

The Journal *of* Wealth Management

“A goals-based asset allocation approach may be better than the alternative to define investment policy, for organizations that must view risk as the inability to make a required payment without selling risk-assets at a potentially difficult time.”

Jean Brunel - Editor's Letter



**Nathan Sosner
and Philip Balzafiore**

**Lot Layering:
*The New Frontier
for Hedge Fund
Partnership
Allocations***



Lot Layering: *The New Frontier* for Hedge Fund Partnership Allocations

NATHAN SOSNER AND PHILIP BALZAFIORE

NATHAN SOSNER is a principal at AQR Capital Management in Greenwich, CT. nathan.sosner@aqr.com

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

KEY FINDINGS

- Lot layering is considered by most tax experts to be the most precise method of partnership allocations, where the precision relates to the alignment between tax and economic outcomes of partners in a partnership.
- We explain how this commonly understood precision is reduced upon redemptions due to the cumbersome and uneconomic basis adjustment method stipulated by Treasury regulations and propose changes to the current regulations that could remedy this problem.
- Despite its unavoidable deficiency caused by the basis adjustment requirements under the current regulations, lot layering both eliminates the need for stuffing allocations and aligns tax and economics more closely than do any of the *aggregation* methods presently used by most hedge funds.

ABSTRACT: *Lot layering may help hedge funds improve the alignment between tax outcomes and the economic experience of their investors. Although lot layering is considered by most tax experts to be the most precise method of partnership allocations, this commonly understood precision is reduced upon redemptions due to the cumbersome and uneconomic basis adjustment method stipulated by Treasury regulations. We propose that changes be made to the current regulations that could remedy this problem. Despite its unavoidable deficiency caused by the basis adjustment requirements under the current regulations, we believe that lot layering aligns tax and economics more closely than do any of the aggregation methods presently used by most hedge funds.*

TOPICS: *Wealth management, legal/regulatory/public policy**

*All articles are now categorized by topics and subtopics. **View at PM-Research.com.**

The purpose of this article is to shed light on the default method for allocating partnership gains and losses, commonly referred to as *layering*, which historically, for practical reasons, has been unavailable to hedge funds that trade significant volumes of securities. Owing to technological advances in data management and computational methods, layering (which in the context of securities partnership hedge funds is more narrowly referred to as *lot layering*) may now help hedge funds improve the alignment between tax outcomes and the economic experience of their investors.

Hedge funds are usually organized as limited partnerships.¹ As partnerships, they are

¹Hedge funds can be also structured as other types of pass-through entities, such as a limited liability company. The analysis in this article applies to

generally not subject to tax at the entity level; rather, their limited partner investors are taxed on items of gain and loss passed through, or allocated, to them by the fund. As a result, for hedge fund investors tax outcomes are determined by two primary factors: (1) realization of tax items (gains, income, losses, and deductions) at the overall fund portfolio level and (2) the allocation of those tax items by the fund to each particular investor. Although the former has received some attention in the investment literature over the past 2 decades, the latter has been generally ignored (an article by Sosner, Balzafiore, and Du, 2018, is an exception). The scant attention to this second factor notwithstanding, the method a hedge fund chooses to allocate its tax items to investors might have a significant effect on their tax outcomes.

Most hedge funds use so-called *aggregation* methods to allocate partnership tax items. Aggregation methods were sanctioned by Treasury regulations exclusively for “securities partnerships” to provide a practical allocation solution for partnerships with frequent turnover in their assets and investor base² (technology available at the time these regulations were issued made it impossible or impractical to track every economic layer for such partnerships). Under aggregation methods, tax items are not tracked on a lot-by-lot basis but rather are apportioned among fund investors using calculations relying only on their participation in the fund’s *overall* (aggregate) cumulative gains or losses. For hedge funds and their tax preparers, the ability to work only with highly aggregated numbers has provided substantial relief of the computational burden associated with tax reporting.

More recently, due to advances in data management technology and computational capabilities of hedge funds and their tax preparers, tax professionals have begun to advocate the use of the more general approach to partnership allocations—lot layering—even for high-turnover hedge funds. In contrast to aggregation methods, when lot layering is used, tax items are not aggregated across all assets. Rather, tax items of specific assets, at a lot level, are tagged to specific investors who economically owned those lots during the period when the tax items accrued.

Lot layering is considered by most tax experts to be the most precise method of partnership allocations, and, as

such entities as well, assuming they elect to be treated as partnerships for tax purposes.

²Reg. § 1.704-3(e)(3) provides the definition of securities partnership, outlines technical rules for alternative aggregation methods, and shows numerical examples of application of such methods.

proof of this, it is also the default method required under § 704 of the Internal Revenue Code (throughout the article, we use “§” and “Reg. §” to reference sections of the Internal Revenue Code and the Treasury Regulations, respectively). Partnership allocations are considered precise when partners’ tax outcomes follow their economic gain and loss and, therefore, replicate a hypothetical separate account investment in the underlying portfolio. In this article, we explain what precision means in the context of lot layering and how such precision can be preserved. We also explain that the commonly presumed precision of lot layering is reduced upon redemptions and associated basis adjustments available under current regulations. Nonetheless, it is our view that, despite some unavoidable dilution in precision, lot layering is the best method of partnership allocations for aligning the tax and economic outcomes of hedge fund investors.

We suggest how the current tax law can be changed to fully remedy the dilution of lot layering precision upon redemptions. Not only does our proposed solution reduce computational complexity, but it also creates a precise alignment for every investor between their allocated tax gains and losses and their economic gain and loss. As of now, however, our solution is purely theoretical, as implementing it would require changes in Treasury regulations.

Our discussion of lot layering would be incomplete without addressing its costs and complexity. Therefore, in the penultimate section we summarize the costs and benefits of lot layering relative to the aggregation methods historically used by most hedge funds.

THE PRECISION OF LOT LAYERING

Information That Makes Lot Layering More Precise Than Aggregation

Exhibit 1 shows the main source of the increased precision of lot layering, compared to aggregation methods. Panel A summarizes the information maintained under both aggregation and lot layering. At the partnership level, both fair market value (FMV) and cost basis are tracked for each lot. The difference between a lot’s FMV and its basis is the unrealized gain of the partnership in the lot. Similarly, at the partner level, the FMV of the partnership interest, or capital account, and the outside cost basis of the partnership interest is tracked for each partner. The difference between a partner’s FMV and outside basis is the partner’s unrealized gain

EXHIBIT 1

Differences in Information Tracked under Aggregation and Lot Layering

Panel A: Information Tracked under Both Aggregation and Lot Layering

Partnership Account			Partner Accounts			
	Basis	FMV	Gain/(Loss)	Outside Basis	Capital Account	
Lot 1	\$300	\$240	(\$60)	Partner 1	\$200	\$180
Lot 2	\$100	\$180	\$80	Partner 2	\$140	\$180
Lot 3	\$100	\$120	\$20	Partner 3	\$160	\$180
	\$500	\$540	\$40		\$500	\$540

Panel B: Additional Information Tracked under Lot Layering Only

Partner Unrealized Gains/(Losses) in Specific Lots				
	Lot 1	Lot 2	Lot 3	Total
Partner 1	(\$40)	\$40	(\$20)	(\$20)
Partner 2	\$10	\$20	\$10	\$40
Partner 3	(\$30)	\$20	\$30	\$20
	(\$60)	\$80	\$20	\$40

in the partnership, which, under aggregation, is referred to as *revaluation account*.

Panel B shows the main differentiating feature of lot layering—whereas under aggregation the unrealized gain and loss information is bunched up, under lot layering unrealized gains are tracked at the granular partner-lot level. In our Exhibit 1 example, the partnership has a \$60 economic loss in Lot 1, as is shown in Panel A. Panel B, which is specific to lot layering, shows how that this \$60 loss comprises a \$40 economic loss of Partner 1, a \$10 economic gain of Partner 2, and a \$30 economic loss of Partner 3. Similarly, Panel A shows that Partner 1 has a \$20 economic loss in the partnership. The information specific to lot layering in Panel B shows that this \$20 loss comprises a \$40 economic loss in Lot 1, a \$40 economic gain in Lot 2, and a \$20 economic loss in Lot 3.

In theory, under lot layering, the detailed breakdown at the partner-lot level shown in Exhibit 1, Panel B, allows for an allocation of tax gains and losses to the partners exactly in line with their economic gains and losses in each underlying lot.

Precision after Disposition of Assets: Remedial Allocations

There is some mathematical complexity to lot layering precision, though. Suppose a lot is disposed of by the partnership at a realized gain. However, this gain results from a large unrealized economic gain in

one period followed by smaller unrealized economic loss in another. Some partners might have a net economic gain in the stock (e.g., those partners who joined the partnership before the large economic gain in the lot); other partners might have an economic loss (e.g., those partners who joined the partnership right before the economic loss in the lot). How best to allocate the realized gain to the partners? The granularity of the unrealized gain and loss ledger maintained under lot layering provides the fairest result: Allocate the total realized gain in the lot in such a way that those partners who prior to the disposition had an unrealized economic gain in the lot (across all periods) receive an allocation of the full amount of that gain, whereas those partners who, prior to the disposition, had an unrealized economic loss receive an allocation of that loss. Because, under lot layering, we know precisely how much each partner gained or lost economically from every lot held by the partnership, it stands to reason that the tax allocations should be made precisely following the economics dollar for dollar.

Such precision would not be possible were the partnership to apply the limiting traditional method of tax allocations. Under the traditional allocation method defined in Reg. § 1.704-3(b), if the *partnership* realizes a gain in a given lot, it can only allocate such gain in that lot to its partners. Similarly, if the partnership realizes a loss in a given lot, it can only allocate such loss in that lot to its partners. This limitation is known as the *ceiling rule*, the application of which means that where a partnership

realizes a net gain from a lot, it cannot allocate any losses from such lot to any of its partners, even those who experienced a true economic loss in the lot. As a result, the economically correct allocation of tax loss to those partners who experienced an economic loss cannot be done. Instead, the gain realized by the partnership would be allocated only between those partners who had gains in the disposed lot, in proportion to their unrealized gains in the lot prior to the disposition. The partners who had an unrealized loss in the lot will receive neither an allocation of gain nor loss. Under this approach, the economic gain partners are effectively underallocated tax gains, and the economic loss partners are underallocated tax losses.

Because under lot layering the precise economic participation of each partner in the unrealized gain or loss of every lot is known, the partnership is permitted to look for ways to correct this misallocation of tax gains and losses among its partners. Such a correction mechanism is specifically available under the “remedial allocation method” defined in Reg. § 1.704-3(d), which, broadly speaking, allows the partnership to eliminate the distortive underallocations caused by the ceiling rule of the traditional method. Under the remedial allocation method, the partnership would allocate a tax loss precisely matching the economic loss to its economic loss partners as a remedial item and simultaneously allocate a matching amount of tax gain to its economic gain partners. These tax allocations would exactly follow the partners’ economic experience in the partnership.

To summarize, when lots are disposed of, the precision of lot layering is preserved through the use of remedial allocations, which allow for an exact match between economic gain and loss and tax gain and loss for each partner, unrestrained by the ceiling rule’s net realized tax gain or loss of the partnership in that lot.

Precision after Redemptions: The § 754 Basis Adjustment

Under current law, redemptions present a challenge to lot layering precision. Redeeming partners generally leave behind their built-in gains and losses in the partnership’s assets. If no additional measures were implemented, when the underlying assets are eventually disposed of, the partnership would have to allocate the economic gains and losses of the redeemers to continuing partners, thereby breaking down the relationship between their economic experience and tax allocations and causing duplication of gains and losses.

To be more explicit, upon redemption, redeeming partners receive a distribution of money or assets and, under § 731, recognize a tax gain or loss equal to the difference between the value of their distribution and their outside basis in the partnership. Effectively, at the time of redemption every redeeming partner *trues up* on the deferred gain or loss accumulated in her partnership interest. To mitigate the duplication of this § 731 gain or loss when underlying partnership assets are sold, a partnership applying lot layering would be well advised to make a so-called § 754 election—an election under § 754 to adjust the basis of the partnership’s assets by the § 731 gain or loss recognized by the redeeming partners.³ If the § 731 amount is a gain, the basis of the partnership assets is increased, thus reducing the built-in gain in assets left behind by the redeemers. If, on the other hand, the § 731 amount is a loss, the basis of the partnership assets is decreased, thus reducing built-in losses in assets left behind by the redeemers. The § 754 basis adjustment is thus intended to alleviate the burden of gain duplication (and symmetrically reduce the benefit of loss duplication) for partnerships.⁴

NO MORE STUFFING

Hedge funds have historically used aggregate allocations sanctioned by Treasury for securities partnerships under Reg. § 1.704-3(e)(3). This allows hedge funds to simplify tax accounting complexities by keeping aggregate revaluation accounts for their partners rather than having to track unrealized gains and losses for every partner, for every asset, at the lot level.

In Exhibit 1, we provide a sense of the bookkeeping complexities under lot layering. For the same reason that aggregation was thought necessary by Treasury for securities partnerships (i.e., administrative impossibility or impracticability of tracking large volumes of asset lots), and as a collateral consequence of this aggregate bookkeeping, such partnerships are generally unable to implement § 754 basis adjustments to properly adjust the cost bases of specific partnership assets to avoid the gain or loss duplication for which § 754 was designed.

³Under § 734(b), if the § 731 amount is a loss exceeding \$250,000, the basis adjustment by the loss is mandatory.

⁴Without the § 754 election gain duplication for a continuing partner is temporary: A higher realized gain today (when the partnership asset with duplicated gain is sold) is offset by a lower § 731 gain upon the later redemption of the continuing partner(s). However, this eventual justice is of little solace to a partner who is planning to be in the partnership for the long haul.

Without a § 754 election to combat gain duplication, hedge funds have historically used special allocations of current-year realized gains (or losses) to redeeming partners, often referred to colloquially as *stuffing* allocations. Conceptually, the idea behind stuffing allocations is fairly straightforward. Suppose a partner redeems at a gain. The partnership would seek to make a special allocation of currently available realized gains in the partnership to the redeemer in the amount of his or her otherwise § 731 gain. In such a way, the redeemer's total recognized gains do not change (rather than having a § 731 gain, the redeemer will now have the same amount of allocated realized gain from the partnership, and no § 731 gain), but realized gains otherwise allocable to the continuing partners are correspondingly reduced. So, the otherwise duplicated gains for the continuing partners become simply a deferral of their own gains. The exact opposite applies if a partner redeems at a loss—he or she is specially allocated existing realized losses, thereby reducing his § 731 loss and current realized losses allocable to the continuing partners.

Although there is no statutory or regulatory authority sanctioning such special allocations (nor are there any rules prohibiting them), they are commonly used by hedge funds in order to reduce the duplication of gains and losses otherwise left behind by redeeming partners.

Uncertain Amount of Stuffing Allocations

At first blush, stuffing allocations seem to neatly diminish the adverse effect of redemptions on the continuing partners. Note, however, that to successfully stuff the redeemers, the partnership must have a sufficient amount of realized gains and losses in the partnership. These realized gains and losses bear little to no relationship to the § 731 gains or losses of the redeemers. As a result, a sufficient amount of gains or losses with which to stuff the redeemer is far from guaranteed. In the case of an insufficient amount of realized gains or losses in the partnership, stuffing might provide only partial relief to the continuing partners. Thus, the practice of stuffing, although in many ways appealing in theory, is by no means certain to achieve its intended result in all practical situations.

Uncertain Character of Stuffing Allocations

Stuffing allocations are made using short-term and long-term gains and losses that the partnership

happened to realize and might bear no relationship to the redeeming partner's tenure in the partnership. As a result, a long-term partner might be allocated a mix of long-term and short-term capital gains, depending purely on what type of realized gains are available in the partnership.⁵

Moreover, if the partnership stuffs on an annual basis, partners redeeming in the middle of the year will have no ability to know the character of capital gains or losses that will be stuffed to them until the end of the year. Because § 754 basis adjustments, which become possible under lot layering, eliminate the need for stuffing allocations, such character uncertainty is eliminated: Long-term redeeming partners recognize long-term § 731 gain or loss on their investment in the partnership, and short-term redeeming partners recognize short-term § 731 gain and loss.

Uncertain Authority of Stuffing Allocations

Over the years, some criticism has been levied at stuffing allocations. Although stuffing is commonly practiced by hedge funds, arguably to resolve disparities between economic gains and losses and tax allocations in the absence of § 754 basis adjustments, there is seemingly no direct authority in the tax law supporting this approach. Nor is there consensus in the legal community as to whether support for stuffing allocations can be inferred from the relevant Treasury regulations.

Exhibit 2 contains excerpts from two articles authored by different tax legal experts published in *Tax Notes*. These nationally recognized tax professionals, having analyzed the same Treasury Regulations, have come to polar opposite conclusions! Whereas Ladin, Lowy, and Woods (2008) took the view that stuffing allocations “further the purpose” (p. 944) of the tax law, Needham (2013) concluded that stuffing allocations “violate one of the most basic tenets” of the tax law (p. 755).

In contrast to stuffing allocations, lot layering, in conjunction with § 754 basis adjustments, provides a statutorily stipulated relief for continuing partners from the gains accumulated by redeeming partners.

⁵Whereas the category—long-term or short-term—of § 731 gain or loss depends only on the redeeming investor's holding period, the category of stuffing allocation gains or losses depends on both the composition of gains and losses realized by the partnership and the method the partnership uses to allocate them. This, for example, might lead to a situation where long-term investors are stuffed with predominantly short-term gains upon their redemption.

EXHIBIT 2

Expert Opinions on whether Stuffing Allocations are Supported by the Tax Law

Stuffing Supported by the Law (Ladin, Lowy, and Woods 2008)

“Hedge Fund Stuffing Allocations represent a reasonable interpretation and application of the Aggregation Approach provided for in the section 704(c) regulations. Stuffing Allocations further the purpose of section 704(c) by preventing a shifting of tax consequences regarding revaluation gain or loss from a departing partner to continuing partners.” (p. 944)

Stuffing not Supported by the Law (Needham 2013)

“... I have been unable to identify any plausible legal support for these allocations, whether under the special aggregation rules governing securities partnerships or the general asset-by-asset rules governing other partnerships. To the contrary, the practice appears to violate one of the most basic tenets of subchapter K, which is that tax follows economics.” (p. 755)

(UNAVOIDABLE) LOSS OF PRECISION CAUSED BY REDEMPTIONS

The Regulatory Glitch

Unfortunately, given the elegant potential for precision of lot layering with remedial allocations and § 754 basis adjustments, it is disappointing to discover that this precision is likely to diminish following redemptions. The elective § 754 basis adjustment, although solving the gain duplication issues discussed previously, has an important limitation. When applying basis adjustments of § 754 to the partnership’s assets, it is required to use the rules of § 755 and Reg. § 1.755-1 to apportion that adjustment among its assets. Suppose a partner redeems at a § 731 gain resulting from gains in some assets and smaller losses in other assets. Although it would be economically sensible to fully increase the cost bases of the lots in which the redeemer had unrealized gains and decrease the cost bases of the lots in which the redeemer had unrealized losses (which would be fully consistent with the remedial method discussed for allocating realized gains and losses among partners), this is not how the mandatory basis adjustment rules under § 755 and Reg. § 1.755-1 work. Reg. § 1.755-1 requires the partnership to ignore the actual (gross) components of the redeemer’s § 731 gain and, instead, spread the basis adjustment among the partnership assets first based on the partnership’s relative unrealized gains in appreciated assets to the extent thereof and then to the extent of any excess based on the relative fair market values of all assets. Let’s be clear: Neither of these two things bear any relation to economic gains and losses experienced by the redeemer (and on which he or she pays tax

under § 731). The cumbersome and uneconomic rules of Reg. § 1.755-1 are summarized in the decision flow chart in Appendix A.

An astute reader might notice that these uneconomic adjustments in asset cost bases will cascade further by creating disparities between partnership unrealized gains and losses in the underlying assets and partner unrealized gains and losses in those assets for tax purposes. As a result, the required basis modifications to these unrealized gains and losses at the partnership level as required by Reg. § 1.755-1 will unavoidably disrupt the exact match between the economic gains and losses experienced by the continuing partners in each asset and their unrealized gains and losses in those assets for tax purposes.

Is there a way to avoid all of these uneconomic adjustments and instead simply update the bases of partnership assets exactly by the actual gains and losses the redeeming partner experienced in each of these assets? The reader might observe that, in so doing, the tax adjustments would exactly match economic reality, and we agree wholeheartedly. However, in its current form, Reg. § 1.755-1 leaves no discretion for a partnership to choose a different method—even if possibly more precise and economically reasonable—beyond its prescribed formula. Hence, suspending any economic intuition, partnerships must follow the regulations to make the basis adjustments.

Economic or not, the Reg. § 1.755-1 method for basis adjustments is the only one currently permissible for avoiding duplication, for tax purposes, of gains and losses left in the partnership by the redeemers. As a result, partnerships using lot layering are, nevertheless, likely to be inclined to make the elective § 754 basis adjustment. After all, half a loaf is better than none!

If We Could Change the Tax Laws Governing the § 754 Basis Adjustments

Now that we've criticized the uneconomic shortcomings of § 754 basis adjustments under Reg. § 1.755-1, the reader may ask whether we have any thoughts on a solution to this problem. Indeed we do, albeit implementing it would require a meaningful change to the language of Reg. § 1.755-1. As a result, a reader interested in a near-term practical solution and not in our musings will be disappointed—under current law, we don't see a better alternative to what we have already described.

Here is our proposed approach. § 755 of the Internal Revenue Code leaves it to Treasury to decide how to adjust the cost basis of the partnership property for built-in gains and losses left behind by redeeming partners. Treasury's current regulation governing the basis adjustment, Reg. § 1.755-1, looks only to the net gain or loss of a redeemer. However, this net gain or loss is composed of unrealized gains and losses in individual assets (and in individual lots of those assets) economically owned by the partner prior to redemption. These unrealized gains and losses should correspond precisely to the economic gains and losses that the partner has experienced in those assets. As a result, rather than netting the redeeming partner's overall economic result into one gain or loss number and then spreading that net amount among partnership's lots based on arbitrary indicia, Reg. § 1.755-1 could allow precise adjustments to individual lots held by the partnership that are perfectly reflective of unrealized gains and losses experienced by the redeemer in those specific lots.

Doing this on a practical level shouldn't be hard. The partnership can use the information already available to it under lot layering to decompose the § 731 amount into unrealized gains and losses of the redeeming partner in specific underlying assets. It can then reflect all gains in assets contributing to the § 731 amount as an increase in basis of those specific assets, and losses in assets contributing to the § 731 amount as a decrease in basis of those specific assets. For example, if a \$50 § 731 gain resulted from an unrealized gain of \$75 in asset A and an unrealized loss of \$25 in asset B, the partnership would increase its cost basis in asset A by \$75 and decrease its cost basis in asset B by \$25. Is there even an argument that this makes much more sense economically than making one net adjustment to the basis of asset A by \$50, as would be required under the current version

of Reg. § 1.755-1? An important benefit of such a precise adjustment as we propose is that after the adjustment, all the partner and partnership tax accounts would look as if the redeeming partner had never been there, because built-in gains and losses of the redeemer in individual lots have been appropriately removed from the partnership assets. The partnership's and partners' unrealized tax gains and losses in the assets would match exactly without any further adjustments to the partners' unrealized gains, both before and after redemptions.

Three main benefits can be derived from this adjustment as we propose it. First, computationally, it is much simpler than the longwinded conditional logic of current Reg. 1.755-1, which we attempt to summarize in the diagram in Appendix A. Second, there is no need to invent reasonable ways to adjust the continuing partners' unrealized gains and losses in each asset to match the partnership's unrealized gains and losses in those assets. Finally, because continuing partners' unrealized gains and losses in the underlying assets are not affected, they continue to precisely reflect the economic gains and losses of the partners in these specific assets. As a result, when a given asset is disposed of any time after a redemption, the continuing partners' realized tax gains and losses will still exactly match their economic gains and losses in the asset (and are no different from what they otherwise would have been absent a redemption).

COSTS AND BENEFITS OF LOT LAYERING

Conceptually, lot layering, combined with a § 754 election, yields several important benefits:

1. *Precision.* With lot layering, economic gain and loss translate directly into unrealized gain or loss specifically and precisely for those investors who benefited from the economic gain or suffered from the economic loss. This is done with high precision on a lot-by-lot level. When partnership assets are later sold, and unrealized gains and losses become realized, they are allocated in precise quantities to those investors who experienced the economic gains and losses in those assets.
2. *Statutory relief to continuing partners.* The § 754 basis adjustment provides relief to continuing partners from the unrealized gains accumulated by redeeming partners that otherwise would remain stored in partnership assets, with the caveat that

the precision of this relief on a lot-by-lot level is diminished because of the constraints imposed by Reg. § 1.755-1. On an overall level, however, across time, all assets, and all investors, the relief is generally complete.

3. *No need for stuffing allocations.* When a partner redeems, no post hoc manual intervention, characteristic of stuffing allocations, is required.

Despite these benefits, lot layering has significant costs (which explains why most hedge funds have not used this method, historically):

1. *Data management burden.* Aggregate allocation methods require maintaining data for K partnership positions and N partnership investors, resulting in data of a K -plus- N dimensionality. In contrast, lot layering relies on storing data at a K -times- N level. As a result, for securities partnerships that, at any point in time, can have holdings in thousands of lots across hundreds of investors, lot layering necessitates maintaining substantial quantities of data. Such quantities of data require data management capabilities that historically have not been contemplated by hedge funds' tax preparers. Automated data storing procedures that commit data to efficient databases are a must. Any attempts to manage such data in spreadsheets, which up until recently have been the workhorse of the industry, are destined to fail.
2. *Computational complexity.* Lot layering requires not only storing large volumes of data but also retrieving and passing data into calculation engines quickly and error-free. Large-scale calculations involving the updating of the unrealized gains and losses tracking ledger on a lot-by-lot and partner-by-partner basis are required for every single transaction. In addition, partnership cost bases in positions need to be adjusted every time there is a redemption from the fund and, as we have seen, the constraints imposed by Reg. § 1.755-1 result in a complex and cumbersome basis adjustment algorithm (which, under current regulations, are difficult to reconcile to other account configurations following redemptions). As a result, all the calculations must be fully automated and executed by efficient and robust computer code.

3. *Error correction challenges.* Similar to automated data management and computations, when a validation test fails, error correction must be fully automated. Given the sheer quantity of data and calculations, a human would be unable to search for errors manually. Fund administrators and their tax preparers, therefore, will have to rely on robust and trustworthy software rather than human grunt work in spotting and correcting errors.
4. *Increase in tax preparation fees.* As with any new technology, lot layering technology requires substantial initial investments by fund administrators and/or tax preparers. The cost of such investment can be expected to be passed on to the funds that are using this newly developed technology, in the form of additional fees or otherwise. Therefore, even for administrators and preparers that overcome the technological challenges outlined, their clients will likely experience increased costs, at least until lot layering technology become robust and wide spread.

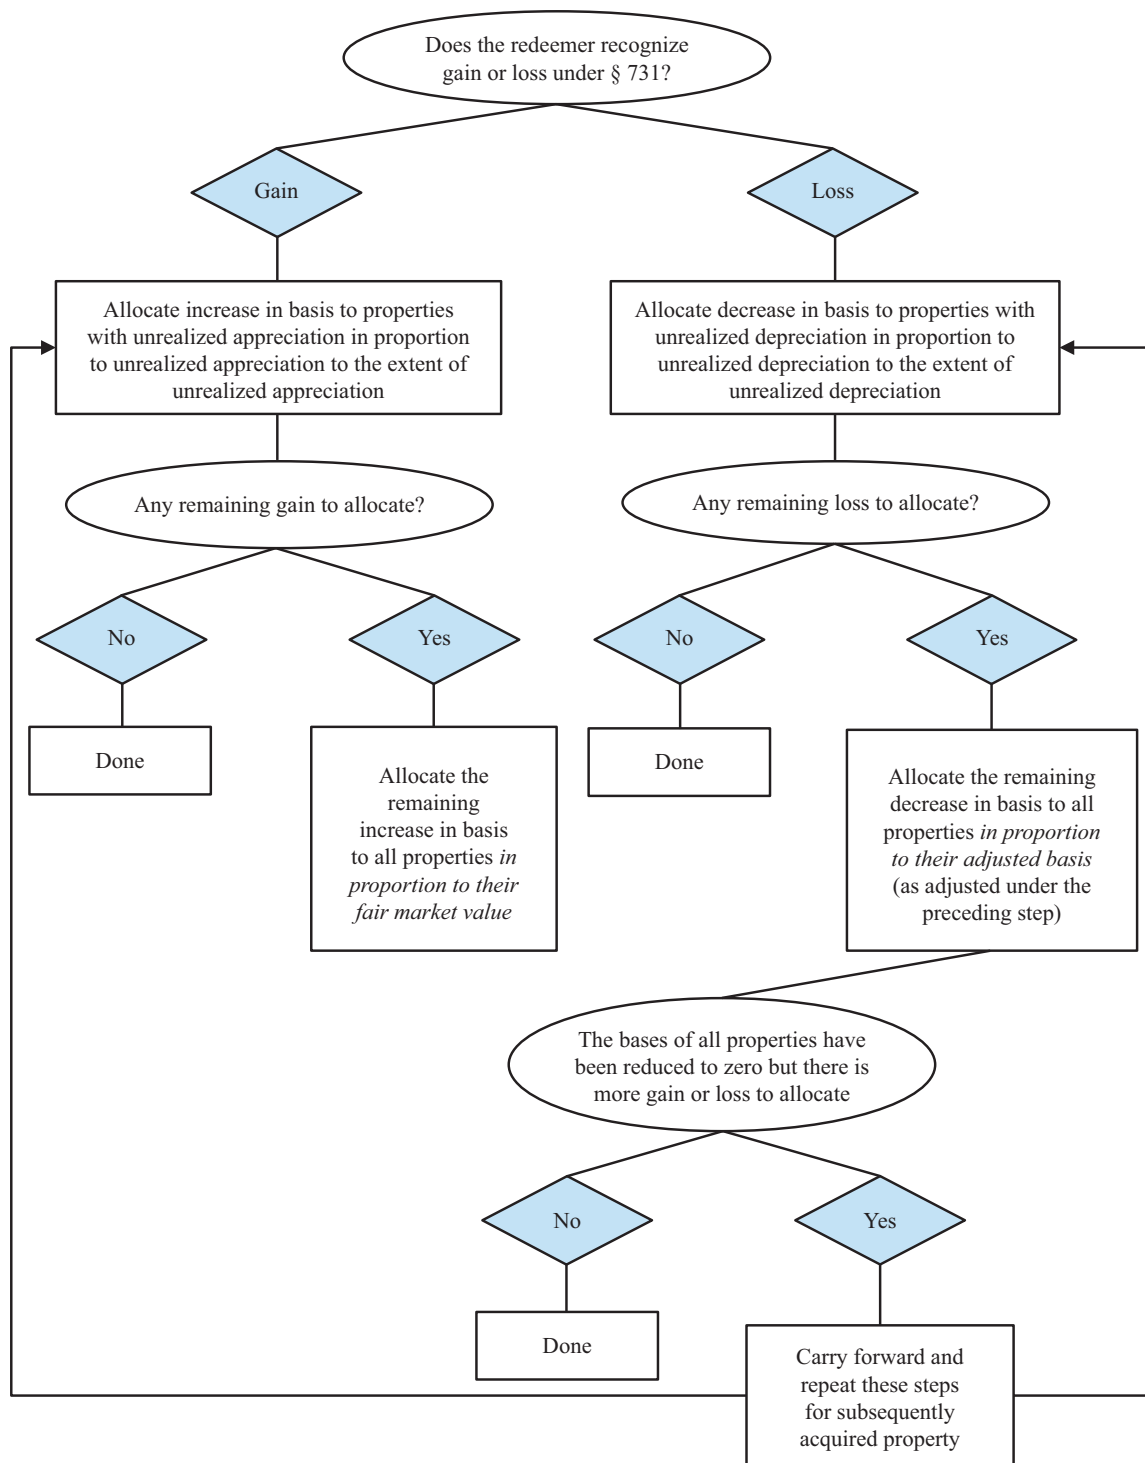
CONCLUSION

Lot layering, combined with remedial allocations and § 754 basis adjustments, is the most precise approach to partnership allocations. We explain how remedial allocations help lot layering achieve its precision. We also explain that some of the precision is lost upon redemptions because of the cumbersome and uneconomic basis adjustment method mandated by current Treasury regulations. We propose possible changes to these regulations that we believe could fully remedy this problem. However, precisely because our solution requires changes in tax law, for now it remains purely theoretical. Despite this unavoidable deficiency under the current regulations, we believe that lot layering nevertheless aligns tax and economics more closely than do aggregation methods. Moreover, lot layering eliminates the need for stuffing allocations upon redemptions, replacing them with § 754 basis adjustments.

The benefits of lot layering, however, come at a cost of significant data management and computation burdens, at least in the present day. In addition, these technology-related burdens can be expected to lead to significantly higher costs charged by tax service providers. For many hedge funds, these costs of maintaining lot layering might outweigh the benefits.

APPENDIX A

REG. § 1.755-1 GENERAL STRUCTURE FOR § 754 BASIS ADJUSTMENT



ACKNOWLEDGMENT

We are grateful for the comments and suggestions by Ronen Israel, Bryan Johnson, Ted Pyne, and Roxana Steblea-Lora.

REFERENCES

Ladin, B. E., J. M. Lowy, and W. S. Woods II. 2008. "Hedge Fund Stuffing Allocations: A Path Through the Maze." *Tax Notes* 121 (November 24): 925–944.

Needham, A. W. 2013. "The Problem with Stuffing Allocations." *Tax Notes* 141 (November 18): 737–756.

Sosner, N., P. Balzafiore, and Z. Du. 2018. "Partnership Allocations and Their Effects on Tax-Aware Fund Investors." *The Journal of Wealth Management* 21 (1): 8–17.

Disclosure

AQR Capital Management is a global investment management firm, which may or may not apply similar investment techniques or methods of analysis as described herein. The views expressed here are those of the authors and not necessarily those of AQR. The reader should conduct his or her own analysis and consult with professional advisors prior to making any investment decisions. AQR is not a tax advisor. This material is intended for informational purposes only and should not be construed as legal or tax advice, nor is it intended to replace the advice of a qualified attorney or tax advisor.

To order reprints of this article, please contact David Rowe at d.rowe@pageantmedia.com or 646-891-2157.

ADDITIONAL READING

Partnership Allocations and Their Effects on Tax-Aware Fund Investors

NATHAN SOSNER, PHILIP BALZAFIORE, AND ZHENDUO DU
The Journal of Wealth Management
<https://jwm.pm-research.com/content/21/1/8>

ABSTRACT: Limited partnerships are attractive investment vehicles for investors because, as limited partners, investors cannot lose more than their invested capital despite the leverage of the partnership's portfolio. Consistent with this, the availability of tax losses to a limited partner is also more limited as compared with a separate account investor. Understanding this limitation is particularly important for investors in tax-aware funds, which tend to allocate

net tax losses. In addition, investors in funds structured as limited partnerships are affected by laws and regulations governing partnership allocations of gains and losses realized by the fund portfolio. This study outlines certain relevant principles of "securities partnership" accounting and shows how these principles apply to investing in tax-aware funds structured as limited partnerships. The authors argue that the laws and practices of partnership accounting align tax results with the economic outcomes of the investors and ensure that new investors do not materially suffer from unrealized gains accumulated in a tax-aware fund. The authors' conclusions are illustrated with a simple stylized model.

Will Hedge Fund Investors Start Asking for Tax Alpha? Can Hedge Fund Managers Deliver It?

ROBERT KIM, EDWARD H. DOUGHERTY, AND MIRIAM KLEIN

The Journal of Wealth Management

<https://jwm.pm-research.com/content/13/4/44>

ABSTRACT: Tax efficiency, a common topic in the mutual fund world, has never been a big priority for hedge funds. Although U.S. taxable investors would benefit if hedge fund managers were more tax aware, the offshore and tax-exempt investors putting up the greater share of the capital haven't required it. But looming higher tax rates, combined with a greater emphasis by hedge funds on client service and retention, are likely to change the dynamic. Tools and techniques for expressing investment ideas in a more tax-efficient manner are available. The task for managers will be to employ them as appropriate while keeping investment strategy firmly in the driver's seat.

The Tax Benefits of Relaxing the Long-Only Constraint: Do They Come from Character or Deferral?

NATHAN SOSNER, STANLEY KRASNER, AND TED PYNE

The Journal of Wealth Management

<https://jwm.pm-research.com/content/21/4/10>

ABSTRACT: In this study, we propose a decomposition of the total tax benefit (or liability) of a strategy into what we define as character and deferral components. Our decomposition is mathematically straightforward and intuitive, and it allows for a quick and informative assessment of tax benefits of different tax-aware strategies without modeling various investor-specific situations. We use this character-deferral decomposition to identify the source of tax benefits resulting from relaxation of the long-only constraint. Our empirical evidence shows that, for tax-aware strategies, relaxing the long-only constraint results in a drastic increase in their tax benefits, in particular owing to an increase in the character benefit. We conclude that tax-aware relaxed-constraint strategies are more attractive to taxable investors than their long-only counterparts.