Harvesting Tax Losses to Minimize Substitution Errors

Frederic Khoury
Dr. Vinay Nair[1]

[1] F. Khoury is an Investment Director at 55ip. V. Nair is the Founder and Chairman of 55ip, and a visiting professor at The MIT Sloan School of Business and The Wharton School at the University of Pennsylvania.

Contents

The rationale for harvesting unrealized losses 2
What happens to pre-tax returns? 2
Minimizing substitution errors 2
Proxy selection 3
Case study 1: IWM 3
Case study 2: GLD 5
Conclusion 6

The rationale for harvesting unrealized losses

A central feature of tax optimization is to be conscious of how the pre-tax returns are split between the short-term realized returns and the long-term unrealized returns. Tax optimization attempts to manage the portfolio and its trading to minimize the short-term realized gains while maximizing the longer term unrealized gains. One of the main techniques used to do this is to "harvest" unrealized short-term losses.[2] This is achieved by selling such positions and converting them to realized short-term losses thus decreasing the short-term capital gains. Assuming such tax optimization does not modify the pre-tax returns, this maximizes the post-tax returns since longer term capital gains are taxed at a lower rate than short-term capital gains.

What happens to pre-tax returns?

However, tax optimization does modify the pre-tax returns compared to an identical investment strategy that is not being optimized for taxes. There are at least two aspects that create a difference between the pre-tax returns of the tax managed and the unmanaged strategy.

First, converting short-term unrealized losses to realized losses requires substituting a loss generating position with a similar but new position. The substitution should of course not violate the wash sales rules.[3] Given that the substitution uses a similar, but not identical exposure, the new position will behave differently from the original intended position. In many cases, this substitution error can be material.

Second, converting unrealized to realized losses requires trading. As a result, there are costs associated with trading and slippage (since decision time and trading time are not identical in most cases).

These features create a difference between the pre-tax returns of tax managed and unmanaged strategies. Tax optimization changes not only the post-tax returns but the pre-tax returns as well! Managing this deviation is central to tax optimization.

Minimizing substitution errors

In tax harvesting algorithms, when taking a loss on an ETF, we replace it with another ETF or a portfolio of ETFs. We call this replacement a proxy. Further, if we harvest the losses from the proxy, we will replace it with a secondary proxy. Each substitution generates substitution error that starts creating differences from the original intended portfolio.

---

[2] Tax aware strategies also try to realize a long-term gain to reset the basis, and increase the chances of generating short term losses in the future. There are other approaches to optimize taxes unrelated to the investment **strategy such as asset location, giving etc.**

[3] Some legalese here that is important: the IRS defines a wash sale as one that occurs when an individual sells or trades a security at a loss, and within 30 days before or after this sale, buys a "substantially identical" stock or security, or acquires a contract or option to do so.

These substitution errors can be material. For example, if the standard deviation of the substitution error is 1% on a monthly basis, the pre-tax returns of the tax managed strategy could differ from the benchmark by over 3.4% in almost one-third of the cases! This is generally larger than long term expected tax benefits. As we'll see in the case studies below, some common ETFs that are part of asset allocation portfolios create such differences. As ETFs become more granular or capture non-equity exposure, these problems are magnified.

Proxy selection

The industry standard to harvest losses on a given ETF is to choose another ETF with the same dollar amount. For example, Wealthfront, a leading robo-advisor in their tax harvesting white paper replaces XLE with VDE. This approach is not optimal as it leads to high tracking error - the correlation between XLE and VDE, for example, is only 73%. Further, a basket of ETFs might be better than a single ETF. Finally, the best proxy can change over time.

Our approach is to be open to a portfolio substitution when its replicating properties are superior. Our basket selection is two-step process. Initially, we will manually select a handful of liquid ETFs that are fundamentally related to the ETF considered. Then we use a proprietary machine learning based algorithm to dynamically select proxy baskets. This enables proxy selection that optimizes substitution error and is dynamic.

In a large sample analysis, we analyzed the top 100 most liquid ETFs, over a period of 10 years. For every single ETF, at every time period we used the past 4 years to choose the best proxy. The industry standard method chooses one ETF, our method chooses one or multiple. We compared the tracking error of the two solutions on the following 2 years. **We found that the tracking error using a naïve single ETF (the one with the highest correlations to the primary position) is, on an average across 100 most liquid ETFs, 1.6 times higher than the optimized substitution baskets we identify.** Here are two specific examples.

Case study 1: IWM

IWM is the Russell 2000 ETF. In February 2016, the proxy generated by our approach was IWO: iShares Russell 2000 Growth (51%) and IWN: iShares Russell 2000 Value (49%). Typical industry proxies are 100% to IJR, VB, VXF, or IJT. In our view, firms that use ETFs based on the same index will be violating wash sales rules. Below, we compare the differences in the substitution error.

As can be seen below, our approach is superior as the correlation of our proxy made of IWO and IWN replaces IWM with a correlation close to 1, and more importantly tracking/substitution error (TE) is more than 9 times lower. This offsets any additional fees in this case. Our proxy has the lowest tracking error at 26 bps versus more than 220 bps for the other ETFs. Historical

analysis shows that this superior replication is not only true in historical in-sample metrics but also in out-of-sample performance.

	IWM	IWO + IWN	IWO	IWN	IJR	VB	VXF	IJT
Name	ISHARES RUSSELL 2000 ETF	Proxy	ISHARES RUSSELL 2000 GROWTH	ISHARES RUSSELL 2000 VALUE	ISHARES CORE S&P SMALL-CAP	VANGUARD SMALL-CAP ETF	VANGUARD EXTENDED MARKET ETF	ISHARES S&P SMALL-CAP 600 GR
Current Average Daily Volume Past 6 Month (in $)	4,097,678,255	162,398,136	125,237,208	201,075,837	207,200,146	100,460,251	42,947,493	28,957,763
Expense Ratio	0.20%	0.25%	0.25%	0.25%	0.07%	0.08%	0.09%	0.25%
Correlation with IWM	100.00%	99.98%	98.44%	97.90%	98.91%	98.30%	97.31%	98.02%
Mean	12.05%	12.05%	14.64%	9.45%	13.64%	12.68%	12.95%	14.90%
Stdev	14.20%	14.13%	15.20%	13.67%	13.27%	12.75%	12.60%	13.42%
TE vs IWM		**0.26%**	**2.78%**	**2.91%**	**2.23%**	**2.88%**	**3.49%**	**2.85%**
Start	Jan-13	Jan-13	Jan-13	Jan-13	Jan-13	Jan-13	Jan-13	Jan-13
End	Dec-15	Dec-15	Dec-15	Dec-15	Dec-15	Dec-15	Dec-15	Dec-15

Case study 2: GLD

GLD is the SPDR Gold Trust ETF. In February 2016, the proxy generated by our approach was GDX: Market Vector Gold Miners and SLV: iShares Silver Trust. The exact mix of the combination changes over time. Typical industry proxies are 100% to GDX, SLV, GDXJ, RING, or PPLT. In our conservative view, firms that use ETFs based on the same index will be violating wash sales rules (e.g. using IAU to replace GLD). Below, we compare the differences in the substitution error.

As can be seen below, our approach is superior as the correlation of our proxy made of GDX and SLV replaces GLD with a correlation of 91.32%, and more importantly tracking/substitution error (TE) is more than 2 times lower. This offsets any additional fees in this case. Our proxy has the lowest tracking error at 652 bps versus more than 1400 bps for the other ETFs. Historical analysis shows that this superior replication is not only true in historical in-sample metrics but also in out-of-sample performance.

	GLD	GDX + SLV	GDX	SLV	GDXJ	RING	PPLT
Name	SPDR GOLD SHARES	Proxy	VANECK VECTORS GOLD MINERS	ISHARES SILVER TRUST	VANECK VECTORS JUNIOR GOLD M	ISHARES MSCI GLOBAL GOLD MIN	ETFS PHYSICAL PLATINUM SHRS
Current Average Daily Volume Past 6 Month (in $)	1,051,581,126	428,017,577	1,722,471,182	153,355,991	767,520,552	4,627,294	5,130,025
Expense Ratio	0.40%	0.28%	0.53%	0.50%	0.56%	0.39%	0.60%
Correlation with GLD		91.32%	84.31%	83.81%	86.22%	83.16%	72.13%
Mean	-14.25%	-14.43%	-32.14%	-23.51%	-35.96%	-33.10%	-16.91%
Stdev	16.00%	14.60%	37.35%	24.45%	44.14%	37.65%	19.27%
TE vs GLD		6.52%	25.36%	14.07%	31.41%	25.91%	13.51%
Start	Jan-13	Jan-13	Jan-13	Jan-13	Jan-13	Jan-13	Jan-13
End	Dec-15	Dec-15	Dec-15	Dec-15	Dec-15	Dec-15	Dec-15

Conclusion

Tax management of investment strategies relies on harvesting unrealized losses to decrease short term gains. When a proxy is used to replace the initial position, it creates a tracking error in the original portfolio. Having a tracking error as low as possible to maintain the initial intent of the portfolio is key to successful tax management. The industry standard has much room for improvement. Our proprietary machine learning based algorithm achieves significant improvement over the industry standard, due to the use of a portfolio approach that jointly optimizes for substitution error, complexity and trading costs, while adjusting over time.

Opinions expressed are current opinions as of the date appearing in this material only. While the data contained herein has been prepared from information that 55 Institutional and 55 Capital believes to be reliable, the firms do not warrant the accuracy or completeness of such information. This communication is for informational purposes only. This is not intended as nor is it an offer, or solicitation of any offer to buy or sell any security, investment or product.

Tax Rates as of 2017