THREE HORIZONS CAPITAL · PRICED VERSUS REALIZED, NOTE THREE

All of the Risk, None of the Premium

A private strategy's liquid twin earned exactly what the index earned, and charged more than twice the volatility to do it.

21.5%
The liquid twin's annualised return over 43 months

22.2%
The index's annualised return over exactly the same window

36%
The twin's peak-to-trough fall, against 10% for the index

Monthly, 2022-12-30 to 2026-07-31 · US dollars · Returns derived from adjusted closing prices, not from a pre-computed return field
Professional investors · No fund, manager, pension system or listed security is named

Executive summary

FINDING 1 A daily-priced basket of the same fee-and-carry economics a private strategy buys returned **21.5%** over 43 months. The index returned **22.2%** over exactly the same window. The basket did it at **30.6% volatility against 13.8%**, and a **36.5% drawdown against 9.8%**.

FINDING 2 The quarterly appraisals of the private vehicles show nothing resembling that fall. They step up most quarters, with occasional small markdowns.

FINDING 3 **The reporting calendar does not explain it.** Observe the same exposure only four times a year and you still see a **30.8%** fall. Quarterly observation hides 5.7% of it, and no more.

FINDING 4 Across every regime window with enough data going back to 2008, this exposure's correlation to equities sat between **0.65 and 0.78**, including through the crisis, while the bond hedge swung from -0.44 to +0.54. This is not a diversifier that stopped working.

FINDING 5 A daily-priced income vehicle run by the same kinds of manager, used here as a control, returned 9.0% at 13.1% volatility. Being priced honestly is not the same as being safe, but it is visible.

> **The mark is quarterly. The risk is not.** And the gap is not a curiosity about one niche strategy. It is a property of every private holding whose risk you measure from its own marks.

1 This is not really a note about one strategy

The strategy in question is small. It raised a few billion dollars last year and a handful of public pension systems hold it in the data we keep. On its own it would not be worth a paper.

It is worth a paper because it is the cleanest laboratory anyone has. Almost every private strategy is hard to study, because there is nothing liquid that holds the same thing. This one is the exception. The cash flow it buys, a minority claim on a manager's fee income and carried interest, is also the cash flow that listed asset managers report as their own earnings. So the same economics exist in two places at once: appraised quarterly in a fund, and priced every day on an exchange.

That accident of structure lets us ask a question that is normally unanswerable. **What would the risk of a private holding look like if somebody had to put a price on it every morning?**

2 What the appraisals say

Read at face value, the reported record describes a calm asset. Net internal rates of return across the vehicles land in a wide but unremarkable band. Distributions relative to capital paid in sit below one times for most of them, which is normal for a strategy whose cash arrives as a slow drip of fee revenue rather than as exit proceeds. Values step up most quarters.

They are not purely one-directional. Looking quarter by quarter rather than at the latest mark, several vehicles show genuine markdowns, and an early version of this analysis was wrong to say otherwise. The point is not that appraised values never fall. It is that they fall rarely, by small amounts, and never in a way that suggests the holding could lose a third of its value in a few months.

3 The same economics, priced every day

We built a basket of the nine largest listed managers whose earnings are that same fee-and-carry stream, weighted by market value and rebalanced monthly, over the 43 months to 2026-07-31.

CHART 1 · SAME DESTINATION, DIFFERENT JOURNEY

The liquid twin of this private strategy ended the period exactly where the index did, having taken a completely different road to get there.

Both rebased to 100 at the start. The basket is nine listed managers whose earnings are the same fee-and-carry economics a GP stakes fund buys a slice of.

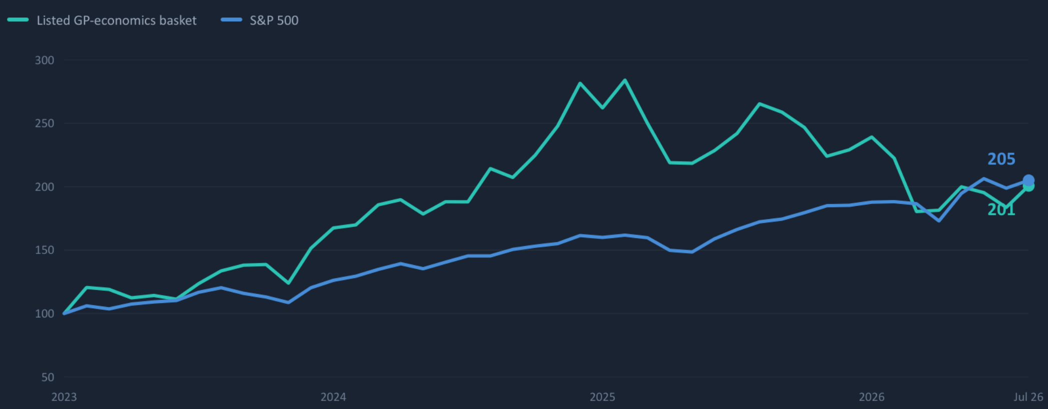

Monthly, 2022-12-30 to 2026-07-31. The basket fell 36% peak to trough on the way; the index fell 10%. Neither is named at constituent level in this note.

Two routes to the same place. The basket rose more than 180% at its peak and gave most of it back.

The basket ends the period at almost exactly the index's level. It got there by rising far higher and then falling much further, twice.

CHART 2 · ALL OF THE RISK, NONE OF THE PREMIUM

Twice the volatility of the index, and not one basis point of extra return for it.

Annualised return against annualised volatility, same window, same currency, same monthly frequency.

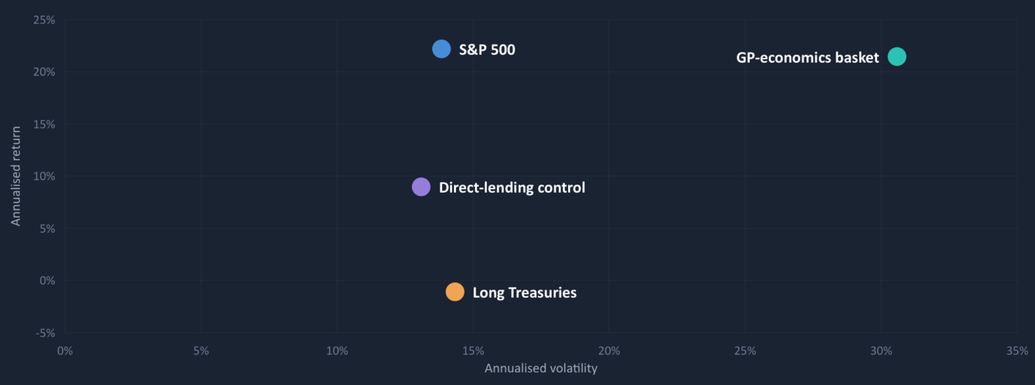

The basket returned 21.5% against the index's 22.2%, at 30.6% volatility against 13.8% and a 36.5% drawdown against 9.8%. A point sitting far to the right without sitting any higher is risk that was taken and not paid for.

Return against volatility. Two points at the same height, one of them a long way to the right.

> This is the finding. Not that the exposure was risky, which anyone might guess, but that over this window **the extra risk earned nothing at all**. An allocator holding it in appraised form saw a smooth line and a respectable internal rate of return, and had no way to see that the same economics, priced openly, had delivered the index return for double the risk.

4 The reporting calendar is not the explanation

The obvious defence is that private vehicles are marked quarterly, so of course the path looks smoother. That defence is testable, and it does not hold.

CHART 3 · THE CALENDAR IS NOT THE EXPLANATION

Look at the same exposure only four times a year and you still see a 31% fall. The reported marks showed nothing of the kind.

Identical series, identical window. The only difference is how often you look at it.

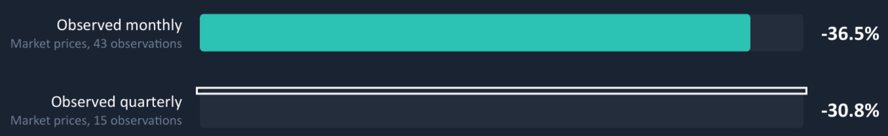

Observing quarterly instead of monthly hides 5.7% of a 36.5% fall. That is the whole effect of the calendar. Whatever else explains the gap between this and a smooth quarterly appraisal, it is not the reporting frequency.

The same series, observed at two frequencies. Looking less often removes remarkably little of the fall.

Take the identical basket and look at it only at quarter ends, exactly as an appraisal cycle would. The drawdown you would have observed is **30.8%**, against 36.5% monthly. Observing less often hides 5.7% of a 36.5% fall, and that is the entire effect of the calendar.

So a quarterly observer of market prices would still have seen roughly a third of the value disappear. The appraisals showed nothing of the kind. Whatever accounts for the difference, **it is the valuation method,**

not the reporting frequency, and that distinction matters because frequency is the part everybody already knows about.

5 There was never a diversifier here to break

CHART 4 · THERE WAS NEVER A DIVERSIFIER TO BREAK

The bond hedge flipped sign across these four eras. The correlation of GP economics to equities did not move at all.

Correlation of monthly returns to the S&P 500, by regime. Windows with fewer than 24 observations are not reported rather than filled in.

REGIME	BOND HEDGE TO EQUITIES	GP ECONOMICS, NAME 1 TO EQUITIES	GP ECONOMICS, NAME 2 TO EQUITIES
Pre-GFC	-0.10 n=60	*too few obs* n=6	*too few obs* n=11
GFC	-0.12 n=24	0.78 n=24	0.66 n=24
Post-GFC normal	-0.44 n=146	0.70 n=146	0.67 n=146
Post-Mar-2022	+0.54 n=53	0.65 n=53	0.65 n=53

The bond hedge moved from -0.44 in the long post-crisis era to +0.54 after March 2022, a genuine regime change. Across every window with enough data, GP economics sat between 0.65 and 0.78. This is not a diversifier that stopped working. On the data we hold, it never was one.

Correlation to equities by regime. Windows without enough observations are left empty rather than filled in.

The bond hedge did what everyone now knows it did: it moved from -0.44 across the long post-crisis era to +0.54 after March 2022. That is a genuine regime change and it is why so many balanced portfolios behaved unexpectedly.

This exposure did not participate in that story. In every window where we have enough observations to report a correlation honestly, it sat between 0.65 and 0.78 against equities, and the highest reading of the three is the crisis window. **It did not stop diversifying in 2022, because on the data we hold it was never diversifying.** That is a sharper caution than a broken-hedge story, because a broken hedge sounds like bad luck and this does not.

6 The strongest objection, and what survives it

The serious reply is this: the holding is genuinely illiquid, so daily volatility is somebody else's problem. Nobody can force a sale. The appraisal is the number that governs, and a path you cannot trade is not a risk you bear.

That objection is largely right, and it is why this note argues for repricing rather than for selling. Illiquidity changes your *behaviour*, which is genuinely valuable, but it does not change the underlying asset. And the distinction stops being academic at exactly three moments.

- **When you size the allocation.** A position sized against appraised volatility is larger than the same position sized against observed volatility. The sizing decision is made once, quietly, and never revisited.
- **When you need liquidity for something else.** The correlation matters most precisely when the rest of the book is falling, and this exposure has never been uncorrelated.
- **When you mark for funding or solvency purposes.** A smoothed number flatters a funding ratio in exactly the periods when the sponsor most needs to know the truth.

7 What to do about it

Not sell. There is usually no way to sell, and the strategy may well be a perfectly good one. The issue is measurement, and measurement is inside your control.

> Most institutional risk systems carry a private sleeve at the volatility its own marks imply. **Feed a smoothed number into a risk model and the model understates total portfolio risk, by the most in the place you have grown the allocation hardest.** That is a modelling input you own, and you can change it this quarter without touching a single holding.

Three practical steps, in ascending order of effort.

- **Find out what volatility your risk model is currently using for each private sleeve.** In most books nobody has looked, and the answer is frequently the appraised series.
- **Build a liquid analogue for each sleeve where one plausibly exists**, and carry its observed volatility and correlation into the risk model instead. Not every sleeve has a clean twin. Some do.
- **Use the analogue as an interim check between appraisal dates.** It will not tell you what your holding is worth. It will tell you whether the last mark is still a reasonable description of the world.

WHAT WE CHECKED, WHAT WE CANNOT SEE, AND WHAT WE GOT WRONG

The comparison is a proxy, and the population does not match. The basket is nine of the largest listed managers. The private strategy buys minority stakes in mid-sized and emerging managers. Our own look-through work confirms no overlap between the two sets, and would not expect any. The basket represents the *type* of cash flow, not the *population*. Read every comparison here as being about the economics, not about the specific holdings.

All returns are derived from adjusted closing prices. While building this we found a pre-computed return field in our own data that does not describe the same period as the monthly row it sits on: across the nine names, **370 of 383 monthly comparisons disagree with the price-derived return by more than half a point**, with a median gap near six points and a worst month of thirty-six. Every series here is therefore rebuilt from prices, and the defect is raised internally rather than quietly worked around. Benchmarks come from a separate table at weekly frequency; both series are cut to one common set of month ends, which is why the window is stated precisely.

Correlations are reported only where there are at least 24 monthly observations. Two of the twelve cells in the regime table are left empty for that reason. An earlier version of this work reported a correlation for a period in which one of the securities had barely started trading.

The appraisal record is thin, and worth being precise about how thin. Of the 431 institutions we track, **31 disclose any private-markets performance data at all**, and the vehicles relevant here are matched on fund naming rather than on a strategy classification, because the strategy field is empty for every one of them. There are known defects in the record too, including at least one commitment captured in the wrong units and one manager appearing under both a legacy and a current name. None of this is a reason to distrust the appraisals. It is a reason to be careful about how much any conclusion drawn from them can be asked to carry.

An earlier version of this analysis said the appraised values only ever rise. That was wrong, and re-reading the quarterly record rather than the latest mark corrected it. Small markdowns do occur. The argument does not depend on their absence.

The Repricing Note

Send us one line describing a private sleeve: what it holds and roughly how big it is. We will come back with a single page. Its appraised volatility against the observed volatility of the closest liquid analogue, the same for drawdown, the correlation to your listed book, and what the difference does to your total portfolio risk number.

No holdings data required, no engagement required, and the method is set out above in enough detail to run it yourself if you would rather.