

Misconceptions about Tax-Aware Long-Short: Misleading Parallels with Direct Indexing, Part I

AQR · 2026-09-21 · English original

Interest in tax-aware long-short investing (TA LS) has grown rapidly among investors, managers, advisers, and the media alike. ^{1 1} Close See More Attention on Tax-Aware Investing Is a Good Thing for the Industry and Investors (AQR Tax Matters, 2026). It is understandable that many draw comparisons to a more familiar and long-established tax-aware approach: direct indexing.

Indeed, direct indexing and TA LS can be thought of as distant cousins. Both are typically equity strategies. Both incorporate tax considerations into trading decisions. And both are commonly managed in separately managed accounts (SMAs). But that is largely where the similarities end.

The comparison becomes problematic when casual observers ^{2 2} Close By "casual observers," we refer to those who learn about these strategies from secondary sources but do not study them sufficiently to understand how they actually work. extrapolate from direct indexing's investment approach to form a mental model of how TA LS works. This can give rise to a number of misconceptions about the design and mechanics of TA LS, leading to a materially incomplete understanding of the strategy. We aim to address these misconceptions in a series of posts, of which this is the first.

Misconception 1: TA LS harvests losing positions and replaces them with close substitutes

The faulty extrapolation goes like this: direct indexing sells a losing position to harvest a tax loss and then replaces it with a close substitute – a stock much like the one it just sold – to preserve benchmark exposure without running afoul of tax rules. ^{3 3} Close On why a genuine strategy must make sense before tax—and the pitfalls of trades built purely around close substitutes—see Separating the Wheat from the Chaff (AQR Tax Matters, 2024). TA LS, the reasoning goes, must do the same, except that it also has short positions to harvest losses on. So, TA LS simply trades losing positions across its long and short books and replaces them with close substitutes.

This all seems to make sense, except that this is not how TA LS works. The two approaches trade for different reasons because they have very different investment objectives.

A direct-indexing strategy seeks to closely track a passive benchmark and deliver its returns more tax-efficiently. By contrast, a TA LS strategy is an active strategy that seeks to materially outperform a benchmark after fees, transaction costs, and financing costs. As a result, the two approaches trade for fundamentally different reasons:

- Direct indexing trades to harvest losses and match a benchmark: Its goal is to track a passive index as closely as possible while harvesting losses along the way. When it sells a position, it looks for a close substitute precisely because its objective is to minimize tracking error and keep the portfolio "hugging" the benchmark.

- TA LS trades to follow the alpha model: Its goal is to pursue pre-tax alpha – to outperform, not to track, a benchmark. When a position is liquidated (potentially at a loss), it is mainly because the alpha model no longer "likes" it. The strategy doesn't seek a lookalike replacement because a close substitute would likely be another low-expected-return position. Instead, it replaces the position with one the model "likes" more: the highest-expected-return position available, subject to constraints, such as industry exposures and position sizes, and penalties, such as transaction costs and shorting fees.

Misconception 2: The long and short portfolios need to be similar/correlated to hedge each other

The idea behind the "close substitute" extrapolation sometimes takes a broader form: the long and short active portfolios must be very similar to each other – effectively close substitutes – in order to hedge one another properly. This way, the investor can harvest losses while limiting the risk of underperforming the benchmark.

Again, this sounds quite reasonable if one applies the rationale of direct indexing. However, this is not what TA LS strategies do. In TA LS, the long portfolio loads positively on factors that the manager expects to be rewarded with excess returns, while the short portfolio loads negatively on those factors. As a result, the long and short portfolios should have opposite exposures to factors associated with expected returns.

For example, if the portfolio manager views quality as a factor rewarded with returns, the long portfolio should hold high-quality stocks, while the short portfolio should hold low-quality stocks. Clearly, portfolios of high-quality and low-quality stocks are not close substitutes. Nonetheless, the manager would still be comfortable holding such different long and short portfolios because those differences reflect the risks for which the manager expects to be rewarded with returns.

Seen this way, the ability to realize losses on both long and short positions is not the purpose of the TA LS investment process, but rather a byproduct of it. Saying that the long-short portfolio is constructed to harvest losses more effectively gets the logic exactly backwards.

Source: <https://www.aqr.com/Insights/Research/Tax-Aware-Investing/Misconceptions-about-Tax-Aware-Long-Short-Misleading-Parallels-with-Direct-Indexing-Part-I>