advisorXLS®

RESEARCH · 05 | MODEL OPERATIONS

Tax-loss harvesting starts with the model you choose

We ran the same harvesting settings on ten model portfolios over the same five years. The model decided more of the after-tax result than any harvesting setting did.

AUTHOR	CATEGORY	PUBLISHED	READ TIME
AdvisorXLS Research	Model Operations	November 27, 2026	5 min read

Most conversations about tax-loss harvesting start with the harvest: how often to scan, what loss to act on, what to hold during the wash-sale window. We spent months testing those choices. The largest difference we found came from somewhere else: the model the harvesting runs on.

Same settings, ten models, very different results

We took ten model portfolios and ran each one twice over the same five years, once with harvesting and once without. The harvesting settings were the same on every model. So were the starting account, the tax rates and the trading costs.

SIMULATED, HYPOTHETICAL, $1,000,000 ACCOUNT	AFTER TAX, HARVESTED VS. UNHARVESTED
Long-only models (3)	-$27,500 to +$34,400
Long/short models with one short sleeve (3)	+$71,100 to +$89,500
Long/short models with another short sleeve (3)	+$5,000 to +$26,700
Long/short model with its own short book (1)	-$607,600

Five years to mid-2026, account sold on the last day, top federal and 13.3% state rates, losses used against the client's other gains. Simulated results are not a guarantee of future results.

Same settings, same years. One model gained close to $90,000 from harvesting, and another lost more than $600,000. Before you ask how well a harvest works, ask what it is working on.

> The spread between models was much larger than the spread between harvesting settings on any one model.

The model decides which gains your client realizes

Every time a model changes its weights, the account sells something. If that sale is at a gain, your client owes tax, and the harvested losses go to offset it. A model that trades often creates a steady supply of gains to offset. A model that rarely changes leaves the losses with nothing to do, and they build up as a carryforward.

So the same harvest can be worth a lot on one model and little on another. On one published run, the model alone turned over 140% of the account a year. The harvested version, which keeps its winners, turned over 34%, so it realized far fewer gains for its losses to offset.

A setting that works on one model can fail on another

We tuned settings on one model and tested them on a second. They often produced a negative tax alpha on the second model. With identical settings over the same six and a half years, a US model earned 0.67% a year in tax alpha and a global model earned 0.18%.

The single change that helped most was also different on almost every model:

- On one long-only model, letting the account sell gains inside a 10% band.
- On two broad index models, sending sale proceeds to the 25 largest names.
- On the long/short models, a different loss trigger, a different substitute fund, or a wider gains band.

There is no one setting that suits every model. Test the settings on the model your client actually owns.

Structure matters as much as holdings

A long/short model behaves differently. A rule that works well on the long side, "never sell a winner", can do damage on the short side. Applied to the shorts, it kept the winning short positions open. One account drifted to shorts worth 2.55 times its value and stayed below a typical broker's maintenance margin on 462 days of the test. With the same rule on the longs only, every long/short model we tested did better.

Size and breadth matter too. A book with 1,277 lines on a $1,000,000 account holds about $783 in each line. At that size, a minimum-dollar loss trigger fires on nearly every lot, and the account trades far more than it needs to.

The model has to earn its place before tax

Harvesting cannot rescue a model that trails its benchmark. Before tax, the long-only models in our study ranged from 3.86% a year ahead of the S&P 500 to 0.44% a year behind it. After tax, only the three long-only models finished ahead of simply holding an S&P 500 fund, and one of those by only $3,800 on $1,000,000.

Choose the model first, on its own merits. Then fit the harvesting to it, and measure what the harvesting adds to that model, not to a benchmark.

How AdvisorXLS handles the model

- **Your model stays in charge.** Harvesting runs on the model you choose: your own, one from the marketplace, a blend, or an ETF rebuilt from its public filings.
- **Results against your model.** Every run shows the benchmark, your model without harvesting, and your model with harvesting, side by side.
- **Rebalancing on your terms.** You choose whether an account rebalances on drift, on a new model version, or on either.
- **Settings tested on your model.** The optimizer tests the settings on the model your client owns, over the periods you choose.

Important information. This article is for financial professionals and is for information only. It is not tax, legal or investment advice, and it is not an offer of any security or service. The results above are simulated and hypothetical. They were produced by AdvisorXLS with the benefit of hindsight, on a $1,000,000 starting value, with assumed top federal tax rates, a 13.3% state rate, trading costs of 2 basis points and taxes paid from outside the account. They do not reflect any client account, actual trading, or advisory fees. Actual results will differ. Hypothetical performance has inherent limits, and advisors who use it in communications must meet the SEC Marketing Rule.

// topics: Model Portfolios Tax-Loss Harvesting Direct Indexing