Learning to Trade I

Visiting Prof. Hans Buehler

hans.buehler@maths.ox.ac.uk

University of Oxford, Institute for Mathematics, 2026

http://deep-hedging.com

Agenda

- Explain real-life considerations when “learning to trade”
- Link to existing derivatives models and infrastructure
- Give pointers for research opportunities

Learning To Trade

- **Buy Side**
  - Hedge Funds, Asset Managers
    - Invest in liquid and bespoke OTC (over-the-counter) instruments to generate returns for investors
    - Typically leverage “working capital” (the actual money the collected from investors) to “trading capital” → requires careful risk management to reduce margin requirements.
    - Fees function of performance and , e.g. 2/20: 2% management fee and 20% upside annually.
  - Corporates
    - Support the company e.g. managing FX risk, finance M&A, share buy backs, etc.
    - Typically not active market participants

- **Sell Side**
  - Provide liquidity e.g. manage an inventory of assets in anticipation of aggregate client demand.
  - Write OTC derivatives contracts for clients and risk-manage (hedge) them.
  - Earn fees – since 2010 Dodd-Frank proprietary trading at banks is severely limited.

UNIVERSITY OF OXFORD

Learning To Trade

- Objective is symmetric:
  - Bank – at any decision point:
    - Offer clients OTC products (which could be blocks of liquid instruments) for some price.
      We have price setting authority for what we sell and will want to know the minimal price we are able to sell a product at.
    - Trade in the liquid market to manage our position.
    - Maximize risk-adjusted return in the sense we want to make money while adhering to our risk management constraints and regulatory obligations.
  - Client – at any decision point:
    - Screen bank OTC products for an attractive price.
      We will need the maximal acceptable price for each such product.
    - Trade in the liquid market to manage our position.
    - Maximize risk-adjusted return in the sense we want to make money while adhering to our risk management constraints, regulatory obligations and our margin requirements.

Classic Quant Finance Primer

Classic Quant Finance

- A *filtration* is a representation of all past events.
  - In our context, assume that s_t is the **state** of the world: all current and past market data, news, etc. In particular, it contains *time*.
    Then our default filtration is the sequence of all observations of s.(*)
  - In machine learning terms, s are all available **features**.
  - We model financial instruments as random variables under the **statistical measure** P.
    The statistical measure is the real-life measure, under which assets can have excess returns.
  - Formally, any process $X = (X_t)_t$ over our filtration can be written as $X_t = X(s_t)$ where X is a measurable function. That simply means that s_t contains all relevant information for us. This does *not* mean that the dynamics of X cannot depend on other, hidden, states.

- In our lecture we use $H = (H^1, \dots, H^n)$ to refer to a range of n tradable "hedging" instruments, each with a price process $H^i = \left(H_t^i\right)_t$.
  - We will comment on this innocuous assumption of existence of "a price" throughout.

- Non-tradable instruments, or portfolios thereof are represented by Z.

(*) formally, a filtration is the sequence of P-completed σ-algebras of s.

Classic Quant Finance

- **The Replication Problem**
  - Assume that we wish to sell a product $-Z_t = -Z(X_T)$ to a client, leaving us with position $+Z_T$.
  - We are able to trade in H at time points $t = 0, \dots, T$.
  - Our cash attracts a symmetric rate r_t.
  - What is a minimum price we should charge?

- Classic Approach:
  - Assume that H is a diffusion which generates cash flows $\mu_{\mathrm{t}} \coloneqq \mu_t(H_t) \in \mathrm{R}^n$;
  - That it has a statistical drift $\alpha_t \coloneqq \alpha_t(H_t) \in \mathrm{R}^n$ in excess of those cash flows; and
  - also assume that $\sigma_t := \sigma_t(H_t) \in R^{n\times d}$ is the (generalized) volatility of H w.r.t to an standard Brownian motion $W_t^P \in \mathrm{R}^d, \mathrm{d} \geq \mathrm{n}$ such that:

$$dH_t = (r_t H_t + \alpha_t - \mu_t)dt + \sigma_t dW_t^P \in \mathrm{R}^n$$

(*) formally, a filtration is the sequence of P-completed σ-algebras of s.

Classic Quant Finance

- **The Replication Problem**
  - It is worth understanding the drift well:

$$dH_t = (r_t H_t + \alpha_t - \mu_t)dt + \sigma_t dW_t^P \in \mathrm{R}^n$$

  - The positive exposure to r_t stems from the same self-financing assumption: in order to buy the asset H_t^i for dt we will need to pay $r_t H_t^i dt$ for borrowing the respective cash.
    All assets which make no payments and settle today must have a drift proportional to r.
  - The negative sign of μ_t is present because if an asset produces a cash flow μ_t^i over dt then in a no-arbitrage world its price must drop by that amount.
  - We may enter into successive self-financing rolling dt-forward contracts $\widetilde{H}$ which satisfy

$$d\widetilde{H}_t \coloneqq dH_t + (\mu_t - r_t H_t)dt = \alpha_t dt + \sigma_t dW_t^P \in \mathrm{R}^n$$

  - α, on the other hand, does not pay anything, and is uncertain, so we cannot trade it away.

(*) formally, a filtration is the sequence of P-completed σ-algebras of s.

Classic Quant Finance

- **The Replication Problem**

$$dH_t = (r_t H_t + \alpha_t - \mu_t)dt + \sigma_t dW_t^P \in \mathrm{R}^n$$

- Assume that our volatility is bounded away from zero such that we may define the measure Q by setting

$$Q_t[A] := \mathrm{E}_t^{\mathrm{P}}[D_t 1_A] \quad \frac{dD_t}{D_t} = \lambda_t' dW_t^P \quad \lambda_t := \sigma_t^{(-1)} \alpha_t$$

- Following Girsanov's theorem under Q we get

$$dH_t = (r_t H_t - \mu_t)dt + \sigma_t dW_t^Q$$

- **Note:** since μ_t represents actual cashflows, the self-financing rolling forward $d\widetilde{H}_t$ is (in theory) a tradable and a martingale under Q:

$$d\widetilde{H}_t := dH_t + (\mu_t - r_t H_t)dt = \sigma_t dW_t^Q$$

Classic Quant Finance

- How does $dH_t = (r_t H_t - \mu_t)dt + \sigma_t dW_t^Q$ help:
- Define

$$Z_t := e^{\int_0^t r_s ds} \mathrm{E}_t^{\mathrm{Q}}\left[e^{-\int_0^T r_s ds} Z(H_T)\right] = e^{\int_0^t r_s ds} \mathrm{E}^{\mathrm{Q}}\left[e^{-\int_0^T r_s ds} Z(H_T) | H_t\right]$$

  - The latter equation is due to the Markov nature of H.
  - The rates handling ensures that Z accrues correctly interest and that $Z_T = Z(H_T)$.

- Then by Ito

$$dZ_t = \partial_H Z_t{}' \sigma_t dW_t^Q + \left\{\partial_t Z_t + \partial_H Z_t'(r_t H_t - \mu_t) + \frac{1}{2}\sum_{i,j} \partial^2_{H^i H^j} Z_t \, \sigma_t^{i\prime} \sigma_t^j\right\} dt$$

- Since $Z_t e^{-\int_0^t r_s ds}$ is a Q-martingale that means

$$\partial_t Z_t + \partial_H Z_t{}'(r_t H_t - \mu_t) + \frac{1}{2} tr \, \partial^2_{HH} Z_t \, \sigma_t' \sigma_t = r_t Z_t$$

Classic Quant Finance

UNIVERSITY OF OXFORD

- Since this holds under Q almost surely, it also holds under P.
  That means the replication problem is solved with $\Delta_t \coloneqq \partial_H Z_t \in \mathrm{R}^d$ such that

$$
\begin{aligned}
& dZ_t - \Delta_t' d\widetilde{H}_t \\
& = \underbrace{\partial_t Z_t}_{\text{Theta}} \, dt + \underbrace{\partial_H Z_t{}'(r_t H_t - \mu_t)}_{\text{Delta Carry}} \, dt + \frac{1}{2} \underbrace{tr \, \partial_{HH}^2 Z_t \; \sigma_t' \sigma_t}_{\text{Gamma Carry}} \, dt = r_t Z_t dt
\end{aligned}
$$

- The initial price is given as Z_0.

- Great theory:
  - We can price derivatives by replication: “the price of a derivative is the cost of its replication”.
  - Avoids us having to estimate hard-to-estimate α as this drops out of the equation.
  - Most importantly all the theta-carry terms should disappear, and P&L “noise” should be zero.

Classic Quant Finance

- However, this approach is pretty optimistic since
  - It does not take into account trading cost, market impact, or risk limits… so the “cost of replication” is not actually captured.
  - It also doesn’t handle that well cases where the hedging instruments are not perpetual assets such as FX or Equity, but derivatives themselves (it can be made work).

- But even then – this is actually *not* what the actual models in quant finance are doing anyway.
- The issue is that we know neither $\mathbb{P}$ nor $\mathbb{Q}$…

[1] Dupire, B. (1994). Pricing with a Smile.

Classic Quant Finance

- We will use a somewhat tedious notation because there are two different time scales involved:
  - The physical time when we fit and use a model is denoted by t.
    Assume we have a real live asset $H = (S_t)_t$ with $S_t \in R^1$.
  - For our training we observe implied market data a time t for future times T.
    - $DF_T^{(t)} := e^{-r_T^{(t)}(T-t)}$ is the discount factor to $T \geq t$ observed in t with $DF_t(t) = 1$,
    - $F_T^{(t)} := S_t e^{\mu_T^{(t)}(T-t)}$ is the forward with $F_t^{(t)} = S_t$, and
    - $Call^{(t)}(T,K) = DF_T^{(t)}\, BS(F_T^{(t)}, T-t, \hat{\sigma}_{T,K}^{(t)})$ is the call price of an option with expiry T and strike K, here parameterized in terms of its *implied volatility* $\hat{\sigma}_{T,K}^{(t)}$.
  - The market state here is $s_t = \left(S_t, r_T^{(t)}, \mu_T^{(t)}, \hat{\sigma}_{T,K}^{(t)}\right)_{T \geq t, K > 0}$.
- When then use these market data in t to fit (train) a model $S^{(t)} = (S_T^{(t)})_{T \geq t}$ under some model risk-neutral measure $\mathrm{Q}^{(t)}$; the model price for some real-life payoff $Z(S_T)$ is then given as

$$DF_T^{(t)} E^{(t)}\left[Z(S_T^{(t)})\right]$$

[1] Dupire, B. (1994). Pricing with a Smile.

Classic Quant Finance

- **Example: Dupire's Local Volatility**[1] for managing derivatives viz a full option market, say S&P500.
  - Dupire's seminal local volatility model[1] for t is defined under its risk-neutral measure as

$$S_T \coloneqq F_T^{(t)} X_T^{(t)} \qquad dX_T^{(t)} = X_T^{(t)} \sigma^{(t)} \left(T, X_T^{(t)}\right) dW_T \qquad X_t^{(t)} \coloneqq 1$$

  - Where the local volatility is given by the famous formula[1]

$$\sigma^{(t)}(T,x)^2 := 2 \frac{\partial_T C^{(t)}(T,x)}{x^2 \partial_{xx}^2 C^{(t)}(T,x)} \qquad C_t(T,x) \coloneqq \frac{Call^{(t)}(T, F_T^{(t)} x)}{DF_T^{(t)} F_T^{(t)}}$$

[1] Dupire, B. (1994). Pricing with a Smile.

Classic Quant Finance

- To price any payoff $Z(S_T)$ we compute the risk-neutral model price:

$$Z_t^{(t)} \equiv Z_t^{(t)}(S_t) \coloneqq DF_T^{(t)}\, \mathrm{E}^{(t)}\left[Z\left(S_T^{(t)}\right)|S_t^{(t)} = S_t\right]$$

  - All input European option prices are recovered.
  - A similar method exists for "stochastic local volatility"[1].
    Both "fit" the market at t very well.
  - That means the prices of liquid instruments are perfectly captured.

[1] Guyon, Julien & Henry-Labordere, Pierre. (2011). The Smile Calibration Problem Solved.

Classic Quant Finance

- However, what happens when t $\rightarrow dt$:

$$Z_t^{(t)} = DF_T^{(t)} \, \mathrm{E}^{(t)} \left[Z\left(S_T^{(t)}\right) | S_t^{(t)} = S_t \right]$$

- In theory, if all market data changed such that the local volatility remains the same, and forward rates and discount rates remain the same, then the model remains the same and we have no excess as we saw before.
- However this is only true if $V_{t+dt}^{(t)} = V_{t\_dt}^{(t+dt)}$ … while, in fact, the whole market has moved: rates, forward drift, all option prices.
- There is nothing inherent in the model to make it even likely that its dynamics correspond to actual market dynamics.

Classic Quant Finance

- Classic models – dynamic behaviour
  - The model is re-fitted every day
  - That gives every day a different “model” X.
  - Nothing in the model makes the dynamics “realistic”.

- Similarly, the “delta” of the Dupire model is driven by only one parameter.
  - No reason to assume this is in some sense correct.
  - No model “vega” (derivative in implied vols) as volatility is a function of spot and not itself a stochastic process in the model.
  - 100% of published models handle only very few stochastic factors – the industry workhorse “Dupire Local Vol” models only the underlying itself as stochastic factor.

Classic Quant Finance

- Our Dupire model is therefore merely a consistent map

$$s_t, Z \to \text{ModelPrice}\,(s_t; Z) \qquad s_t = \left(S_t, r_T^{(t)}, \mu_T^{(t)}, \hat{\sigma}_{T,K}^{(t)}\right)_{T \geq t, K>0}$$

- We compute “greeks” with respect to non-model parameters
  - E.g. we compute “vega” as derivative in $\hat{\sigma}$ for Local Volatility.
  - In fact, we will usually compute a whole host of greeks which doesn’t make sense in a “complete market”
- Greeks” are not good enough
  - Every trading desk has their overwrites and heuristics to make the presumed replicating model work.
  - Capture wrongly captured model dynamics, most notably “vol skew” and other effects
  - Account for transaction cost and other market frictions such as liquidity
  - Express view on the market
  - … few papers assess properly out-of-sample performance of automated hedging with the proposed derivative models for a broad range of instruments.

Learning to Trade

- … few papers assess properly out-of-sample performance of automated hedging with the proposed derivative models for a broad range of instruments.

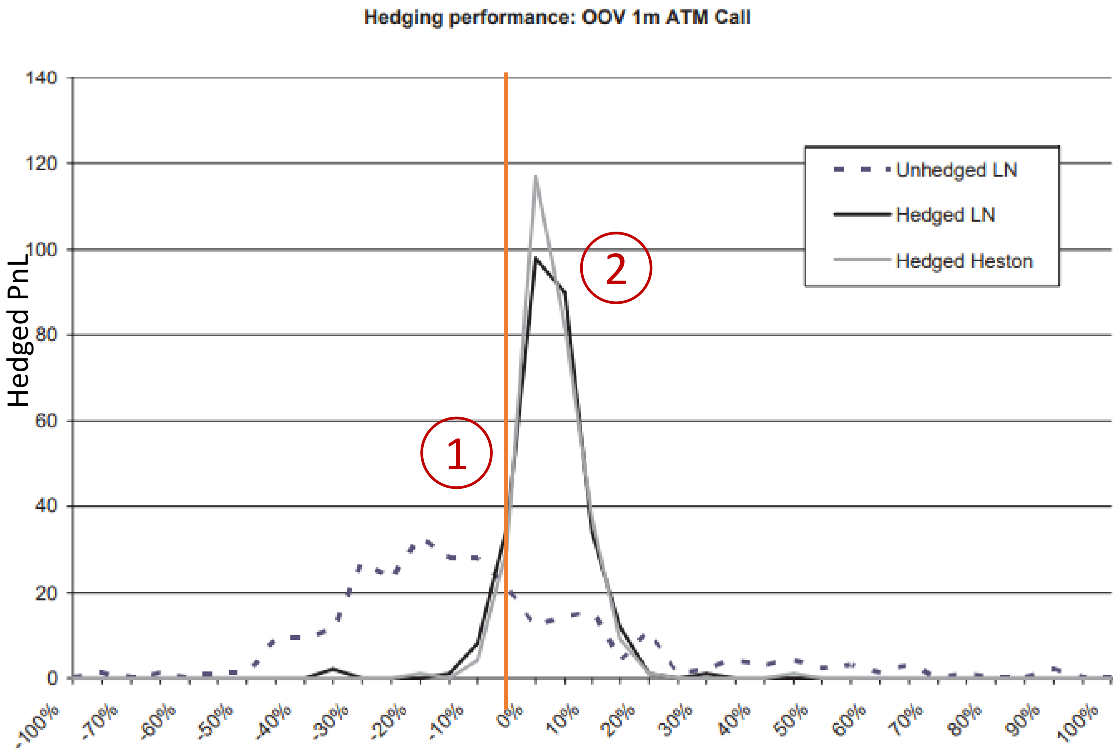

Analysis of hedging performance for options on variance with classic models, 2008 [1]

[1] Volatility Markets: Consistent Modelling, Hedging and Practical Implementation, Buehler 2008
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1118245

Learning to Trade

- Today's derivative models
  - Focus on "fitting the market"
  - We compute "greeks" with respect to non-model parameters
  - "Greeks" are not good enough
  - Theoretically dubious
  - Very few papers assess properly out-of-sample performance of automated hedging with the proposed derivative models

Hard to automate

- This is a historical effect: when derivatives analytics where developed
  - Data and compute were sparse and expensive.
    At the time of Black, Scholes, Merton a normal assumption was entirely reasonable
  - Trading was a slower affair
  - Fewer electronic platforms to trade automatically with

Learning to Trade

- Machine Learning and AI
  - Culture of using data to drive decisions
  - Much more data available
  - Availability of large-scale specialized compute
  - Modern “automatic adjunction differentiation” engines such as PyTorch, jax
  - Requires programming; numerical math; cloud tech

- AI
  - Do something good enough but at scale – detect cats and dogs in videos
  - Do something better than humans – play chess

Learning to Trade

- We want to let machine trade (mostly) by themselves

- Today’s derivative models … aren’t quite working

- Let’s go back to first principles

- We will look at two main ideas
  - Do something good enough but at scale – “Statistical Hedging”
  - Do something better – “Deep Hedging”

Framework

Framework

The Market

- The market is observed under the **statistical measure** P.

- Everything is in discrete time → markets cannot be complete.

- We call s_t the **state** of the market. It represents all known current and past information such as prices, bid/asks, twitter feeds, etc.*
  - That means that any observable random variable R_t can be written as $R(s_t)$.
  - Our trading activity may affect the distribution of s_{t+1} e.g. in markets with **impact**.

(*) That means that the σ-algebra F_t of our underlying filtration at time t is generated by s_t.

Framework

- Hedging Instruments
  - A key contribution of our work is providing a hedging framework for **derivatives as hedging instruments** e.g. swaps, futures, options.

  - We allow for **“floating” instrument definitions**, i.e. at each time step a potentially different set of instrument is available to trade,
    - For example the current on-the-run swaps, available listed options with some minimum liquidity etc.
    - This is a strong departure from most of the hedging literature, which tends to focus on “perpetual” instruments such as stocks and and FX, or on “fixed” instruments such as a specific bond.

  - For notional simplicity we will fix a total number of n instruments $H = (H^1, \dots, H^n)$.
    - The framework allows for different instruments per time step; to ease notation here we just consider all of them - we can and will restrict their trading to ensure a realistic setting, for example disallowing trading of instruments which are not listed yet.

Framework

- Hedging Instruments
  - Any instrument we may trade has to have a **book value** given for example by an official closing price, a weighted bid/ask, or a classic valuation model. E.g. for options, prices are given by implied volatilities. We denote the **book values** at time t by

$$H_t^i \equiv H^i(s_t) = \mathrm{ModelPrice}\ (s_t; H)$$

  - For simplicity, assume that all prices are expressed in our natural *numeraire*.
    - E.g. if a model-price $\tilde{X}_t$ is in USD, and if our numeraire $B_t^\$$ is the value of a USD bank account with a fixed notional, then $X_t \coloneqq \tilde{X}_t / B_t^\$$.
    - This is not a real restriction: it is straight forward to incorporate an actual cash account as, say, H^0, and even charge different rates depending on whether we are in credit or debit (a sort of market impact). It makes notation tedious though.

- We use $Z_t \equiv Z(s_t) = \mathrm{ModelPrice}_t(Z)$ to refer to the book value of a generic short OTC instrument position. For some applications will assume a book value in form of a valuation model is available; later we will also determine a fair price.

Framework

- Trading cost
  - Trading $a \in \mathrm{R}^n$ units of H in t will cost $c_t(a) \geq 0$ in excess of the bool price H_t.
  - The function c is normalized to $c_t(0) = 0$, non-negative, increasing in all directions, and convex.
  - *Convexity excludes fixed trading cost.*
  - Trading is naturally limited to where $c < \infty$.

- Example of Trading Frictions
  - Assume we trading vanilla options with mid-prices H_t^i.
  - Denote by Δ^i their Black & Scholes Delta, and by V^i their Black & Scholes Vega,
  - Example non-trivial trading cost for parameters $\gamma_\Delta, \gamma_V^{1,2} \geq 0$:

$$\hat{c}_t(a) \coloneqq \gamma_\Delta |a'\Delta| + \gamma_V^1 |a'V| + \gamma_V^2 |a'V|^2$$

  - A maximum Vega capacity of $V_{\max}$ is incorporated as follows

$$c_t(a) \coloneqq \hat{c}_t(a) + \infty 1_{|a'V| > V_{\max}}$$

  - Open question: what are consistent trading cost for options?

Parameter Hedging

Markowiz-Type Hedging of Daily Derivatives Risk

Parameter Hedging [1]

- Remember that risk management models have at least three roles
    1. Provide an optimal hedge to minimize risk vs. implementation cost
    2. Provide a “risk price”, i.e. a price irrespective of client of trader’s market view.
    3. Provide the ability to run stress scenarios for adverse scenarios
- Let us start with only the first property

- Consider the actual situation on a trading floor:
    - Every derivative in our books, and any derivative we might want to trade with our clients *has* a classic model “book value” (banking regulation) e.g. Local Volatility. Formally at time t every payoff is assigned a book value

$$s_t, Z \to \mathrm{ModelPrice}\ (s_t; Z)$$

    - Even a liquid instrument which is quoted on exchange has a book value: usually the mid-price of the actual bid/ask.
    - Given those model prices we wish to establish the best hedging strategy for “one day”.

Statistical Hedging, Buehler 2017 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2913250

Parameter Hedging

- We select a reasonable subset $x_t \subset s_t$ of **features** which we think are relevant factors for risk management. For example we exclude past data.

- For $X_t \in \{Z_t; H_t^1, \dots, H_t^n\}$ recall that $X_t \equiv X(s_t)$. We may use Taylor/Ito similar to before; let

$$\Theta_t \coloneqq \partial_t X_t \quad \Delta_t \coloneqq \partial_x X_t \quad \Gamma_t \coloneqq \partial_{xx}^2 X_t$$

- Then

$$dX_t = \underbrace{\Theta_t dt}_{Theta\ P\&L} + \underbrace{\Delta_t dx_t}_{Delta\ P\&L} + \underbrace{\tfrac{1}{2} tr \Gamma_t d\langle x \rangle_t}_{Gamma\ P\&L} + O(x_t^3)$$

- Recall: dx_there extends far beyond “spot”: even if the underlying model (e.g. Dupire) only has spot as “model state” we still compute derivatives to all other parameters. E.g. in the Dupire case: interest rates, forward rates, implied volatilities of all options.

- In other words “Δ_t” contains first order derivatives such as Vega etc.

Parameter Hedging

- For $X_t \in \{Z_t; H_t^1, \dots, H_t^n\}$ and $X_t \equiv X(s_t)$:

$$dX_t = \underbrace{\Theta_t dt}_{Theta\ P\&L} + \underbrace{\Delta_t dx_t}_{Delta\ P\&L} + \underbrace{\tfrac{1}{2} tr \Gamma_t dx_t^2}_{Gamma\ \{\&L} + O(dx_t^3)$$

- Our real P&L **gains** over $(t, t + dt)$ when trading $a \in R^n$ are given by:

$$G_t^a \coloneqq dZ_t + a' dH_t - c_t(a)$$

  - Idea is to minimize exposure to each “Greek” (derivative)
  - As it stands, hard to tell which of the terms are the most important to hedge… $\partial_{S\&P\ spot} G^a$ vs. $\partial_{S\&P\ vol}\, G^a$ vs $\partial_{USD\ 3m\ Swap\ Rate} G^a$ …. ?
  - *Parameter Hedging* formalizes this approach with data.

Parameter Hedging

- Assume we estimated $d\tilde{x}_t = \mu_t dt + \sigma_t dW_t$ with correlated dW_t under the *statistical* measure $\mathbb{P}$. We therefore estimate

$$d\tilde{X}_t \approx \underbrace{\left\{\Theta_t + \Delta_t\mu_t + \frac{1}{2}tr\Gamma_t\sigma_t^2\right\}}_{Drift} dt + \underbrace{\Delta_t\sigma_t}_{Risk} dW_t$$

- This formula provides us with a normal approximation of the **returns of any instrument**.
- Contrary to our earlier example, we do not assume that the drift reduces to interest rates.
- This gives us an estimated normal **gains process** where both drift $\mu(a)$ and variance $\Sigma(a)$ depend on the action:

$$\tilde{G}_t^a \coloneqq d\tilde{Z}_t + a' d\tilde{H}_t - c_t(a) \sim N(\mu(a), \Sigma(a))$$

Parameter Hedging

- Stylistically we may write this as a classic portfolio optimization problem:

$$\tilde{G}_t^a := \left(\text{drift}(Z) + a'\text{drift}(H)\right)dt + (\text{risk}(Z) + a'\text{risk}(H))dW_t - c_t(a)$$

- Lends itself rather obviously to optimization programs of the form

$$\max_a : U\left(\tilde{G}_t^a\right)$$

- What should U be?

Parameter Hedging

- What kind of objectives $U(G^a)$ should we maximize:

$$\tilde{G}_t^a := \big(\mathrm{drift}(Z) + a'\mathrm{drift}(H)\big)dt + (\mathrm{risk}(Z) + a'\mathrm{risk}(H))dW_t - c_t(a)$$

- Markowitz seminal work [1] suggests using
  - Markowitz Mean-Variance $U(X) = E[X] - \frac{1}{2}\lambda E[(X - EX)^2]$
  - Markowitz Mean-Volatility$U(X) = E[X] - \lambda E[(X - EX)^2]^{\frac{1}{2}}$
  - General form: $U(X) = E[X] - \frac{q}{p}\lambda E[(X - EX)^p]^{\frac{1}{q}}$ for $q \in \{1, p\}$. $\lambda \geq 0$ represents our **risk aversion**.
- Mean-Volatility has an easy interpretation:
  - The value $U(X)$ for $\lambda = 2$ denotes the level of returns X will exceed in 98% of cases.*

[1] Harry Markowitz. Portfolio selection. The Journal of Finance, 7(1):77–91, 1952.

Parameter Hedging

- Classic closed-form solution for *mean-variance* in the case $c = 0$:

$$a = \frac{\frac{1}{\lambda}\mathrm{drift}[H] - \mathrm{Covar}[Z, H]}{\mathrm{Var}[Z]}$$

  - We estimated normalized returns for our book and all hedging instruments
  - We then define the optimal hedge as “regression” of all hedging instruments vs our portfolio.
  - We essentially used our classic quant finance models as “prior”.
- Markowiz Optimization is a very well-developed field
- Used traditionally for “linear” asset allocation where a normal assumption is natural
  - Equity
  - FX
  - Long-dated debt
- Now usable for derivatives.

(*) A normal with mean μ and volatility σ is has probability 96% to be within $[\mu - 2\sigma, \mu + 2\sigma]$. Useful: https://en.wikipedia.org/wiki/Standard_deviation

Parameter Hedging

- Recall that we have assumed that $d\tilde{x}_t = \mu_t dt + \sigma_t dW_t$.
- In practise, we will want much more powerful predictors of mean and covariance of our market data.
  - Mean and volatility will depend on current market environment.
  - Estimating mean and volatility for a large vector x requires robust methods.

Parameter Hedging

- Assume that implied volatilities are given by parametric function of the type

$$\hat{\sigma}_{T,K}^{(t)} = \mathrm{atm}_T^{(t)} \left(1 + \mathrm{skew}_T^{(t)} z + \mathrm{cvxy}_T^{(t)} z^2\right) \qquad z := \log \frac{K}{F_T^{(t)}}$$

  with parameterized interpolation in expiries.
  - The point is that instead of a variable grid of implied volatilities we use the *same number of parameters* as variables.
  - This also allows us to price every option, not just those whose grid points we considered.
- Then generate hidden state h_t from the sequence of past s_t using a time series model such as GRU, LSTM, S5 or LLMs , e.g. (*)

$$h_{t+dt} = G(h_t, x_t) h_t + (1 - G(h_t, x_t)) N(h_t)$$

- Write $(\mu_t, \sigma_t) = F(h_t)$ in terms of some feed forward network F.
- Train for example conditional log-likelihood

$$\ell \equiv \frac{1}{2} \log|\sigma_t' \sigma_t| + \frac{1}{2} (dx_t - \mu_t)' (\sigma_t' \sigma_t)^{-1} (dx_t - \mu_t)'$$

(*) we drop the dependency of networks on their weights θ to declutter notation.

Parameter Hedging

- If a vol parametrization is not given, a natural choice is a set-invariant auto-encoder
- At each time step we observe options identified by $\mathrm{w}_j^{(j)} := \left(k_j^{(t)}, \tau_j^{(t)}, c_j^{(t)}\right)$ for $j = 1, \dots, d^{(t)}$ with
  - Expires $T_j^{(t)} := t + \tau_j^{(t)}$
  - Strikes $K_j^{(t)} := F_{T_T^{(t)}}^{(t)} k_j^{(t)}$
  - Call/Put indicator $c_j^{(t)}$
  - We also observe prices $H_j^{(t)}$ and implied volatilities $\hat{\sigma}_j^{(t)}$ for each option.
- We encode

$$c_t := \mathrm{E}\left(w_j^{(t)}, H_j^{(t)}, \hat{\sigma}_j^{(t)}\right)_j$$

- And decode, for any option $w = (k, \tau, c)$ its price as $D(w; c_t)$ so we wish to minimize

$$H_j^{(t)} - D(w_j^{(t)}; c_t)$$

(*) we drop the dependency of networks on their weights θ to declutter notation.

Parameter Hedging

- The encoder

$$c_t := \mathrm{E}\left(w_j^{(t)}, H_j^{(t)}, \hat{\sigma}_j^{(t)}\right)_j$$

needs to be "set invariant" in the sense that the order and number of observed option prices does not matter.

- Classic Deep Set architecture for feed forward networks F, G:

$$E(y_1, \dots, y_m) := F\left(\frac{\sum_j G(y_j)}{m}\right)$$

- Attention-like

$$E(y_1, \dots, y_m) := F\left(\frac{1}{m}\sum_j F\left(\frac{\sum_{j,k} e^{-K(y_i)Q(y_k)} G(y_k)}{\sum_{j,k} e^{-K(y_i)Q(y_k)}}\right)\right)$$

[1] Deep Sets, Zaheert et al 2018, https://arxiv.org/abs/1703.06114

Parameter Hedging

- If we have the actual historic hedging instrument values H_{t+dt} in $t + dt$ in our sample set, we can apply the same log-likelihood to the normal

$$d\widetilde{H}_t \approx \underbrace{\left\{\Theta_t + \Delta_t \mu_t(s_t) + \tfrac{1}{2} tr \Gamma_t \sigma_t(s_t)^2\right\}}_{Drift} dt + \underbrace{\Delta_t \sigma_t(s_t)}_{Risk} dW_t$$

  and vs the actual dH_t.
- In theory this captures second order effects as hedging instruments will not usually be linear functions of x_t.
- However, in the cases we tried impact was not high (mostly due to the fact that overall sample size is not large vs. model capacity):
  - S&P 500 has 1000's of options every day. Use, say, 100.
  - Twenty years of daily data ~ 5000 data points.
  - Can add other indices, etc but total number of data points when using daily data is loiw.

Parameter Hedging

- More advanced generative models which generate distributions of market parameters:
  - "Quant GANs: Deep Generation of Financial Time Series with TCNs" 2019 Magnus Wiese, Robert Knobloch, Ralf Korn, Peter Kretschmer https://arxiv.org/abs/1907.06673
  - VolGAN Cont, Vuletic @ Oxford 2023 discussed how to build models for the entire implied volatility surface of a given asset using generative methods https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4617536
  - "Forecasting implied volatility surface with generative diffusion models" Chen Jin, Ankush Agarwal, 2025 extends beyond GANs https://arxiv.org/abs/2511.07571

Parameter Hedging

Parameter Hedging: with normal assumption solve

$$\max_{a}: U(\, G^{a}\,)$$

- Trivial and fast implementation.
  - Will naturally take care of second order cross-parameter dynamic
  - This is already a much more robust approach than used in most financial institutions.
  - Can easily incorporate short-term transaction cost, liquidity constraints, limts (in form of more complicated complex functions U)
- However
  - Normal assumption questionable for derivatives with non-linear payoffs,
  - Real life the drift/covariance estimators change depending on market scenarios → requirement for robust return models for market states → see
- Addresses the hedging problem *locally* under normality.

Statistical Hedging

Statistical Hedging of non-normal Derivatives Risk

Statistical Hedging

- We assumed our portfolio returns are normal.
- In this situation Markowiz optimization is a good choice.
- Key advantage: can use greeks for complex derivatives computed with existing compute infrastructure, and multiply with returns of my market data → easy to adapt in any institution.
- Solves the hedging problem locally … if everything is normal enough.
- However, it is questionable to approximate a portfolio of derivatives by essentially symmetric returns…

Statistical Hedging

- Fails close to non-linearities
  - Example Down KO at 90%, Up KO at 110% →
  - Spot is at 91%, plotting historic 1D returns (left) and 4D (right).

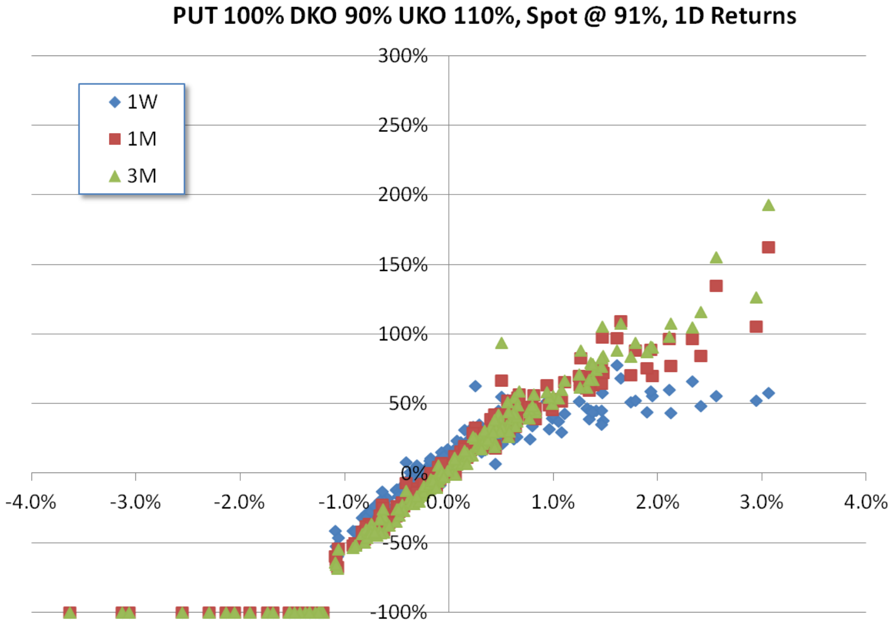

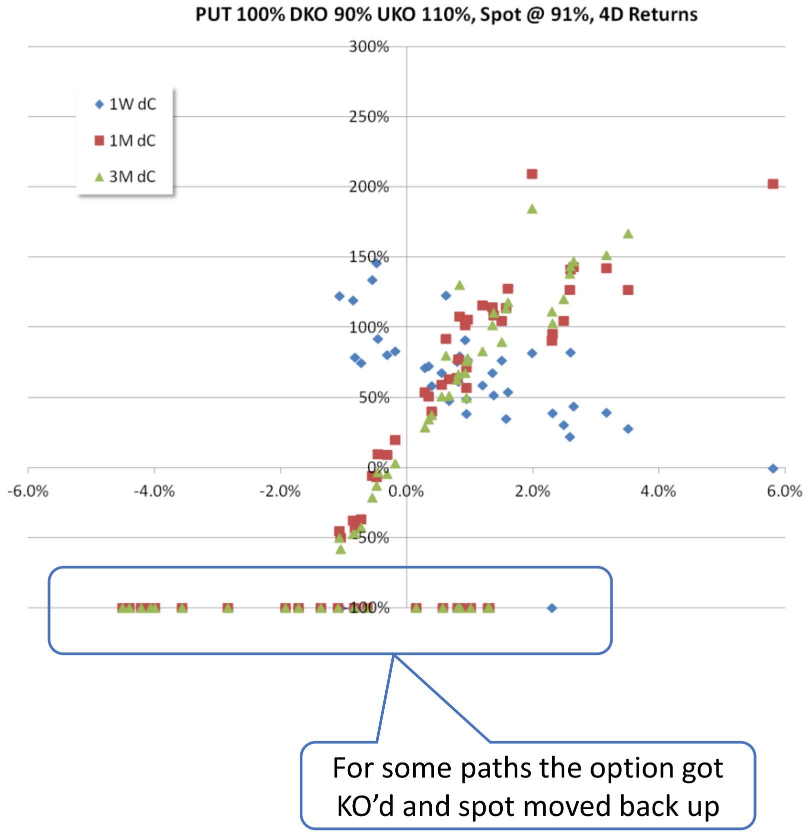

Statistical Hedging

- Towards full Hedging [1]:
  - Assume we have sufficient historic market data.
  - Re-value the return of our portfolio and all hedging instruments using such historic data.
  - Provides actual historic samples for dZ_t and dH_t.
  - This can then also be applied to longer periods than dt, for example over a week. In this case we can also hedge during that week with “model greeks”.
  - *United States patent application number 20210174441 17/110253 Filed 10-06-2021*

- Then more recent [2] uses this idea on top of a market generator for implied volatility and show performance vs naïve Greek hedging.

- What do we optimize?

---
[1] Statistical Hedging, Buehler 2017 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2913250
[2] Data-driven hedging with generative models, Cont & Vuletic 2025 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5282525

Statistical Hedging

- When we do this … why would we penalize gains of G^a as much as loses?
  - Markowiz volatility $U(X) = E[X] - \lambda\sqrt{E[(X - EX)^2]}$
  - Markowiz semi-volatility $U(X) = E[X] - \lambda\sqrt{E[2\min\{0, X - EX\}^2]}$

Statistical Hedging

- Attractive from a practical point of view but they are not *monotone.*
  - That means that it is possible that $X_1 > X_2$ but $U(X_1) < U(X_2)$ … in which case our optimizer would falsely return X_2.

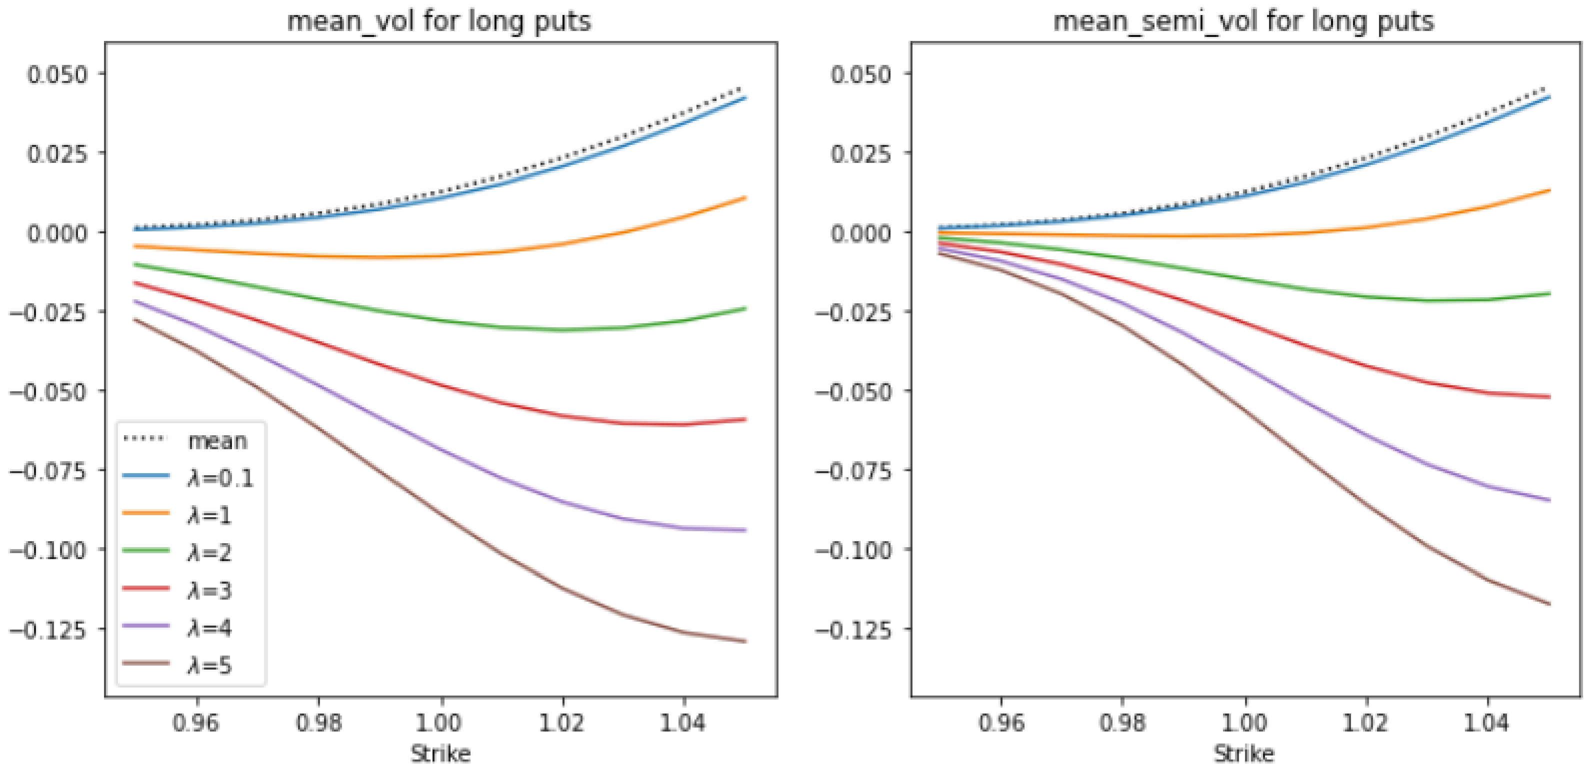

"Values" of long puts are decreasing for increasing strikes for some risk aversions (which cannot be right).

Statistical Hedging

- Another intuitive measure is the *confidence level* which we will also loosely call VaR*

$$\mathrm{VaR}(X) \coloneqq P[X \leq \cdot]^{-1}(1-\alpha)$$

  - Assume that α is your confidence level, e.g. 90%.
  - Then X will exceed $\mathrm{VaR}(X)$ in 90% of cases.
  - This generalizes the mean-volatility intuition.
  - It is well-known, however, that VaR is not concave*, and therefore not *risk averse*.
    - This happens because VaR does not consider the size of the loss beyond the confidence level. Therefore we can have a variable X_1 which in 95% of cases is just slightly above X_2 (and therefore is better), but whose loss in the 5% case way exceeds that of another variable X_2.

(*) Actual Value at Risk, VaR, is defined as –U(X).

Statistical Hedging

- This is rectified with *Expected Shortfall* or CVaR* which computes the average loss below VaR.

$$U(X) \coloneqq \mathrm{CVaR}(X) \coloneqq E[X | X \leq \mathrm{VaR}(X)]$$

- This metric has a number of attractive properties
  - It is **monotone**, i.e. if $X_1 \geq X_2$ then $U(X_1) \geq U(X_2)$ … “more is better”
  - It is **concave**. This means it is risk-averse wrt $E[\cdot]$ i.e. $U(X) \leq U(E[X])$.**
  - It is **cash-invariant** in the sense that $U(X + y) = U(X) + y$ for any real y.
    This means in particular that <u>finding any hedge is invariant of current wealth</u>.

—
(*) Classic Expected Shortfall or, mostly equivalently, CVaR, is –U(X)
(**) Note that the relation $E_Q[U(X)] \leq U\left(E_Q[X]\right)$ is not necessarily true if the measure Q is different than the measure U was defined with.

Statistical Hedging

- Motivating Cash Invariance
  - Assume $\widetilde{U}$ is concave and monotone (a "pre-kernel" in [1])
  - Assume now that we allow "writing off" any part W of a portfolio X for the benefit its worst outcome, e.g.

$$U(X) \coloneqq \sup_{W>-\infty} \widetilde{U}(X-W) + \inf W$$

  - Since $\widetilde{U}$ was monotone $-W \le -\inf W$ hence

$$U(X) \coloneqq \sup_{w\in R} \widetilde{U}(X-w) + w$$

  - The functional U is cash-invariant [1]

[1] Statistical Hedging; Buehler 2013 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2913250

Statistical Hedging

- We call a **monotone**, **concave**, and **cash-invariant** functional U which is normalized to $U(0) = 0$ a **monetary utility.**
  - Then $-U(X)$ is a (normalized) **convex risk measure**.
- A monetary utility is called **law-invariant** with respect to a measure Q if $U(X) = U(Y)$ for all variables which have the same law under Q.
- A Q-law invariant monetary utility is **risk averse** in the sense that

$$U\left(E_Q[X]\right) \geq U(X)$$

- We call U **coherent** if $U(nX) = nU(X)$ for positive n.
  Coherence is not usually a desirable property.

Statistical Hedging

- **Optimized Certainty Equivalent** [1]: let u be an increasing, concave C^1 utility function normalized to $u(0) = 0$ and $u'(0) = 1$.*
- Let Q be a measure. Then the following functional is a Q-law invariant monetary utility:

$$U(X) \coloneqq \sup_{y \in R} : E_Q[u(X + y)] - y$$

- For our setting this translates to

$$\sup_{a \in R^n} U(G^a) = \sup_{a \in R^n, y \in R} E[u(dZ_t + a' dH_t + y) - y] - c_t(a)$$

- From a machine learning perspective, this means that we can optimize a complex risk-adjusted return using what is still an expectation and therefore suitable for batching (as opposed to mean-variance, for example, where variance is defined across the training set).

[1] Aharon Ben-Tal and Marc Teboulle. An old-new concept of convex risk measures: The optimized certainty equivalent. Mathematical Finance, 17(3):449–476, July 2007.

Statistical Hedging

- Exponential utility $u(x) := (1 - e^{-\lambda x})/\lambda$
  - Exponential utility is the generalization of mean-variance.
    Indeed, if $X = \mu + \sigma Y$ for a normal Y then $E[U(X)] = \mu - \frac{1}{2}\lambda\sigma^2$.
  - Annoyingly, the exponential is *very* averse vs large losses. Indeed a short position in a Black Scholes stock has infinite negative utility.
- Truncated exponential utility: use quadratic on the downside.
- CVaR: $u(x) = (1 + \lambda)\min\{0, x\}$
- Handerson and Hobson [1] proposed $u(x) = (1 + \lambda x + \sqrt{1 + \lambda^2 x^2})/\lambda$
- "Quadratic" utility: quadratic function cut off and shifted to satisfy $u'(0) = 1$.

---
[1] V. Henderson and D. Hobson. Utility indifference pricing: An overview. 2004.
https://warwick.ac.uk/fac/sci/statistics/staff/academic-research/henderson/publications/indifference_survey.pdf

Statistical Hedging

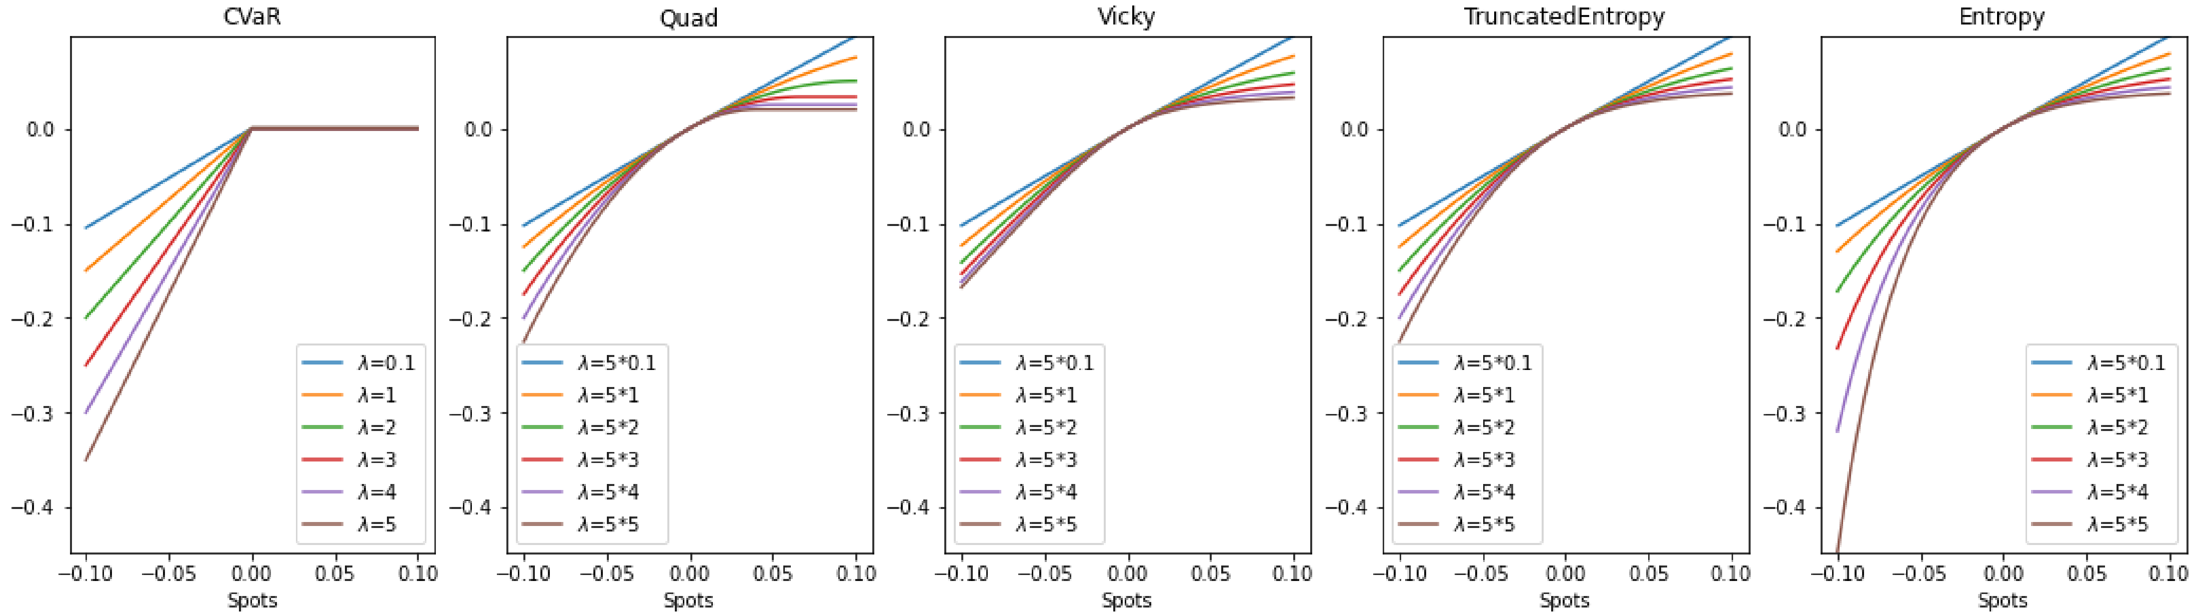

Statistical Hedging

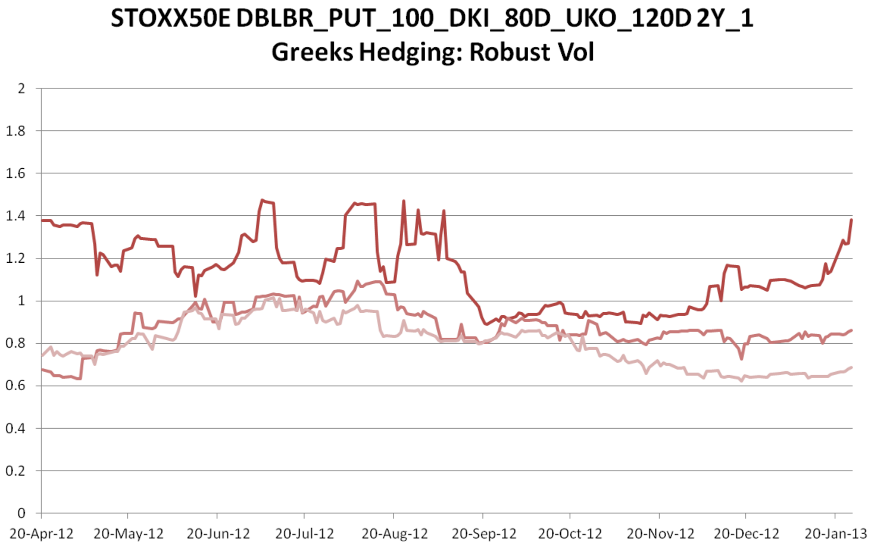

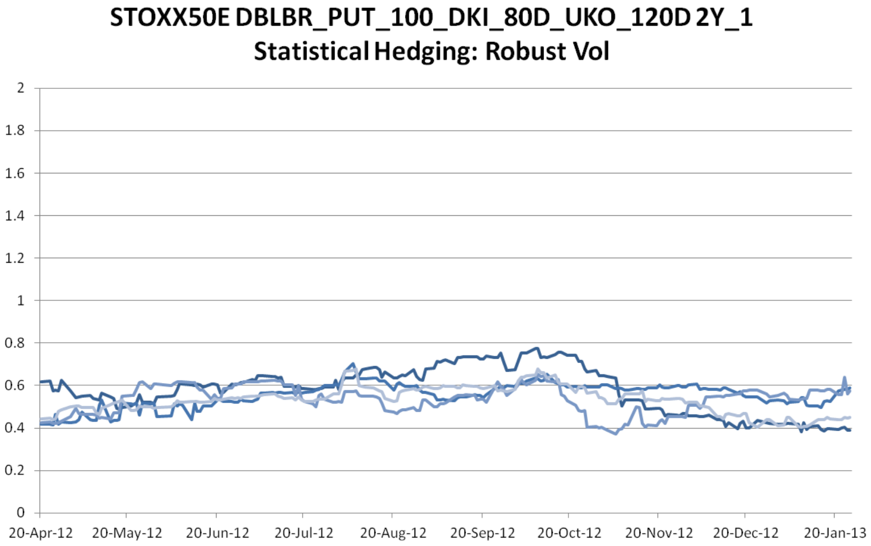

Less vol is better

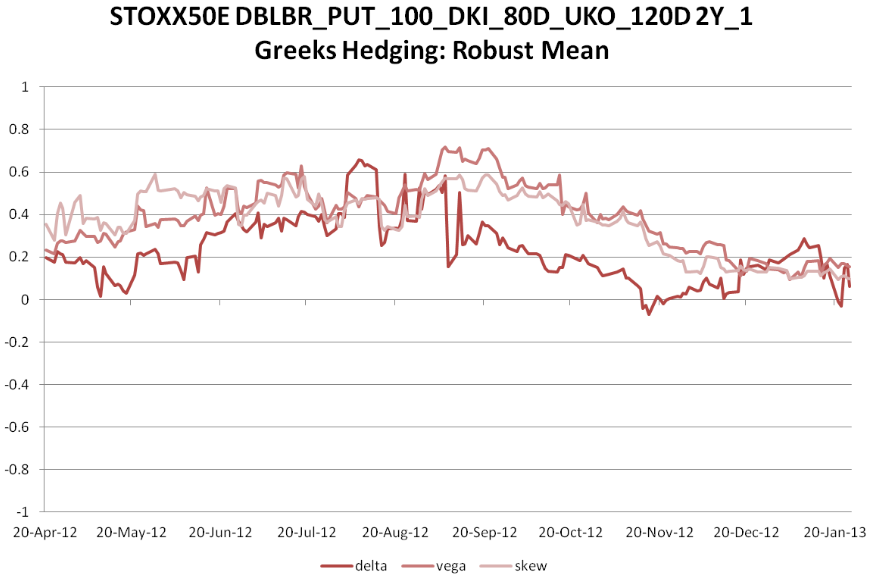

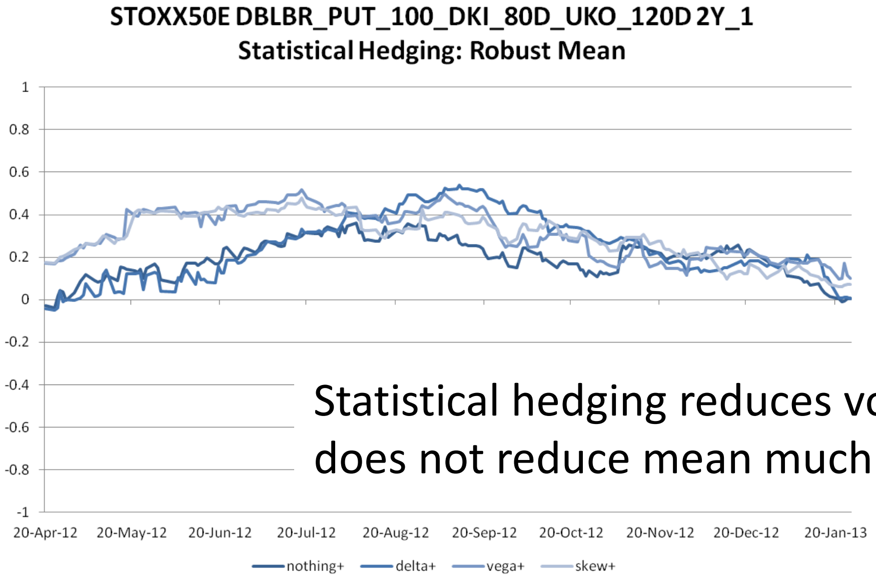

Mean should be zero

Statistical hedging reduces volatility but does not reduce mean much (CVaR 90%)

Statistical Hedging

Statistical Hedging

- Returns for G^a are based on historic scenarios.
- Due to non-normality non-trivial objective functions are preferred
- Sound choice is the concept of **optimized certainty equivalents** as they are numerically very efficiently solvable using standard cone optimizers (use cvxpy for experimentation)
- One-step optimization problem becomes:

$$\sup_{a \in R^n} U(G^a) = \sup_{a \in R^n, y \in R} E[u(dZ_t + a' dH_t + y) - y] - c_t(a)$$

Statistical Hedging

Data-driven hedging with generative models [1]

- Use a market scenario generator generate dZ and dH.
- In [1] use VolGAN, essentially a GAN model for joint implied volatility and stock dynamics.
- Then solve the statistical hedging problem. In [1] the authors hedge a 1M long straddle and compare delta-vega hedging with statistical hedging ("regression")

Covid-19 period	Method	Statistics					
		Mean	Median	Std	5% VaR	2.5% VaR	1% VaR
Included	Unhedged position	0.16	0.29	60.58	78.95	102.07	155.29
	Delta hedging	-1.23	-0.44	32.70	19.49	36.33	58.63
	Delta-vega hedging	0.98	0.00	29.70	10.90	19.70	43.72
	Data-driven hedging	0.55	-0.16	32.98	12.79	23.42	50.79
Excluded	Delta hedging	-1.46	-0.36	8.40	13.22	22.84	36.66
	Delta-vega hedging	-0.37	0.00	9.34	9.58	16.67	34.81
	Data-driven hedging	-1.05	-0.18	8.15	10.55	17.32	33.85

[1] Data-driven hedging with generative models, Cont & Vuletic 2025 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5282525

Statistical Hedging

Statistical Hedging

- Used at scale in JP Morgan for Flow Derivatives (2022)
  https://www.risk.net/awards/7928696/equity-derivatives-house-of-the-year-jp-morgan

- Good
  - Provides base line hedging strategy for “any” portfolio of derivatives
  - Very versatile, robust, and “model-free”.
  - Conceptually trivial, but represents a major progress towards automated trading:
    - No arbitrarily defined greeks required
    - All major (local) dynamics naturally captured
  - Easily applicable natively in markets with lots of data e.g. intraday trading of derivatives

- But
  - Price of an instrument given by “classic model”. What if that is wrong – as it likely is?
  - Hedge only locally optimal.
  - Generating sufficient scenarios requires historic data and for OTC products the ability to re-value the product.*

—

(*) From pure compute complexity running market scenarios is of similar complexity to running a large number of Greeks. However, most institutions will have the former readily available while the latter requires additional implementation.

Please ask questions