

Learning to Trade I

Statistical Hedging


TU Munich
Visiting Prof. Hans Buehler
Prof. Blanka Horvath

TU München 2022

<http://deep-hedging.com>

Learning to Trade

- “Trading” is a business
 - Sell derivatives to “franchise” clients who have a genuine requirement
 - Institutional clients: managing future cashflows or liabilities, commodity hedging, M&A activity, cash management, loans
 - Retail clients: directional trading, yield enhancement, mortgages
- Hedge the resulting franchise exposure in the market
 - Exchanges, multi-dealer platforms, even decentralized.
 - Hedge excess risk with “arbitrageurs”.
- Idea generation:
 - Propose ideas for both groups of clients based on our knowledge of their business



Learning to Trade

- Today's derivative models
 - Focus on “fitting the market”
 - Interpolate desired hedging instruments (spot, forwards, option prices)
 - Expand model with additional stylistic risk factors to capture higher order effects
 - Examples
 - Local Volatility (fits all forwards, discount factors, FX or Equity option prices)
 - Hull-White fits ATM all swap rates, and a strip of swaptions e.g. ATM
 - Great paper “The Smile Calibration Solved” [1]
 - No statistical claim on those mechanics being particularly realistic

Learning to Trade

- Today's derivative models
 - Focus on “fitting the market”
 - We compute “greeks” with respect to non-model parameters
 - E.g. we compute “vega” for Local Volatility ... but this is a one-factor process
$$\frac{dS_t}{S_t} = \mu_t dt + \sigma(t, S_t) dW_t$$
 - In fact, we will usually compute a whole host of greeks which doesn't make sense in a “complete market”

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
 - 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


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
 - Assumption of unique martingale measure in the first place ...
... and that it is somewhat related to the stylistic model we are using (e.g. local vol).
 - Absence of market frictions
 - Absence of risk aversion

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

Learning to Trade



Early example of analysis of hedging performance


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

[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
- This is a historical effect: when derivatives analytics where developed
 - Data and compute where 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



Hard to
automate

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 TensorFlow and PyTorch
 - 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

Deep Hedging Framework

The Market

- The market is observed under the **statistical measure** P .
- Everything is in discrete time $\rightarrow$ markets not complete.
- We call s_t the **state** of the market. It represents all known 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 .

Deep Hedging 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.

Deep Hedging Framework

Hedging Instruments

- For notional simplicity we will fix a total number of n instruments. We can restrict their trading to ensure realistic behaviour. See [1] for a more detailed discussion.
- Any instrument we may trade has to have a mark-to-**model price** given for example by an official closing price, a weighted bid/ask, or a classic valuation model.
We denote at time t the **model prices** by $H_t^i \equiv H^i(s_t)$ of n hedging instruments.
- 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 $X_t := \tilde{X}_t / B_t^\$$.
- The model prices of (sum of) the instruments in our existing portfolio is $Z_t \equiv Z(s_t)$.

Framework

- We split s_t into
 - True random drivers x_t which we think are relevant risks.
Mathematically, these are those that have a quadratic variation.
 - Examples: spot prices, FX, option prices, swap rates
 - States τ_t which do not have quadratic variation, and are absolutely continuous with respect to time, i.e. $\tau_t = \int^t \tau'_t dt$.
In particular, the vector τ contains time itself.
 - Examples: interest payments, dividend payments, realized volatility
- See [1] for background

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

[1] Delta-Hedging Works: On Market Completeness in Diffusion Models, Buehler 2009, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1464865

Framework

- Trading cost
 - Trading $a \in \mathbb{R}^n$ units of H at t will cost $c_t(a)$ in excess of the model price H_t .
 - The function c is normalized to $c_t(0) = 0$, non-negative, and convex.
 - Convexity excludes fixed trading cost.
 - Trading is 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:

$$\hat{c}_t(a) := \gamma_\Delta |a \cdot \Delta| + \gamma_V^1 |a \cdot V| + \gamma_V^2 |a \cdot V|^2$$

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

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

Lesson 1.1: Parameter Hedging

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

Parameter Hedging

- For $X_t \in \{Z_t; H_t^1, \dots, H_t^n\}$ recall $X_t \equiv X(s_t)$ we may use Taylor and obtain

$$dX_t = \underbrace{\partial_\tau X_t d\tau_t}_{\text{Theta}} + \underbrace{\partial_x X_t dx_t}_{\text{Delta}} + \underbrace{\frac{1}{2} \partial_{xx}^2 X_t d\langle x \rangle_t}_{\text{Gamma}} + O(x_t^3)$$

- Important: even if the underlying model (e.g. Black Scholes) only has spot as “model state” we still compute derivatives to all other parameters (in the Black & Scholes case: interest rates, forward rates, volatility).

Parameter Hedging

Greek Hedging

- For $X_t \in \{Z_t; H_t^1, \dots, H_t^n\}$ recall $X_t \equiv X(s_t)$ we may use Taylor and obtain

$$dX_t = \underbrace{\partial_\tau X_t d\tau_t}_{\text{Theta}} + \underbrace{\partial_x X_t dx_t}_{\text{Delta}} + \underbrace{\frac{1}{2} \partial_{xx}^2 X_t d\langle x \rangle_t}_{\text{Gamma}} + O(x_t^3)$$

- Our gains process is:

$$G^a := dZ_t + a \cdot dH_t - c_t(a)$$

- We can write the Taylor expansion for dZ and dH including higher terms.
- Idea is to minimize exposure to each such “Greek” (derivative)
- As it stands, hard to tell which of the terms is most important to hedge...
 $\partial_{S\&P \text{ spot}} G^a$ vs. $\partial_{S\&P \text{ vol}} G^a$ vs $\partial_{USD \text{ 3m Swap Rate}} G^a$?
- Relies on heuristics. *Parameter Hedging* formalizes this approach with data

Parameter Hedging

- Let $dx_t = \mu_t dt + \sigma_t dW_t$ under the *statistical* measure P . We sort above with some abuse of notation into

$$dX_t \approx \underbrace{\left\{ \partial_\tau X_t \tau'_t + \partial_x X_t \mu_t + \frac{1}{2} \text{tr} \partial_{xx}^2 X_t \sigma^2 \right\}}_{\text{Drift}} dt + \underbrace{\partial_x X_t \sigma_t}_{\text{Risk}} dW_t$$

- This formula provides us with a normal approximation of the **returns of any instrument**.
- Therefore we have an estimate of the distribution of the portfolio for any trading **action** a .
- This gives us the **gains process**

$$G^a := dZ_t + a \cdot dH_t - c_t(a)$$

- Meaningful data-driven weighting scheme for our Greeks when compared to Greek Hedging.

Parameter Hedging

- We now have an approximation of the statistical distribution of

$$G_t^a := dZ_t + a \cdot dH_t - c_t(a)$$

- Lends itself rather obviously to optimization programs of the form

$$\max_a U(G^a)$$

- What should U be?

Parameter Hedging

- Classic choices for $dG^a := dZ_t + a \cdot dH_t - c_t(a)$ with

$$dG_t^a \approx dX_t \approx \underbrace{\left\{ \partial_\tau X_t \tau'_t + \partial_x X_t \mu_t + \frac{1}{2} \text{tr} \partial_{xx}^2 X_t \sigma^2 \right\}}_{\text{Drift}} dt + \underbrace{\partial_x X_t \sigma_t}_{\text{Risk}} dW_t - \underbrace{c_t(a)}_{\text{Cost}}$$

- Markowitz seminal work [1] suggests using
 - Markowitz Mean-Variance $E[X] - \frac{1}{2} \lambda E[(X - EX)^2]$
 - Markowitz Mean-Volatility $E[X] - \lambda E[(X - EX)^2]^{\frac{1}{2}}$
 - 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**.
- Close form for Mean-Variance for $c \equiv 0$

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

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

Parameter Hedging

- Mean-Volatility in particular has an easy interpretation:
 - The value $U(X)$ for $\lambda = 2$ denotes the level of returns X will achieve in 98% of cases.*
- However, Mean-Volatility is also **coherent** i.e. $U(aX) = aU(X)$.
 - This means that we do not care about the size (notional) of our risk outside trading cost.

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

Parameter Hedging

- Markoviz Optimization is a very well developed field
- Used traditionally for “linear” asset allocation where a normal assumption is natural
 - Equity
 - FX
 - Long-dated debt
- Commercial covariance estimators are available for such assets
- Main challenge is modelling *cost* and ... the drift, μ .
Called **alpha** when it comes to a directional forecast.

Parameter Hedging

Parameter Hedging: with normal assumption solve

$$\max_a: U(G^a)$$

- Trivial and fast implementation.
- No machine learning required.
- A number of observation:
 - Our approach will naturally take care of second order cross-parameter dynamics
 - This is already a much more robust approach than used in most financial institutions.
 - Classic Greek Hedging is equivalent to some arbitrary choices of $\mu, \sigma \dots$ c.f. [1]

[1] In “Volatility Markets: Consistent Modelling, Hedging and Practical Implementation” we discussed the use of L1/L2 optimizers for managing derivatives risk without “statistics” https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1118245

Lesson 1.2: Statistical Hedging

Statistical Hedging of Daily Derivatives Risk

Statistical Hedging


- We assumed our portfolio returns are normal
 - In this situation there is not much more to do than Markovitz optimization.
 - However, it is questionable to approximate a portfolio of derivatives by essentially symmetric returns.
- Much more powerful as proposed in [1]:
 - Estimate returns of *today's* hedging instrument and portfolio using historic scenarios.
 - Key is to use the *relative-same* instruments not “the same” fixed instruments:
 - Keep time-to-maturity constant in business time
 - Keep moneyness for strikes and barriers constant relative to market levels
 - Ensure that past events such as barrier breaches are kept
 - *We are still using our classic derivative model values !*
 - We can even extend the “one day” horizon by longer periods such as a week with intermittent (mechanical) hedging.
 - First discussed 2012, 2013 on the Global Derivatives conferences.

Statistical Hedging

- When we do this ... why would we penalize gains of G^a as much as losses?
 - Markoviz semi-variance $U(X) = E[X] - \frac{1}{2}\lambda E[2 \min\{0, X - EX\}^2]$
 - Markoviz semi-volatility $U(X) = E[X] - \lambda E[2 \min\{0, X - EX\}^2]^{\frac{1}{2}}$

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.

Statistical Hedging

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

$$\text{VaR}(X) := P[X \leq \cdot]^{-1}(1 - \alpha)$$

- Assume that α is your confidence level, e.g. 90%.
- Then X will with 90% VaR $U(X)$.
- 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 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.

$$U(X) := \text{CVaR}(X) := E[X | X \leq \text{VaR}(X)]$$

- 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 finding any hedge is invariant of current wealth

(*) Classic Expected Shortfall or, mostly equivalently, CVaR, is $-U(X)$

(**) Note that the relation $E_Q[U(X)] \leq U(E_Q[X])$ is not necessarily true if the measure Q is different than the measure U was defined with.

Statistical Hedging

- Motivating Cash Invariance

- Assume $\tilde{U}$ 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.

$$U(X) := \sup_{W > -\infty} \tilde{U}(X - W) + \inf W$$

- Since $\tilde{U}$ was monotone $-W \leq -\inf W$ hence

$$U(X) := \sup_{w \in \mathbb{R}} \tilde{U}(X - w) + w$$

- The functional U is cash-invariant [1]

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

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

- 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:

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

- We call the OCE *strict* if it is strictly increasing and strictly concave.
- We note that cash invariance implies

$$U(G^a) = U(Z_{t+dt} + a \cdot H_{t+dt}) - Z_t - a \cdot H_t - c_t(a)$$

[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

- **Optimized Certainty Equivalent** [1]: let u be a strictly increasing, strictly 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:

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

- We note that cash invariance implies

$$U(G^a) = U(Z_{t+dt} + a \cdot H_{t+dt}) - Z_t - a \cdot H_t - c_t(a)$$

[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.

(*) strict monotonicity and concavity are only required for some of the results here. See [1].


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

- Practical Comment

For most practical applications, we look at delta-hedged returns of instruments. In this case, monotonicity seems to matter a lot less. So far we found only pathological examples for when this becomes an issue.

That means semi mean-variance and semi mean-volatility remain a popular return metrics.

Statistical Hedging

Statistical Hedging

- Returns for G^a are based on historic scenarios using the “relative-same” instruments.
- 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)

$$U(G^a) = U(Z_{t+dt} + a \cdot H_{t+dt}) - Z_t - a \cdot H_t - c_t(a)$$

Statistical Hedging

Statistical Hedging

- Used at scale in JP Morgan for Flow Derivatives
<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
- But
 - Price of an instrument given by “classic model”. What if that is wrong – as it likely is?
 - Hedge only locally optimal.

There isn't any machine learning ... ?

Please ask questions