it increases the eventual gain on the position by one dollar.

The model tests this identity directly against a counterfactual, the same $25,000,000 held in a passive index fund, growing at the same 7%, with no harvesting at all.

AT YEAR 10	DIRECT INDEX SLEEVE	INDEX ETF	DIFFERENCE
Market value	$49,178,784	$49,178,784	$0
Cost basis	$15,625,000	$25,000,000	($9,375,000)
Embedded gain	$33,553,784	$24,178,784	$9,375,000

The difference in embedded gain is $9,375,000. The cumulative losses harvested were $9,375,000. These are the same number, and they are the same number under any growth path, because basis reduction is dollar-for-dollar with loss realization.

The sleeve ends the period with 68% of its market value as embedded gain, against 49% for the passive alternative. The strategy has, quite deliberately, manufactured a concentrated low-basis position.

> It is not a return enhancement. It is an interest-free loan from the tax authority, drawn down over the early years of the sleeve and repaid in full at exit, unless the exit is structured so that repayment never comes due.

◆

6 · THE TWO EXITS

6.1 Exit A: Full liquidation

If the client sells the sleeve at year ten, the deferred gain is recognized.

Present value of harvest benefits over the life	$1,941,045
Extra tax at exit versus the ETF counterfactual (nominal)	$2,231,250
Extra tax at exit (present value, 5% discount)	$1,369,794
Net advantage of direct indexing	**$571,251**

Note that the nominal extra tax of $2,231,250 is precisely equal to the total nominal tax benefit harvested along the way. Every dollar comes back. What the client keeps is the difference between receiving those benefits early and paying them back late.

$571,000 on $25,000,000 is 2.3% of the sleeve, realized over a decade, roughly 23 basis points a year before fees, against a strategy that typically charges 20 to 40. On a liquidation exit, direct indexing is close to a wash once the manager is paid.

That conclusion should not be softened. It follows directly from §5 and it applies to any client whose eventual intention is to sell.

6.2 Exit B: Contribution to a donor advised fund

The client in this scenario was already charitably inclined and making substantial annual gifts. The direct indexing sleeve did not create that intent; it changed what the gifts were funded with.

A contribution of long-term appreciated securities to a donor advised fund produces two effects. The donor deducts the full fair market value, subject to the 30% of adjusted gross income limit applicable to appreciated property, with a five-year carryforward. And the embedded gain is never recognized by anyone, since the fund is tax-exempt and sells without tax.

Applied to the terminal sleeve, the entire $33,553,784 of embedded gain escapes recognition permanently.

Present value of harvest benefits over the life	$1,941,045
Extra tax at exit	$0
Net advantage of direct indexing	**$1,941,045**

The strategy is worth 3.4 times what it is worth under a liquidation exit. The $1,370,000 difference is created by nothing except the choice of exit.

There is a structural elegance here worth naming. The strategy's defining flaw, that it manufactures a large low-basis position, is exactly what makes it an ideal funding source for charitable giving, because appreciated securities are the most tax-efficient possible gift asset. The sleeve's terminal state is a liability under one exit and an asset under the other.

This also connects directly to the lot-selection framework in the companion gifting memo: the terminal sleeve consists of hundreds of individual lots with widely varying basis, and gifting should be executed at the lot level rather than pro rata.

6.3 Partial donation

Most clients sit between the two cases. The relationship is close to linear:

SHARE OF SLEEVE DONATED	NET ADVANTAGE
0%	$571,251
25%	$913,700
50%	$1,256,148
75%	$1,598,597
100%	$1,941,045

Each incremental 25% of the sleeve directed to charity adds roughly $342,000 of present value. A client can therefore size the sleeve against realistically expected lifetime giving rather than treating the decision as binary.

The practical implication for an advisory process is direct: the charitable intent conversation should precede the manager selection conversation, because it determines the size of the prize by a factor of more than three.

7 · SENSITIVITIES

Discount rate. The liquidation case is a timing arbitrage, so it is sensitive to the discount rate applied to tax cash flows. A higher rate increases the value of early benefits relative to the deferred cost, and a lower rate compresses it toward zero. The workbook computes this grid live across a 3% to 7% range. The charitable case is far less sensitive, because it has no offsetting terminal cost to discount.

Growth rate: no effect on the liquidation result. This is the model's most counterintuitive output. Because the extra exit gain equals cumulative losses harvested regardless of path, and because both the sleeve and the counterfactual grow identically, the growth assumption cancels. A view on market direction is not a reason to adopt or reject this strategy.

The rate at which losses are valued. The base case values harvested losses at the same 23.8% applied to the exit gain. Where a client can apply losses against income taxed at a higher rate, the benefit rises materially, because deductions are taken at a high rate and repaid at a low one.

This deserves a caution rather than a recommendation. Capital losses offset capital gains first, and only $3,000 per year against ordinary income. The rate arbitrage is therefore available only where the offsetting income is itself gain-character, or where specific entity structures route ordinary-character cash flow away from the individual return. Both are client-specific determinations for tax counsel, not modeling assumptions an investment team should adopt on its own authority. The base case here deliberately claims none of this.

Tracking error. Not modeled. A harvested sleeve accumulates active risk against its benchmark as constrained positions build up. Over a ten-year life this is a real cost that offsets some portion of the tax benefit, and a fuller treatment would price it.

8 · THREE STANDING RULES

8.1 Never fund a sleeve by realizing gain while a non-recognition route exists

The funding decision is worth more than several years of harvesting. Before selling, exhaust: substitution powers in existing grantor trusts, unfunded cash in related entities, new contributions, in-kind transfers, and pending liquidity events. In this scenario the substitution was worth $2,261,000, nearly four times the ten-year net present value of the strategy under a liquidation exit.

8.2 Do not adopt direct indexing without an identified exit

A sleeve without a non-recognition exit is a deferral vehicle. That may still be worth doing, but it should be presented and sized as such. The failure mode is a client who adopts the strategy for its harvesting, holds it for a decade, and then discovers that unwinding it costs back everything it produced.

8.3 Substitute the lots you least want to sell

Legacy basis is irrelevant to a cash-funded sleeve's harvesting capacity. The correct screen for which lots to move is portfolio construction, which positions do you want out of the personal portfolio and never want to sell, not tax characteristics.

9 · IMPLEMENTATION NOTES

Benchmark selection. The modeled sleeve tracks the S&P 500. A broader index offers more constituent names and therefore more harvesting opportunities, at the cost of higher trading friction in small positions. This is a real tradeoff and would merit its own analysis before implementation.

Wash sale discipline. Harvesting requires repurchasing correlated but not substantially identical exposure within the 30-day window. This constraint is what generates tracking error and it tightens as the sleeve ages.

Vintage accounting. If contributions are made over time, each behaves as its own vintage with its own harvest curve and its own clock. A sleeve receiving regular contributions has a materially longer productive life than the single-funding case modeled here.

Long-only versus long/short. Extension strategies generate larger and longer-lived loss harvesting than long-only. That is a separate strategic question, addressed in a companion write-up, and it does not change any conclusion here: the lifecycle framework and the central identity of §5 apply to both.

Sequencing with the charitable program. If the exit is charitable, the gifting should be planned alongside the sleeve rather than after it, including whether a donor advised fund is established early enough to receive contributions in high-income years.

10 · APPENDIX: MODEL STRUCTURE

The companion workbook contains six sheets. All inputs sit on a single Assumptions page; every other sheet derives from it.

- **Entry** – substitution versus sell-to-fund
- **Harvest** – loss schedule, tax benefits, and the basis and gain path over the sleeve's life
- **Exits** – the ETF counterfactual and both exit paths, including a live cell that tests the §5 identity rather than asserting it
- **Sensitivity** – discount rate, loss valuation rate, and share donated

All figures in this memo are quoted from that workbook, which is authoritative where the two differ.

This analysis was developed as a modeled scenario. The structure it describes was evaluated but not implemented; a change in the client's tax position removed the need for the losses before execution. Nothing here should be read as tax advice, and the trust eligibility determinations described in §3.3 were made by tax counsel.