Unpacking Private Equity Performance *

Prof. Gregory Brown
Research Director
gregwbrown@unc.edu

William Volckmann
Research Associate
wmvolckmann@unc.edu

September 26, 2023

Abstract

Performance analysis of private equity funds is challenging because fund ownership does not trade in a liquid market with observable prices. Instead performance analysis, especially during a funds life, must rely on observed cash flows to and from the fund and quarterly net asset value (NAV) estimates. Further complicating the analysis are the increasingly common practices of funds using subscription lines of credit (fund-level debt) and recycling capital. Even the variation in the timing of capital deployment across funds has important implications for common performance measures used to evaluate funds such as internal rate of return (IRR) and multiple on invested capital (MOIC). In this paper, we analyze a set of simulated funds to better understand how fund performance analysis is affected by these common issues. Overall, our analysis suggests that intermediate IRRs – i.e., values likely observed during fundraising periods for subsequent funds – are strongly affected by subscription lines and deployment pacing. Intermediate MOICs are only weakly affected by subscription lines, but strongly affected by capital deployment pacing. Both IRRs and MOICs are strongly affected by recycle deal accounting methodology. Consequently, LPs need to be very cognizant of these issues when measuring and utilizing fund performance measures during the life of a fund as well as when assessing ultimate performance at the end of fund life.

1 Introduction

Private equity (PE) fund limited partners (LPs) are constantly faced with the prospect of evaluating the performance of their existing and prospective general partners (GPs) that manage the fund assets. A standard part of any evaluation process is examining the performance of prior

*Comments are welcome. Please contact Professor Gregory Brown, Sarah Graham Kenan Distinguished Professor of Finance, CB 3440, UNC at Chapel Hill Kenan-Flagler Business School, Chapel Hill, NC, USA; +1-919-962-9250, gregwbrown@unc.edu. We thank Barry Griffiths, Pierre-Yves Mathonet, David Robinson, and Avi Turetsky for helpful comments.

funds managed by the GP. Given the relatively long life of PE funds (typically in excess of 10 years) and the relatively short PE fundraising cycle (roughly every 2-4 years), LPs will typically seek to evaluate funds that are still far from fully realized. In these cases, the performance evaluations will depend critically on cash flows to and from the fund between the fund inception date and the performance evaluation date, as well as the estimates of net asset values (NAVs) reported by the GP to the LPs.[1]

Yet, there are many factors that can determine the timing and magnitude of cash flows besides simply the investments in the underlying portfolio companies. For example, the GP will charge fees to the LPs. GPs also increasingly use subscription lines of credit (fund-level debt) to delay or smooth capital calls from the LPs.[2] Many funds are allowed to reinvest profits obtained from an early exit into a new deal, commonly referred to as "recycling," and the cash flows can be treated in various ways when calculating fund performance. But even very basic issues such as how quickly a GP deploys capital in a fund (i.e., deal pacing) can have potential impacts on interim fund performance. In this paper, we look in detail at how these three features of PE funds (recycling, deal pacing, and subscription lines) affect the two most common performance measures: the multiple of invested capital (MOIC) and the internal rate of return (IRR).[3] We show that each of these features can have large impacts on fund performance measures – especially early in a fund's life when it is likely to be fundraising – and provide some recommendations to LPs to assist them in unpacking performance metrics to better understand how various factors can influence return analysis.[4] Our goal is not to pass judgment on any particular strategy or reporting technique, but simply to improve understanding and transparency.

PE funds typically have five years to deploy capital, and in practice there is a wide range of observed deal activity over the investment period. Some funds have been observed to deploy all of their capital within the first two years of the investment period and others undertake more even pacing, or even back-load deployment toward the end of the investment period. Many factors

[1] For a discussion of NAV reporting in private equity see Brown and Nykyforovych-Borysoff (2023).

[2] For example, Albertus and Denes (2020) document that in a historical sample of funds in the Burgiss Manager Universe, those funds using subscription lines have 1.9% higher IRRs and increase their probability of raising a next fund by 20.6%.

[3] We do not consider public market equivalent measures (PMEs) such as those described in Kaplan and Schoar (2005) because in our deterministic setting the MOICs and PMEs provide identical insights.

[4] For a detailed analysis of the relation between performance measures and subsequent fundraising see Barber and Yasuda (2017) and the citations therein.

can determine pacing including deal- and industry-specific factors as well as broader economic conditions and the PE fundraising environment. In some cases, GPs are known to have their own preferred strategy for deployment pacing. Regardless, the pace at which committed capital is deployed by a fund is important for at least two reasons. First, depending on the fee structure, it could be beneficial to LPs for the fund to front-load deal activity because then fees are lower as a proportion of deployed capital earlier in the fund's life. This derives from the common practice of fees typically being fixed as a percentage of total committed capital (not deployed capital) early in a fund's life. A concern for LPs would then be that pressure to "put capital to work" could potentially lead to the selection of sub-par investments (Mendoza, 2019). Second, and more relevant to our analysis, the pacing of deployments could affect both intermediate and final performance metrics of a fund. The private equity industry is sometimes criticized for embellishing intermediate performance (Kazemi, 2021), and deployment pacing may be one method of doing so.[5]

Subscription lines are fund-level loans that allow fund managers to delay capital calls to investors. Historically, subscription lines were used to smooth capital calls (e.g., combine calls within a month or quarter) and provide flexibility to act quickly on a deal if necessary, for example, within the time window provided for LPs to meet a capital call. Subscription lines were traditionally limited to 90 days or less, but over the last decade many funds have lengthened their allowed loan terms, sometimes for up to 3 years (Prequin, 2019). Subscription lines potentially benefit LPs because investment opportunities can be taken advantage of while committed capital has not yet been called; but they also impose a potential cost on investors, namely the interest paid for subscription line loans that enters net profit calculations. In particular, if the fund's cost of borrowing is greater than the LP's cost of borrowing, there will be an economic loss to the LP from this unnecessarily high interest rate.[6] Subscription lines also affect time-sensitive performance metrics, like a fund's IRR, because the time between capital calls and distributions is compressed. We show that even fairly small delays in capital calls can have meaningful effects on IRRs early in a fund's life. Large delays in capital calls stemming from use of subscription lines can have huge near-term effects and moderate long-term effects on IRRs.

[5]Some evidence suggests GPs may also strategically impact the estimates of NAVs reported to LPs durign fundraising. See Brown et al. (2019) and citations therein.

[6]There are additional portfolio considerations and we discuss these later.

In a recycle deal, some of a fund's realized gains and capital invested are re-invested in an additional deal rather than returned to investors. From a performance measurement standpoint, there is no standard practice for how to consider the exit deal cash flows that are then recycled. For example, when reporting performance, some GPs count the recycled capital as additional called capital and some don't – and this simple difference can have meaningful consequences. More specifically, there are two common ways of accounting for recycled capital. At an intuitive level, the first way considers the realized gains (net of carry and other expenses) to have been distributed back to investors and then immediately called back into the fund as needed for the recycle deal (though the capital does not actually leave and get returned to the fund), thereby constituting a new investment from a performance analysis standpoint. The second way considers realized gains to have not left the fund before the recycled portion is reinvested, thereby effectively constituting a continuation of the original investment (i.e., no increase in contributed capital). The implications on performance measures can be large. This has caused concerns among some and accordingly has been reported to the Securities and Exchange Commission. Some analysts have developed novel performance metrics to overcome issues involving measurement in the face of complex management strategies; for example, the MaxMult metric calculates the highest multiple possible for *any* set of recycling rules, thereby conforming the effect of recycling into a single comparable measurement (Griffiths, 2016). At a minimum, without transparency in accounting methodology, the validity of cross-fund comparisons (whether they are "apples to apples") is thrown into question (Garcia, 2022).

Our analysis shows that intermediate and final IRR are strongly affected by subscription line maturity, and intermediate IRR is significantly affected by deployment pacing. Intermediate and final MOIC are strongly affected by recycle deal accounting methodology, and intermediate MOIC is strongly affected by deployment pacing and subscription line maturity. Consequently, LPs need to be very cognizant of these issues when measuring and utilizing fund performance measures.

2 A Simulated Private Equity Fund

All of the analysis in this paper examines a stylized model of a private fund so that we can precisely isolate the effects we seek to examine. Specifically, the simulated funds in this paper share

the following assumptions. We consider a fund that has $100 in committed capital. The fund can take actions (capital calls, investments, and distributions) every 6 months.[7] All fund investments in portfolio companies have a gross multiple on invested capital (MOIC) of 3.0, although the holding period of each investment will vary in some of our analysis. The model is entirely deterministic, so the growth rate and holding period of each deployment are known with certainty. Before year 5.5 (i.e., the end of the investment period), the GP charges a flat fee of 2% of committed capital, although the start of actual fee collection from LPs could be delayed with a subscription line. After the investment period has ended, the fee is multiplied by the end of period cost basis (i.e., the sum of capital actively in deployment as a proportion of committed capital).[8] In the final period of the fund, all deployed, but as-of-yet uncalled, capital is called (including fees and interest payments).

The GP earns carry as 20 percent of the realization net of the investment amount and pay down of fees and interest payments (net LP profit is then the other 80 percent).[9] Likewise, the net LP (investor) MOIC is net distributions (i.e., net profit plus returned capital) as a proportion of capital called, and net LP (investor) IRR is calculated using the net cash flows actually experienced by the LP.

We vary fund parameters along the following three dimensions (more details are discussed below). First, we vary the pace of deployment, i.e., how rapidly funds are invested, to see how performance metrics differ over time and at exit (if at all). Second, we vary subscription line maturities to see how the compressed time frame between investor outflows and inflows, as well as the accumulation of interest payments on subscription line loans, affect performance metrics. Because these factors can have a large effect on performance metrics during the life of the fund, we especially focus on effects during the common fundraising period when fund age is 2-5 years. Third, we vary the size of recycle deals and use two accounting methods to see how simple accounting differences lead to differences in outcomes and reported performance metrics.

[7] Looking at quarterly or monthly intervals does not have any meaningful impact on the analysis and so we simplify the time line into 6-month increments.

[8] To be precise, we consider a separate management fee and partnership fee which is often done in practice. During the investment period the management fee is 1.5% of committed capital and the partnership fee is 0.5% of committed capital and collected every 6 months. After the investment period the partnership fee is halved and both fees are charged only on deployed capital.

[9] Adding a hurdle rate for earning carry does not change the qualitative analysis and conclusions.

2.1 Deployment Pacing

In this section we consider how different pacing of capital deployment affects performance measures. We assume that the fund will invest \$85 of its \$100 committed capital in portfolio companies, reserving the remaining \$15 for fees.[10] We consider three alternative paces of capital deployment: *fast*, *even*, and *slow*, as detailed in Table 1. For simplicity, we assume that each 6-month period is a single deal and all deals are held for 3.5 years with a gross MOIC of 3.0.

TABLE 1: Paces of Deployment

Year	1	1.5	2	2.5	3	3.5	4	4.5	5	5.5
Fast	\$30	\$20	\$10	\$10	\$2.5	\$2.5	\$2.5	\$2.5	\$2.5	\$2.5
Even	\$8.5	\$8.5	\$8.5	\$8.5	\$8.5	\$8.5	\$8.5	\$8.5	\$8.5	\$8.5
Slow	\$2.5	\$2.5	\$2.5	\$2.5	\$2.5	\$2.5	\$10	\$10	\$20	\$30

These *fast*, *even*, and *slow* pacings are roughly consistent with the paid-in capital data collected by Burgiss for a large sample of private equity funds.[11] All three pacing schemes generate the same gross realization of $\$85 \times 3 = \$255$, but the pacing differences will affect both IRRs and MOICs due to discounting differences as well as post-investment period fee calculations.

The left panel of Figure 1 shows that deployment pacing can have a large effect on IRRs especially early in a fund's life when it is likely to be fundraising (shaded area). The fund with *fast* deployment (red line) has the highest net IRR throughout as well as the only IRR that is always positive. The fund with *even* deployment (blue line) generates a negative IRR in the first six months, but the IRR turns positive in the second year and steadily approaches (but never matches) the IRR of the *fast* deployment. Finally, the fund with *slow* deployment (green line) generates a negative IRR through the fund's third year and only slowly approaches its ultimate IRR of about 23%, the lowest of the three.

[10] We obtain nearly identical qualitative results if we assume a total of \$80 or \$90 is invested instead of \$85 or even if we assume that \$100 is deployed, i.e., that fees are accessed on top of full deployment of committed capital. This is just a scaling issue that washes out for IRRs and MOICs.

[11] Specifically, we calculate paid-in capital as of the end of the first year for all equity funds in the Burgiss Manager Universe. We then sort funds into deciles and examine the cash flow patterns of funds at the fastest-to-deploy decile, the median fund, and slowest-to-deploy decile. Using these data we generate smoothed deployment schedules shown in Table 1 that have similar cumulative paid-in capital values.

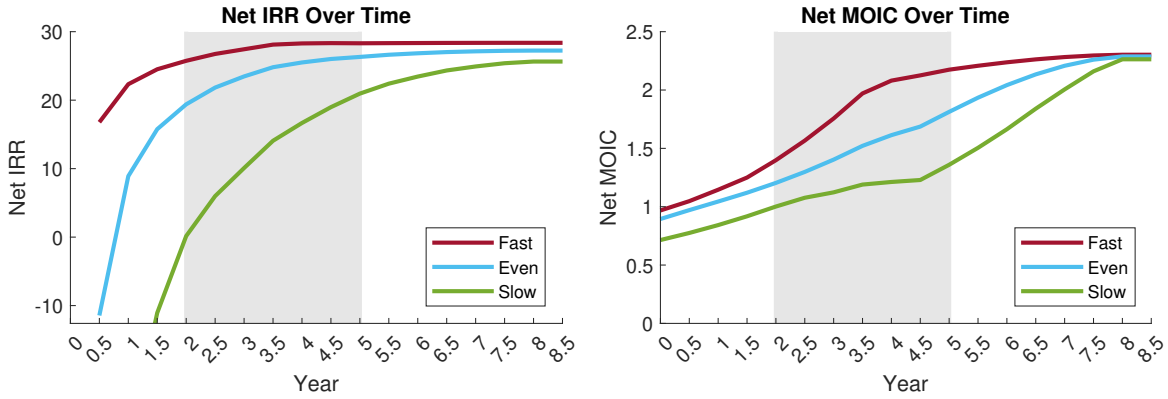

FIGURE 1: Fast pacing has a permanently higher net IRR. The initial net IRR for slow pacing is -70 percent but was cut off so as do not dominate the plot. The effect on net MOIC is less pronounced. The shaded areas indicate fundraising periods.

To understand the intuition for this result, consider the initial \$30 investment by the fund with fast deployment. Having an annual accretion rate of about 37 percent (by virtue of a gross MOIC of 3 after 3.5 years) generates a contribution to the net asset value (NAV) of \$35.10 when the fund is 6-months old. This exceeds the initial investment plus the \$2 in fees accrued through then. Now consider the fund with even deployment: after 6 months the initial investment has a NAV of just $\$9.95 (= \$8.5 \times 1.17)$, which is less than the initial investment plus \$2 in fees, thus a negative IRR at that date. Similarly, the fund with slow deployment generates an even more negative IRR. In effect, early IRR is higher with fast-paced deployment because the fees are smaller in proportion to the deployments.

We now turn to understanding why the net IRRs do not converge to the same values at the end of the funds' lives. This result derives from management fees after the investment period being proportional to fund asset cost basis. With fast pacing, fund asset value is relatively low after year 5 since the majority of invested capital is quickly realized, thus lower fees. In contrast, the fund with slow pacing will have an asset value that is relatively high after year 5 since the majority of invested capital isn't realized until well after the investment period has ended. This then generates relatively higher fees. Of course the fund with even pacing is always between the fast and slow paced funds. To put it more succinctly, final net IRR is higher with fast-paced deployment because the fees are smaller in absolute terms. With the set of deployments considered in this paper, the

final difference in net IRR between fast and slow amounts to around 2.7 percentage points.[12]

When we consider fund MOICs, similar logic applies to explain the ordering of performance. Fast deployment always outpaces slow deployment as shown previously in Figure 1. The final difference in net MOICs between fast and slow deployment rates is just 0.04 of a multiple, but the difference can be stark midway through the fund and actually reaches a maximum dispersion during the likely fundraising period (shaded area). The increase in the MOIC denominator (cumulative contributions) of the fast-paced fund slows down as subsequent deployments become smaller, which leads to a large intermediate difference between NAV (numerator) and contributions. The same is not true of slower-paced funds whose denominators grow at an even or increasing pace. But by year 5.5, nearly all capital calls have been made for any pace of deployment, so thereafter the larger late realizations of the slower-paced funds close the NAV gap. To better understand the actual differences in cash flows, investor profit and investor fees are shown in Figure 2. Cumulative investor profit is about $1.50 lower for the fund with slow deployment and this is driven entirely by the difference in fees.

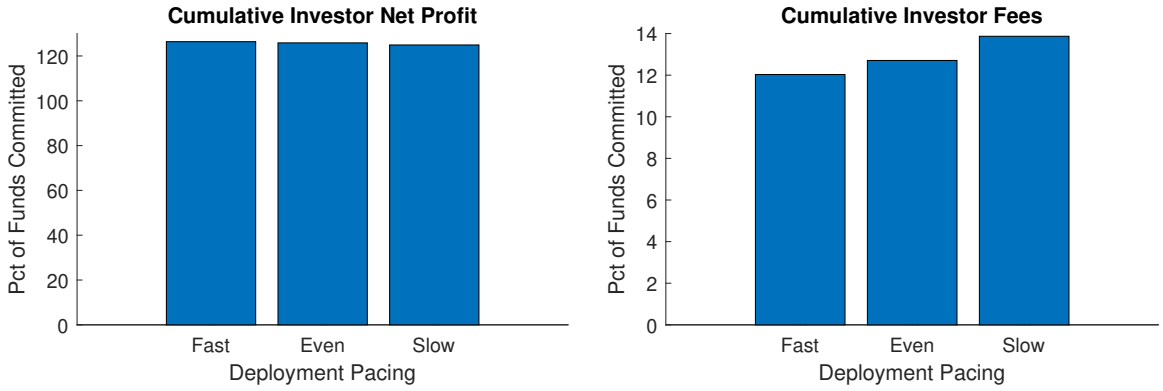

FIGURE 2: The difference in investor profit is driven by difference in fees, but the difference is very small in context.

The key take-aways of this section are that differences in capital deployment pacing can have i) large effects on IRRs especially early in a fund's life and ii) large effects on MOICs especially in the middle years of a fund, and iii) meaningful effects on final IRRs but not MOICs.

[12]One might also be concerned that the interaction of the slower pacing and the discounting process means that investments have less of an ability to make up for the earlier fee impact, however this is a negligible effect.

2.2 Subscription Lines

In this section, we consider the effects on IRR and MOICs when a fund delays capital calls through the use of subscription lines of credit. Use of subscription lines has become increasingly common with the median fund effectively delaying capital calls by 3 months and the upper decile of funds delaying by close to 1 year. While often considered a strategy utilized by funds just during the invesment period, recent data provided by Burgiss indicates that a quarter of funds have used subscription lines past the fifth year of fund life.[13]

We conduct our analysis of subscription lines by examining a fund that makes even capital deployments throughout its investment period (i.e., the deployment sequence is $8.5 semi-annually in through year 5). All investments generate a gross MOIC of 3.0, but the initial $8.5 investment is held for 3.5 years whereas other deals are held for 5 years. We consider subscription lines with 0.5 years, 1 year, and 1.5 years, where each line has an annual interest rate of 5 percent.[14]

Because delaying a capital call shortens the time span between investor outflows and inflows, a subscription line should increase the net IRR (whenever it is positive and the interest rate is sufficiently low compared to the investment return). We also expect that longer delays lead to larger effects but that the effects decrease with fund age. Figure 3 shows the exact effects on IRRs for the four scenarios we consider (including no usage of subscription lines). As expected, when the subscription line delays capital calls by 6 months (red line), the fund IRR is consistently higher and the difference declines as the fund ages to about 2%. As the delay increases to one year (blue line), the effect becomes more pronounced especially early in the fund's life. In the extreme case of 1.5 year delay, the IRR early in a fund's life spikes to over 100% but rapidly declines to around 35% by year 5. We emphasize the very large possible effects on IRRs of subscription lines during the period GPs are likely to be fundraising for their next fund (shaded area). The righthand graph in Figure 3 plots the final difference in IRRs as a function of the term of the subscription line borrowing. The increase in delay is associated with a growing difference in the final net IRR – with a change in IRR from no delay to a 1.5 year delay of about 8 percentage points.

[13] See, for example, Warren (2023).

[14] Of course, higher interest rates, such as those experienced more recently, will provide a drag on returns which will offset some of the subsequently documented boost in interim performance.

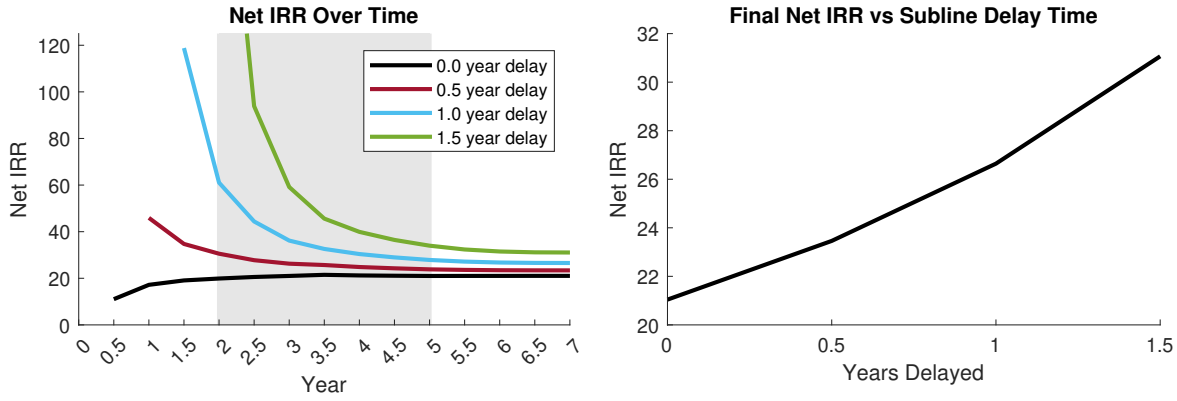

FIGURE 3: Large delays lead to very large increases in IRR. The 1.5-year delay has an initial IRR of about 148 percent, but is cut off so as to not dominate the vertical axis. The shaded area indicates fundraising periods.

The net MOICs paint a very different picture of the effect of subscription lines. This is because MOICs are not affected by the compression of time between capital calls and investment growth. Instead, the MOIC is mostly affected by the interest payments that accompany subscription lines which ultimately increases the amount of capital called for fees but without any affect on the value of investment realizations. Nonetheless, the net MOIC with subscription lines will sometimes be higher early on in the investment cycle before the full brunt of capital calls can manifest. This effect is illustrated in Figure 4 where the MOICs are typically slightly higher with subscription lines up until about year 3 of fund life. When we focus in on the fundraising period (shaded area) we see that the effects of subscription lines on MOICs can be positive or negative but are generally small except for early on for the most extreme 1.5 year delay. Yet while the difference in net MOIC between no use of a subscription line and a 0.5 year delay is always small, the difference in final net MOIC between no delay and a 1.5 year delay is almost 0.15 of a multiple.

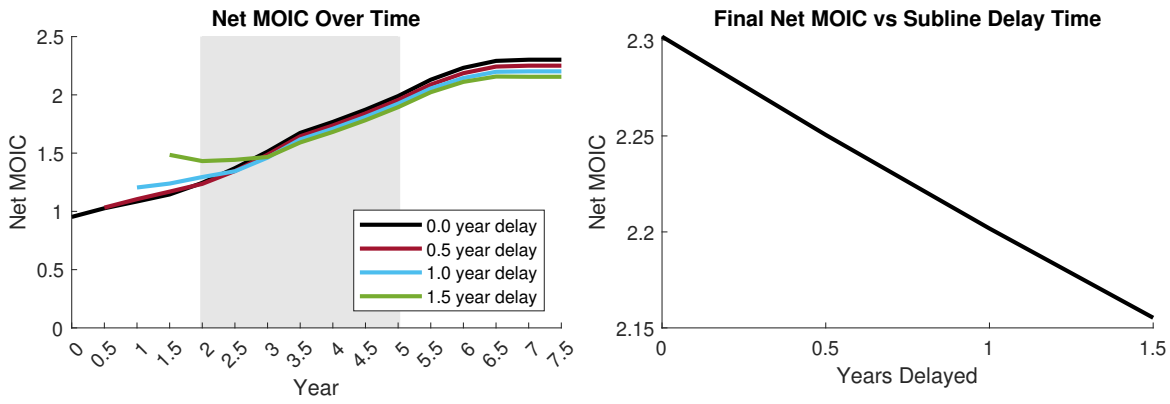

FIGURE 4: Subscription lines can have a positive effect on net MOIC early on, but ultimately a negative effect due to interest payments. The shaded area indicates fundraising periods.

The differences in investor net profit and interest payments are shown in Figure 5. As predicted by the net MOIC but not the IRR, investor net profit is lower with subscription lines on account of interest payments. That said, it could be that investors make productive use of delayed capital which, provided the return is high enough relative to the subscription loan interest rate, could overcome the differences in investor net profit when also taken into consideration. On the other hand, in some jurisdictions LPs are required to hold a cash reserve for uncalled commitments. Consequently, the overall economic cost or benefit from subscription lines will be specific to each investor.

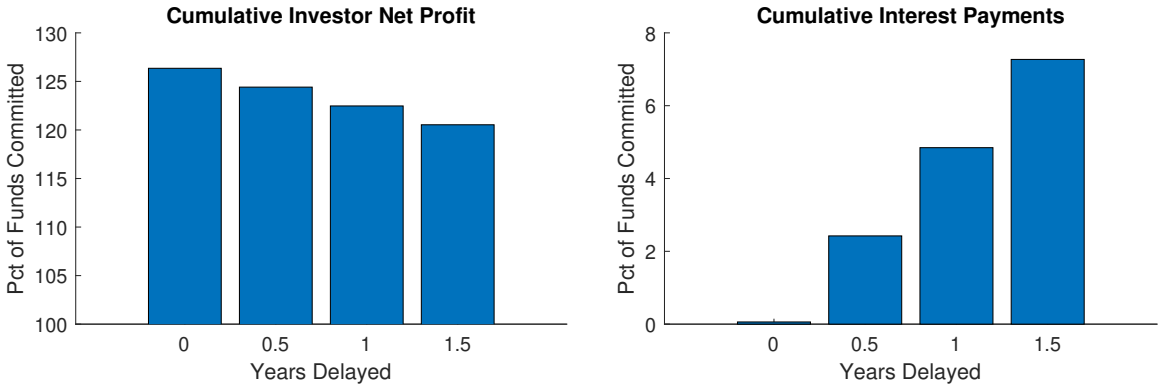

FIGURE 5: Investor profit is lower with subscription lines because of the interest payments.

2.3 Recycling of Capital

Recycling has become a very common feature of PE funds. A recent survey of 50 large GPs found that every fund recently raised allowed for some type of recycling of capital.[15] In this section we consider how the amount of recycling, as well as how recycling deals are accounted for, affects fund performance measurement. In each scenario the same non-recycle deployments are made with the same pacing, specifically an initial investment (at year 0) of $20 is made with a holding period of 3.5 years. This is followed by three subsequent investments of $20 in years 0.5, 1, and 1.5 with holding periods of 5 years each; and a final investment of $5 in year 3 with a holding period of 5 years.[16]

The investment of $20 made in year 0.5 becomes a gross realization of $60 in period 5.5, the final period in which investments can be made. We consider four different scenarios in which this realization will be recycled in amounts as suggested by Burgiss paid-in capital data—$45 (max), $30 (high), $15 (low) and $0 (none)—and compare the outcomes.[17] Furthermore, we calculate outcomes by using two different accounting methods—Method A and Method B—to see how otherwise identical investments might produce different performance metrics.

With Method A, the recycle deal is considered a new investment, fundamentally no different than any of the other deployments. To that end, the new recycle investment requires a new capital call: it is as if the realization of $60 had been distributed back to investors (net fees and carry), and a portion of it is called back into the fund. With Method B on the other hand, the recycle deal is not considered a new investment: it is as if the realization of $60 never left the fund and a portion of it was immediately funnelled back into an investment; that portion is not considered part of the gross realization because only the remainder (net fees and carry) was distributed back to investors. Because the recycled capital never left the fund, the recycled capital does not need to be called back into the fund.

For example, consider a simplified scenario with no fees where an investment of $20 generates

[15] See, 2022 Private Equity Fundraising: Key Trends and Market Survey by Paul, Weiss, Rifkind, Wharton, & Garrison.

[16] While front-loaded investments tend to increase the opportunity for recycling capital, we also examine an alternative with even deployment pacing and common holding periods. Results are presented in Appendix A.

[17] These values are derived as follows: In the Burgiss paid-in capital data we examined final paid-in capital as a percent of committed capital. The max of $45 was roughly the top 1% breakpoint of funds in excess of 100%. Likewise, $30 (high) and $15 (low) represent presumed recycling for roughly the 75th and 25th percentiles of funds with paid-in capital in excess of committed capital.

a gross realization of \$60, and \$45 of that realization is recycled. With both Method A and Method B, \$65 is deployed. With Method A, total gross realization is $\$20 \times 3 + \$45 \times 3 = \$195$, whereas Method B has total gross realization of $(\$20 \times 3 - \$45) + \$45 \times 3 = \$150$. The most salient difference between Method A and B is fact that Method B has \$45 less capital called and \$45 less gross realization.[18]

Using the full model with a \$45 recycle, the difference in methods is illustrated in Figure 6. The capital call in year 5.5 for only Method A stands out starkly. Comparatively, the higher gross realization of Method A in year 5.5 is less stark. The cumulative difference over time is shown in Figure 7. Notice that the cumulative difference in capital called widens slightly over time since Method B has fewer fees (by around \$2-\$4 depending on recycle size) due to having fewer new investments (in an accounting sense). Also note that total capital called with Method A is well in excess of the total committed capital, but the same cannot be said of Method B.

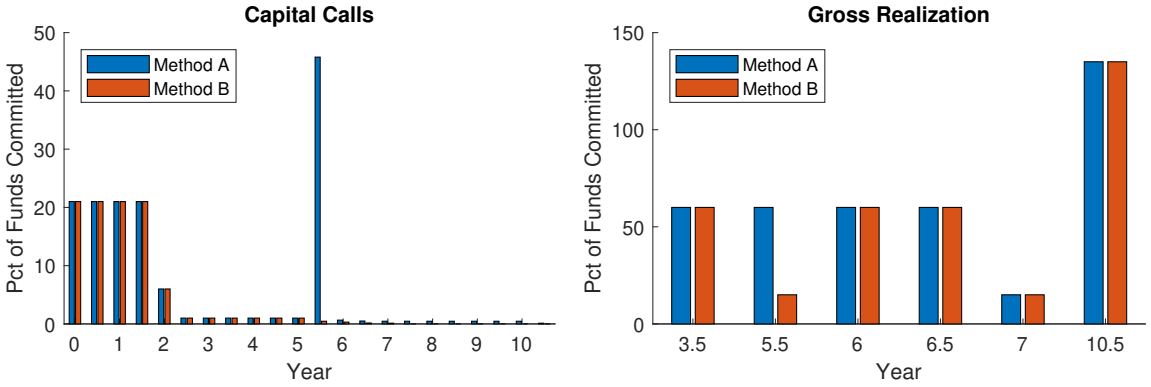

FIGURE 6: The most glaring difference between Method A and Method B is the capital call for the recycle deal in year 5.5.

[18]It is worth noting that recycling provisions can vary across funds and thus, in practice, some types of recycling are feasible and others are not. For example, when we describe recycling \$45 of a \$60 distribution, some funds would be allowed to do this while others would not. Specifically, some sub-documents allow the recycling of all distributions as long as the amount of distributions is less than capital called to date. Thus in this case, the fund could recycle \$45 as long as it had called more than \$45 by this point regardless of where the \$45 was deployed. However, other sub-documents only allow funds to recycle the cost basis of the specific exited investment, which would be \$20 in this case. We abstract away from these real-world constraints in our analysis.

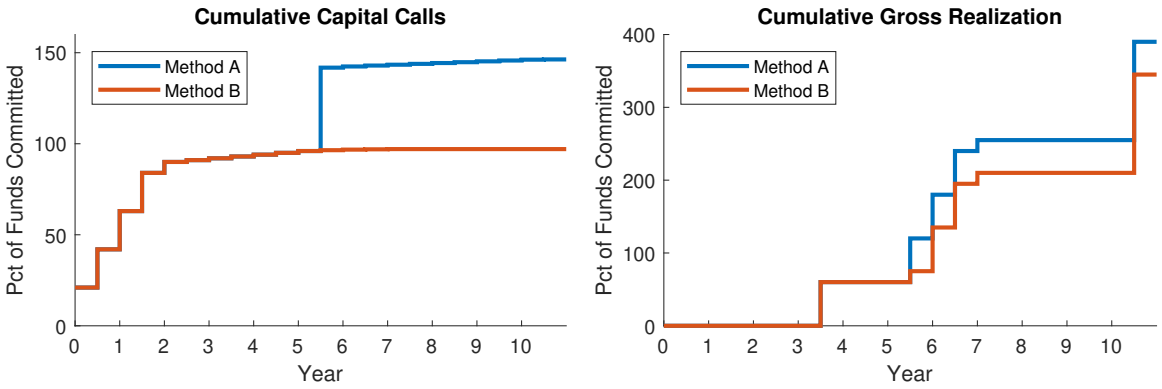

FIGURE 7: The difference in capital calls is larger in proportional terms than the difference in gross realizations.

Because the differences in fees are relatively small, we continue to focus on capital calls and realizations (or distributions) in discussing two performance metrics: the net IRR and the net MOIC. Again refer to Figure 7: for Method B, note that investor inflows (distributions) are lower by a small proportion relative to investor outflows (capital calls). Accordingly, one would expect net IRR to be higher with Method B. Likewise in considering net MOIC, the numerator (distributions) are lower by a small percentage relative to the denominator (capital calls). One would expect a higher net MOIC with Method B as well.

Figures 8 and 9 confirm this intuition, and in fact the effect is more pronounced for larger recycle deals.[19] The largest difference is only around 1 percentage point for net IRR, but the difference in net MOIC is substantial with a maximum difference of about 1 full multiple. In fact, the max recycle with Method B gives a net MOIC above 3, even though each individual deployment is assumed to have a gross MOIC of only 3. It is, of course, counter-intuitive that the final net MOIC could be larger than the MOIC of any individual investment made within the fund, especially since this MOIC is calculated net fees; but the accounting of Method B effectively allows for two distinct deployments—and therefore two distinct realizations—based on a single capital call.

[19] Again, note that these differences are not primarily driven by fees: even with zero fees, the differences have the same basic pattern and practically the same magnitudes.

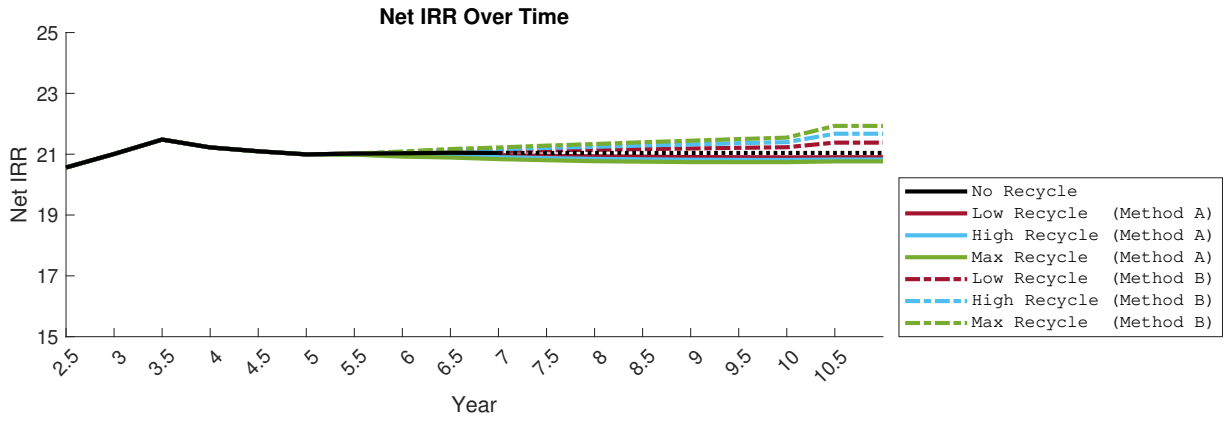

FIGURE 8: Method B generates higher net IRR, but the magnitude only spans about 0.5-1.0 percentage points.

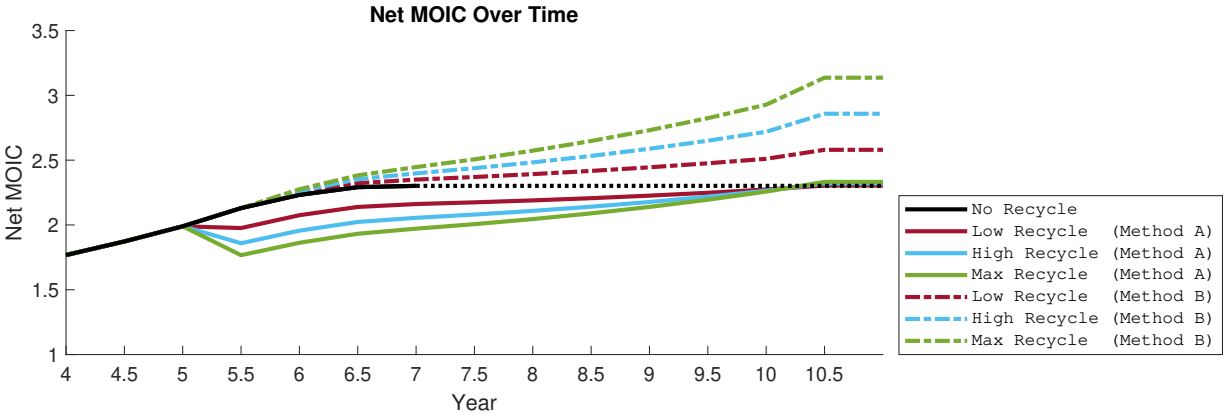

FIGURE 9: Method B generates higher net MOIC, and the difference is economically significant.

It is also worth noting that measuring investor profit as simply a percent of capital committed is higher with Method A, as shown in the left panel Figure 10. To see why, again consider the simple example with \$45 recycle deal and no fees. With Method A, gross profit is $\$195 - \$65 = \$130$, of which \$104 (80 percent) is distributed as net profit back to investors. With Method B, gross profit is $\$150 - \$65 = \$85$, of which \$68 (80 percent) is distributed as net profit back to investors. Note that net MOIC, defined as net distribution (i.e. net profit plus returned capital) as a proportion of called capital, with Method A is $(104 + 20 + 45)/(20 + 45) = 2.6$, whereas net MOIC with Method B is $(68 + 20)/(20) = 4.4$. Simply put, Method A puts more of the investors' committed capital to work, which results in more "profit" for the investors. So in this sense, the additional net profit of Method A more than offsets the additional fees of Method A.

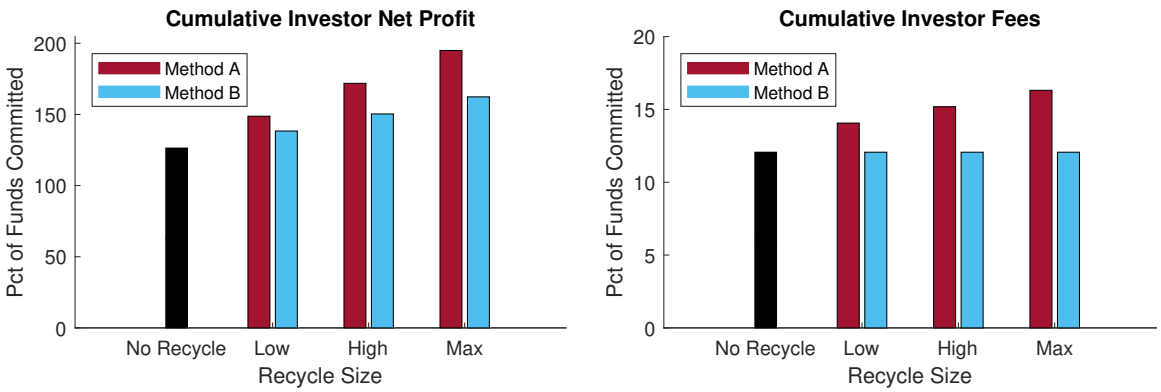

FIGURE 10: Method A gives higher investor profit in excess of extra fees.

However, the actual economics of dollars earned by the investors tells a different story. In both methods A and B, the same investments are made and the same returns are achieved at the portfolio company level, so it is only the way that fees and profits are being calculated that differ.

3 Conclusion

We make the following general, rule-of-thumb conclusions about the sensitivity of performance metrics when considering each of the preceding dimensions in isolation as summarized in Table 2.

TABLE 2: Rule-of-Thumb Stability Chart

	Recycling	Pacing	Subscription Lines
Intermediate Net IRR	stable	unstable	unstable
Final Net IRR	stable	stable	unstable
Intermediate Net MOIC	unstable	unstable	stable after investment period
Final Net MOIC	Method A: stable Method B: unstable	stable	stable
Net Profit	large differences	small differences	small differences
Fees and Interest	small differences	small differences	small differences

During the likely fundraising period for the next fund, both net IRR and net MOIC are very sensitive to the pace at which a fund deploys capital, especially insofar as slow deployment can

give extremely low or even negative net IRR and MOIC much less than 1. Final values are only affected slightly, as are investor profit and fees. Unlike with recycle deals or subscription lines, we cannot simply imagine a scenario in which pacing isn't present; IRR and MOIC in general should therefore be treated with extreme caution during the investment period of a fund.

When a fund utilizes subscription lines, net IRR is very sensitive over the life cycle of the fund and can massively exaggerate performance during the investment/fundraising period. Net MOIC can also be exaggerated early in the investment period, but ends up lower overall by the time a fund is about 3 or 4 years of age. Net IRR should therefore be treated with extreme caution at any stage of a fund that uses subscription lines, whereas MOIC is more stable, and especially so after the investment period.

When a fund utilizes recycling, net IRR remains a stable metric during the life cycle of a fund. On the other hand, net MOIC can become highly sensitive over the life cycle and can greatly exaggerate the underlying success of the fund's investments when recycle deals are not accounted for as new investments. Investor fees are slightly larger when a recycle deal is accounted for as a new investment, and that effect grows when recycle deals are larger, but the increase in investor profit grows more so. MOIC should therefore be treated with extreme caution at any stage of a fund with recycling and ascertaining the accounting method is critically important.

Although this analysis informs us of important features regarding fund-level performance, it begs the question of performance analysis of portfolio companies. In light of recent research suggesting a decline in performance persistence at the fund-level, these factors suggest yet another reason for more granular due diligence. Understanding the specific drivers of value creation at the deal-level resolves many of the complications encountered with fund-level performance analysis.

References

Albertus, James F and Matthew Denes (2020) "Private equity fund debt: Capital flows, performance, and agency costs," *SSRN working paper 3410076*.

Barber, Brad M and Ayako Yasuda (2017) "Interim fund performance and fundraising in private equity," *Journal of Financial Economics*, 124 (1), 172–194.

Brown, Gregory W, Oleg R Gredil, and Steven N Kaplan (2019) "Do private equity funds manipulate reported returns?" *Journal of Financial Economics*, 132 (2), 267–297.

Brown, Gregory W. and Maria Nykyforovych-Borysoff (2023) *Net Asset Value in Private Equity*, 1–4: Springer International Publishing, 10.1007/978-3-030-38738-9_9-2.

Garcia, Luis (2022) "Recycled Fund Capital Can Pad Return Multiples, Colmore Tells SEC," *The Wall Street Journal*.

Griffiths, Barry E. (2016) "MaxMult:A New Performance Metric for Private Equity," *Landmark Partners*.

Kaplan, Steven N and Antoinette Schoar (2005) "Private equity performance: Returns, persistence, and capital flows," *The Journal of Finance*, 60 (4), 1791–1823.

Kazemi, Hossein (2021) "Editor's Letter," *The Journal of Alternative Investments*, 23 (4), 1–3, 10.3905/jai.2021.23.4.001.

Mendoza, Carmela (2019) "Faster deployment doesn't help investments perform well," *Private Equity International*.

Prequin (2019) "Preqin Special Report: Subscription Credit Facilities,"Technical report, Preqin Holding Limited, `https://docs.preqin.com/reports/Preqin-Special-Report-Subscription-Credit-Facilities-June-2019.pdf`.

Warren, Patrick (2023) "Subscription Lines of Credit, Part One: The Rise (and Rise) of Sub Lines,"Technical report, Burgiss, `https://www.burgiss.com/applied-research-blog/subscription-lines-of-credit-part-one-the-rise-and-rise-of-sub-lines`.

A Appendix: Recycle Deals with Neutral Pacing and Holding

Because of the uneven deployment and difference in holding periods in the baseline model, we also consider even deployments and identical holding periods for all deployments to see if the outcomes change.

Specifically, we now suppose that \$10.625 is invested in years 0 through 3.5 (again at 6-month increments) and all deals have a holding period of 3.5 years. The "low" recycling scenario will have a recycle deal of \$15 in year 4.5; the "high" recycling scenario will have a recycle deal of \$15 in both years 4.5 and 5; and the "max" recycling scenario will have recycle deals of \$15 in years 4.5, 5, and 5.5.

As shown in Figure 11, final net IRR with Method A is above its no-recycle level, unlike in the baseline model. The main reason is that the (unweighted) average holding time in the base model with no recycling is $(3.5 + 5 + 5 + 5 + 5)/5 = 4.7$ years, and therefore adding another new investment with a 5-year holding period increases the average the holding time (equivalently reducing the accretion rate), thereby "slowing down" the rate of return.

In this section, however, adding new 3.5-year investment via recycle has no effect on the average holding time because all holding times are 3.5. The net IRR ultimately ends up higher here because fees as a proportion of capital calls are lower since some of them are collected after the investment period and therefore are only a fraction of their investment period levels. There is little qualitative difference for the final MOICs, but MOICs prior to year 8 can differ by as much as about 0.3. Of course, year 4 is included as part of what we assume to be the likely fundraising period for the GP's next fund.

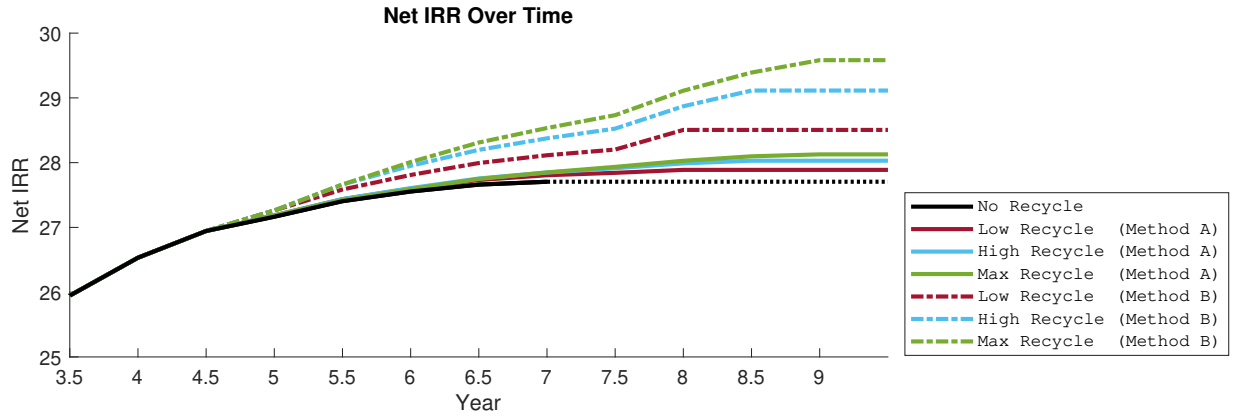

FIGURE 11: Final IRR with Method A is not under the no-recycle level when holding periods and pacing are even.

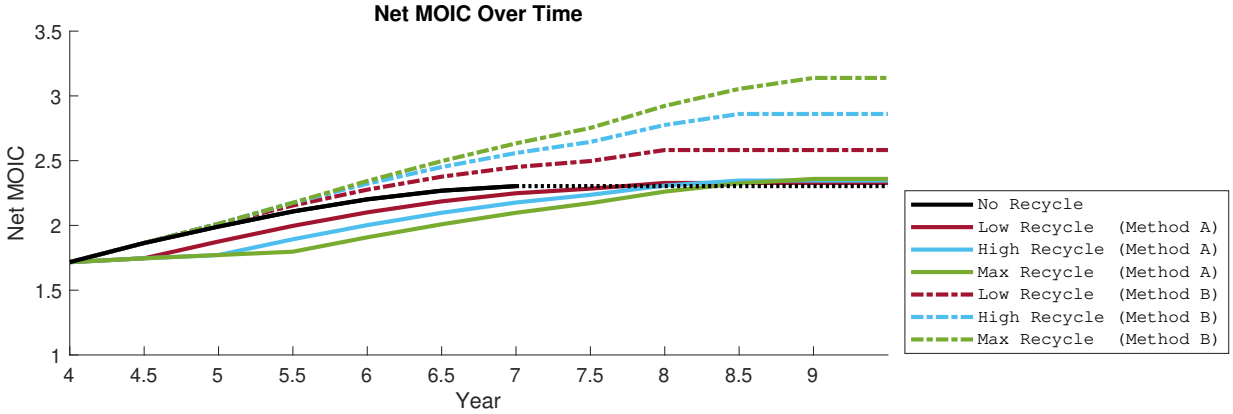

FIGURE 12: Not much changes qualitatively with net MOIC when holding periods and pacing are even.