Deep Hedging: from Theory to Practice

From Greeks to Hedging under Market Frictions

Imperial Frontiers in Quantitative Finance
April 24th, 2019

Dr. Hans Buehler * Global Head Equities Analytics, Automation and Optimization & Investor Services QR
Baranidharan Mohan * EMEA Analytics QR
Dr. Ben Wood * EMEA Head Equities Modelling QR

Joint research project with Josef Teichmann ETH Zurich

** Opinions expressed in this paper are those of the authors, and do not necessarily reflect the view of JP Morgan.*

J.P.Morgan
QUANTITATIVE RESEARCH

Disclaimer

J.P.Morgan

J.P.Morgan

QUANTITATIVE RESEARCH

Summary

- **Greek Hedging** is a legacy approach once justified by lack of data and computational power
- **Statistical Hedging** brings data-driven risk management but still relies on classic models for pricing
- **Deep Hedging** defines a new data-driven "AI" reinforcement learning risk and pricing concept for derivatives.
  Its challenges are
  - Realistic and robust simulation of markets
  - Efficient modern Reinforcement Learning techniques for rapid evolution

Framework

J.P.Morgan

Portfolio

- We are given an portfolio Z of securities and derivatives, all of which are assumed to expire before some terminal maturity T. Negative values represent losses.
  - We assume interest rates are deterministic, hence we may consider discounted variables.
  - We assume that FX transactions are cost-free, hence we may assume w.l.g. that all assets are denominated in the same currency.
  - The portfolio and all subsequent instruments are considered as ``total return" assets. The total return of Z, i.e. the sum of all cash flows of Z at T is denoted by Z_T.
- We set $T=t_m>0$ and denote by $0=t_0<\ldots<t_{m-1}<T$ possible intermediate hedging days.

Framework

J.P.Morgan

Mark-to-Model

- At each *t* we observe relevant market data such as spots, implied vols, rates, etc. We denote this set of market data by $\boldsymbol{X}=(X^1,\dots,X^K)$.
- We assume that for each element of our portfolio we have a way to compute a "mark-to-model" value from $\boldsymbol{X}$.
  - Volume-weighted mid-prices for equity, FX
  - Classic derivative pricing models such as Stochastic-Local Vol for derivatives.
- Combined, this yields a (mark-to-) **model value** Z_t for our portfolio.
  - This is not a tradable quantity.
  - This meta model will yield a range of classic Greeks in the form of first or higher order derivatives.

Quant Finance as an Interpolation Problem

- Classic derivative models are neither equivalent to the statistical measure Q, nor have they been designed to behave realistically. Their primary objective is **interpolation** between observable market data in $\boldsymbol{X}$.
  A "good" model is measured by:
  - Quality of fit to reference market data in $\boldsymbol{X}$, e.g. implied volatilities.
  - Speed of calibration and execution
  - Stylized dynamics such as stochastic volatility or stochastic interest rates.

Framework

J.P.Morgan

Hedging

- We are given a range of liquid **hedging instruments** $\boldsymbol{H}=H^1,\dots,H^n$ such as options, swaps, futures, ETFs, stocks, FX etc.
- The **mid-price** at time t is denoted by $\boldsymbol{H}_t$, which is a **model value** computed from $\boldsymbol{X}_t$, for example the volume-weighted mid-price for an equity. It is not a tradable quantity.
- The actual price for trading $\boldsymbol{a}=(a^1,\dots a^n)$ is given by

$$\boldsymbol{H}_t(\boldsymbol{a}) := \boldsymbol{a}\,\boldsymbol{H}_t + c_t(\boldsymbol{a})$$

  in terms of a non-negative and normalized cost function c_t. We usually assume c_t is convex, but there are valid examples it is *not*, e.g. fixed fees per trade.
    - Cost can depend on past trading activity to model impact. Research topic: consistent impact model for option prices.
- The formal mark-to-model P&L over the period dt due to trading a in t is given as

$$\boldsymbol{a}\, d\boldsymbol{H}_t - c_t(\boldsymbol{a})\,.$$

  We note that this does not take into account unwind cost.

Framework

J.P.Morgan

Liquidity

- Not all instruments are tradable at all times:
  - An exchange traded option $(S_r - k\, S_t)^+$ for $r>t$ may only be traded at times $u \in [t,r)$, i.e. when S_i is known and therefore the strike is fixed.

- We denote by A_t the convex, non-empty set of **admissible actions** at time *t*.
  - The set A_t may depend all observable market data and our historic trading decisions. For example,
    - Short-sell restrictions
    - Available liquidity as a function of past trading activity
    - Risk limits for our overall position (e.g. maximum Vega exposure)
  - We call $\pi=(\boldsymbol{a}_0,\ldots,\boldsymbol{a}_{m-1})$ with $\boldsymbol{a}_t \in A_t$ a **trading policy**; we will usually omit the "$\in A_t$" unless necessarily.
  - $\delta_t := \boldsymbol{a}_0 + \ldots + \boldsymbol{a}_{t-1}$ is our current position in *H.*
  - At time *t* our mark-to-market is $M_t \coloneqq Z_t + \boldsymbol{\delta}_t \boldsymbol{H}_t$.

The End of the Greek Era

J.P.Morgan

QUANTITATIVE RESEARCH

From Greeks to Statistical Hedging

- Before we focus on Deep Hedging, we discuss alternative approaches.
- Each of these will hedge only over a given period *dt*.

The Greek Era

J.P.Morgan

QUANTITATIVE RESEARCH

Risk-Neutral Hedging

- We denote by $\boldsymbol{\delta}^{\$} := (X_1 \cdot \partial_{X_1}, \dots, X_K \cdot \partial_{X_K})$ the first order "Cash Delta" derivative operator to relevant observable market parameters. We denote by $\Gamma^{\$}$ the respective second order "cash Gamma" derivatives and by Θ "Theta".
- The risk-neutral paradigm stipulates to reduce all Delta exposure to zero, i.e. to minimize regardless of cost in each t

$$\sup_{\boldsymbol{a}_t}: -\left\|\boldsymbol{\delta}^{\$} M_t + \boldsymbol{a}_t \boldsymbol{\delta}^{\$} \boldsymbol{H}_t\right\| \qquad (M_t := Z_t + \Delta_t \boldsymbol{H}_t)$$

NB: for the L^2 norm this can be solved using the Pseudo-inverse of $\boldsymbol{\delta}^{\$} \boldsymbol{H}_t$.

- This formulation does not take into account trading cost.
  Ad-hoc heuristics

$$\sup_{\boldsymbol{a}_t}: -\lambda \left\|\boldsymbol{\delta}^{\$} M_t + \boldsymbol{a}_t \boldsymbol{\delta}^{\$} \boldsymbol{H}_t\right\|_2 - c_t(a_t)$$

Pros

- Fast
- Needs only today's market data.

Cons

- Inconsistent:
  - No sense of carry
  - Difficult to add cost to this approach as the "Cash Deltas" have only nominally a \$ interpretation: how do we know whether spending 100k\$ to hedge a nominal 1m\$ vega position is worth it?
- Unrealistic: does not account for "Skew Delta".

The Greek Era

Parameter Hedging

– We assume that we are estimating a simple normal model

$$\frac{d\boldsymbol{X}_t}{\boldsymbol{X}_t} = \boldsymbol{\mu}_t(\boldsymbol{X}_t)dt + \boldsymbol{\Sigma}_t(\boldsymbol{X}_t)d\boldsymbol{W}_t$$

for our market parameters, e.g. simply by replaying historic data.
This model gives rise to the operator

$$d(\cdot) = \Delta^{\$}\boldsymbol{\Sigma}_t d\boldsymbol{W}_t + \left\{\Delta^{\$}\boldsymbol{\mu}_t + \frac{1}{2}\boldsymbol{\Sigma}_t\Gamma^{\$}\boldsymbol{\Sigma}_t + \Theta\right\}dt$$

– Applying Markoviz' mean-variance approach to $M_t \coloneqq Z_t + \Delta_t \boldsymbol{H}_t$ yields the intuitive "carry" expression

$$\sup_{\boldsymbol{a}_t}: \left\{\boldsymbol{a}_t\Delta^{\$}\boldsymbol{\mu}_t + \frac{1}{2}(\boldsymbol{a}_t\boldsymbol{\Sigma}_t)'\Gamma^{\$}(\boldsymbol{a}_t\boldsymbol{\Sigma}_t) + \Theta\right\}dt - \frac{\lambda}{2}(\boldsymbol{a}_t\Delta^{\$}\boldsymbol{\Sigma}_t)'(\boldsymbol{a}_t\Delta^{\$}\boldsymbol{\Sigma}_t)dt - c_t(\boldsymbol{a}_t)$$

Carry term combined both of Gamma-Theta carry and drift

Risk term driven by the covariance matrix of ***X***

Implementation cost

The Greek Era

J.P.Morgan

QUANTITATIVE RESEARCH

Parameter Hedging

$$\sup_{\boldsymbol{a}_t}:\left\{\boldsymbol{a}_t\boldsymbol{\delta}^{\$}\boldsymbol{\mu}_t+\frac{1}{2}(\boldsymbol{a}_t\boldsymbol{\Sigma}_t)'\Gamma^{\$}(\boldsymbol{a}_t\boldsymbol{\Sigma}_t)+\Theta\right\}dt-\frac{\lambda}{2}(\boldsymbol{a}_t\boldsymbol{\delta}^{\$}\boldsymbol{\Sigma}_t)'(\boldsymbol{a}_t\boldsymbol{\delta}^{\$}\boldsymbol{\Sigma}_t)dt-c_t(\boldsymbol{a}_t)$$

Carry term combined both of Gamma-Theta carry and drift

Risk term driven by the covariance matrix of $\boldsymbol{X}$

Implementation cost

Pros

– Data-driven combination of carry and risk
– Captures well-known effects such as “skew delta”
– Fast

Cons

– Normal Approximation:
  - Does not capture strong non-linearities such as short-term barriers
  - I.e. can lead to strictly worse “hedged” portfolios
– Hedges are only locally optimal / what is the optimal horizon *dt* vs. the absolute cost term
– Still requires mark-to-model values for Greeks and pricing

Statistical Hedging I – quadratic case

- In our previous example, we effectively approximated the return of a derivative over *dt* as a normal.
- In *Statistical Hedging* [2013] we proposed replacing this approximation with genuine historic returns of “the same” derivatives.
  - For each historic day, use a derivative with the *then*-same moneyness and time-to-maturity, and compute that derivative’s return over *then*-dt.
    - Do not use today’s fixed derivative terms.
    - Compute returns of fixed instruments.
    - For path-dependent options, keep past states consistent (e.g. past barrier breaches).
  - This yields genuine historic returns of both *Z* and *H* from *t* to *t+dt*.
- In practise, we will want to adjust the drift term to take into account views on relevant carry: e.g. just because the S&P went up for the last years does not mean we wish to capture this with our model → topic of model uncertainty

Beyond Greeks

J.P.Morgan

QUANTITATIVE RESEARCH

Statistical Hedging I – quadratic case

– Given returns dM_t and $d\boldsymbol{H}_t$ we may now solve the Markoviz problem

$$\sup_{\boldsymbol{a}_t}: \mathrm{E}[dM_t + \boldsymbol{a}_t d\boldsymbol{H}_t]dt - \frac{\lambda}{2}\mathrm{Var}[dM_t + \boldsymbol{a}_t d\boldsymbol{H}_t] - c_t(\boldsymbol{a}_t)$$

Carry term derived from historical performane

Risk term driven by the covariance matrix of ***X***

Implementation cost

– We note that if all variables are well approximated by our previous normal representation, then the two approaches coincide.

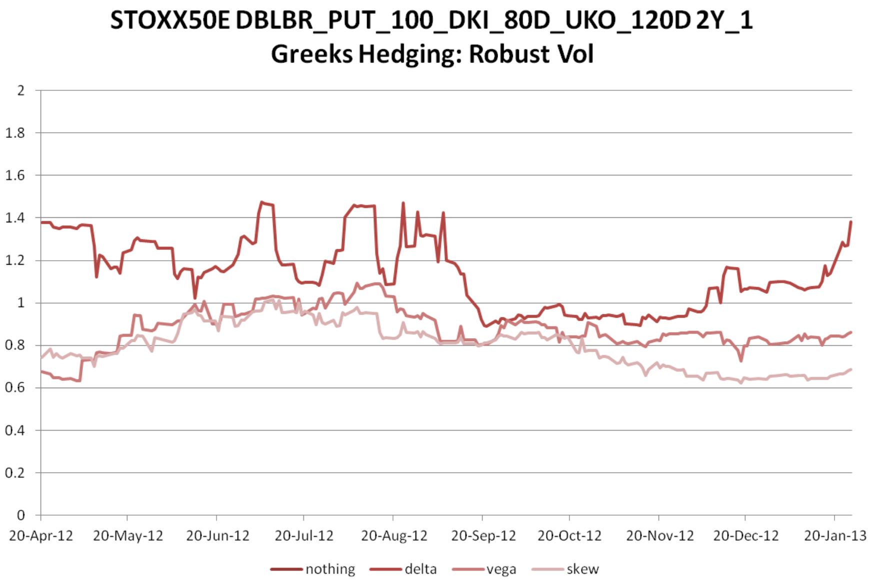

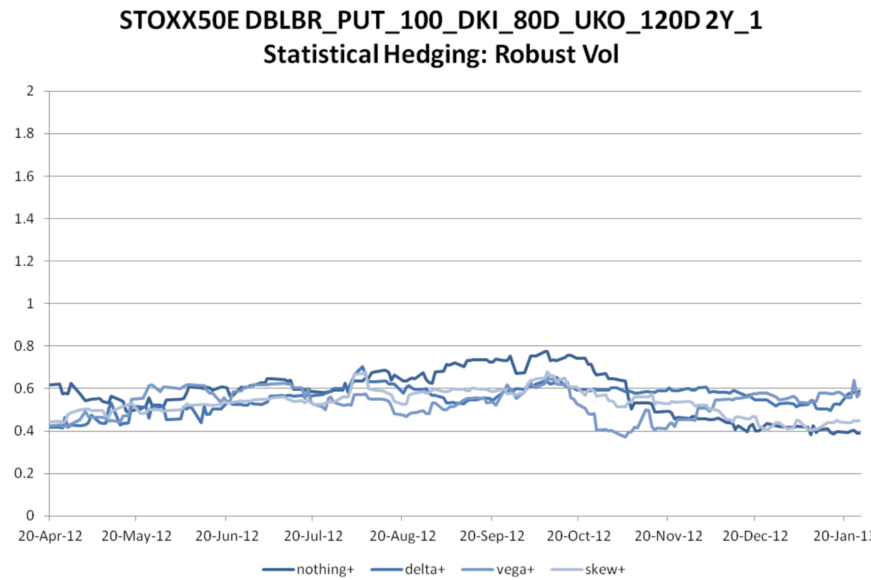

Beyond Greeks

J.P.Morgan

QUANTITATIVE RESEARCH

What is wrong with Markoviz for Derivatives?

Mean-Dispersion measures are not monotone

$$U(A) := \mathrm{E}[A] - \frac{\lambda}{2}\|A - \mathrm{E}[A]\|_p$$

… it makes very little sense to optimize over a non-monotonic objective.

Graph shows the fair risk-adjusted price for taking on a short position in a digital call.
The lower the strike, the higher the required compensation.

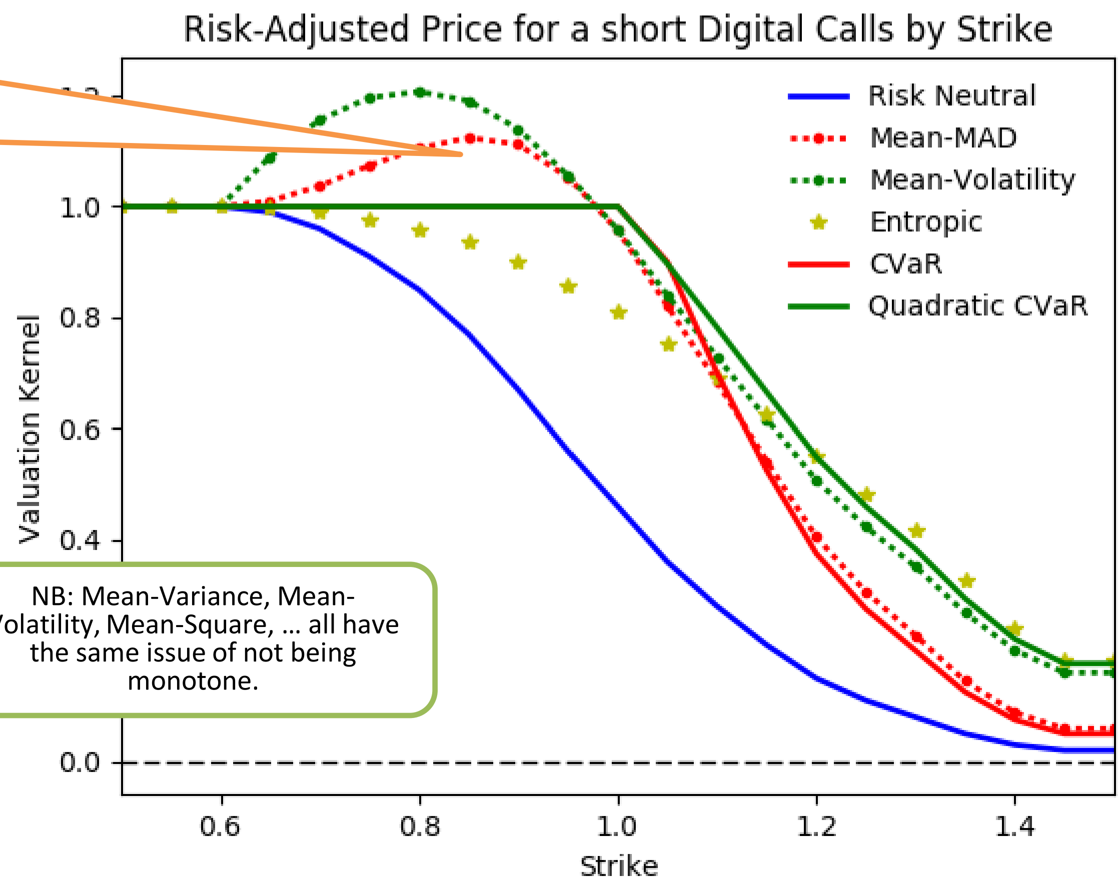

NB: Mean-Variance, Mean-Volatility, Mean-Square, … all have the same issue of not being monotone.

Risk aversion levels manually calibrated to roughly fit between CVaR and Mean-Dispersion measure: Mean-MAD 1, Mean-Vol 1, CVaR 1.5, Quadratic VaR 0.5.

Beyond Greeks

J.P.Morgan

What is wrong with Markoviz for Derivatives?

Using entropy vs. mean-variance to hedge a net long butterfly position maturing tomorrow. Data generated with BS model, 50% vol. The same vanillas which compose the butterfly are available to trade for free.

Entropy

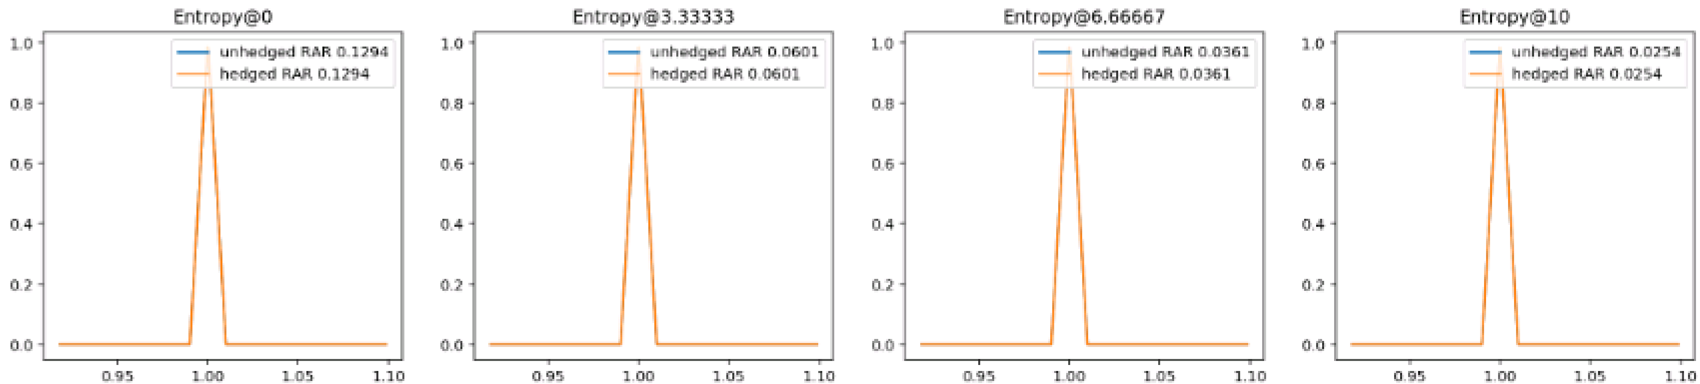

MeanVar

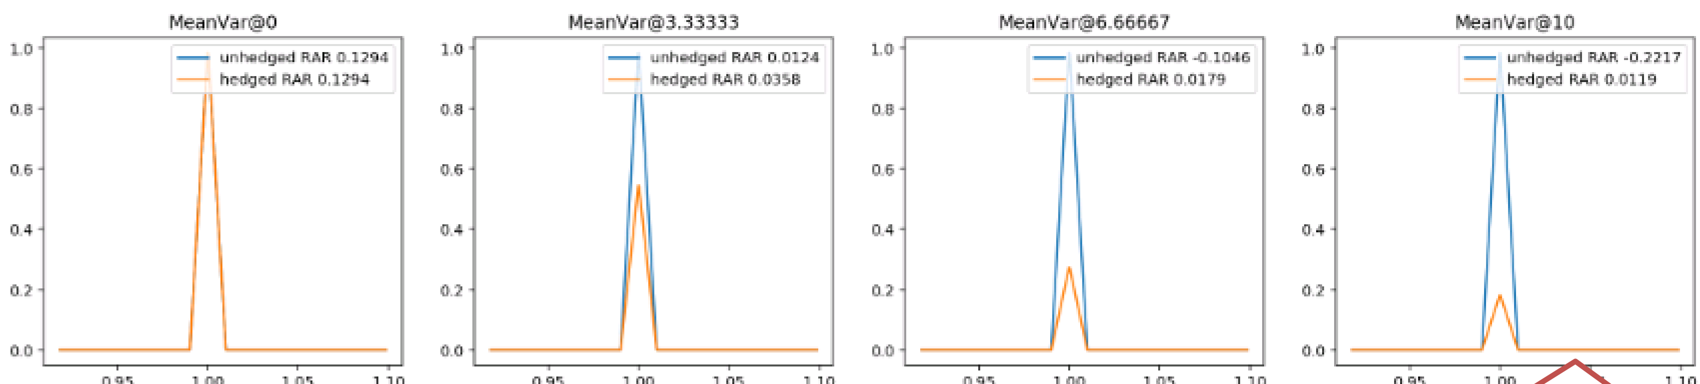

Increased Risk Aversion

Mean-variance sells a net long position for positive cost to reduce perceived "variance" risk

Statistical Hedging II – Convex Risk Measures

- Mean-Variance is not monotone; we therefore move to a systematic approach to measure carry/risk: we call *U* a **risk-adjusted return** if
  - *U* is normalized to *U(0) = 0*.
  - *U* is monotone, i.e. *A≥B* then *U(A) ≥U(B)*.
  - *U* is concave.
  - *U* is cash-invariant, *U(A+c) = U(A) + c* for all constants *c*.
- We note –*U* is a **convex risk measure**.
- A classic example is the **entropy** risk-adjusted return $U(A) \coloneqq -\frac{1}{\lambda} \log \mathrm{E}[\exp(-\lambda A)]$. Another important example is (the negative of) CVaR.
- Boundary cases for most reasonable *U*'s: $\mathrm{E}[A] \geq U(A) \geq -\inf(A)$
- We now solve the local problem

$$\sup_{\boldsymbol{a}_t} : U(dM_t + \boldsymbol{a}_t d\boldsymbol{H}_t) - c_t(\boldsymbol{a}_t)$$

which we may also write in terms of the associated dispersion risk measure $\rho(A) \coloneqq -U(A - E[A])$ as

$$\sup_{\boldsymbol{a}_t} : \mathrm{E}[dM_t + \boldsymbol{a}_t d\boldsymbol{H}_t] - \rho(dM_t + \boldsymbol{a}_t d\boldsymbol{H}_t) - c_t(\boldsymbol{a}_t)$$

This resembles the previous carry-risk-cost representation.

Beyond Greeks

Cost vs. Risk Aversion – Quadratic CVaR Example

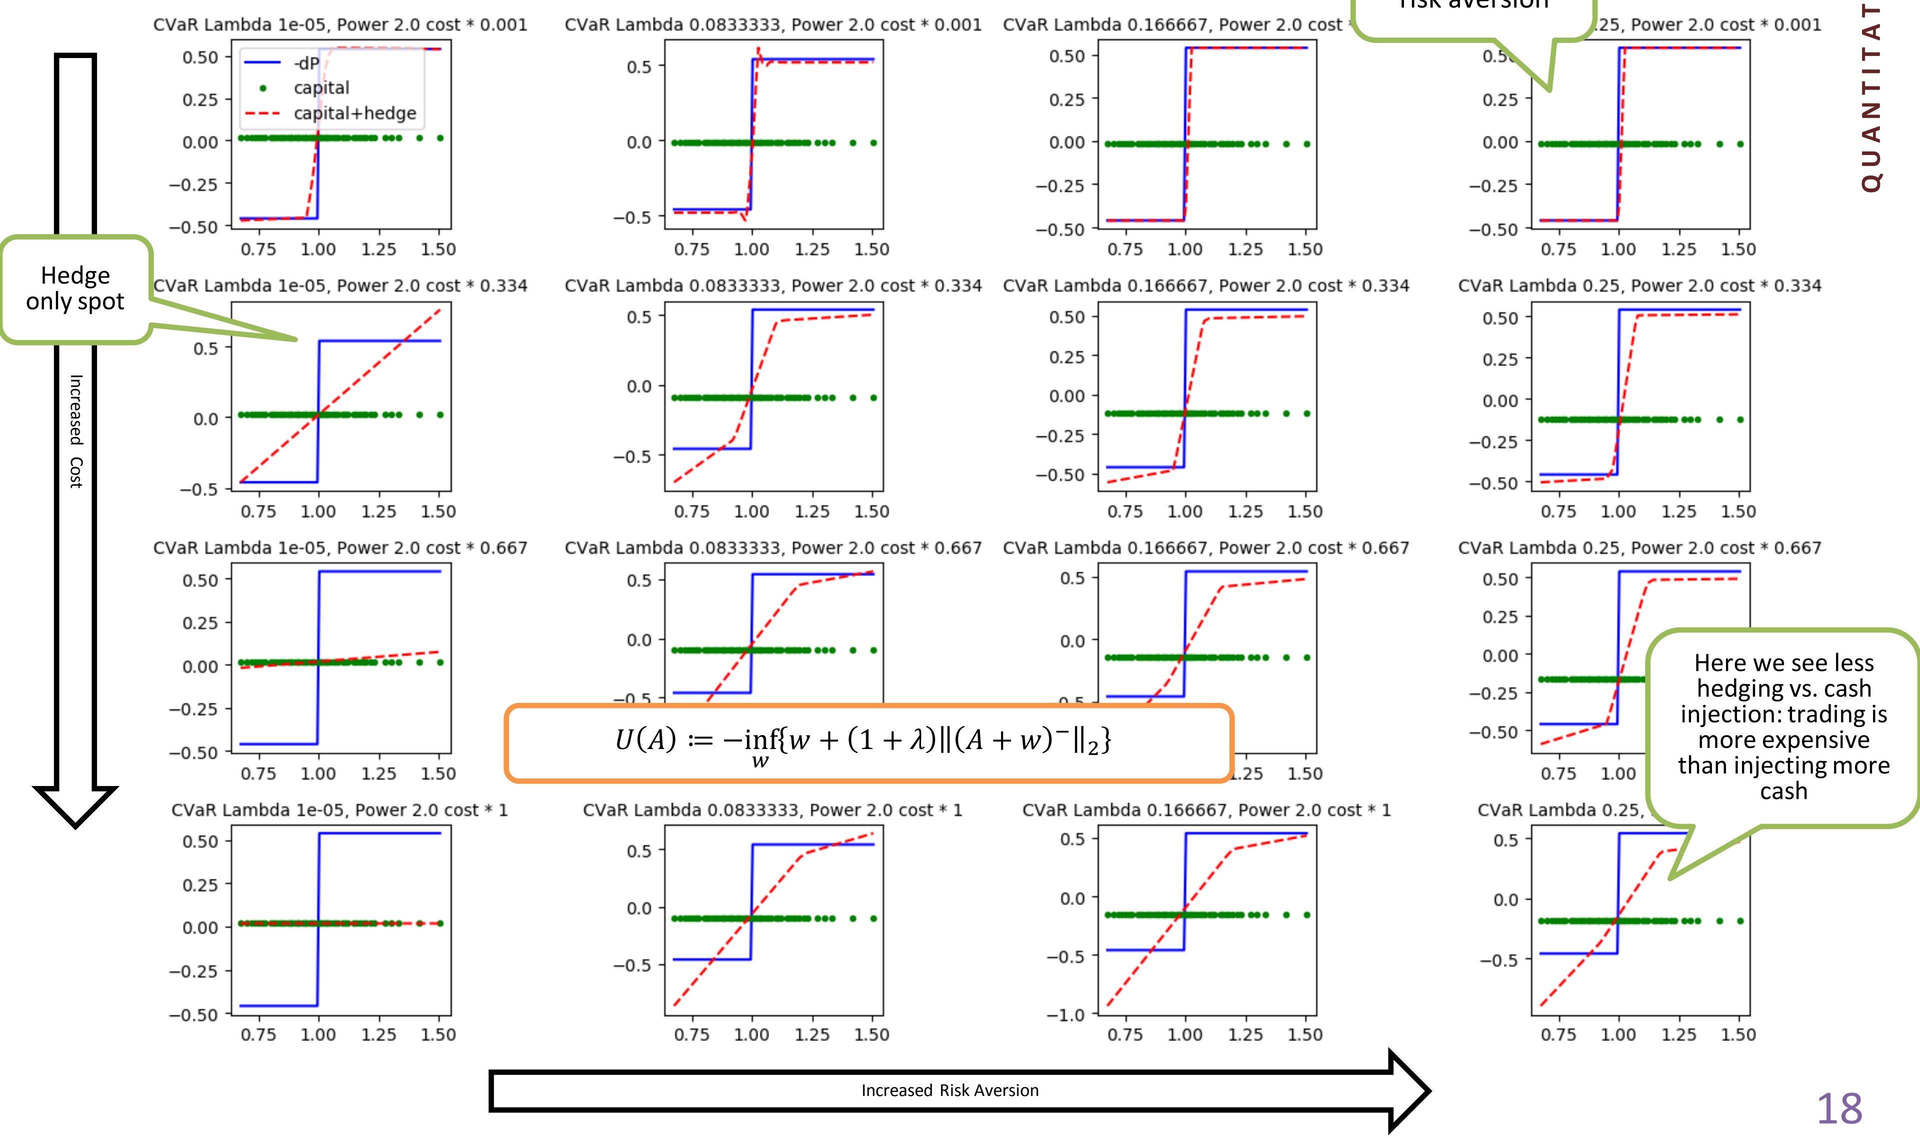

$$U(A) := -\inf_{w}\{w + (1+\lambda)\|(A+w)^-\|_2\}$$

Beyond Greeks

J.P.Morgan

QUANTITATIVE RESEARCH

Statistical Hedging

$$\sup_{\boldsymbol{a}_t}: U(dM_t + \boldsymbol{a}_t d\boldsymbol{H}_t) - c_t(\boldsymbol{a}_t)$$

Pros

- Data-driven combination of carry and risk which captures non-linearities
- Captures well-known effects such as "skew delta"
- Monotone, convex hedged portfolios.

Cons

- Hedges are only locally optimal / what is the optimal horizon *dt* vs. the absolute cost term
- Still requires mark-to-model values for return computation and pricing
- Compute intensive

The End of the Greek Era

J.P.Morgan

QUANTITATIVE RESEARCH

Summary

- **Classic Greek Hedging** is unsuitable for data-driven risk management
- **Parametric Hedging** works for very smooth portfolios
- **Statistical Hedging** expands this to strong non-linear features

Challenges

- Future hedging cost for the now-optimal portfolio are not taken into account: tomorrow’s mark-to-model value is assumed to be realizable.
- Pricing relies on classic model prices, possibly with some ad-hoc adjustment due to immediate hedging cost.

Deep Hedging

- Simulate the market to maturity and then solve the generic problem

$$\sup_{\boldsymbol{a}_0 \dots \boldsymbol{a}_{T-1}} : \mathrm{U}\left(Z_T + \sum_{t=0}^{T-1} \boldsymbol{\delta}_t d\boldsymbol{H}_t - \sum_{t=0}^{T-1} c_t(\boldsymbol{a}_t) \right)$$

Deep Hedging

Deep Hedging

- Simulate the market to maturity and then solve the generic problem

$$\sup_{\boldsymbol{a}_0 \ldots \boldsymbol{a}_{T-1}} : \mathrm{U}\left(Z_T + \sum_{t=0}^{T-1} \boldsymbol{\delta}_t d\boldsymbol{H}_t - \sum_{t=0}^{T-1} c_t(\boldsymbol{a}_t) \right)$$

- Classic “hedging under market frictions” problem.

Challenges

- Theory
- Numerical implementation
- Market dynamics

Theory

J.P.Morgan

QUANTITATIVE RESEARCH

Deep Hedging

$$v^*(Z) := \sup_{\pi=\boldsymbol{a}_0 \dots \boldsymbol{a}_{T-1}} : \mathrm{U}\left(Z_T + \sum_{t=0}^{T-1} \boldsymbol{\delta}_t d\boldsymbol{H}_t - \sum_{t=0}^{T-1} c_t(\boldsymbol{a}_t) \right)$$

Statistical Arbitrage

- We say the market has **statistical arbitrage** if $v^*(0) > 0$.
  - This *does not imply* presence of strict arbitrage.
    Example: Black & Scholes model with positive drift.
  - Strictly speaking, presence of strict arbitrage also does not imply presence of statistical arbitrage.
    Example: market has 100 scenarios, in 6 of which the asset returns 0%. In all other scenarios the asset returns 1%. Under CVaR@95%, $v^*(0) = 0$.
- Statistical Arbitrage is real – it means there are opportunities in the market, depending on one's risk aversion.
  - Example: realized-implied vol carry; rates curve carry etc.
- However, it can pollute the question of risk management: just as in classic cash portfolio optimization the estimation of "alpha" is much more involved.
  - We may therefore separate the estimation of carry from the estimation of the higher moments of the distribution of our instruments.
- We call a trade a **static arbitrage** opportunity if the return is non-negative under *any* market dynamics (e.g. violation of butterfly or calendar arb).

Deep Hedging

$$v^*(Z) \coloneqq \sup_{\pi = \boldsymbol{a}_0 \dots \boldsymbol{a}_{T-1}} : \mathrm{U}\left(Z_T + \sum_{t=0}^{T-1} \boldsymbol{\delta}_t d\boldsymbol{H}_t - \sum_{t=0}^{T-1} c_t(\boldsymbol{a}_t) \right)$$

Pricing

– Consider a current position of *Z*.
  The price of selling a derivative *Y* to a customer is given by the marginal cost

$$p(Y) \coloneqq v^*(Z - Y) - v^*(Z)$$

  - Of course, $v^*(Z - Y + p(Y)) = 0$.
  - Reflects naturally a bid/ask spread.
  - The model-price is no longer used.

Deep Hedging

$$v^*(Z) := \sup_{\pi=\boldsymbol{a}_0\ldots\boldsymbol{a}_{T-1}} : \mathrm{U}\left(Z_T + \sum_{t=0}^{T-1} \boldsymbol{\delta}_t d\boldsymbol{H}_t - \sum_{t=0}^{T-1} c_t(\boldsymbol{a}_t)\right)$$

Hamilton-Jacobi-Bellman

- One of the biggest short comings of the approach presented here is that Z is fixed and not a part of the "state" of the market.
- The Bellman form of the problem can be written as follows:
  - Denote by M^t all future cash flows of our portfolio on and after t, and by m_t the cashflow arising from holding M^t at t.
    This gives the classic HJB reward form

$$V^*(M^t|\boldsymbol{S}_t) = \sup_{a_t} : U(\, V^*(M^{t+1} + \boldsymbol{a}_t H^{t+1})|\boldsymbol{S}_t) - c(\boldsymbol{a}_t|\boldsymbol{S}_t) + m_t$$

  - Research topics:
    - Find a representation for a portfolio of derivatives which is efficient for this to be applicable.
    - Under what conditions does this equation have a fixed point for all combined states $(M^t, \boldsymbol{S}_t)$.

Deep Hedging

$$v^*(Z) := \sup_{\pi=\boldsymbol{a}_0 \ldots \boldsymbol{a}_{T-1}} : \mathrm{U}\left(Z_T + \sum_{t=0}^{T-1} \boldsymbol{\delta}_t d\boldsymbol{H}_t - \sum_{t=0}^{T-1} c_t(\boldsymbol{a}_t) \right)$$

Pricing

- Consider a current position of Z.
  The price of selling a derivative Y to a customer is given by the marginal cost

$$p(Y) := v^*(Z-Y) - v^*(Z)$$

  - Of course, $v^*(Z - Y + p(Y)) = 0$.
  - Reflects naturally a bid/ask spread.
  - The model-price is no longer used.

Numerical Implementation

J.P.Morgan

QUANTITATIVE RESEARCH

Deep Hedging

$$\sup_{\pi=\boldsymbol{a}_0\ldots\boldsymbol{a}_{T-1}} : \mathrm{U}\left(Z_T + \sum_{t=0}^{T-1} \boldsymbol{\delta}_t d\boldsymbol{H}_t - \sum_{t=0}^{T-1} c_t(\boldsymbol{a}_t) \right)$$

Structure of the Problem

- The problem above is convex in π.
- We will solve this over a fixed set of paths, and therefore fixed set of terminal payoffs.
- Path-dependency of the problem enters due to transaction cost and from path-dependent restrictions on liquidity.

Reinforcement Learning

- Parameterize $\boldsymbol{a}_t$ as a neural network, with the result of the previous step feeding into the next step.
- This is called “model-based policy search” in the ML literature.
- Theoretical result [DH’18]: neural networks approximate any policy arbitrarily well with increasing depth and width.
- Practical choices:
  - Each $\boldsymbol{a}_t$ has its own network → rather deep network
  - Share network across t → LSTM to capture path
- Example code in Karas is just about a page of code https://people.math.ethz.ch/~jteichma/deep_portfolio_optimization_keras.html
- Efficient, scalable, model-independent implementation.

Numerical Implementation

J.P.Morgan

QUANTITATIVE RESEARCH

Deep Hedging

- We show the efficiency and impact of the algorithm for classic derivative model dynamics.
- Used to validate convergence of the numerical scheme against (the few) analytically available results.
- Even here, we are able to compute previously inaccessible problems for entire portfolios of derivatives

Comparison with theory: vanilla option with Heston dynamics

- Hedge an ATM 30-day call with spot and var swap
- No costs, no limits, 50% CVaR value function

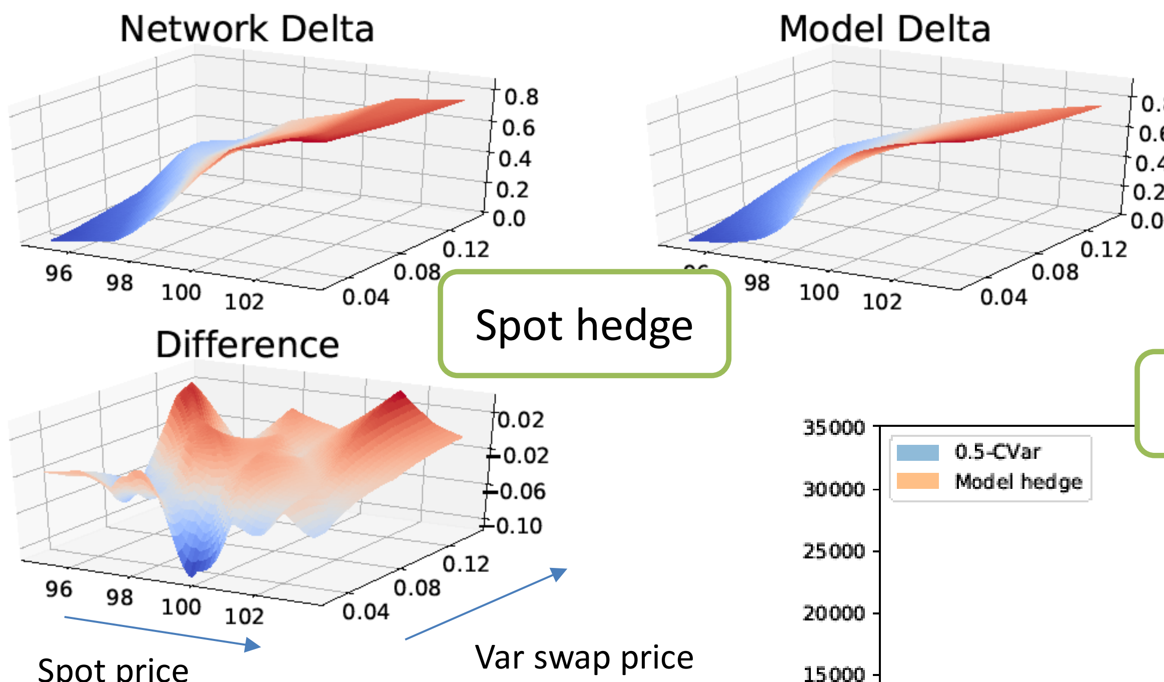

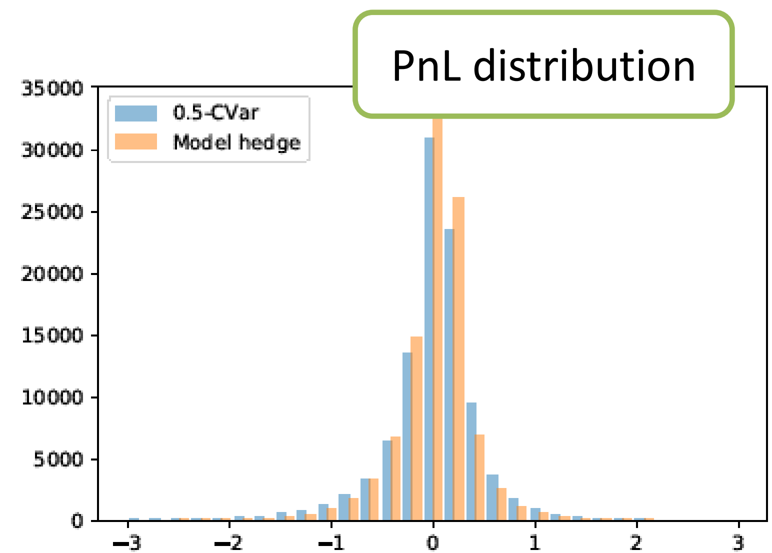

J.P.Morgan

QUANTITATIVE RESEARCH

Comparison with theory: Heston dynamics with cost

- We have asymptotic results for small transaction cost for classic one-factor models such as Black Scholes.
- There are no analytic results for higher order models. We show that the same asymptotics hold using our numerical scheme

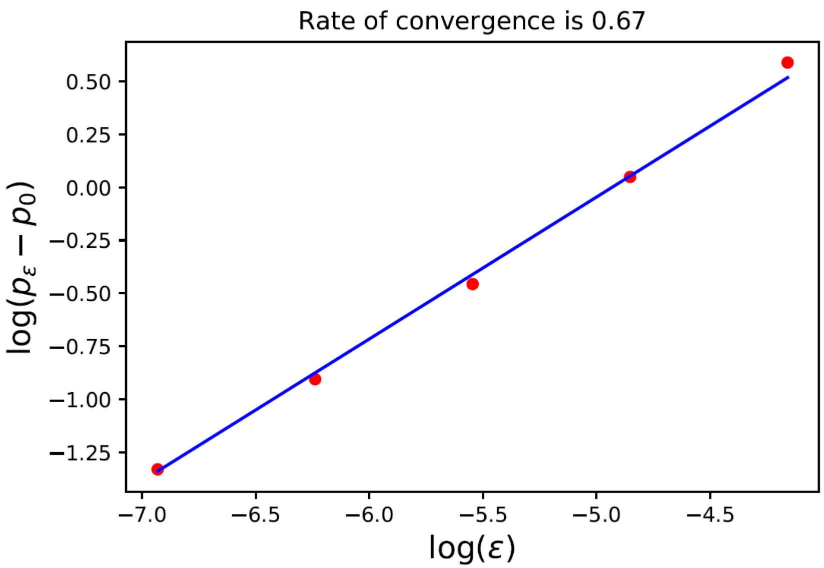

Black-Scholes for which we know theoretical results

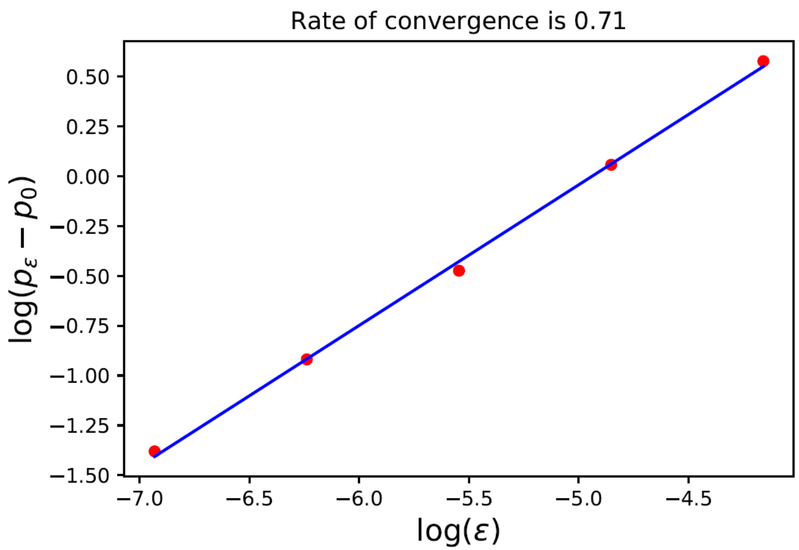

Heston
No theoretical results but we plot the 1-factor expected behaviour

J.P.Morgan

QUANTITATIVE RESEARCH

Vanilla option: impact of transaction costs

- 30-day ATM call option
- Plot shows spot hedge ("Delta")
- No limits
- Entropic value
- Risk aversion 10
- Black-Scholes simulator

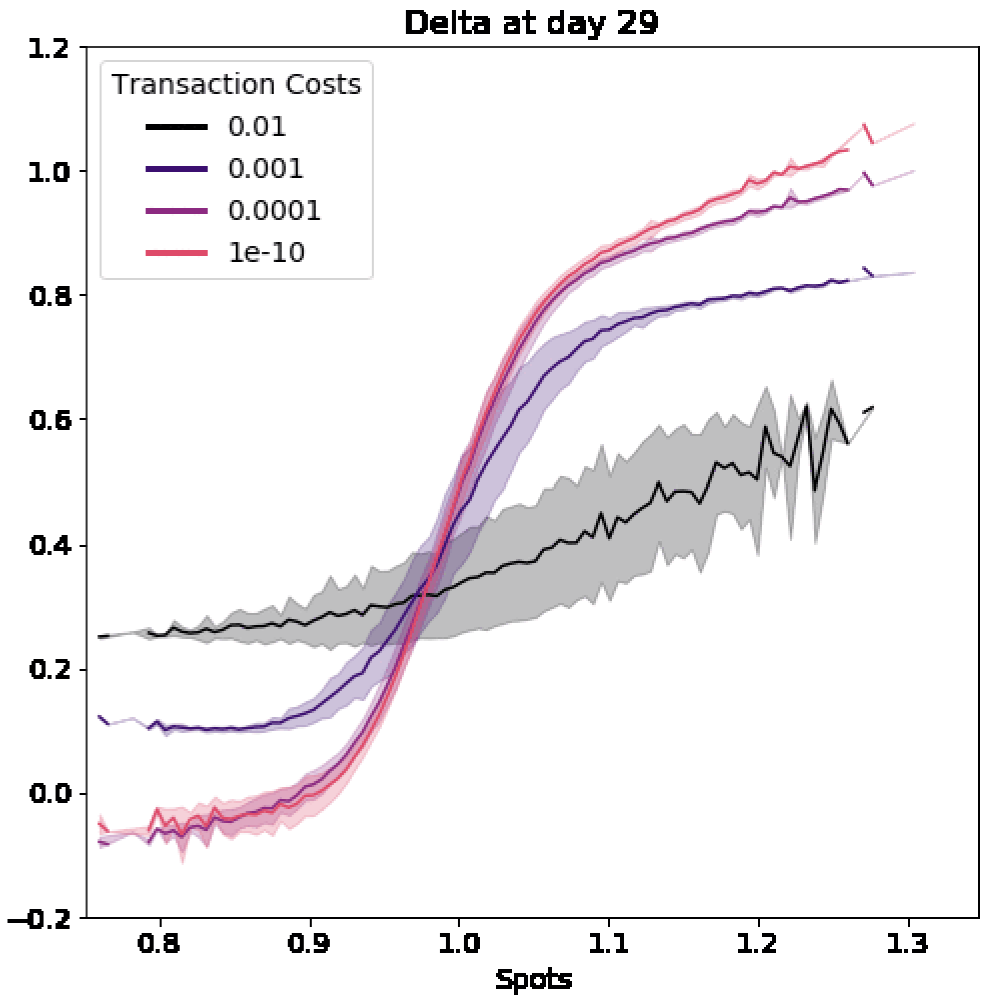

Deep Hedging

Forward starting option: impact of transaction costs

- 30-day ATM call starting in 15 days
- Heston simulator, no limits, entropic value, risk aversion 50
- Hedge with spot only

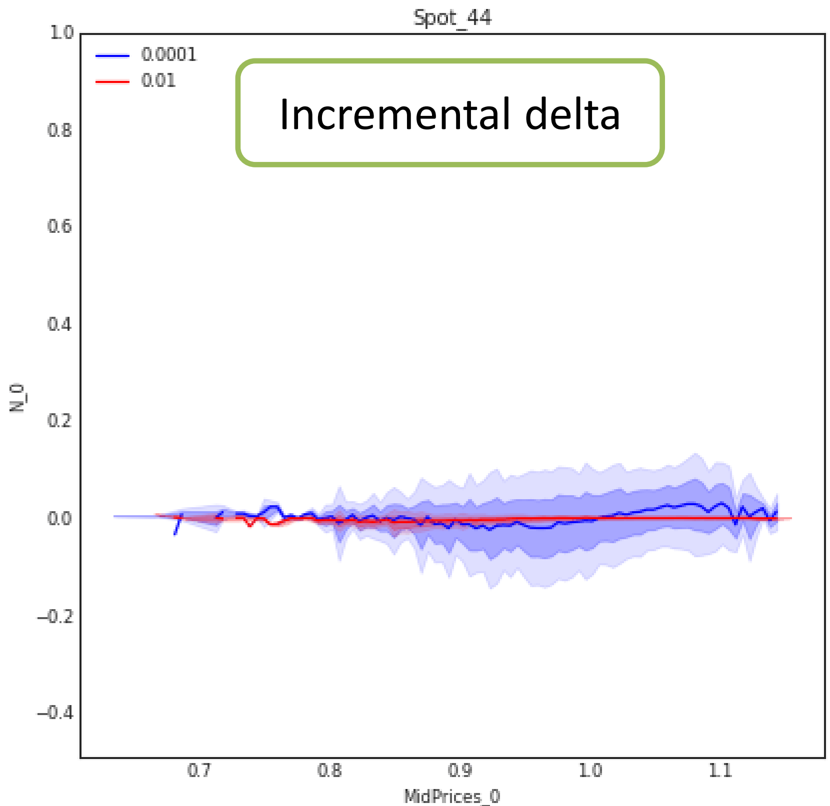

J.P.Morgan

QUANTITATIVE RESEARCH

Forward starting option: impact of transaction costs

- 30-day ATM call starting in 15 days
- Heston simulator, no limits, entropic value, risk aversion 50
- Hedge with spot only

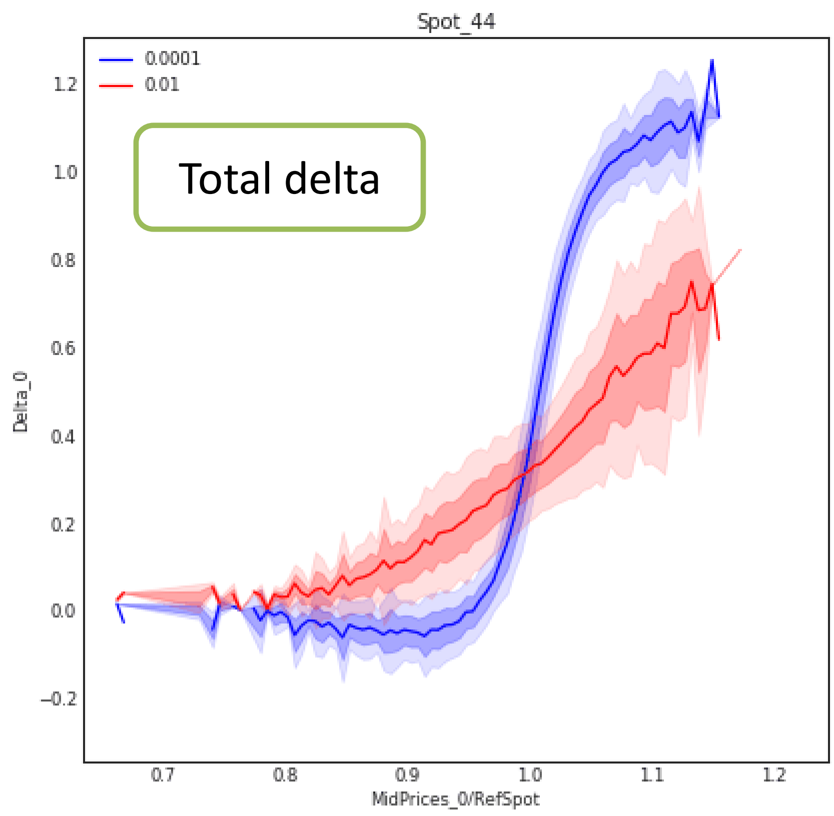

Deep Hedging

Popular retail payoff:

- Client is short a down-and-in put paid at maturity
- Upper knockout barrier
- Fixed coupons until KO

Market

- 0.1% transaction costs
- No limits
- Risk aversion 20
- Entropic value
- Monthly hedging
- Local volatility simulator

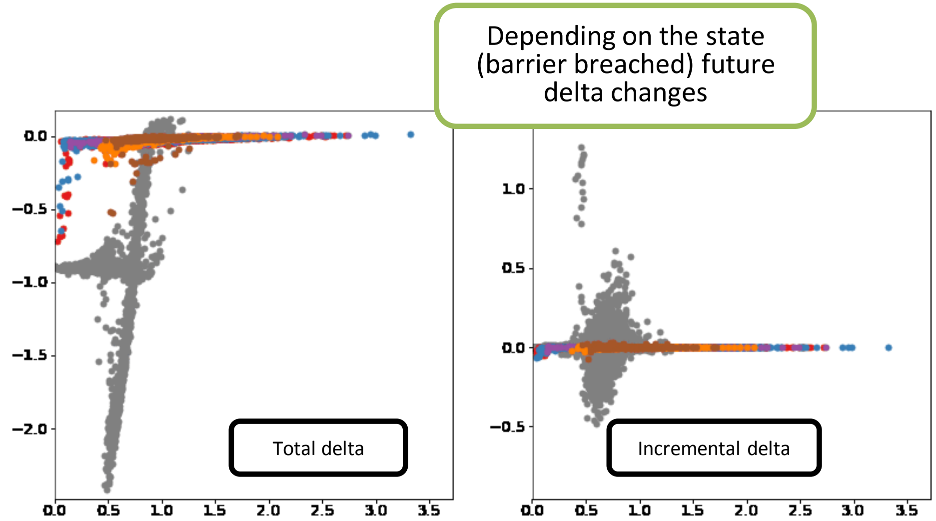

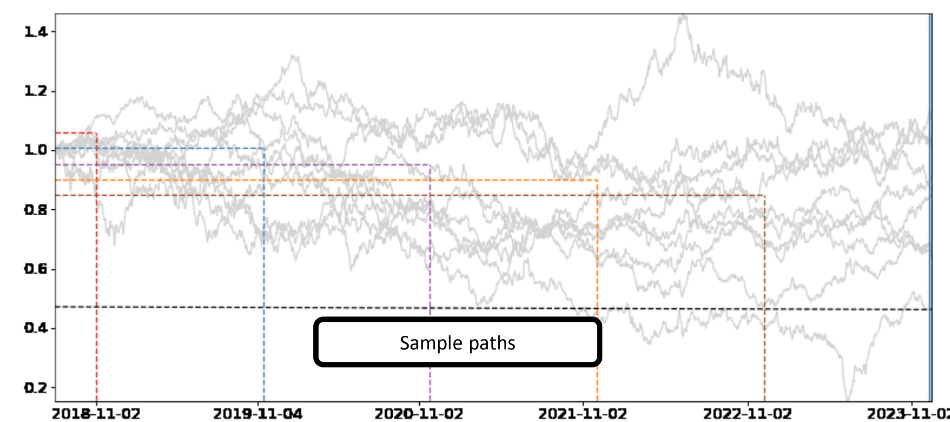

Deep Hedging

J.P.Morgan

QUANTITATIVE RESEARCH

Deep Hedging with Classic Derivatives Generators

- Useful for assessing liquidity and cost impact at scale
  - Straightforward with TensorFlow
  - Inherently parallelizable since risk-adjusted returns tend to be expectation based.
  - Speed independent of number of instruments in portfolio
  - Asset-class agnostic: discounting, FX, own callability etc all no problem
  - Client-callability: master thesis under way @ ETH

…. shouldn't we use the real market to train our model?

Market Dynamics

Market Dynamics

J.P.Morgan

QUANTITATIVE RESEARCH

Challenges

- Sparse data set vs. large number of instruments.
- Non-stationarity and robustness
  Idea: solve a robust version of our problem,

$$\inf_{Q} : \sup_{\pi = \boldsymbol{a}_0 \ldots \boldsymbol{a}_{T-1}} : U_Q\left[Z_T + \sum_{t=0}^{T-1} \boldsymbol{\delta}_t \, d\boldsymbol{H}_t - \sum_{t=0}^{T-1} c_t(\boldsymbol{a}_t)\right]$$

- Avoid static arbitrage
  A possible approach is parameterizing the implied vol in "discrete local vol" (see also Wissel) which is an arb-free parameterization.

Model Challenges

- Statistical arbitrage.
  - Robustify the estimator.
  - Find the closest risk-neutral measure without drift.

GAMeD

J.P.Morgan

QUANTITATIVE RESEARCH

Generative Adversarial Market Dynamics

– We want to solve

$$p_{\mathrm{Q}}(\mathrm{Z}) \coloneqq \sup_{\pi=\boldsymbol{a}_0 \dots \boldsymbol{a}_{T-1}} : U_{\mathrm{Q}}\left[Z_T + \sum_{t=0}^{T-1} \boldsymbol{\delta}_t d\boldsymbol{H}_t - \sum_{t=0}^{T-1} c_t(\boldsymbol{a}_t)\right]$$

– However, we are not sure about Q since we only observe samples $\hat{Q}$. Classic robust approach [Follmer, Schied: Stochastic Finance: An Introduction in Discrete Time, 2011]: define set of reasonable measures "close" to $\hat{Q}$ and a distance *d* to $\hat{Q}$ and solve for $\alpha \uparrow \infty$:

$$\hat{p}_F(Z) \coloneqq \sup_{\pi} \inf_{Q} \{\alpha U_Q[G(\pi_t; Z_T)]^+ + d(\mathrm{Q}, \hat{\mathrm{Q}})\}$$

– Machine Learning interpretation [Generative Adversarial Networks, Goodfellow 2014]:
  - **Generator** *Q* tries to fit the target distribution and take away money
  - **Adversary** π tries to make money

– We use unconditional Wasserstein distance W_1 as our metric using a fast stochastic algorithm c.f, [Stochastic Optimization for Large-scale Optimal Transport, 2016, Genevay / Cuturi /Peyré /Bach]

GAMeD

STOXX50 Dataset

- We obtain historical spot and option prices for last 10 years of data ~ 2000 historical data points
- Option Grid with relative strikes $K = \{0.8, 0.85, 0.9, 0.95, 1.0, 1.05, 1.1, 1.15, 1.2\}$, and maturities of $T = \{20, 40, 60, 80, 100, 120\}$ days
- $\boldsymbol{H}_t$ = 109-dimensional vector = 1 Spot + 54 Calls + 54 Puts
- At every time step, $\boldsymbol{H}_t$ = generated prices for the grid of options + spot
- Each simulation time step = 20 days to match the maturities

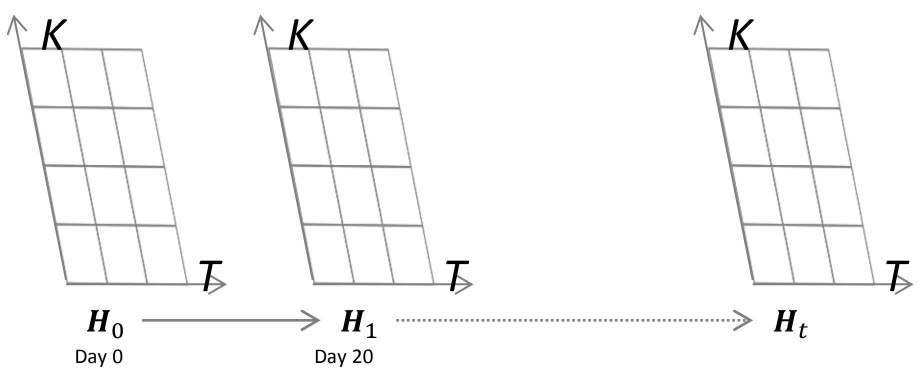

GAMeD

Setup and Training

- Generator network: 2 layer LSTM of size 64, f ~62K parameters
- Hedger network: 2 layer LSTM of size 64, π ~112K parameters
- Use 2% transaction costs to regularize the hedger network
- Training:
  - Train with batch sizes of 32K to minimize noise, RMS Prop with a learning rate of 1e-4
  - Dropout of 25% between state-to-state LSTM connections to regularize training
  - 50 K updates with a batch size of 32K takes 10hrs to train on Tesla V100 GPU

Results

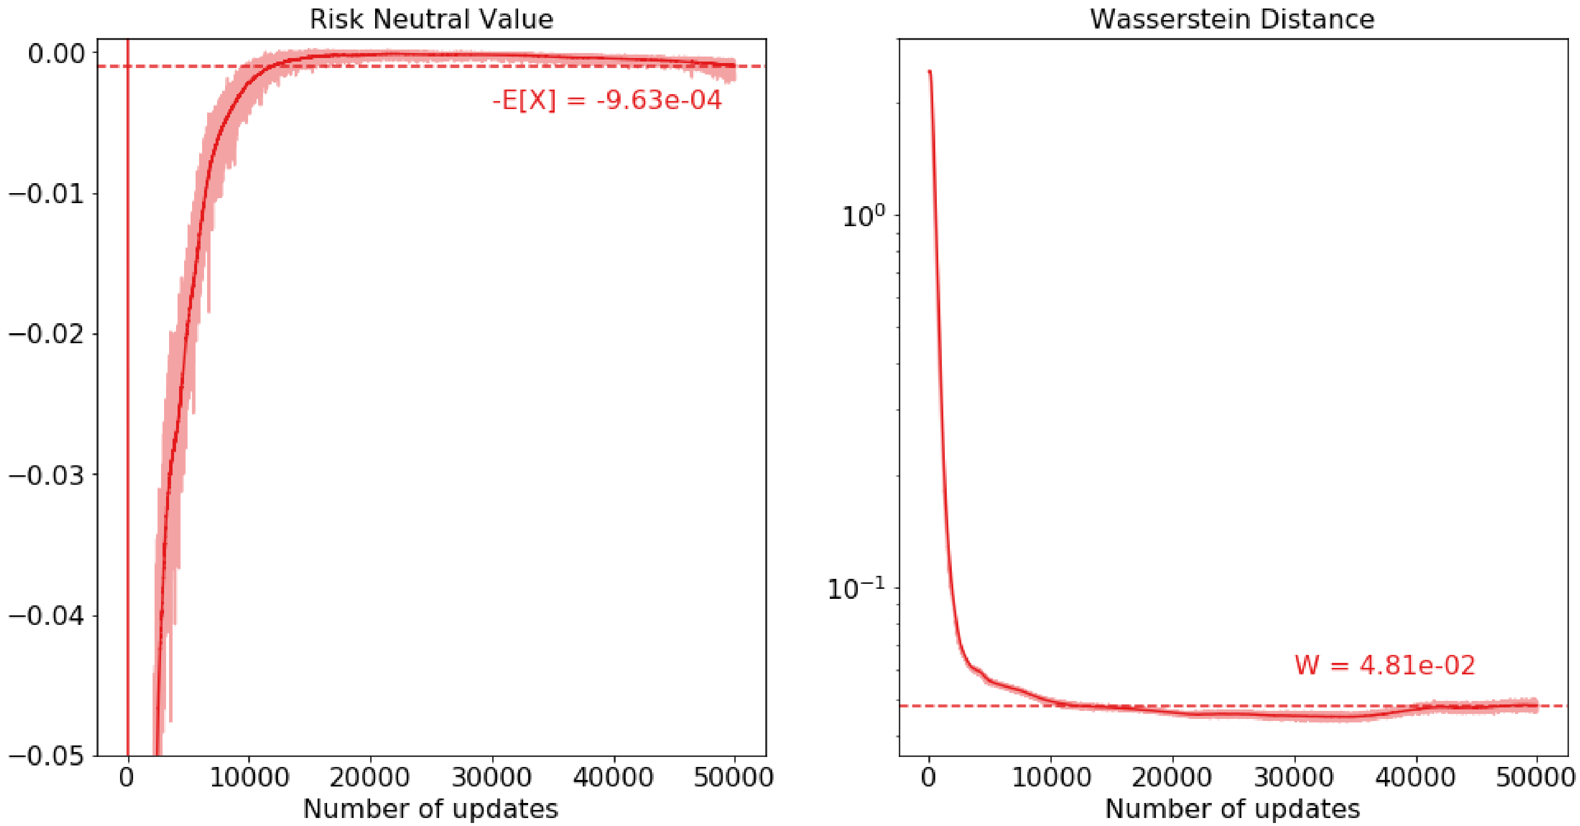
Risk Neutral Value
0.00
−0.01
−0.02
−0.03
−0.04
−0.05
-E[X] = -9.63e-04
0
10000
20000
30000
40000
50000
Number of updates
Wasserstein Distance
10^0
10^{-1}
W = 4.81e-02
0
10000
20000
30000
40000
50000
Number of updates

J.P.Morgan

QUANTITATIVE RESEARCH

Results

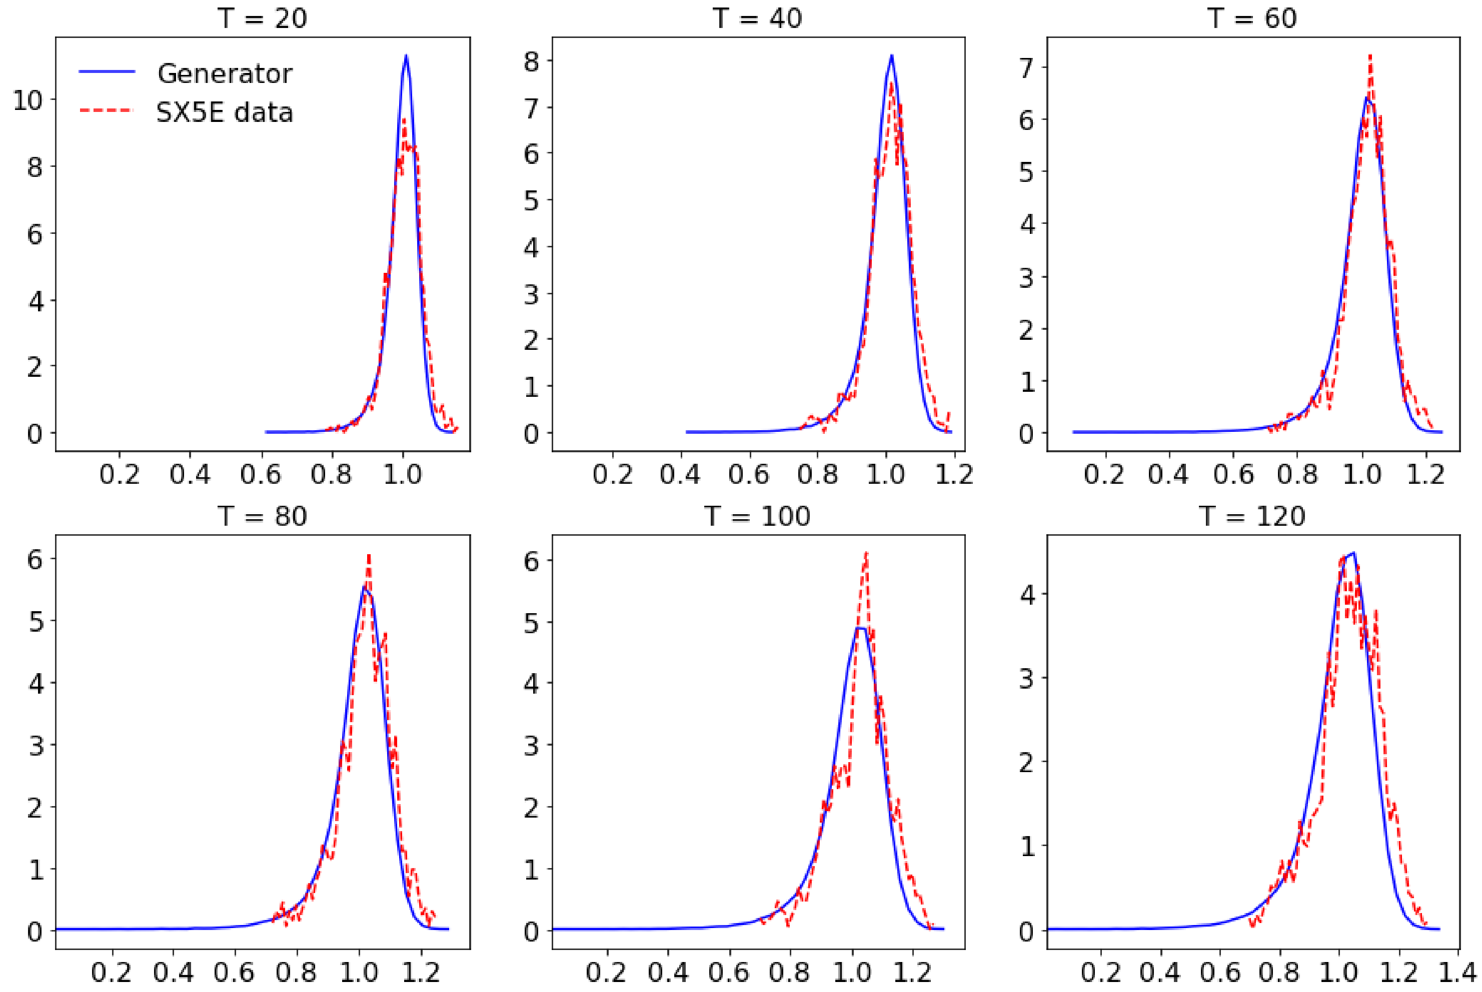

Histogram of generated spot prices compared to historical data

GAMeD

J.P.Morgan

Results

Histogram of generated call option prices at t = 3 compared to historical data

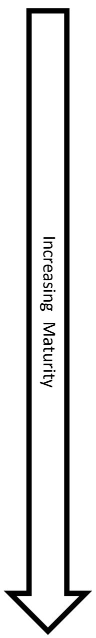

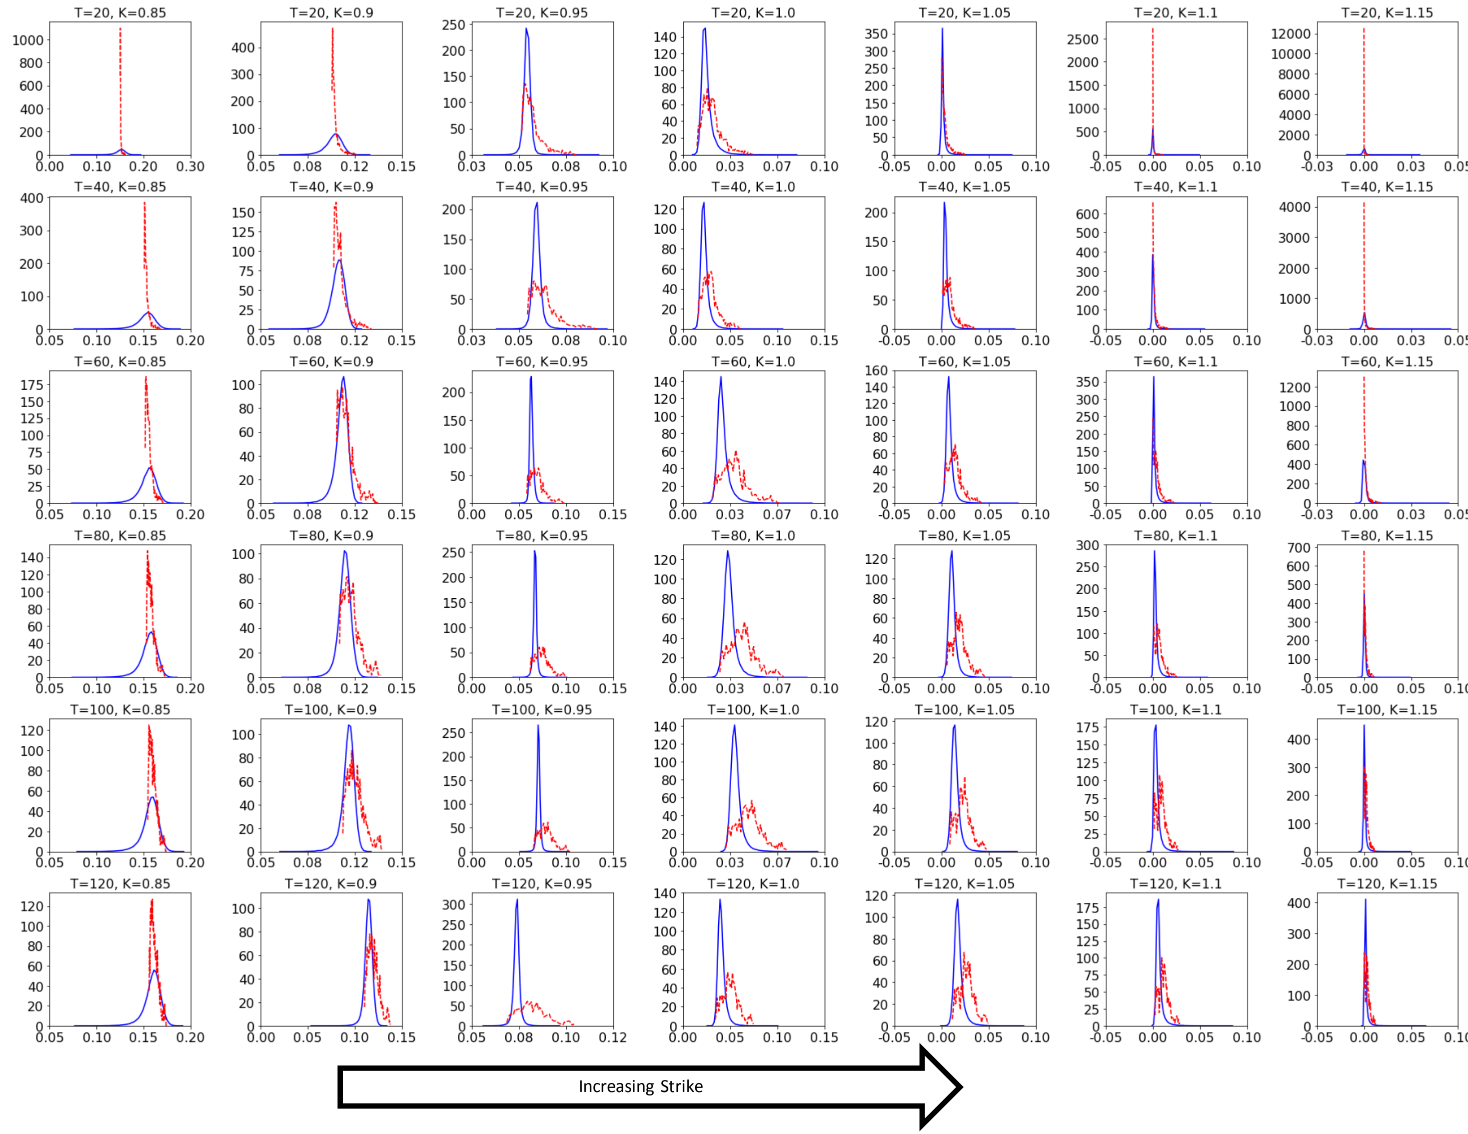

GAMeD

J.P.Morgan

QUANTITATIVE RESEARCH

Results

Graphs show the time series of option prices from simulated by LSTM based generator

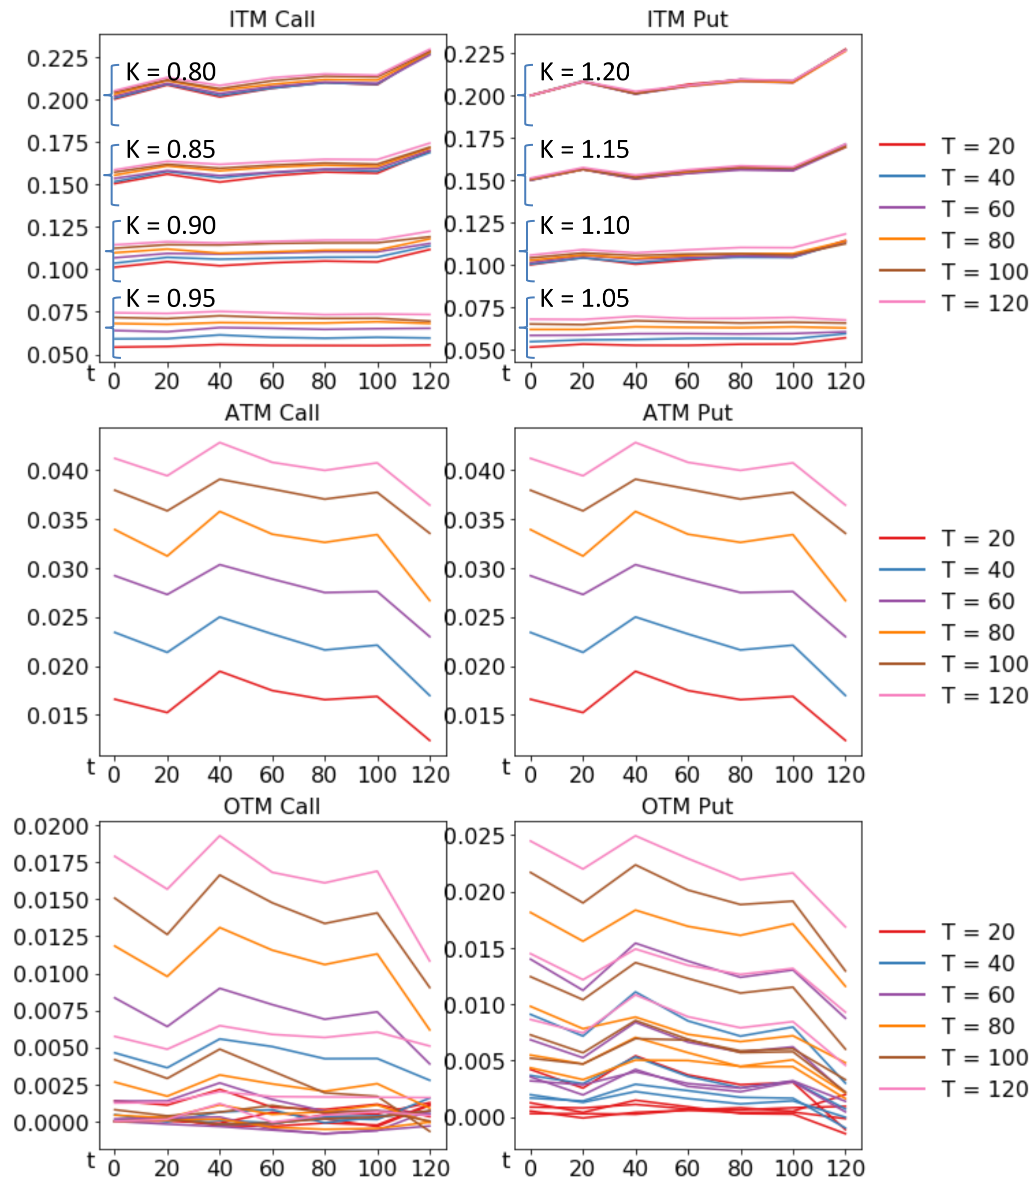

What Next

J.P.Morgan

Next steps

- Define conditional version of GAMeD
- Use distance metric which reflects use case, i.e. DH itself
- Move to recursive version of DH.
- Formalize theoretical relationship between adversarial learning and robust statistics

Thank you very much for your attention