Evergreen vs. Drawdown Funds: Risk, Returns and Cash Flows *

Prof. Greg Brown
IPC & Kenan-Flagler Business School
gregwbrown@unc.edu

William Volckmann
IPC
wmvolckmann@unc.edu

June 11, 2025

Abstract

Evergreen funds are a rapidly growing segment of the private capital universe. While open-ended structures have existed for decades in some asset classes such as real estate, the more recent products have included a broader set of assets such as private equity and credit. In addition, many new funds are designed to appeal to individual investors (e.g., low minimum investments). Currently, little is known about the portfolio properties of evergreen funds versus traditional closed-end, drawdown funds. In this analysis, we consider the differences between evergreen and closed-end funds in a simulation exercise with a focus on expected returns, risk, and portfolio cash flows. While the average risk-adjusted performance is similar across fund types, we show that investing in a series of closed-end funds is likely to lead to substantially more variation and unpredictability in cash flows and portfolio risk levels than an investment in evergreen funds. This suggests that investors without a large diversified portfolio of closed-end funds could face unwelcome challenges undertaking portfolio allocation and rebalancing as well as slightly lower expected returns. We also propose some additional questions that investors should examine when trying to decide what structure is most appropriate for their investment needs.

*The authors thank the Institute for Private Capital and its members for support of this project.

1 Introduction

There is a long history of open-ended fund structures in some private asset classes with real estate being the most common example. Typically these funds have held a blend of public and private assets to manage liquidity risk and provide some flexibility to arbitrage between public and private market valuations. Yet, these funds have sometimes faced liquidity challenges such as real estate ODCE funds during the global financial crisis (GFC) and more recently gating of redemptions at Blackstone's BREIT.[1] Currently, there is fast-growing interest in "evergreen" funds in other asset classes such as private equity, private credit and infrastructure. For example, Pitchbook recently estimated there have been over 200 product launches with an estimated $400 billion in assets under management since 2019 (Carmean et al., 2024).

Evergreen funds allow new investors to buy into the fund at net asset value, and existing investors can withdraw capital at regular intervals (provided liquidity is available). Some observers have termed these funds "semi-liquid" investments. A typical evergreen fund might hold 10-20% in liquid assets, although some evergreen funds hold almost no liquid assets. The remaining value of the fund is allocated to illiquid private investments on an ongoing basis with the fund general partner (GP) managing new investment opportunities, the portfolio of existing holdings (and exits), and the target level of liquidity. In a typical structure, evergreen funds will have a management fee and an incentive fee (carried interest). Management fees in evergreen funds are usually based on fund net asset value (rather than committed capital during the investment period as is common in a traditional drawdown fund), and incentive fees are based on rolling total returns (rather than on realized returns since inception), but ultimately fees are similar in both fund structures (Blostein et al., 2024).[2]

Much of the demand for evergreen structures has come from individuals through the so-called "wealth" channel. These investors likely have different goals and constraints than the large institutional investors that have dominated the growth in private closed-end funds in recent decades. Specifically, individuals are less likely to have highly diversified portfolios of private funds (man-

[1] Blackstone's private credit fund, BCRED, also reached it's 5% redemption limit in 2022, but apparently did not limit withdrawals.

[2] While certainly worthy of additional examination, the primary focus of our analysis in this paper is not fee differences between fund types.

aged by dedicated professional staff), are more conscious of liquidity constraints, have shorter investment horizons, and are generally less tolerant of substantial variation in cash flows from private investments.[3]

In terms of individual investors, higher net worth households typically invest a larger proportion of their wealth in alternative assets, primarily in buyout and venture capital equity funds (Gabaix et al., 2024). This is partly explained by the high minimum investment required by a typical private equity fund, but also by legal requirements that private fund investors be accredited or qualified investors. Evergreen funds can provide an alternative because they often allow investors access to private markets without the large investment minimums that are typically required of drawdown funds. However, changes in legal constraints are under consideration. Recent discussions within the Securities and Exchange Commission have focused on increasing access to private funds for retail investors, decrying "misguided efforts to shield retail investors from risk" via legal barriers, and noting that a shrinking public market (in terms of publicly listed companies) and the growth of private markets practically necessitates broader access to private markets (Myers, 2024).

Murphy and Hadas (2024) discuss trade-offs between drawdown funds and evergreen funds. They conclude that evergreen funds have an advantage in liquidity (despite still being only semi-liquid), and evergreen funds also have the advantage of being relatively simple for LPs since the fund manager manages private market exposure and capital reinvesting. On the other hand, drawdown funds are easier to access (despite having typically higher investment minimums) because the universe of drawdown funds is so much larger, and the authors submit that drawdown funds have an advantage in performance since evergreen funds typically hold some capital as cash.

On the other hand, Wood et al. (2024) note that an evergreen fund puts LP capital to work immediately, whereas drawdown funds typically involve years of uncalled capital which is often held as cash (or near-cash). This logic can be extended into a "reinvestment risk" faced by LPs in drawdown funds whose realized capital (fund distributions) once again must be called into a

[3]These differences in preferences can be because of taxes and liquidity constraints as well as the desire to not hold large cash buffers or rely on potentially expensive lending facilities. See, for example, Sikes and Verrecchia (2012).

new fund, whereas evergreen fund managers handle reinvestment automatically. Eissler (2024) similarly concludes from a simple simulation of an evergreen fund and a drawdown fund that an investor with an extant private equity program might be best off continuing with drawdowns, but a newer investor might be better off using an evergreen fund to generate immediate capital deployment into a diversified portfolio.

Because evergreen funds have become prominent only recently, there is a dearth of research investigating how some of the aforementioned factors—namely, liquidity management and capital deployment—ultimately manifest for investors in different fund types. To that end, our goal in this paper is to understand differences in return, risk, and cash flows between fund types and how different investor types might consider trade-offs via stylized simulations of fund types. We furthermore outline portfolio management perspectives, namely fundamental differences in creating and managing a portfolio of private investments.

2 Comparing Drawdown and Evergreen Funds

Given the newness of evergreen-style funds, it is not possible to conduct a detailed historical comparison of returns, risk, and cash flows. Instead, we turn to a simulation exercise to better understand how differences in fund structures might affect variables of interest. The benefit of this is that we can create a laboratory with complete clarity with respect to our modeled inputs and outputs. The primary downside is that the analysis is hypothetical and potentially constrained by simplifying assumptions necessary to generate the simulations. With these caveats in mind, we simulate three feasible strategies for an accredited investor seeking exposure to the investments of a particular GP:

1. A single investment in a closed-end drawdown fund. Figure 1 depicts the exposure (i.e., percent of committed capital invested) for an investor who makes a commitment to a new fund (at year 0) to a fund with a 12 year life. We assume the fund deploys capital over a 5-year investment period and returns capital (starting in year 3) to the investor until fully liquidated. In our analysis, this is a just straw-man strategy since it obviously results in the investor having very low exposure to the investment portfolio in most years (with the bulk

of the committed capital assumed to remain in cash).

FIGURE 1: Investment Profile for Single Drawdown Fund

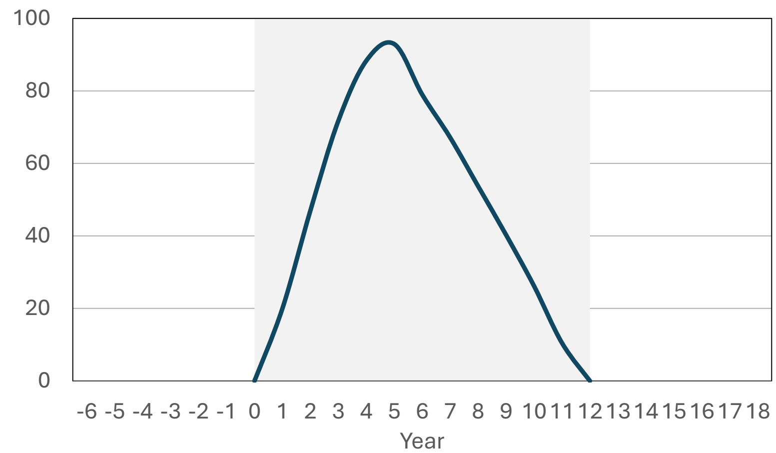

2. A single investment in an evergreen fund (at year 0) that is liquidated at the end of year 12. For our base case we assume that the fund targets being 90% invested and maintains a 10% cash buffer (but examine other levels of targeted liquidity buffer as well). Figure 2 depicts the investment profile for this strategy. Clearly, the investor will benefit (on average) from being immediately and continuously invested over the 12-year holding period, though there will be a performance drag (on average) from the cash buffer.

FIGURE 2: Investment Profile for an Evergreen Fund

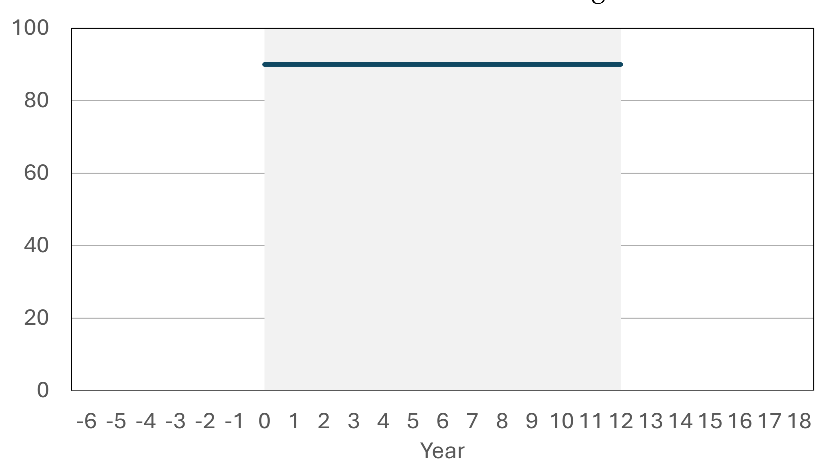

3. A rolling series of (properly-sized) investments in closed-end drawdown funds like the one described above. Specifically, we assume that the investor makes a new commitment every 3 years and started this process 6 years ago (in year -6).[4] Thus, as shown in Figure 3 by

[4]Or alternatively, we could think about the investor purchasing interests in the two seasoned funds in the secondary

the thin dotted lines, the investor starts the evaluation period we consider (years 0 through 12) with two seasoned funds and a new commitment to a third fund. Then the investor makes new commitments at the beginning of years 3, 6, and 9. The dark solid line in Figure 3 depicts the investment profile for the combined series of drawdown funds. While the exposure varies from year to year during years 0 through 12, the average exposure to the investment portfolio is around 90% (with the remainder assumed to be in cash). We discuss the appropriate sizing of commitments below.

FIGURE 3: Investment Profile for Series of Drawdown Funds

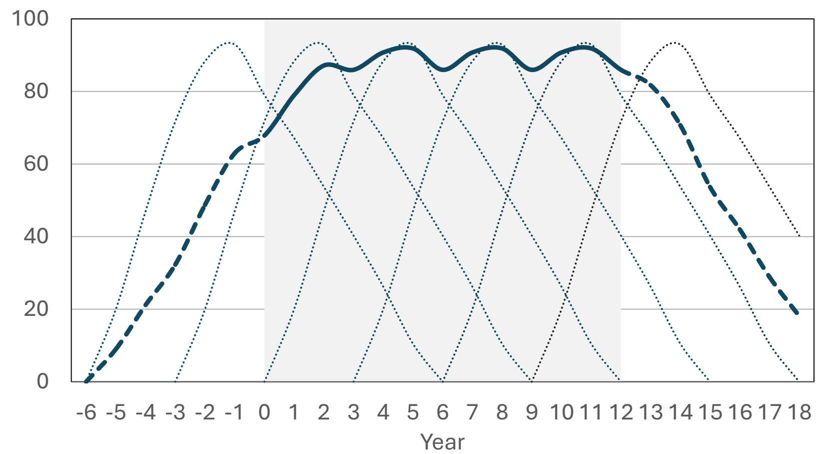

The specifics of the simulations are provided in the appendix, and we just provide an intuitive overview here. We consider an investor who seeks to have $90 in capital committed to private assets at year 0 with $10 remaining in cash. We assume that the underlying portfolio of assets in both drawdown and evergreen funds is the same even though the structures are different. Each year, the funds generate returns on their portfolio of assets with two components: a random market-wide component that is common across funds as well an independent random "alpha" component. We calibrate the random distributions of market returns and alphas to match net-of-fee historical performance similar to that observed for buyout funds in the MSCI-Burgiss data with the assumption that the market beta of the fund is 1.0 (similar to what recent academic literature documents for buyout funds). Specifically the average annual market return is 12% and the average alpha is 3%.

We also model capital calls and distributions as random variables with ranges similar to those

market at $t = 0$ and then subsequently liquidates all funds at the end of year 12.

observed in the MSCI-Burgiss data, but we force funds to be fully liquidated by year 12.[5] For the drawdown funds, we assume that the investor earns a cash (risk-free) interest rate on committed but uncalled capital (including capital distributed from an earlier fund and not yet called by a new fund). Likewise, the GP of the evergreen fund earns the risk-free rate on cash it holds as a liquidity buffer. Obviously fees can vary across funds, LPs and fund structures, but we don't seek to model those effects here (but provide a qualitative discussion below). We conduct 10,000 random simulations for each type of fund/strategy and evaluate portfolio returns, risk, and cash flows from year 0 to year 12.[6]

3 Simulation Results

For each type of simulation, we report in Table 1 the mean and standard deviation of annualized returns and percent invested, as well as the mean fund alpha and beta. Recall that the return for evergreen funds and the series of drawdown funds can be thought of as the LP committing capital at year 0 and fully liquidating its funds at fair value on the secondaries market at the end of year 12. Also, recall that the values we calculate for the series of drawdown funds are calculated from years 0, . . ., 12, so that the evaluation period is the same across the three scenarios.[7]

TABLE 1: Performance and Risk Metrics

Metric	Single Drawdown	Evergreen	Series of Drawdowns
Annualized Return - Mean	6.5%	13.5%	12.7%
Annualized Return - StDev	2.2%	5.0%	4.3%
Percent Invested - Mean	28.7%	90.0%	90.0%
Percent Invested - StDev	2.7%	0.0%	19.5%
Fund Alpha - Mean	0.9%	3.6%	2.9%
Fund Beta - Mean	0.31	0.91	0.91

Based on 10,000 simulations each.

[5]Fully liquidating at the end of year 12 is a modeling choice just for convenience to contain the time frame we model. Of course, in reality many funds still have non-trivial NAVs after year 12. However, if we extend the cap on fund life to be longer (e.g., 15 years), it does not change the main results and conclusions in a meaningful way.

[6]Specifically for the series of drawdown funds, we ignore the performance before year 0 and after year 12 which is akin to assuming the investor buys a seasoned portfolio of the first two funds at year 0 at fair value and liquidates all fund holdings at the end of year 12 at fair value.

[7]That is, we look at the portfolio with a series of drawdown funds for periods in which there are at least 3+ active and overlapping staggered funds—when the series is ”seasoned” up until year 12.

Unsurprisingly, the average annual return for single drawdown fund is the lowest at 6.5%, since the investor with a single drawdown fund is just 28.7% invested in private markets on average over the 12-year evaluation period. Evergreen funds have the highest average return at 13.5% compared to 12.7% for the series of drawdown funds. The modestly higher return for evergreen funds is a bit surprising since it holds the same assets and amount of cash on average as the series of drawn funds. Specifically (and by design), the percent invested in underlying investment (versus held in cash) for evergreen funds and the series of drawdown funds are both 90% on average, and the average fund beta for both the evergreen funds and the series of drawdowns is about 0.91, closely reflecting the average 90/10 allocations to portfolio companies and cash. That said, the series of drawdowns exhibit a wider variation in percent invested since contributions and distributions are not deliberately tailored by the GP to ensure that the LP maintains its desired 90/10 balance; that onus is on the LP, unlike with the evergreen fund.

Indeed, the reason why (on average) the evergreen fund has higher returns than the series of drawdowns, despite investing in the same assets, is its greater consistency over time in maintaining the 90/10 balance. To see why, consider the following simple numerical examples. Suppose an investor has \$1 in capital, faces an expected 15% private fund return, and for simplicity a zero risk-free rate. In the first example, the investor wants its portfolio to be exactly 50% cash and 50% fund value in all periods. The expected value after one period is $0.50(1 + 0.15) + 0.50 = \$1.075$, and then 50% of this is reinvested at the private fund rate, giving $0.50(1.075)(1 + 0.15) + .50(1.075) \approx \$1.156$ after the second period. Now suppose instead the investor only wants the 50/50 balance across periods *on average*—for the sake of argument, suppose the investor goes with 75% private fund and 25% cash for the first period allocation, and 25% private fund and 75% cash for the second period allocation. Their expected return after one period is $0.75(1 + 0.15) + 0.25 = \$1.1125$, then 25% of this is reinvested at the private fund rate, giving $0.25(1.1125)(1 + 0.15) + 0.75(1.1125) \approx \$1.154$. The same result follows if the 75/25 order is reversed.

The difference after only two periods is obviously small, but compounding this difference over 10 periods yields a difference of 1.25 percentage points in total return. Again, not large, but not negligible either, like the 13.5% versus 12.7% seen in the simulation results. The takeaway is that for an average 90/10 balance, the total return is maximized when the balance is as close to 90/10 as possible over time; and furthermore, that total return deviates more from that maximum the more the balance deviates from the 90/10 target.[8] Since the LP controls neither capital calls nor distributions with a series of funds, the LP cannot maintain that constant ratio when investing in a series of funds, and the return suffers accordingly relative to the evergreen fund and its constant ratio. This is also reflected in the higher evergreen alpha of 3.6% compared to the 2.9% of the series of drawdowns.

Histograms for total returns, alphas, and betas are shown in Figure 1. The top row of graphs shows results for the evergreen funds and the bottom row shows results for the series of drawdown funds. Consistent with the results in Table 1, the evergreen fund is observed to have both a higher mean total return as well as more variation in total returns. As with returns, evergreen alphas have a wider spread, but betas are a different story: evergreen betas have less variation than for the series of drawdown funds, and the series of funds has a noticeable right skew that is not evident in evergreen betas.

It may seem paradoxical that the evergreen fund has more variation in total return and alpha, but lower variation in beta – especially since, as discussed above, we can trace the slightly higher return for evergreen funds to lower variation in the percent of assets invested in the portfolio assets (i.e., not cash). However, these facts are all tied together. The intuition is as follows: In the series of drawdown funds the variation in beta is driven by the capital calls and distributions of the series of drawdown funds. Sometimes the LP in the series of funds will get a large distribution and have no option but to hold it in cash while waiting for the capital to be called by another fund (thus lowering the overall portfolio beta). But sometimes the opposite will occur and the LP will have to meet a large capital call without enough cash on hand and thus have to borrow (and thereby levering the portfolio slightly and increasing the overall portfolio beta). In fact, in our simulation, the series of drawdown funds has a negative cash position (is borrowing) in about 10% of years.

The differences in cash flows is one of the most important differences between the evergreen funds and the series of drawdown funds. To illustrate this, LP net cash flows are shown in Figure

[8] The underlying mathematics are practically identical to those showing that a mean-preserving spread will reduce the utility of an agent with concave utility. Which is to say, the result for the evergreen fund versus the series of drawdown funds can be proved generally when the evergreen fund has a lower variation in cash holdings than the series of drawdown funds. When using an 80/20 balance, the evergreen fund has an average total return of 13.4% compared to the series return of 11.7%.

FIGURE 4: Performance and Risk Metric Distributions for Evergreen Fund and Series of Fund Simulations

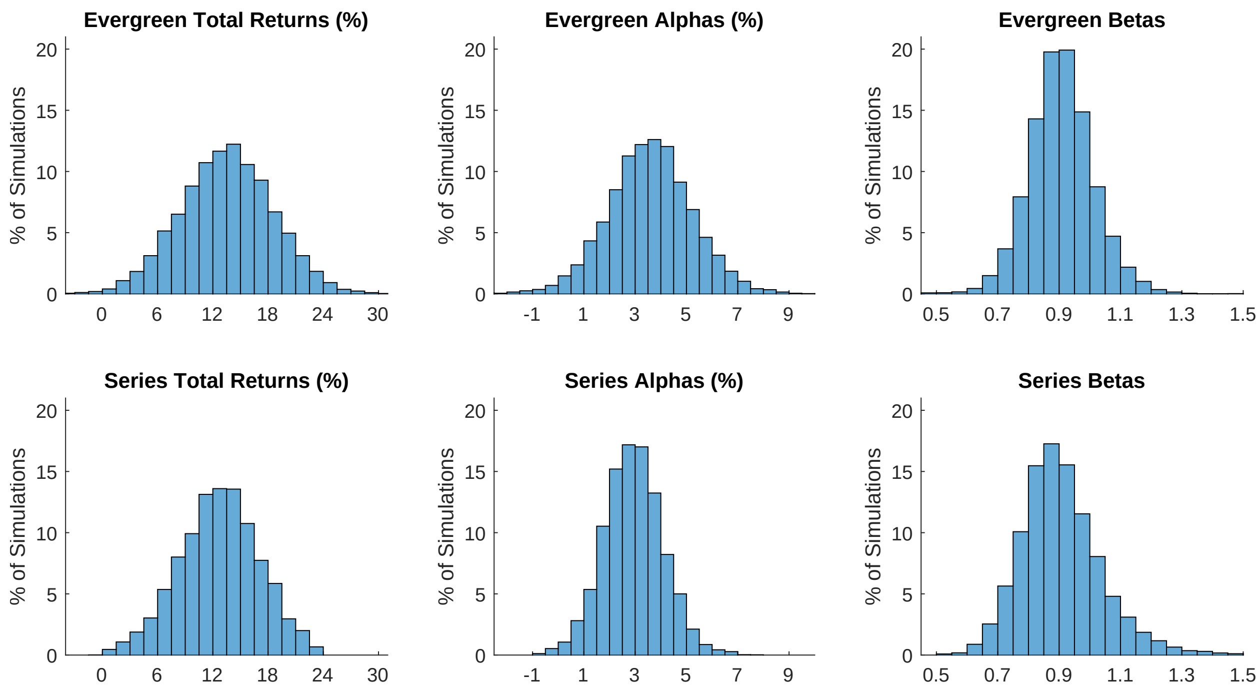

This figure shows histograms of total annualized return, alpha, and beta for 10,000 simulations of an evergreen fund (top) and a series of drawdown funds (bottom).

5. The evergreen fund has very simple cash flows: the LP invests \$100 at year 0, and the GP takes care of everything else until the LP decides to sell its share at the end of year 12. Naturally, there is a wide range of possible outcomes at year 12 based on market returns and realized alpha for the fund, but there are no cash flows in the interim years. In contrast, the series of drawdown funds involves considerable cash flow activity throughout its life. The initial cash flow and the sale in year 12 are not shown since they dominate the y-axis, but all interim periods show fluctuations over time with considerable heterogeneity within and across periods. The cash flows can be quite large, sometimes as much as 146% of the initial commitment in magnitude in a single period. We also see cyclical features with a 3-year frequency based on the assumed schedule of capital commitments (e.g., higher average cash flows in years 3, 6, and 9).

Overall, our simulations provide some interesting takeaways:

1. Approximating an investment in an evergreen fund with a series drawdown funds is complicated and will result in comparatively large variation in interim cash flows and portfolio

FIGURE 5: Cash Flow Distributions for Evergreen Fund and Series of Fund Simulations

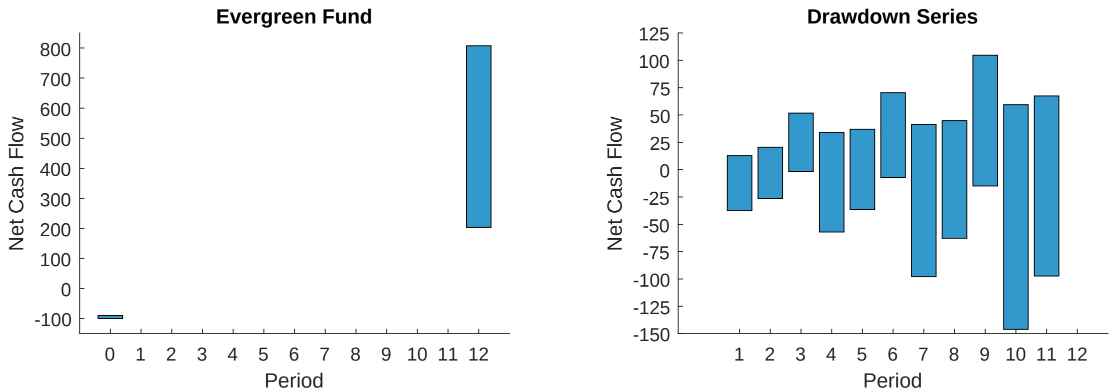

This figure shows the 90th and 10th quantiles of net cash flows in each period for 10,000 simulations of an evergreen fund (left) and a series of drawdown funds (right). The histogram for draw downs excludes the initial outflow and final inflow so as to not dominate the y-scale.

allocations to the underlying fund assets.

2. The additional variation in cash flows and allocations for the series of drawdown funds drives *higher* variation in portfolio betas and slightly lower average returns. Perhaps counterintuitively, these also result in slightly lower variation in total returns and alphas relative to the evergreen fund.

3. Together these results suggest that investors in the series of drawdown funds should be able to effectively manage the inherent cash flow and allocation variation. We expect large institutional investors with highly diversified portfolios of private assets (including the ability to co-invest and do directs) to have this advantage. On the flip side, investors with limited ability to manage the cash flow and allocation variation should benefit from the evergreen structure. We expect that this describes many individual (e.g., "wealth") investors.

While this final point is likely intuitive to most industry professionals contemplating this question, it is nonetheless interesting to see that it can also be shown in a rigorous simulation – and has implications for both average returns and portfolio risk.

4 Other Real-World Considerations

Our simulation exercise is very stylized and ignores a range of other real-world issues. Here we only note some of these issues and hope to explore them in subsequent research.

The assumption that the manager of an evergreen fund can maintain exactly a 10% ratio of cash is obviously unrealistic. Data on the typical variation in the liquidity ratios of evergreen funds would be informative and allow us to estimate the effect of this variation on average returns and market betas. Of particular interest would be how variation in cash holdings relates to market cycles.

Drawdown fund investments and evergreen investments might have different strategies and different holdings, perhaps driven by liquidity concerns. For example, an evergreen fund might hold a larger number of smaller assets in order to make individual entries and exits less disruptive to the fund's pool of liquidity. Another concern might be that, following large fund inflows or a big exit, an evergreen manager would feel pressured to put that (liquid) capital to work quickly which could result in inferior investment selection.

Also related to liquidity, we assume similar borrowing and lending rates for the investor in drawdown funds when typically borrowing costs are 1-3% higher than lending rates. Borrowing costs would weigh on performance for the drawdown strategy when the LP has to borrow to meet capital calls and would weigh on evergreen fund performance when the fund must borrow to meet redemptions. The potential costs associated with liquidity and borrowing could also affect commitment and portfolio strategy beyond the dynamics we have simulated.

There are potentially many other reasons why the assets may differ between drawdown and evergreen funds and optimal allocations (i.e., degree of diversification) may differ even when the investment opportunity set is the same. These differences could arise because of differing portfolio constraints (such as liquidity issues discussed above) between evergreen and drawdown funds, but also because of potential agency conflicts between LPs and GPs. Differences in incentive structures could drive differences in how deals with higher return expectations or risk are allocated between evergreen and drawdown funds, or even how deal flow is managed (e.g., warehousing). Once assets are owned there could be conflicting incentives for how to manage the assets in terms of value-creation strategy, exit timing, exit type, etc. Given that evergreen funds will generate performance fees based on rolling returns with high-water marks, there is the potential for short-term misalignment between LPs and GPs, as well as effectively higher fee structures as described in Ben-David et al. (2020). In addition, because performance metrics for drawdown and evergreen funds are necessarily different, there could be internal agency issues at a GP (e.g., among partners) that are also relevant. How GPs allocate assets internally (e.g., between companion funds) could also be subject to both internal and external agency conflicts.

Finally, GPs are likely to face more pressure both internally and externally regarding asset valuation in evergreeen funds. In drawdown funds, net asset values (NAVs) are reported quarterly, but often with substantial discretion on the part of GPs. Reported NAVs have relatively few implications for investors in drawdown funds since they cannot transact at these values. In contrast, NAVs may have tremendous importance to investors in evergreen funds, because new investors will enter, and old investors can exit, at these values. This means that, for example, when a new investor enters a fund at a NAV that is too conservative (aggressive) it will advantage (disadvantage) new investors relative to existing investors. As a consequence, GPs will need to be very confident that NAVs reflect "true" fair value in order to prevent investors from gaming systematic misvaluations.

5 Conclusion

We construct stylized simulations of a limited partner investing in both a series of drawdown funds as well as an evergreen fund to see how differences in liquidity management can affect investor outcomes. We find that an investor in a series of drawdown funds experiences difficulty maintaining their preferred liquidity ratio because managers of drawdown funds are not targeting a specific liquidity ratio in the same way that an evergreen fund manager does: large or small contributions and distributions can massively disrupt an investor's liquidity ratio, even to the point of requiring an investor to borrow. While an investor can roughly maintain their preferred liquidity ratio on average, the increased variation in that liquidity ratio from year to year reduces average cumulative return and average alpha relative to the more consistent liquidity ratio that is deliberately targeted by the evergreen fund. Nonetheless, there exists a range of other factors that we identify, but do not model, which may also be drivers of differential return and risk between evergreen and drawdown funds, but are beyond the scope of this analysis.

References

Ben-David, Itzhak, Justin Birru, and Andrea Rossi (2020) “The Performance of Hedge Fund Performance Fees,” *NBER Working Papers 27454*.

Blostein, Alexander, Michael Vinci, Luke Bianculli, Anthony Corbin, Aditya Sharma, and Mohak Batra (2024) “Americas Asset Managers: State of Retail Alternatives: Recent trends continue apace as product landscape expands,” *Goldman Sachs Equity Research*.

Carmean, Zane, Hilary Wiek, Nick Cicero, and Brian Rooke (2024) “The Allocator’s Atlas: Private Wealth and the Growth of Evergreen Funds,” *Pitchbook*.

Eissler, Ralph (2024) “Comparing Evergreen and Traditional Fund Returns in Private Equity,” *Neuberger Berman Insights*.

Gabaix, Xavier, Ralph SJ Koijen, Federico Mainardi, Sangmin Oh, and Motohiro Yogo (2024) “Asset demand of US households,”Technical report, National Bureau of Economic Research.

Murphy, Daniel and Juliana Hadas (2024) “Choose Your Vehicle: A Closer Look at Private Market Fund Structures,” *Goldman Sachs Featured Insights*.

Myers, Bill (2024) “SEC committee flips narrative on retail investors in private funds,” *Private Funds CFO*.

Sikes, Stephanie and Robert Verrecchia (2012) “Capital Gains Taxes and Expected Rates of Return,” *The Acccounting Review*, 87 (3), 1067–1086.

Wood, Alisa Amarosa, Paula Campbell Roberts, and Justin Park (2024) “An Evergreen Vehicle Can Be an Important Tool in Private Equity Asset Allocation,” *KKR Wealth Insights*.

Appendix

A Simulation Details

A.1 Single Drawdown Fund Simulation

The simulation process for a single drawdown fund is as follows. We consider 12 periods (indexed by t) of fund life, where each period is thought of as one year. At the beginning of the first year ($t = 0$), an LP has \$100 in liquid capital that has just been committed but has yet to be called into the fund, and is meanwhile invested at an assumed risk-free rate with guaranteed 3% annual return. For simplicity of exposition, we refer to any capital that is not invested in a private fund as *cash*. One year later at $t = 1$, the LP starts making contributions to the fund: of the \$100 committed capital still available, somewhere between 0% and 40% of that, randomly drawn with uniform distribution, is contributed to the fund, and the remainder plus any of the risk-free returns are kept as cash and invested at the risk-free rate. The capital invested in the fund is assumed to grow based on a market return, distributed normally with a mean of 12% and 18% standard deviation; plus a private market alpha, also distributed normally but with a mean of 3% and 5% standard deviation. When $t = 2$ arrives, the LP contributes (randomly and with uniform distribution) somewhere between 20% and 40% of remaining committed capital, again keeping the remainder of cash invested at the risk-free rate. The same pattern follows at $t = 3, 4, 5$, with 20% to 60% of remaining committed capital invested at $t = 3$, 40% to 80% at $t = 4$, and any remaining committed capital at $t = 5$.

Distributions begin in period $t = 3$. It is assumed that any distributions are immediately re-invested either to fund a capital call or as cash. We take distributions in a given period t to be a random and uniformly distributed proportion of that period's net asset value (NAV), gradually scaling up from 0-20% in $t = 3$ to 100% of NAV in period $t = 12$ to ensure complete liquidation of the fund; the full distribution schedule is shown in Table A1.

To briefly summarize the dynamics of the simulation, let rm_t denote the market rate in period t, let α_t denote the private market alpha in period t, and let rf denote the (time invariant) risk-free rate. Then the allocation variables of interest can be written as

$$\text{NAV}_t = (1 + rm_t + \alpha_t)\text{NAV}_{t-1} + \text{Contributions}_t - \text{Distributions}_t, \tag{1}$$

$$\text{Cash}_t = (1 + rf)\text{Cash}_{t-1} + \text{Distributions}_t - \text{Contributions}_t, \tag{2}$$

$$\text{NAV}_0 = 0, \tag{3}$$

$$\text{Cash}_0 = 100, \tag{4}$$

which illustrate that cash outflows become private fund inflows, and vice versa. This process is repeated 10,000 times to create the set of simulated drawdown funds, each with a unique set of draws for market returns, private fund alphas, and schedules for contributions and distributions.

TABLE A1: Distribution Schedule

Period	0	1	2	3	4	5	6	7	8	9	10	11	12
Lower Bound	0%	0%	0%	0%	10%	15%	20%	20%	25%	30%	40%	60%	100%
Upper Bound	0%	0%	0%	20%	30%	35%	40%	40%	45%	50%	60%	100%	100%

Distributions percentages are taken as a percentage of net asset value.

A.2 Evergreen Fund Simulation

The evergreen fund simulation process shares the same risk-free and market return structure, as well as the private market alpha structure, as that of the single drawdown fund process. Unlike the single drawdown fund, however, we assume that an LP immediately enters the evergreen fund with a \$100 contribution at $t = 0$; there is no delay between commitment and capital call as there is with the drawdown fund. Also unlike with the drawdown fund, the LP does not hold on to any of the initial \$100 nor any of its return as cash; all capital remains in the fund, and the fund manager balances the fund cash to private investment ratio within its portfolio. In other words, the LP outsources cash management to the evergreen fund's GP.

In terms of dynamics, the evergreen fund is modeled as a portfolio for which the GP strictly maintains 10% of portfolio value in cash and the other 90% in private funds. Since the fund managers allocate both cash and private holdings, the NAV is inclusive of both and is equivalent to the portfolio value. We can summarize the variables of interest at each period t as

$$\text{NAV}_t = (1 + rf)\text{Fund Cash}_{t-1} + (1 + rm_t + \alpha_t)\text{Fund Holdings}_{t-1}, \tag{5}$$

$$\text{Fund Cash}_t := 0.10 \times \text{NAV}_t, \tag{6}$$

$$\text{Fund Holdings}_t := 0.90 \times \text{NAV}_t, \tag{7}$$

$$\text{NAV}_0 = 100. \tag{8}$$

Equations (6) and (7) illustrate that the GP is actively re-balancing 10% of the total portfolio value into cash and 90% into the private funds at the beginning of each period. Of course, in reality a GP will not perfectly manage a 90/10 balance at all times, but this idealized assumption will nonetheless shed light on the consequences of an evergreen GP closely managing such a balance, at least relative to the alternative. We again repeat this process 10,000 times to generate the set of simulated evergreen funds.

A.3 Series of Drawdown Funds Simulation

Modeling the series of drawdown funds is relatively complex compared to the preceding simulations. On one hand, we simply repeat overlapping and staggered single drawdown funds, largely as described above in Subsection A.1. The added complication is that we want the series to reflect certain properties of the evergreen fund investment in order to align comparisons: an investor into the portfolio of a series of drawdown funds should start with a NAV of $90 along with an average cash ratio of 10% in the first three periods before making commitments to any new funds, and furthermore the 10% cash to 90% NAV balance should on average be maintained over the investor's involvement.

We start by seeding the series of funds with an initial investment and introduce a new fund every 3 periods. The size of subsequent commitments are, on average, calibrated so that the 10/90 ratio is maintained on average over the life of the fund, and so adjusted dynamically based on performance up to the point of each new commitment. For example, suppose that an investor wants to make a $100 commitment at period $t = 3$ when returns to date are exactly as expected (i.e. exactly 15% private fund returns in all periods upon reaching $t = 3$); if the investor instead only experienced 80% of that expected return at the beginning of $t = 3$, then the investor will instead commit $80 to the new private fund in period 3. This dynamic adjustment helps prevent private fund commitments from becoming disproportionately large or small relative to portfolio value, and it is also realistic insofar as LPs are likely to adjust commitment size relative to their historical performance. To that end, the average percent of cash held during the salient periods is 10.02% in the 10,000 simulated series (a more detailed illustration of cash/fund balance is discussed below).

The single drawdown fund simulation uses a single commitment of $100. For the series of drawdown funds, on the other hand, we introduce a new fund every 3 periods and need to adjust the size of the commitment dynamically as follows. A hypothetical investor, starting with $100 cash, makes a commitment to a new drawdown fund in periods $t = -6, -3, 0, 3, 6$, and 9, in respective base amounts of $99, $149, $224, $337, $507, and $763 (an increase of roughly 14.6% annually). These numbers might seem ad hoc at first, but were chosen because they lead to an average cash ratio of 10.02% across periods $t = 0, \ldots, 12$, that is, during the periods in which the investor is active and 3+ staggered, overlapping funds are active.

Upon reaching period $t = 0$, the entire fund is normalized so that the fund NAV is $90, matching the private fund value of the evergreen fund; it is at this point that the investor under analysis becomes involved. To at first roughly match the evergreen cash balance of 10%, qualifying simulations must have a cash balance averaging 8% to 12% in periods $t = 0, 1, 2$, before a commitment to a new fund is made; this balance cannot be made exactly because of the randomness involved in distributions and capital calls that the LP cannot control (which is a feature of the simulation, and perhaps a bug of investing in a staggered series of funds as discussed in the main text).

From periods $t = 3, \ldots, 12$, the LP adjusts commitments based on total return-to-date, but otherwise there are no restrictions on the actual ratio of cash to NAV in any given simulation; that said, the average balance across simulations is still 90% invested and 10% in cash.

Due to a handful of extreme outliers, we throw away and re-simulate any simulations that exhibit an alpha, beta, or portfolio value more extreme than the 0.5% largest and smallest observed in the unfiltered simulation. Numerically this means discarding simulations with alpha beyond the interval $(0.00\%, 7.56\%)$, beta beyond the interval $(0.52, 1.65)$, and portfolio value beyond $(87.24, 1147.37)$. As shown for alphas and betas in Figure 4, these values are indeed extreme relative to the mass.