Tax Loss Harvesting Optionality

Ben Davis, Ph.D.

April 25, 2023

Capital Gain Taxes 101

- Realized capital gains are taxed
- Gains are offset by Losses
- Example
    1. Purchase 245 shares of SPY (an S&P 500 ETF) at a price of $600/share
    2. Portfolio cost basis = 245 shares x $600/share = $147,000
    3. Suppose SPY falls 8.33% to $550/share
    4. Now portfolio market value = 245 shares x $550/share = $134,750
    5. Unrealized loss =market value – cost basis = -$12,250
    6. Sell all SPY shares, realizing short-term capital loss = -$12,250
    7. The loss offsets capital gains elsewhere in portfolio, cutting investor tax bill

	Realized Gain or Loss	Capital Gains Tax (40.8%)
SPY	-$12,250	-$4,998
Elsewhere	$20,000	$8,160
Total Tax Due		$3,162

- *Investors have the right, but not the obligation, to sell assets held at a loss, resulting in valuable tax credits*

Tax Loss Harvesting Benefit as Put Exercise Payoff

- The put "hockey stick" shows how the exercise payoff changes with the underlying asset price

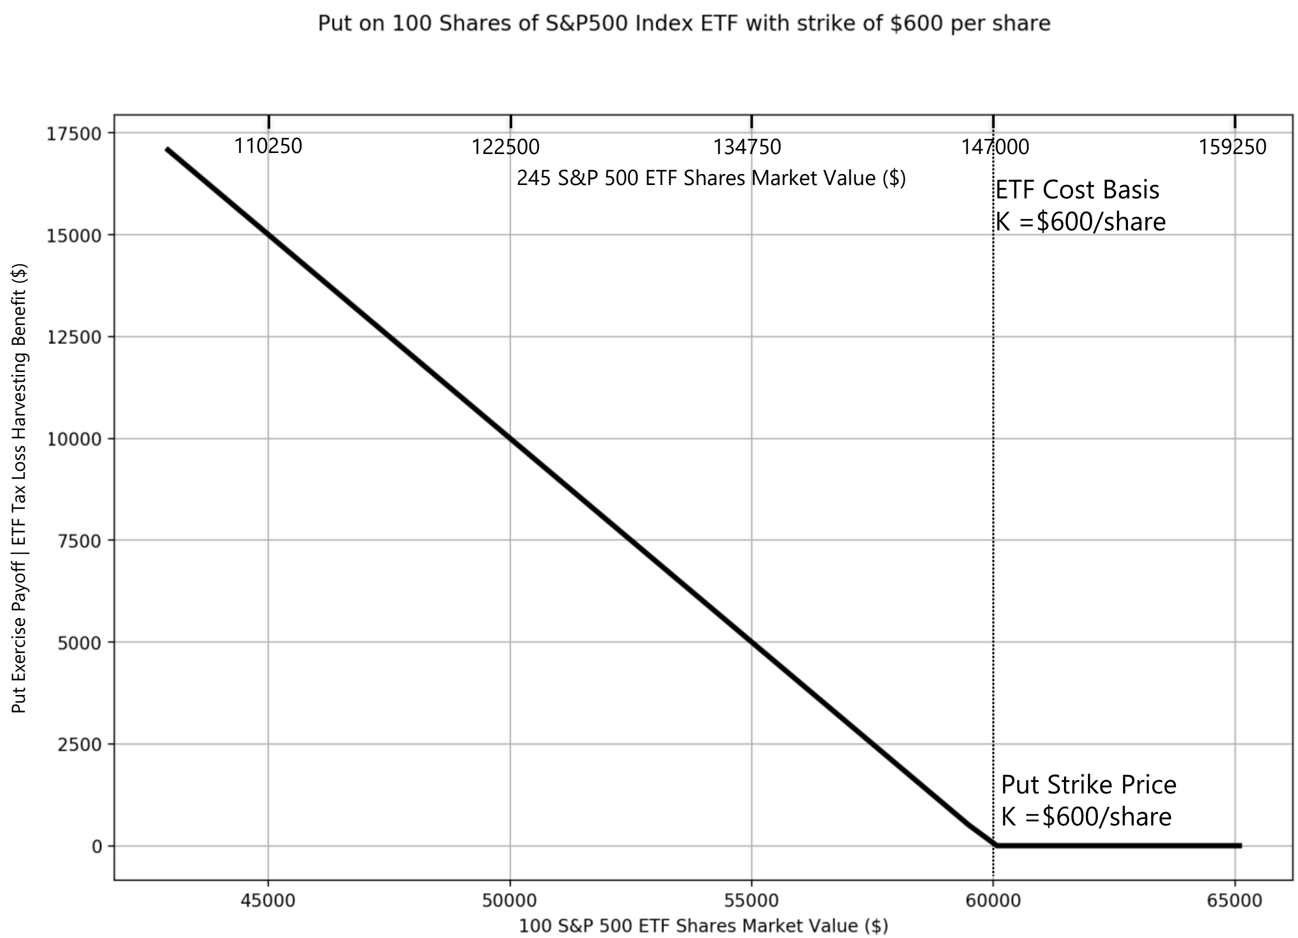

Optimal Option Exercise

- For **a perpetual American put** on a non-dividend paying equity, the **optimal policy** is to exercise when the underlying stock price falls below

$$P^* = \left[\frac{r}{r + \sigma^2/2}\right] K$$

  where r and σ are the continuous time risk-free rate and volatility of the underlying stock price (see John Hull, *Options, Futures and Other Derivatives)*

- The key points are
  - The optimal policy is loss-depth trigger-based
  - There is a simple analytical formula for the trigger price
  - The trigger price depends on the stock volatility
    - When volatility goes up, the trigger price moves down (deeper trigger)
    - When volatility goes down, the trigger price moves up (shallower trigger)
- As we will describe, the tax loss harvesting option is more exotic than this

IRS Wash Sale Rule

- When buying and selling securities, a taxable investor must adhere to the IRS publication 550 wash sale rules to claim deductions for realized capital losses

> *"A wash sale occurs when you sell a security at a loss and within 30 days before or after the sale you buy substantially identical stock or securities."*

- Once a security is sold at a loss, to claim a capital gains tax deduction, the investor must wait 30 days before buying it back.
- Israelov and Lu (2022) argue that the IRS wash-sale rule creates a barrier to re-investment such that investor should be selective about when to harvest a loss.
- They present a trigger-based loss harvesting policy based on loss-depth determined by stock volatility and supported by extensive Monte Carlo simulation.
- We extend this line of reasoning by applying stochastic process theory to prove that such a trigger-based policy is optimal for a stylized, continuous time model of tax loss harvesting.

The Stylized Model

- The ETF pays no dividends
- The ETF price exhibits geometric Brownian motion; e.g the price returns are continuous-time log normal of constant volatility and with mean log return of zero

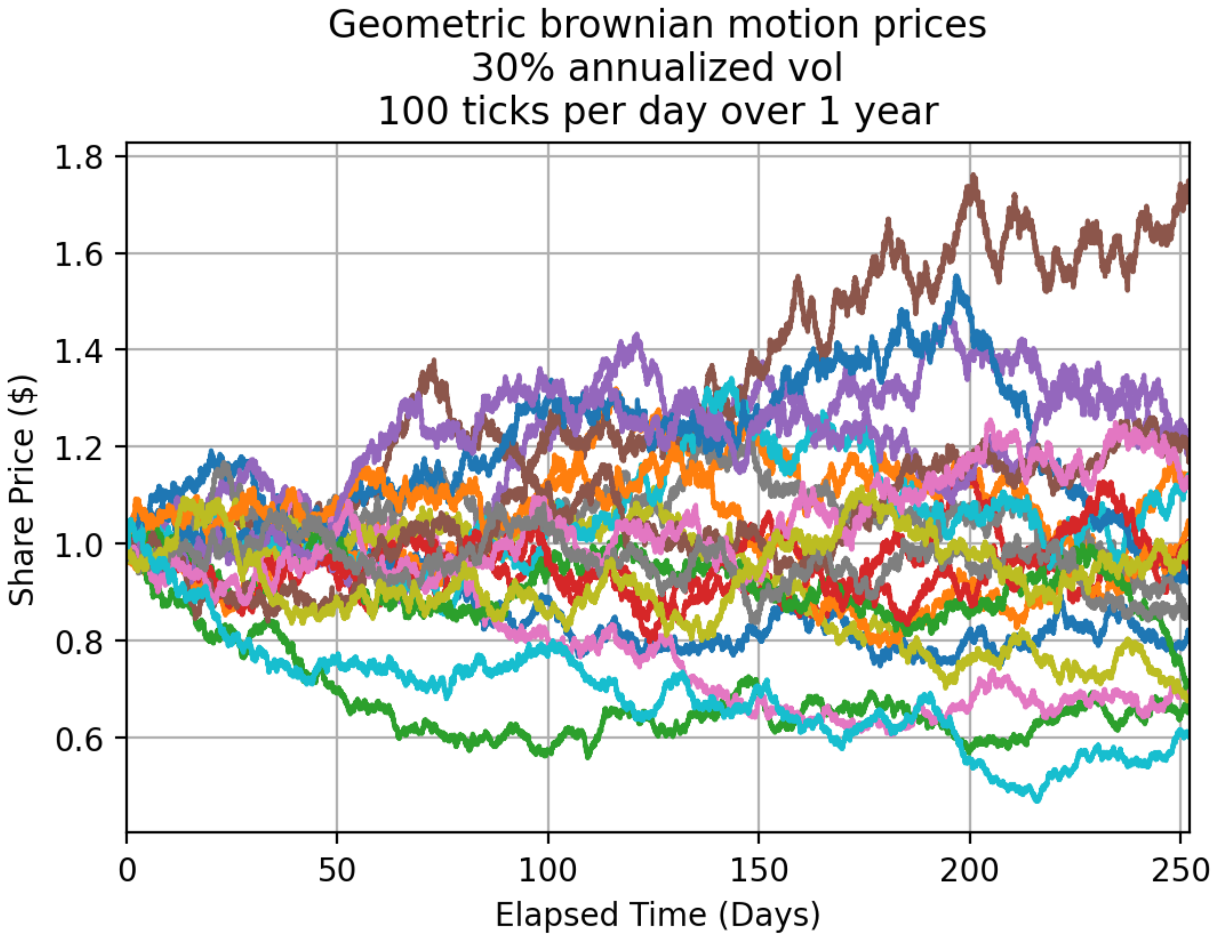

The Stylized Model – Set up

- We continuously monitor the price and sell when the unrealized losses hit a given trigger level, as a percentage of the cost basis
- We hold cash for 21 trading days (the "wash-sale period"), and then we fully re-invest in the ETF

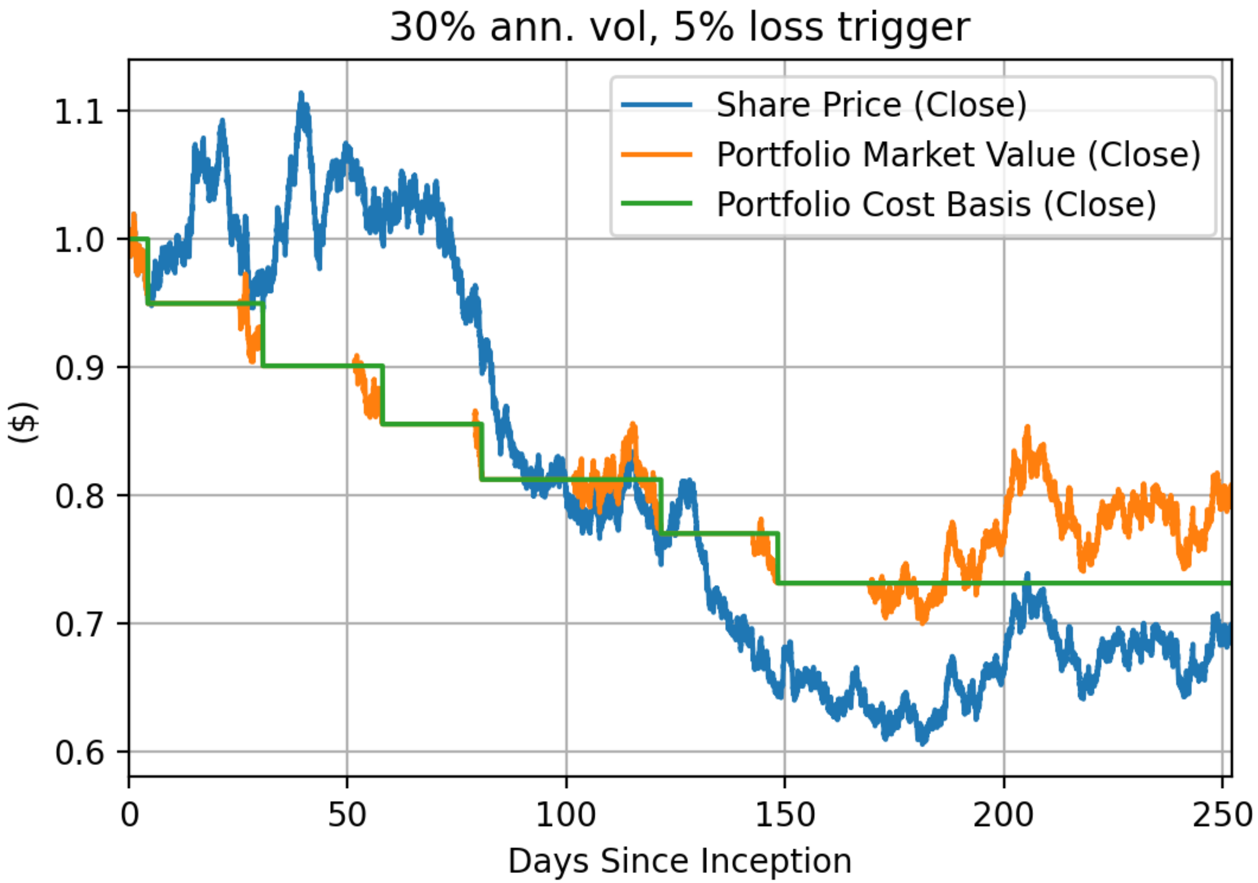

The Stylized Model – Set up

Example simulated portfolio history: ETF 30% ann. vol with 5% loss trigger

Days Since Inception	Share Price (Close)	Stock Daily Return (Close)	Portfolio Invested in Equities as of Prior Day Close	Portfolio Market Value (Close)	Portfolio Cost Basis (Close)	Portfolio Unrealized Losses (Close, Pre-Trade)	Sell at Close Flag	Porfolio Single Day % Realized Losses (Close)	No-Buy Restriction for the Day	Buy at Close Flag
0	$1.0000	0.00%	FALSE	$1.0000	$1.0000	0.00%	FALSE	0.00%	FALSE	FALSE
1	$1.0206	2.06%	TRUE	$1.0206	$1.0000	2.06%	FALSE	0.00%	FALSE	FALSE
2	$1.0260	0.52%	TRUE	$1.0260	$1.0000	2.60%	FALSE	0.00%	FALSE	FALSE
…	…	…	…	…	…	…	…	…	…	…
81	$0.9612	0.26%	TRUE	$0.9612	$1.0000	-3.88%	FALSE	0.00%	FALSE	FALSE
82	$0.9242	-3.85%	TRUE	$0.9242	$0.9242	-7.58%	TRUE	-7.58%	FALSE	FALSE
83	$0.9046	-2.12%	FALSE	$0.9242	$0.9242	0.00%	FALSE	0.00%	TRUE	FALSE
…	…	…	…	…	…	…	…	…	…	…
101	$0.8410	0.60%	FALSE	$0.9242	$0.9242	0.00%	FALSE	0.00%	TRUE	FALSE
102	$0.8278	-1.57%	FALSE	$0.9242	$0.9242	0.00%	FALSE	0.00%	TRUE	FALSE
103	$0.8506	2.75%	FALSE	$0.9242	$0.9242	0.00%	FALSE	0.00%	FALSE	TRUE
104	$0.8607	1.19%	TRUE	$0.9352	$0.9242	1.19%	FALSE	0.00%	FALSE	FALSE
105	$0.8932	3.77%	TRUE	$0.9705	$0.9242	5.01%	FALSE	0.00%	FALSE	FALSE
106	$0.8664	-3.00%	TRUE	$0.9414	$0.9242	1.86%	FALSE	0.00%	FALSE	FALSE
107	$0.8757	1.08%	TRUE	$0.9516	$0.9242	2.96%	FALSE	0.00%	FALSE	FALSE

The Stylized Model

- Consider the distribution of simulated portfolio realized loss history (colorful stairs)
- Compute the expected value of losses at each time T (heavy black line)
- Convenient to work with log losses (e.g. $\log(1 - \frac{5}{100})$ corresponds to a 5% loss)
- Stochastic process theory: There is an analytical formula the expected log losses!

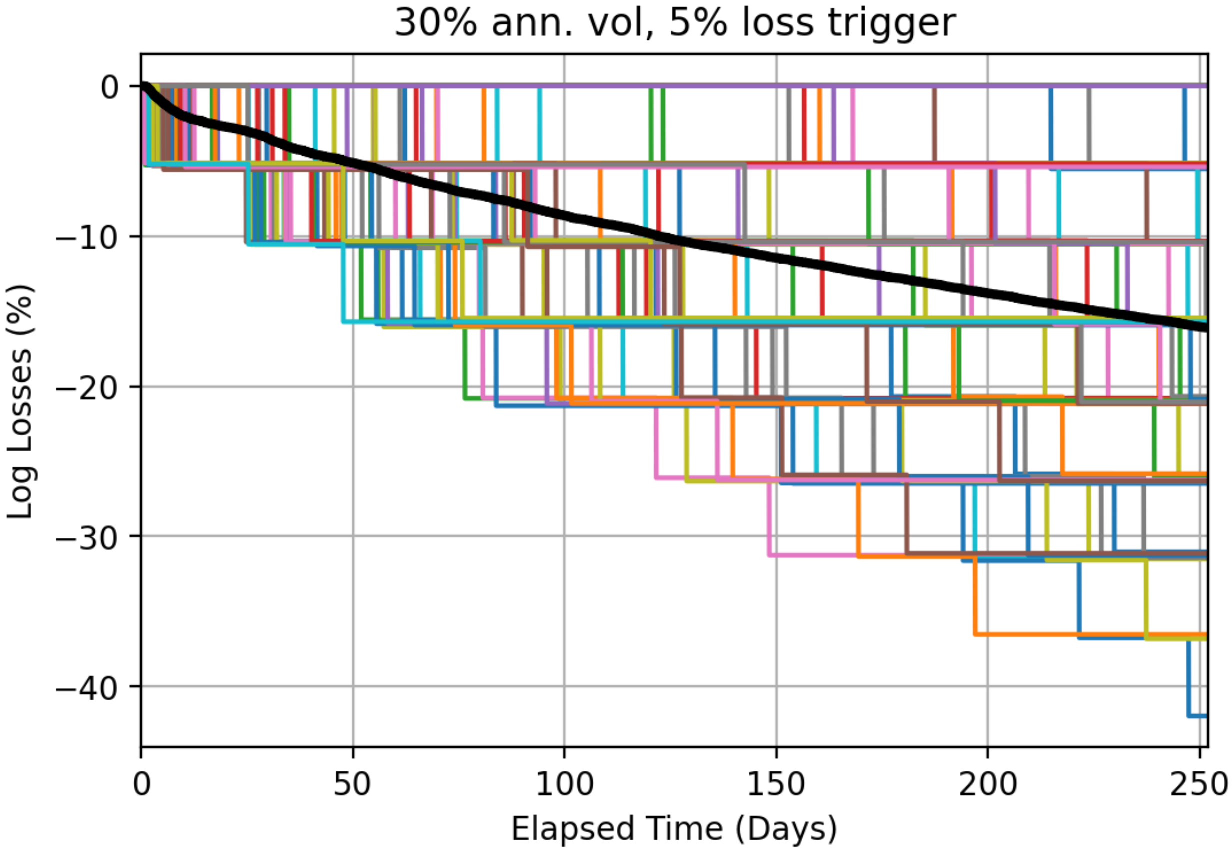

The Stylized Model

- Different triggers have different rates of expected loss accumulation

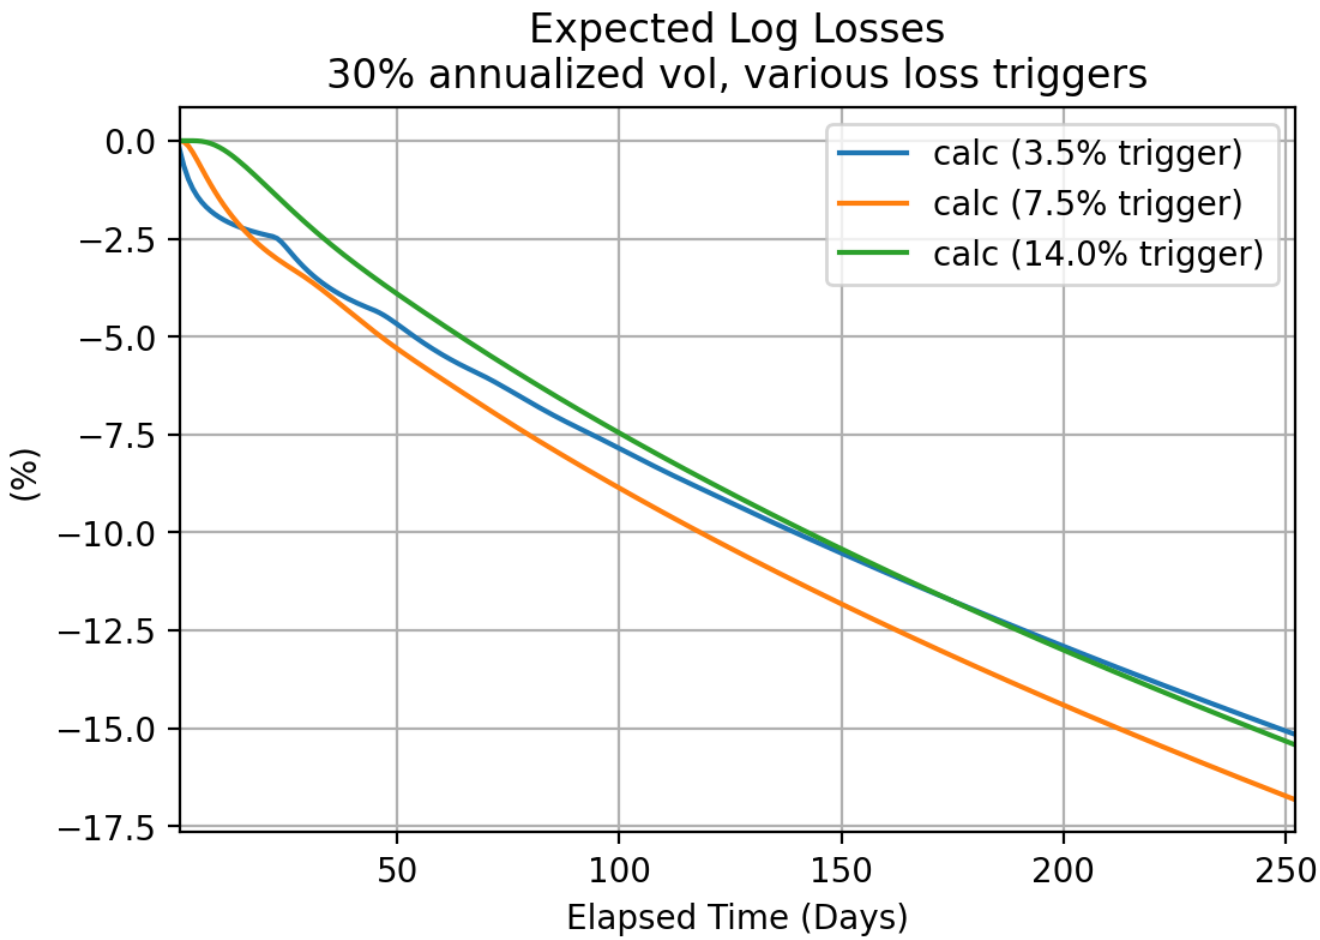

The Stylized Model

- We find the trigger depth with the largest expected log losses using the analytical formula for expected log losses

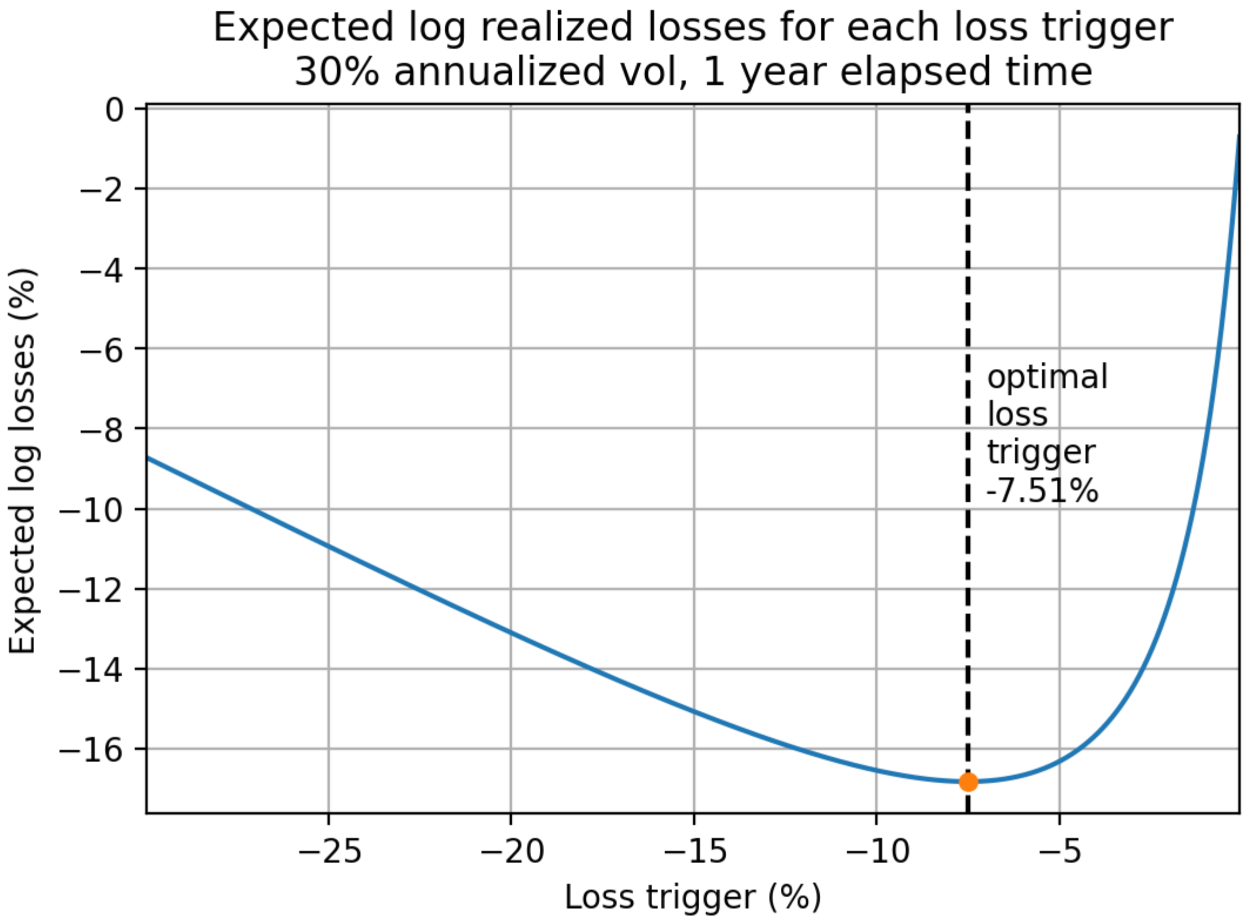

The analytical formula

- Let $\boldsymbol{\sigma}$ be the volatility of the ETF log daily returns (e.g. stdev{log(1+r)}, where r represents the distribution of daily ETF price returns
- Let $\boldsymbol{\lambda = \|\log(1+L)\|}$where the trigger level *L* is expressed on a decimalized basis as a percentage of the cost basis; e.g. L=-0.05 represents a 5% loss trigger
- Let $\mathbf{z} = \frac{\lambda}{\sigma}$ represent the trigger level expressed as standard deviations of ETF log daily returns
- The expected log losses at elapsed time time T are given by

$$E[LogLoss(T)] = \lambda \cdot \sum_{m=0}^{M} \tilde{h}_m(T) \cdot m$$

where

$$M = \text{ceil}(T/21)$$

$$\tilde{h}_m(T) = \begin{cases} \text{erf}\left(\frac{(m+1)z}{\sqrt{2(T-21m)}}\right) - \text{erf}\left(\frac{mz}{\sqrt{2(T-21(m-1))}}\right) & 21m < T \\ 1 - \text{erf}\left(\frac{mz}{\sqrt{2(T-21(m-1))}}\right) & 21(m-1) < T \leq 21m \\ 0 & T \leq 21(m-1) \end{cases}$$

First hitting times

- The *first hitting time distribution* is a foundational formula in stochastic processes and answers the question: What is the probability of hitting the trigger the first time between times T_0 and T_1?
- Answer: $\int_{T_0}^{T_1} f(t)\,dt$ where $f(t) = \frac{z}{\sqrt{2\pi}} t^{-3/2} e^{-\frac{z^2}{2t}}$
- This integral can be computed using *u*-substitution with $u = t^{-1/2}$ and the established *error function* $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-x^2}\,dx$
- The result is $\int_{T_0}^{T_1} f(t)\,dt = \left(1 - \operatorname{erf}\left(\frac{z}{\sqrt{2T_1}}\right)\right) - \left(1 - \operatorname{erf}\left(\frac{z}{\sqrt{2T_0}}\right)\right)$
- Predicted frequency of first hitting times versus results of 10,000 path Monte Carlo simulation; 30% annualized vol, 5% loss trigger

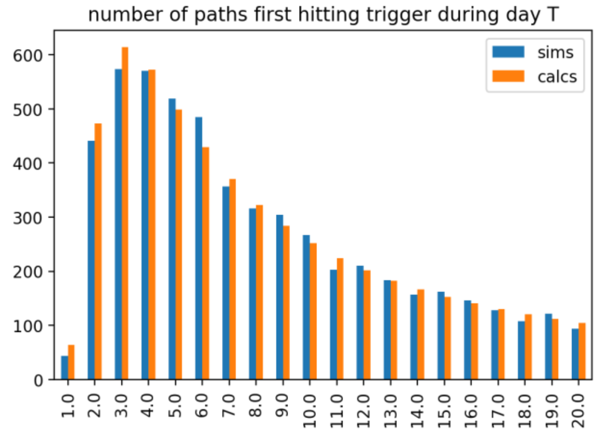

Renewal Processes

- Regardless of whether the return generating process is log-normal or not, the CDF of the first hitting time is the probability of first hitting the trigger on or before elapsed time T
- $F_1(T) = \int_0^T f(t)\, dt$
- For our log-normal process,

$$F(T) = 1 - \mathrm{erf}\left(\frac{z}{\sqrt{2T}}\right)$$

- For a given sample path, consider the function that counts the number of trigger hits by elapsed time T, $N(T)$.
- $\mathbb{E}(N(T))$ is the expected number of trigger hits at time T
- The probability of hitting the trigger precisely *m* times at time T, $h_n(T) = Prob(N(T) = n)$
- The expected log losses are then $\lambda \cdot \mathbb{E}(N(T)) = \lambda \cdot \sum_{n=0}^{\infty} n \cdot h_n(T)$
- The hit count probabilities h_n can be computed from the the first hitting time CDF F, as we will now show

Two Key Observations

- Observation #1 : (No wash-sale period) The probability that hit *n* for trigger depth λ occurs within elapsed time T is equal to the probability of the first hit for trigger depth n$\cdot\lambda$ occurs within elapsed time T

$$F_n(\lambda;T) = F(n \cdot \lambda;T)$$

- In the case of lognormal returns, we know the formula for the right-hand side

$$F_n(\lambda;T) = 1 - \mathrm{erf}\left(\frac{na}{\sqrt{2T}}\right)$$

- Observation #2: Let $h_n(T)$ be probability that precisely n trigger hits will occur by elapsed time T. Then

$$F_n(T) = h_n(T) + h_{n+1}(T) + \cdots$$

and consequently, we can solve for the h's

$$h_n(T) = F_n(T) - F_{n+1}(T)$$

In the case of lognormal returns, we obtain

$$h_n(T) = \mathrm{erf}\left(\frac{(n+1)z}{\sqrt{2T}}\right) - \mathrm{erf}\left(\frac{nz}{\sqrt{2T}}\right)$$

Precisely *m* trigger hits probability formula (No wash sales)

- No wash sales, Predicted frequency of precisely *m* trigger hits at 252 days versus results of 10,000 path Monte Carlo simulation; 30% annualized vol, 5% loss trigger

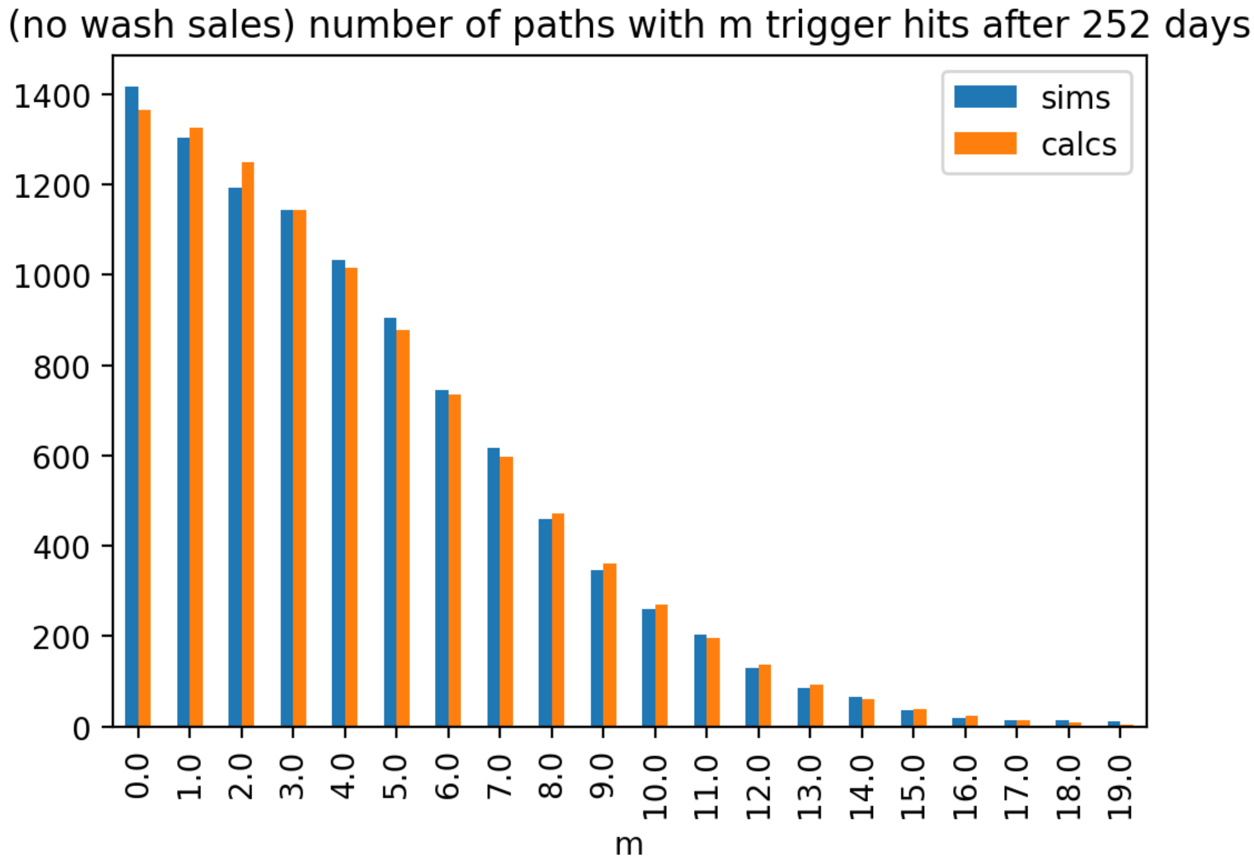

Incorporating wash sale restrictions

> It is helpful to assume that each trigger hit occurs after a 21 day waiting (wash-sale) period

> There is a small adjustment needed for the first trigger hit, since the investor in reality has no waiting period after inception

> We will solve the problem where there is a 21 day waiting period after inception; then we will adjust the solution by sliding the solution back by 21 days

> Okay, recall that, in complete generality, $F_n(T)$ is the CDF of the time of the nth trigger hit $T_n = X_1 + \cdots + X_n$

> Then, with wash-sale restriction, the random variable of the time to the nth trigger hit is

$$\begin{aligned} \tilde{T}_n &= \tilde{X}_1 + \cdots + \tilde{X}_n \\ &= (21 + X_1) + \cdots + (21 + X_n) \\ &= 21n + X_1 + \cdots + X_n \\ &= 21n + T_n \end{aligned}$$

> Thus,

$$\tilde{F}_n(T) = F_n(T - 21n)$$

where F_n is zero for negative input values

Incorporating wash sale restrictions

- The key observation #2 still holds, namely,

$$\begin{aligned}\tilde{h}_n(T) &= \tilde{F}_n(T) - \tilde{F}_{n+1}(T) \\ &= F_n(T - 21n) - F_{n+1}\big(T - 21(n+1)\big)\end{aligned}$$

- In the case of log-normal returns,

$$\tilde{h}_n(T) = \operatorname{erf}\left(\frac{(n+1)z}{\sqrt{2\big(T - 21(n+1)\big)}}\right) - \operatorname{erf}\left(\frac{nz}{\sqrt{2(T - 21n)}}\right)$$

- Finally, slide the entire formula 21 days back, to adjust for the missing 21 day waiting period at the start (replace T with T+21 in the formula)

$$\tilde{h}_n(T) = \operatorname{erf}\left(\frac{(n+1)z}{\sqrt{2(T - 21n)}}\right) - \operatorname{erf}\left(\frac{nz}{\sqrt{2(T - 21(n-1))}}\right)$$

- When 21(n-1) < T < 21n, the first term becomes 1, because it is 1-F_n+1, and F_n+1 becomes zero in this range. Similarly, when T < 21(n-1) you can think of both terms becoming 1 because both F's become zero in this range.

Precisely m trigger hits probability

> Predicted frequency of precisely *m* trigger hits at 252 days versus results of 10,000 path Monte Carlo simulation; 30% annualized vol, 5% loss trigger

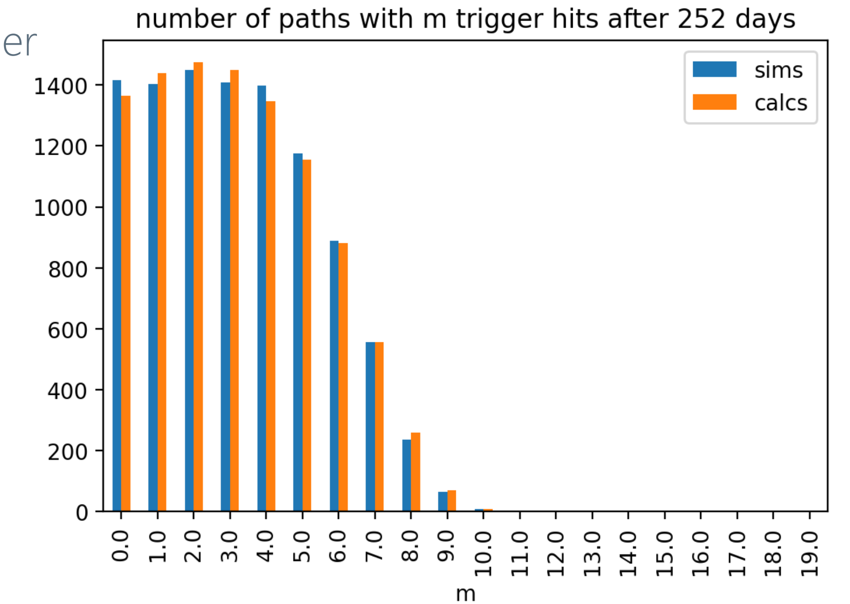

> Predicted frequency of *m* trigger hits over time

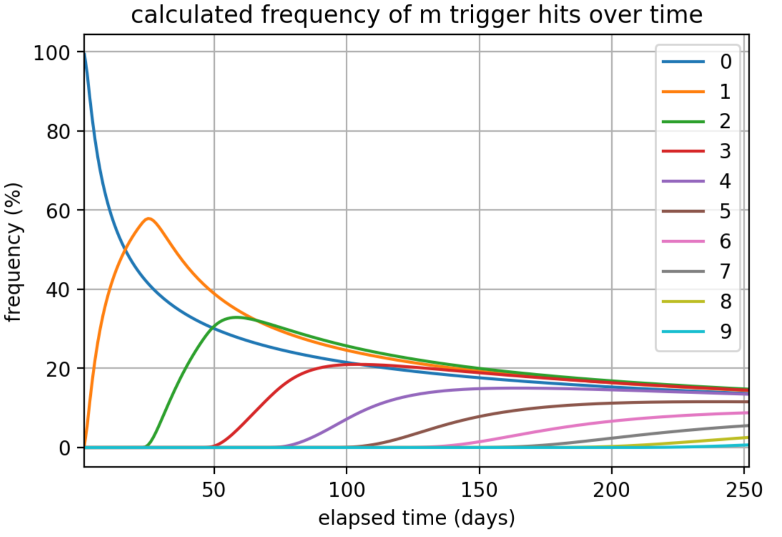

Finding the optimal loss trigger

> Even though we know the analytical formula for the expected log losses for each loss trigger, finding the minimum using calculus has eluded me

$$\frac{dE[LogLoss(T,\sigma)(\lambda)]}{d\lambda} = 0$$

> The minimum in the plot below was found numerically

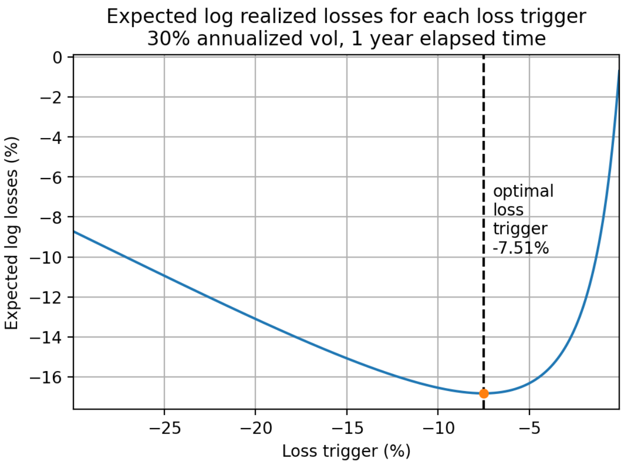

Finding the optimal loss trigger

> The optimal loss trigger depends on the time horizon T

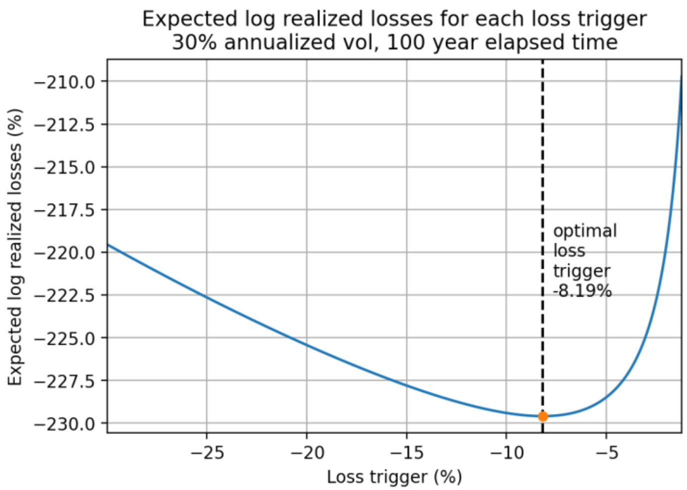

> We can make some further progress by considering symmetries of the simulated portfolio histories

Scaling Laws

- Scaling log prices or time by a constant is a symmetry on simulated portfolio histories
- To state fully, replace duration of wash sale period "21 days" with variable τ
- Notice $\tilde{h}_n(\sigma, \tau, \lambda)(T) = \tilde{h}_n(k\sigma, \tau, k\lambda)(T)$, since $\tilde{h}_n$ is a function only of $z = \frac{\lambda}{\sigma}$
- Thus, $\boxed{\lambda_{opt}(k\sigma, \tau)(T) = k \cdot \lambda_{opt}(\sigma, \tau)(T)}$
- For time scaling, we can verify by substitution into the formula for $\tilde{h}_n$ that

$$\tilde{h}_n(\sigma, \tau, \lambda)(T) = \tilde{h}_n(\sigma, k\tau, \sqrt{k}\lambda)(kT)$$

- So, $\boxed{\lambda_{opt}(\sigma, k\tau)(T) = \sqrt{k}\lambda_{opt}(\sigma, \tau)(kT)}$
- The implication is, if we know $\lambda_{opt}(\sigma, \tau)(T)$ for a single reference pair of values for (σ, τ), then we know it for all values

$$\boxed{\lambda_{opt}(\sigma, \tau)(T) = \sigma \cdot \sqrt{\tau} \cdot \lambda_{opt}(1,1)(\sqrt{\tau} \cdot T)}$$

- We will continue to work for the case of 30% annualized vol and 21 day wash-sale period
- For these fixed values, how does the optimal trigger vary with time horizon?

Optimal Trigger at Various Horizons

- Plotted below is he optimal trigger at various horizons, from less than a day to 1000 years (note logarithmic scale on time axis)

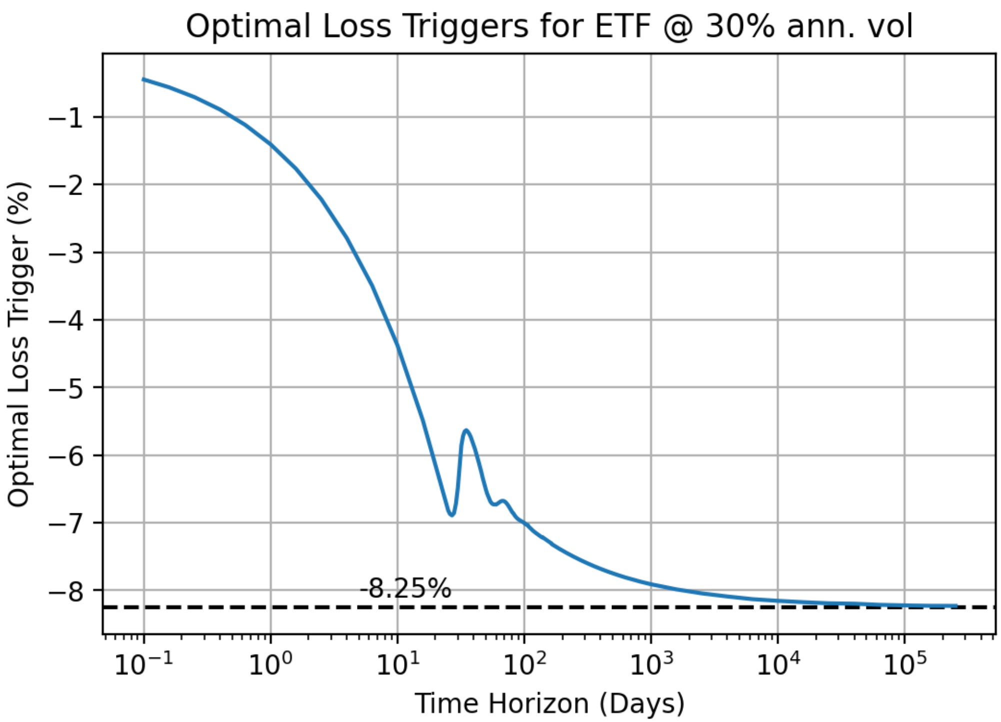

- The scaling laws guarantee that this plot is the same, regardless of the time units (days, months, etc.) used to generate the plot

$$\lambda_{opt}\left(\sigma\sqrt{21}, \frac{\tau}{21} = 1\right)(T/21) = \lambda_{opt}(\sigma, \tau)(T)$$

Long-Run Optimal Conjecture

> The long-run optimal trigger of -8.25% corresponds to a log loss trigger of

$$\lambda_{opt,LR} = \ln\left(1 - \frac{8.25}{100}\right) = -8.61\%$$

> On the other hand, for an ETF of a given annualized vol, mean zero lognormal returns have vol over a monthly horizon given by

$$\text{vol} \quad \rightarrow \quad \sigma = \sqrt{\ln\left(\frac{1 + \sqrt{1 + 4\left(\frac{\text{vol}}{\sqrt{12}}\right)^2}}{2}\right)}$$

> For vol=30%, this given monthly horizon lognormal vol of $\sigma = 8.61\%$

> Long-run optimality conjecture: $\boldsymbol{\lambda_{\mathrm{opt,LR}} = -\sigma\sqrt{\tau}}$ where $\boldsymbol{\sigma}$ is lognormal vol over time duration d (e.g. d could be 1 day, or 21 days, or 1 month) and $\boldsymbol{\tau}$ is the duration of the wash sale period measured in multiples of d (e.g., if d is 21 days, then for example the IRS wash-sale period is $\boldsymbol{\tau = 1}$; if d is 1 day, $\boldsymbol{\tau = 21}$).