The Journal *of* Wealth Management

20th
anniversary
PART ONE

Partnership Allocations and Their Effects on Tax-Aware Fund Investors

NATHAN SOSNER, PHILIP BALZAFIORE, AND ZHENDUO DU

IIJ

AQR

Partnership Allocations and Their Effects on Tax-Aware Fund Investors

NATHAN SOSNER, PHILIP BALZAFIORE, AND ZHENDUO DU

NATHAN SOSNER
is a managing director
at AQR Capital
Management in
Greenwich.
nathan.sosner@aqr.com

PHILIP BALZAFIORE
is a managing director
at AQR Capital
Management in
Greenwich.
philip.balzafiore@aqr.com

ZHENDUO DU
is an associate at AQR
Capital Management
in Greenwich.
zhenduo.du@aqr.com

Many tax-aware equity strategies seek to achieve tax benefits by deferring capital gains while continuing to realize, or even accelerating, capital losses. While long-only loss-harvesting strategies, such as those described by Stein and Narasimhan [1999] and Arnott, Berkin, and Ye [2001], among others, can be implemented in separate accounts, tax-aware strategies that use shorting, such as those studied by Sialm and Sosner [2018], are generally implemented in funds structured as limited partnerships. Limited partnerships are attractive to investors because, as limited partners, investors cannot economically lose more than their invested capital despite the leverage of the partnership's portfolio. As a result, the economic investment loss exposure of a limited partner in a fund can be far more limited than that of a separate account investor. Consistent with this, the availability of tax losses to a limited partner is correspondingly more limited than that of a separate account investor. Whereas a separate account investor includes all taxable gains and losses of the investment strategy directly in her tax returns, the tax outcomes of a limited partner in a fund managing the same strategy can be different. To shed light on the differences in taxation, this study outlines certain relevant principles of "securities partnership" tax accounting and shows how these prin-

ciples apply to investing in tax-aware funds structured as limited partnerships.

RULES GOVERNING GAIN AND LOSS ALLOCATIONS AND LOSS DEDUCTIBILITY

This section focuses on two topics: 1) allocations of realized gains and losses to partners and 2) limitations on the utilization of net allocated capital losses by such partners. (We denote sections of the Treasury regulations as "Reg." All other code section references are to the Internal Revenue Code.)

Allocation of Realized Gains and Losses

There are two acceptable methods for allocating realized gains and losses to partners in securities partnerships:[1] lot layering and aggregation. Both methods allocate realized gains and losses based on partners' relative shares of *unrealized* gains and losses. Lot layering is the more precise and more complex of the two methods. The aggregate method is a simplified proxy for lot layering that is only available to securities partnerships. It was made available by the U.S. Treasury because the sheer volume of transactions that securities partnerships tend to carry out renders the more precise lot layering method difficult to implement in practice. For brevity of

exposition, this study only discusses the less precise (and more commonly practiced by investment funds) aggregate method of allocation and shows that, even under this method, allocations of realized gains do not necessarily systematically disadvantage newer partners relative to the older partners in limited partnerships managing tax-aware strategies.

There are two approaches to aggregate allocations both of which are outlined in Reg. §§ 1.704-3(e)(3)(iv) and (v), respectively: *partial netting* and *full netting*.[2] In both approaches, *book* (economic) gains and losses are netted. The difference between the two arises from the netting of *tax* gains and losses. Tax gains and losses can be allocated either on a gross basis (gains and losses are allocated separately) or on a net basis (gains and losses are first netted within the category prior to allocation). Partial netting implements gross allocations and full netting implements net allocations of tax gains and losses in each category. ("Partial" netting therefore refers to the fact that only book gains and losses are netted, while tax gains and losses are not netted but rather allocated separately.) Comparison of these two approaches is outside the scope of this study.

In this article, we will discuss a variation of full netting whereupon gains and losses are netted within category—long term or short term—but are not netted across categories—long-term gains against short-term losses and vice versa—for the purpose of calculating tax allocations.[3] Per § 702(a) and Reg. § 1.704-3(e)(3)(vi) long-term and short-term gains and losses are allocated to partners as separate items. We proceed with an illustration of such a full netting approach. For simplicity, we assume that there are no performance allocations to the partnership's investment manager and that all investors receiving book and tax allocations are limited partners. Let's first define the concept of a *revaluation account*: An investor's revaluation account is the cumulative difference between her book (economic) and tax allocations and is updated periodically (annually, in this example) by the difference between book and tax allocations. For example, if a partner is allocated $10 of book gain, $5 of long-term tax gain, and $8 of short-term loss, the partner's revaluation account increases by $13 ($10 − [$5 − $8]). For funds that tend to allocate short-term tax losses in excess of long-term tax gains, such negative tax allocations can increase partners' revaluation accounts rather rapidly.

Calculation of tax allocations begins with the allocation of book income. *Net* book gain or loss is allocated to each partner in proportion to his or her interest in the partnership. Each partner's revaluation account is then updated by these book allocations—a net book gain increases the revaluation account, whereas a net book loss decreases it.

Once each partner's revaluation account is updated for book allocations, allocations of realized tax gains and losses in each category are performed using the following rules. Net realized tax *gains* in each category, if any, are first allocated based on each partner's relative share in the total of *positive* revaluation accounts, if any, but only up to the size of each such partner's positive revaluation account,[4] and any remaining net realized tax gains in each category are allocated to the partners based simply on their current relative percentage interests in the partnership. Similarly, net realized tax *losses* in each character category, if any, are first allocated based on each partner's relative share in the total of *negative* revaluation accounts,[5] but only up to the size of each such partner's negative revaluation account. Any remaining net realized tax losses are allocated to the partners based on their current relative percentage interests in the partnership.[6]

In the context of tax-aware funds, which have the propensity to accrue unrealized gains and realize gains as long-term and losses as short-term, our approach to aggregation—netting gains and losses within category (long-term or short-term) but not across categories—is not only reasonable but is perhaps *the most* reasonable. Reg. § 1.704(e)(3)(i) requires that a securities partnership aggregate gains and losses using a reasonable approach that is "consistent with the purpose of § 704(c)": § 704(c) seeks to allocate realized tax gains to those partners who benefited from economic gains and realized tax losses to those partners who suffered economic losses. Most if not all partners in tax-aware funds, which tend to accumulate unrealized gains, will have positive revaluation accounts, that is, unrealized gains in their partnership interests. As a result, allocation of long-term gains in proportion to partners' revaluation accounts is consistent with the long-term accumulation of unrealized gains in their partnership interests. On the other hand, short-term losses are not reflective of the specific experience of partners accumulating unrealized gains and are thus allocated independently of unrealized gains simply in proportion to partnership interests.[7]

Relevant Implications of Allocation of Realized Gains and Losses

- In a partnership managing a tax-aware strategy, which has the propensity to accrue unrealized gains and realize gains as long-term and losses as short-term, a reasonable aggregation method is to net realized tax gains and losses within categories but not across categories—long-term and short-term gains and losses are allocated separately.
- In a partnership managing a tax-loss-harvesting strategy, partners are much more likely to have positive revaluation accounts, reflecting unrealized gains.
- If all partners in a securities partnership have positive revaluation accounts and in a given period the partnership realizes net long-term gains and net short-term losses, under our approach the net long-term gains are allocated based on the partners' respective shares in the total positive revaluation accounts and the net short-term losses are allocated based on their current partnership interest percentages.
- Recently joining partners, who have not yet accumulated large positive revaluation accounts through net tax loss allocations, are allocated under this approach only a small portion of long-term tax gains realized by the partnership but are allocated realized short-term losses based on their partnership interest percentages, which is roughly reflective of their economic experience in the partnership.

The final point is an important one. The fact that a new partner receives a disproportionately small share of realized long-term gain to some extent replicates the experience of an investor in a separate account where tax efficiency (or tax benefit) is expected to be higher in the first few years when the cost basis of most positions is closer to their market value. However, the source of tax efficiency is very different. In a separate account, the investment portfolio tends to realize less gain and more losses. By contrast, in a limited partnership that uses a character-sensitive aggregate allocation approach, even if its investment portfolio realizes significant gains, only a relatively small share of those gains is allocated to the new investor due to her relatively small positive revaluation account.

In addition, there is a limitation on the quantity of tax losses allocated by a limited partnership that a partner can use as a deduction against capital gains from other investments. We explain this limitation next.

Limitation on Deduction of Allocated Tax Losses

§ 704(d) limits a partner's deduction of allocated partnership tax losses to his or her "outside basis" (adjusted cost) in the partnership at the end of the year. The partner's outside basis is adjusted under § 705 for prior capital contributions and distributions and allocations of partnership income and loss items. A partner's allocated tax losses in excess of her outside basis at the end of the year are disallowed and are carried forward for potential future use. These disallowed excess tax losses may be used in a future year to offset unrelated taxable gains if (and only if) the partner's outside basis increases (generally through additional contributions, allocations of income or gains, or an increase in allocated liabilities, as we explain in the next paragraph). The category (short-term or long-term) of such carryforward losses is preserved in future years.

Under § 752, an increase in a partner's share of partnership liabilities is considered to be a "contribution of money" by the partner to the partnership, while a decrease in her share of partnership liabilities is considered a "distribution of money." This means that for purposes of the § 704(d) limitation on allocated losses, an increase in allocated liabilities increases a partner's outside basis and thus increases her ability to deduct allocated tax losses, whereas a decrease in allocated liabilities decreases her outside basis and therefore the amount of allocated tax losses that can be potentially deducted.

An additional limitation on allocated partnership losses—above and beyond the § 704(d) limitation—is imposed on partners who are individuals by the "at-risk rules" in § 465. § 465(a)(1)(A) states that an *individual* taxpayer cannot deduct more tax losses than the amount she has "at risk."

What does this mean for an individual investor in a fund structured as a limited partnership? A partnership may invest in a levered strategy and thus suffer losses that exceed the amount of limited partner invested capital. However, an investor who is a limited partner can only lose up to invested capital and is not liable for any further losses incurred by the partnership. § 465 thus limits the

deduction of tax losses to the partner's maximum potential economic loss: For every dollar contributed to the fund plus every dollar of allocated taxable income that hasn't been distributed (plus every dollar incrementally borrowed by the partnership for which an investor is *personally liable*), an investor can deduct at most a dollar of tax losses allocated to him or her by the fund.

It is important to make the distinction between a *partner's cost basis* in the partnership (outside basis) and her share of the *partnership's cost basis* in assets and securities held by the partnership (inside basis). The following example provides a simple illustration of inside and outside basis concepts.

Suppose that John and Mary invest $50 and $100, respectively, in XYZ fund structured as a limited partnership. XYZ fund borrows another $150 and invests the total of $300 in 3 stocks—1, 2 and 3, $100 in each. The fund has $100 inside basis in stock 1, $100 inside basis in stock 2, and $100 inside basis in stock 3. If the $150 loan were allocated $50 to John and $100 to Mary, respectively, their outside bases in the fund would be the total of their contributed capital plus their respective shares of liabilities—$100 outside basis for John and $200 outside basis for Mary. However, they are limited partners and the fund creditors have no recourse to them, only to the general partner. As a result, whether the liabilities are allocated to John and Mary or not (and thus whether their outside bases increase accordingly), their at-risk amounts remain only $50 and $100, respectively. The logic is that if the value of all the assets in the fund goes to zero, and the fund thus loses $300 economically, John can only lose $50 and Mary $100. According to the "at-risk" rules described previously, therefore, John can only deduct $50 of aggregate losses allocated to him by the fund and Mary can only deduct $100.

For simplicity, henceforth we assume that a limited partner is not allocated any fund liabilities and thus her outside basis defined under § 705 and her amount considered at risk under § 465 are equal. This will allow us to discuss a partner's outside basis limitation while at the same time implying the partner's at-risk amount as being the same thing. For a limited partner, additional cash contributions and allocations of taxable income and gains increase her outside basis (and her at-risk-amount) dollar for dollar, while cash redemptions and allocations of losses decrease it.[8] Thus, barring additional contributions, the outside basis of a limited partner in a leveraged partnership that on average tends

to allocate losses might eventually be reduced to zero. At that point, any further net losses allocated to the limited partner (that is, losses of one category, say short-term, allocated in excess of gains allocated of a different category, say long-term) become *suspended*—the partner is not allowed to use them as deductions unless and until her basis in the partnership becomes positive.

An example helps clarify these rules. Let's continue with the example of John and Mary. Suppose that each year the fund allocates to John $25 of long-term capital gains and $50 of short-term capital losses. Because gain allocations increase the outside basis and loss allocations decrease the outside basis, John's outside basis in the fund at the end of year 1 is $25 and at the end of year 2 it is $0. In year 3, John receives the same allocations of gains and losses but his outside basis is $0. For John, the $25 long-term gain allocation increases his outside basis in the fund to $25, and $25 of the $50 allocated short-term loss then decreases his basis back to $0 (under the tax rules outside basis cannot be reduced below zero). As a result, John can offset $25 of short-term gains from other strategies in his investment portfolio with $25 of the short-term losses allocated by the fund; the remaining $25 of short-term loss allocation becomes suspended until John's outside basis in the fund becomes positive. In year 4, John's situation is similar: He is allocated $25 of long-term gain and $50 of short-term loss; $25 of the short-term loss allocation can be used to offset short-term gains elsewhere in the portfolio and his suspended short-term capital loss in the fund increases by an additional $25 to $50. In other words, while one might think of John as theoretically having a negative basis in his partnership interest, his outside basis under the tax law is zero. The difference between the two is his suspended loss.

Suspended losses become "unlocked" and therefore deductible when outside basis increases from zero through the allocation of income or net capital gains or additional contributions to the partnership. For example, suppose that in year 4 John sells an unrelated long-term investment with a $50 cost basis and $100 market value. He now has $100 in cash and a realized long-term gain of $50. If John were to contribute the $100 of proceeds to the fund, his outside basis would increase to $100 and his $50 of suspended short-term capital losses would be "unlocked" and deductible (thereby reducing his outside basis in the fund from the initial $100, caused by the contribution, to $50).

The net result of these two transactions for John is as follows: $50 of realized long-term gain from his liquidation of the unrelated investment; an increase of his economic interest in the partnership by $100; an increase in his outside basis in the fund by net $50; and a release of the $50 short-term loss previously allocated but suspended. Note that if John does not have short-term gains from other investments, the unlocked short-term loss allocation will be used to offset long-term gains, including those from having sold the unrelated investment, and any losses in excess of long-term gains will be carried forward in their character under a different set of rules (i.e., they are no longer "suspended"). However, in the presence of other short-term gains, the $50 of previously suspended short-term loss will be used most efficiently to offset those gains.

Relevant Implications of Limitation on Deduction of Allocated Tax Losses

- Partners' deductible losses are limited by their outside basis in the partnership—a partner cannot deduct losses in excess of her outside basis.
- While outside basis might include allocated liabilities if properly structured, for a partner who is an individual, deductible losses are also limited by the amount of her investment considered at risk. Non-recourse liabilities are specifically excluded from the at-risk amount.
- Suspended losses can be used by the partner when the partner's outside basis, or the at-risk amount, in the partnership becomes positive.
- When a partner's outside basis, or at-risk amount for an individual, in the partnership is zero, in the absence of additional contributions she can only deduct losses to the extent of allocated realized gains. This is because gain allocations increase both outside basis and the at-risk amount.
- Additional contributions to the partnership by a partner with a zero outside basis, or at-risk amount, allow the partner to deduct suspended losses to the extent of the cash contribution (or the tax basis of contributed property). "Unlocking" of suspended losses reduces the partner's outside basis, or her at-risk amount, dollar for dollar, which generally counteracts the thing that increased it in the first instance.

CALCULATION OF PARTNERSHIP ALLOCATIONS WITH MULTIPLE PARTNERS

Calculation Methodology

This study is intended to illustrate general outcomes of partnership allocations under only the most basic of assumptions, and as such, we rely on the following stylized facts. First, the strategy's pre-tax return and realized gains and losses are a constant fraction of the fund's net asset value (NAV) and are unaffected by contributions to the fund. Second, investors can contribute to the fund only at year-end. Third, we assume that the investor is an individual and thus is subject to the "at-risk" rules as discussed in the previous section.

The assumptions about the strategy's pre-tax return and realized gains and losses are based on the long–short tax-aware strategy results in Sialm and Sosner [2018]. The pre-tax annual return is 4.3%, the realized annual long-term capital gain is 13.8% of the NAV, and the realized annual short-term capital loss is 20.2% of the NAV. Finally, we assume that all of the tax savings achieved by the strategy are not reinvested back into the strategy but rather are used for other investments.

In our calculations, we will also account for the fact that a partner's amount at risk (and outside basis) is reduced by the allocation of short-term losses in excess of long-term gains, such that eventually the amount at risk is reduced to zero and further allocations of short-term losses exceeding long-term gains become suspended.

Partnership Allocations

We first calculate the effect of a contribution by a new partner on future realized gain allocations of the fund. Note that we assume that the fund generates pre-tax income while at the same time experiencing net tax losses. As a result, all partners' revaluation accounts will be positive, leading to allocation of gains based on relative revaluation accounts, while losses will be allocated based on relative partnership interest percentages. Allocation of losses based on partnership interest percentages is straightforward; for example, a 50% partner will be allocated 50% of losses every year (not all of these losses will be deductible under the outside basis and at-risk limitations, but they are allocated nonetheless). The realized gain allocation pattern, however, is more involved.

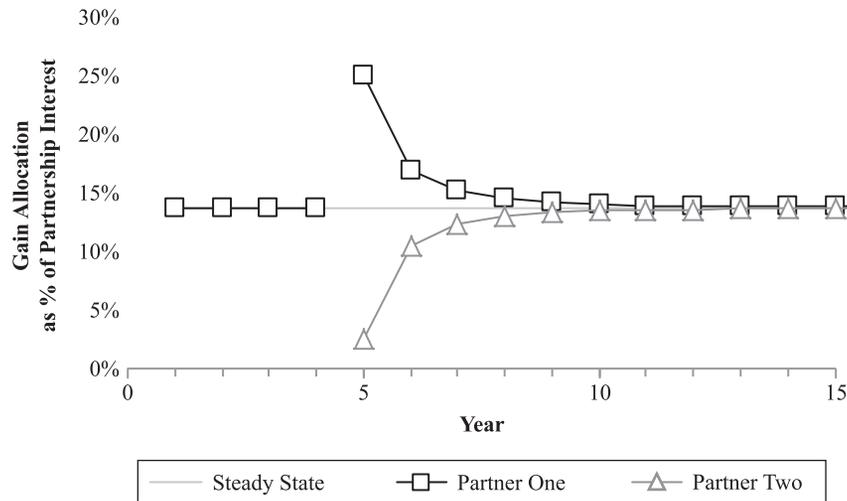**A New Partner Joins in Year 5 as a 50% Partner**

Exhibit 1 shows the pattern of long-term gain allocations under the assumption that a fund is launched with one limited partner, Partner One. Another 50% partner, Partner Two, joins the fund on the first day of year 5. Based on our assumptions for gain allocations, each year Partner One is allocated 13.8% of the value of her partnership interest in long-term gains and 20.2% in short-term losses. We will refer to these allocations as "steady state." When Partner Two joins as an equal partner, both partners continue to be allocated 20.2% of the value of their respective partnership interests in short-term losses, but their long-term gain allocations diverge substantially from the steady state.

What causes the pattern we observe in Exhibit 1? Over time Partner One accumulates a large positive revaluation account. In fact, each year her revaluation account is increased by 10.7% of her partnership interest value (Book income − [Gain allocation − Loss allocation] = 4.3 − [13.8 − 20.2] = 10.7%). Partner Two comes in with a zero revaluation account at the beginning of year 5 and sees an increase in her revaluation account of only 4.3% (book income) before tax gain or loss allocations are considered.[9] As a result, Partner One is allocated the majority of the long-term gains realized by the fund in year 5.

Because the revaluation account is updated by the difference between the book and tax allocations and Partner Two is initially allocated a smaller portion of gains than Partner One, going forward her revaluation account grows faster than the revaluation account of Partner One. Eventually, the two revaluation accounts converge around years 9 to 10, at which point each partners' gain allocations return to steady state.

Exhibit 2 shows the evolution of the partners' at-risk amounts (and outside bases) and their suspended losses. Due to allocations of short-term losses in excess of long-term gains, the partners' at risk amounts eventually reach zero, at which point they begin to accumulate suspended losses. The entry of Partner Two increases the allocation of gains to Partner One above the steady state level but does not change the allocation of losses. As a result, while in the steady state Partner One reaches a zero at-risk amount after 13 years, entry of the new partner extends that period to 15 years.

What happens if Partner Two joins the partnership at an earlier or a later year? We expect the shift in gain allocation to Partner One to increase as Partner Two joins the fund in progressively later years. This is because the revaluation account of Partner One becomes larger over time. However, each additional year adds proportionately less to an already large revaluation account.

Exhibit 3 illustrates Partner Two joining in different years and shows gain allocations to Partners One and Partner Two in the year when Partner Two first joins the fund. The first-year allocation to Partner One indeed increases but at a slowing rate. The two circled points in Exhibit 3 correspond to the year 5 allocations in Exhibit 1. Similarly, the convergence to the steady
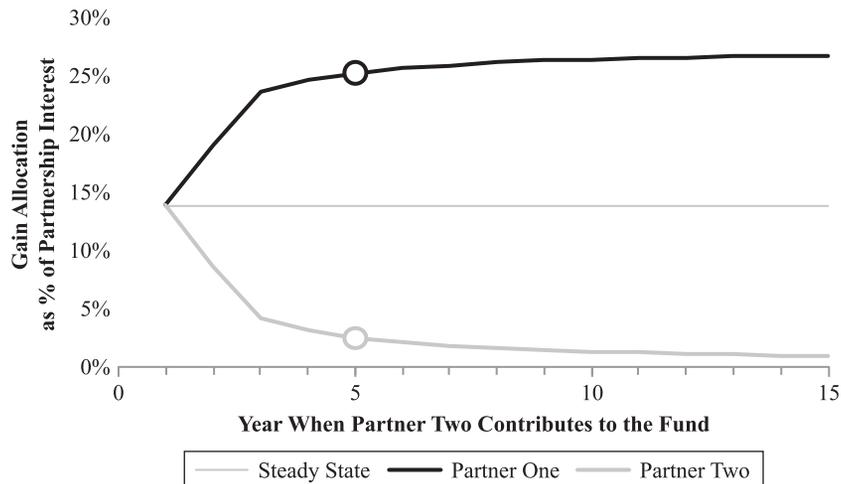
E X H I B I T 2

Amount at Risk (AAR) and Suspended Loss (SL) when a New Partner Joins in Year 5 as a 50% Partner
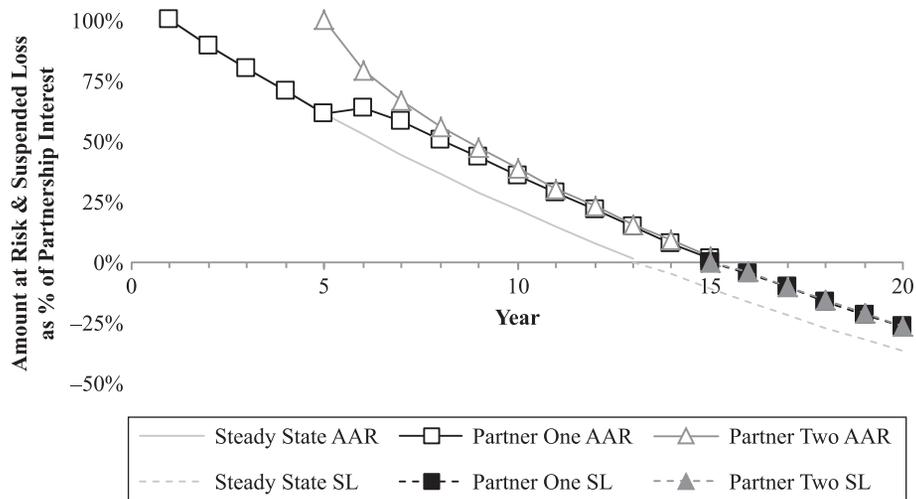
E X H I B I T 3

**A New Partner Joins in Different Years as a 50% Partner**

state gain allocations occurs faster (slower) if Partner Two joins at an earlier (later) year. For example, as can be seen in Exhibit 1, following the contribution of Partner Two in year 5 it takes 4 to 5 years for the gain allocations to converge to their steady state level. In calculations not shown here, we find that if the contribution of Partner Two occurs in year 3, the convergence to the steady state level takes only one year, and if contribution entry occurs in year 7, the convergence takes approximately 7 to 8 years.

What about the size of the Partner Two's contribution to the fund? As the size of Partner Two's contribution increases, Partner One is affected proportionately more; however, there is little effect on Partner Two. The intuition behind these outcomes is straightforward: A larger contribution by a new partner results in larger realized gains for the fund, and because the old partner in the beginning attracts most of those gains (due to the accumulated revaluation account), her gain allocation relative to the value of her partnership interest increases.
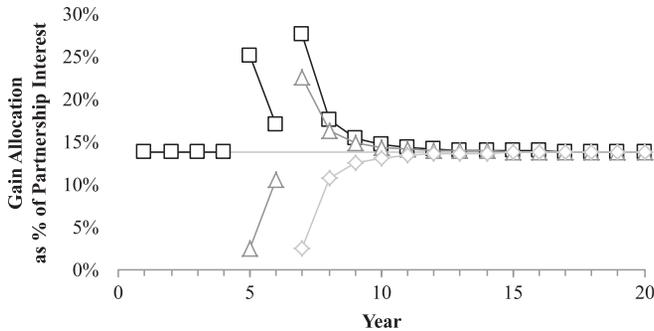
Exhibit 4
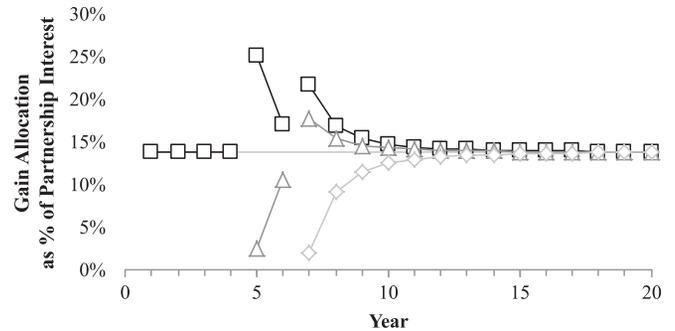## A New Partner Joins in Year 5 with Different Levels of Partnership Interest
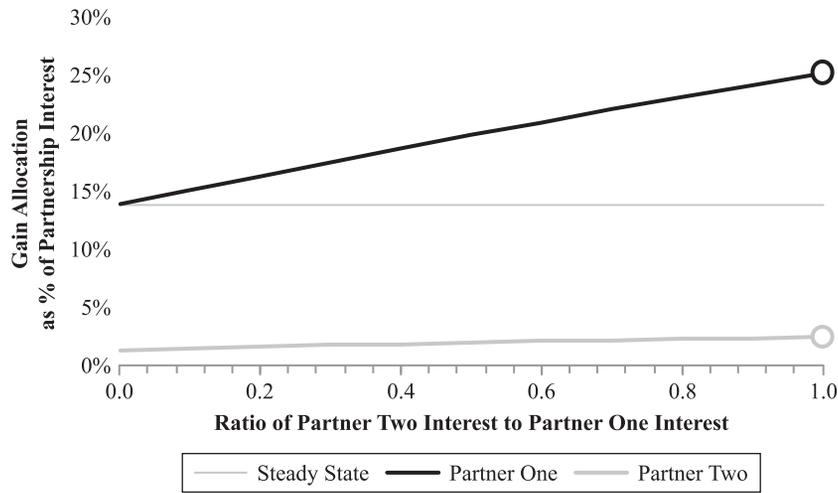
Exhibit 5
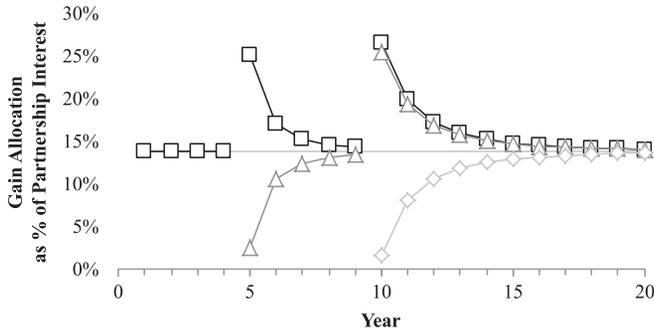## Successive Entry of New Partners with Various Contribution Sizes

Panel A: One New Partner Joins in Year 5 as a 50% Partner, Another New Partner Joins in Year 7 as a 50% Partner
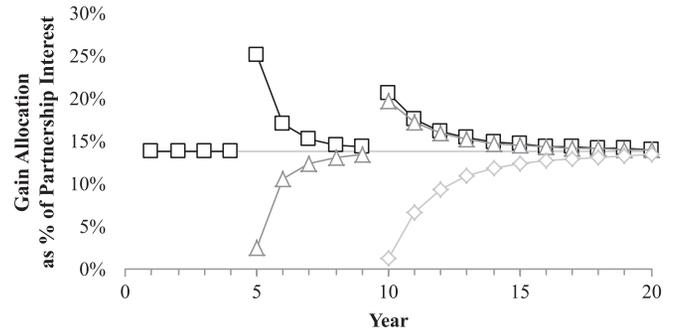
Panel B: One New Partner Joins in Year 5 as a 50% Partner, Another New Partner Joins in Year 7 as a 33% Partner

Panel C: One New Partner Joins in Year 5 as a 50% Partner, Another New Partner Joins in Year 10 as a 50% Partner

Panel D: One New Partner Joins in Year 5 as a 50% Partner, Another New Partner Joins in Year 10 as a 33% Partner

This result is shown in Exhibit 4, where Partner Two joins in year 5 with a progressively larger contribution size. The two circled points in Exhibit 4 correspond to the year 5 allocations in Exhibit 1, when Partner Two joins as an equal partner. In calculations not shown here, we find that the relative size of the contribution by Partner Two affects the speed of convergence to the steady state gain allocations. In particular, larger contributions help the allocations to converge faster while smaller contributions prolong the convergence process. If Partner Two is equal in size to Partner One, the convergence takes about 4 to 5 years as we saw in Exhibit 1. If Partner Two is 100 times smaller than Partner One, the convergence takes 6 to 7 years. And if Partner Two is 3 times larger than Partner One or more, the convergence occurs after just one year.

With these results in mind, it is possible to anticipate how a successive entry of new partners with various contribution sizes will affect partnership allocations. Panel A of Exhibit 5 shows the entry of 50% Partner Two in year 5 and 50% Partner Three in year 7. Panel B changes Partner Three's contribution from 50% to one-third. The calculations in Panels C and D assume that Partner Three enters in year 10 with 50% and a one-third interest, respectively.

A few previously observed patterns are reinforced in Exhibit 5. First, the timing of new contributions matters: A later entry of a new partner causes larger increases in gain allocations to the old partners due to their larger relative revaluation accounts—compare Panels A and B with Panels C and D. Second, the size of new contributions matters: A larger contribution by a new partner causes a larger increase in gain allocations to the old partners— compare Panels A and C with Panels B and D. Finally, with time a new partner's revaluation account as a fraction of partnership interest converges to that of old partners, such that the later an additional new partner joins, the more similar effect that a new contribution will have on all existing partners. Whereas Panels A and B show some difference between the allocations of partners One and Two upon the entry of Partner Three, those differences almost disappear in Panels C and D, where Partner Three enters three years later than in Panels A and B.

To summarize, upon entry new partners tend to benefit from a disproportionately small allocation of realized taxable gains while at the same time being allocated losses in proportion to their partnership interest percentage. As a result, new partners are not materially adversely affected by the unrealized gains accumulated by the old partners in the fund.

CONCLUSION

Investors in commingled funds might be concerned about the potential impact of other investors on their tax liabilities. For funds that are "securities partnerships" and structured as limited partnerships, these concerns are alleviated by the laws and regulations that evolved over decades with the goal of aligning tax results with economic outcomes. Even in tax-aware funds that tend to allocate net tax losses, new partners do not materially suffer from unrealized gains accumulated in the fund prior to their admission if an appropriate partnership allocation method is implemented. Our stylized model illustrates why this is the case.

ENDNOTES

[1]Reg. § 1.704–3(e)(3)(iii) provides the definition of securities partnerships and in particular of investment partnerships that are the main interest of this study. It states that a partnership is an investment partnership if 1) On the date of each capital account restatement, the partnership holds qualified financial assets that constitute at least 90% of the fair market value of the partnership's non-cash assets and 2) The partnership reasonably expects, as of the end of the first taxable year in which the partnership adopts an aggregate approach under this paragraph (e)(3), to make revaluations at least annually.

[2]An interested reader can find numerical examples of application of these two approaches in Reg. § 1.704–3(e)(3)(ix).

[3]As we discuss later in this section, such an approach to full netting is not only a reasonable approach as the tax law requires, but is perhaps *the most* reasonable approach in the context of tax-aware funds that tend to defer capital gains and accelerate capital losses.

[4]Net tax gains are first allocated based on the partner's relative share in the total of positive revaluation accounts because a positive revaluation account means that the partner has unrealized gain in her partnership interest, and thus taxable realized gain allocations go in the direction of unrealized gains with the goal of "closing the gap" between book and tax accounts.

[5]Mirroring the logic of tax gain allocations, tax losses go in the direction of unrealized losses.

[6]Securities partnerships often make special allocations to redeeming partners. Hedge funds use such allocations to best align the economic and tax outcomes of the redeeming

and remaining partners by ensuring that the revaluation accounts of redeeming partners are closed out upon redemption. A discussion of such special allocations is outside of the scope of this article.

[7]We would like to point out that such character-sensitive full netting is particularly pertinent for tax-aware funds. Funds that tend to realize most of their gains within a year may not see a meaningful difference in their ability to satisfy the purpose of § 704(c) whether they use this approach or any other approach to aggregation, for example, netting gains and losses not only within long-term and short-term categories but also across these categories.

[8]Contributions and redemptions in-kind have more complex consequences for at-risk amount calculations and are outside of the scope of this study.

[9]The earlier section on the allocation of realized gains and losses explains the logic of partnership allocations.

REFERENCES

Arnott, R.D., A.L. Berkin, and J. Ye. "Loss Harvesting: What's It Worth to the Taxable Investor?" *The Journal of Wealth Management*, Vol. 3, No. 4 (Spring 2001), pp. 10–18.

Sialm, C., and N. Sosner. "Taxes, Shorting, and Active Management." *Financial Analysts Journal*, Vol. 74, No. 1 (First Quarter 2018), pp. 88–107.

Stein, D.M., and P. Narasimhan. "Of Passive and Active Equity Portfolios in the Presence of Taxes." *Journal of Private Portfolio Management*, Vol. 2, No. 2 (Fall 1999), pp. 55–63.

Disclosure

To order reprints of this article, please contact David Rowe at drowe@iijournals.com or 212-224-3045.