Simulating spot and equity option markets using rough path signatures

Hans Buehler*, Ralf Korn†, Alexandre Pachoud*,
Magnus Wiese*†, Ben Wood*

* J.P. Morgan, London
† TU Kaiserslautern

March 12, 2021

[0] Opinions expressed in this presentation are those of the authors, and do not necessarily reflect the view of J.P. Morgan.

Deep Hedging

Problem setting

- Deep Hedging (DH) aims to automate hedging a portfolio of derivatives under market frictions.
- Calibration of the DH agent requires *large amounts* of realistic daily market data for an underlying which is not available.

Deep Hedging

Problem setting

- Deep Hedging (DH) aims to automate hedging a portfolio of derivatives under market frictions.
- Calibration of the DH agent requires *large amounts* of realistic daily market data for an underlying which is not available.

We aim to mitigate the problem of limited market data by calibrating a realistic market simulator $G : \mathbb{R}^{p\cdot(nm+1)+q} \to \mathbb{R}^{nm+1}$

$$(R_{t+1}, \boldsymbol{\sigma}_{t+1}) = G(Z_{t+1}, R_t, \ldots, R_{t-p+1}, \boldsymbol{\sigma}_t, \ldots, \boldsymbol{\sigma}_{t-p+1})$$

where $p \in \mathbb{N}$, $Z_{t+1} \sim \mathcal{N}(0, I_q)$, R_t is the 1-dimensional spot log-return at time t, $\boldsymbol{\sigma}_t$ is a nm-dimensional grid (n strikes, m maturities) of **discrete local volatilites** (DLVs) at time t, and G is **unknown**.

Deep Hedging

Problem setting

- Deep Hedging (DH) aims to automate hedging a portfolio of derivatives under market frictions.
- Calibration of the DH agent requires *large amounts* of realistic daily market data for an underlying which is not available.

We aim to mitigate the problem of limited market data by calibrating a realistic market simulator $G : \mathbb{R}^{p\cdot(nm+1)+q} \to \mathbb{R}^{nm+1}$

$$(R_{t+1}, \boldsymbol{\sigma}_{t+1}) = G(Z_{t+1}, R_t, \ldots, R_{t-p+1}, \boldsymbol{\sigma}_t, \ldots, \boldsymbol{\sigma}_{t-p+1})$$

where $p \in \mathbb{N}$, $Z_{t+1} \sim \mathcal{N}(0, I_q)$, R_t is the 1-dimensional spot log-return at time t, $\boldsymbol{\sigma}_t$ is a nm-dimensional grid (n strikes, m maturities) of **discrete local volatilites** (DLVs) at time t, and G is **unknown**.
How do we calibrate G?

Orthonormal path compression with signature cumulants

- Direct calibration of G suffers from the curse of dimensionality; especially for dense grids $nm \gg 1$
- Instead, the grids of DLVs are compressed to an efficient lower-dimensional representation $\gamma_t := F(\boldsymbol{\sigma}_t), F : \mathbb{R}^{nm} \to \mathbb{R}^l$ by using autoencoders
- Additionally, the components $(\gamma_t^1, \ldots, \gamma_t^l)$ are regularized to be statistically independent through space and time by leveraging a **signature cumulants-based penalty** (Bonnier et al., 2020)

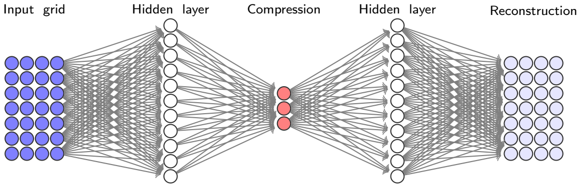

Figure: Illustration of a shallow DLV autoencoder $(28-11-3-11-28)$.

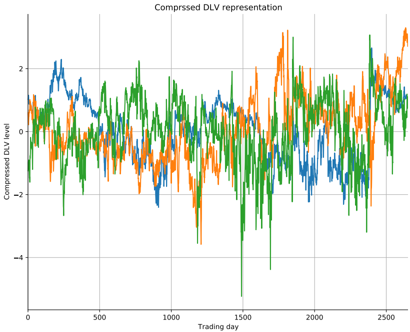

Figure: Compressed DLV times series $(\gamma_t)_{t=0}^{T}$

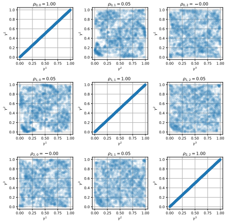

Figure: Copula of γ_t

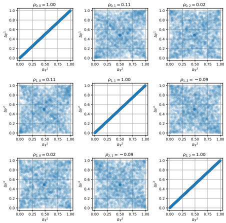

Figure: Copula of $\Delta\gamma_t$

Calibrating the market simulator with signatures

The market simulator is calibrated by minimizing the **conditional Sig-Wasserstein-1** *loss* (Hao Ni et al., 2020):

$$\mathcal{L}_\mu(\nu) := \mathbb{E}_{X_{0,t}\sim\mu}\left(\|\mathbb{E}_\mu(S(X_{t,t+T})|X_{0,t}) - \mathbb{E}_\nu(S(X_{t,t+T})|X_{0,t})\|_2\right)$$

where $X_{0,T} = (X_t)_{t\in[0,T]}$ is a linear interpolation of the discrete-time process $(R_t, \gamma_t)_{t\in\{0,\dots,T\}}$, μ is the empirical law, ν the law induced by the market simulator, and the approximation of

- $\mathbb{E}_\mu(S(X_{t,t+T})|X_{0,t})$ boils down to a linear regression which can be estimated **robustly** via ridge regression / LASSO and cross-validation,
- $\mathbb{E}_\nu(S(X_{t,t+T})|X_{0,t})$ is a Monte Carlo approximation evaluated through generating samples $X_{t,T} \sim \nu(\cdot|X_{0,t})$.

Comparison of empirical distribution and invariant distribution induced by G

$$(R_{t+1}, \gamma_{t+1}) = G(Z_{t+1}, R_{t-p+1:t}, \gamma_{t-p+1:t})$$

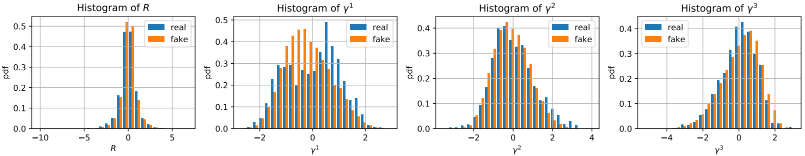

Figure: Marginal distribution of (R_t, γ_t)

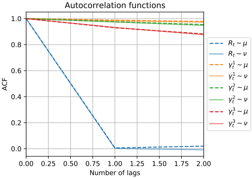

Figure: Autocorrelation function of (R_t, γ_t)

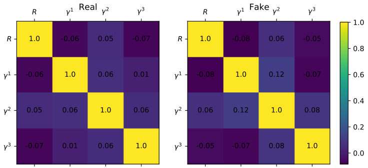

Figure: Cross-correls of (R_t, γ_t)

Comparison of empirical distribution and invariant distribution induced by G

$$(R_{t+1}, \gamma_{t+1}) = G(Z_{t+1}, R_{t-p+1:t}, \gamma_{t-p+1:t})$$

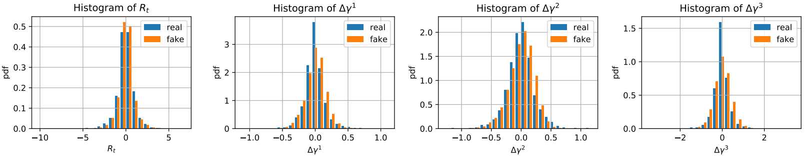

Figure: Marginal distribution of $(R_t, \Delta\gamma_t)$

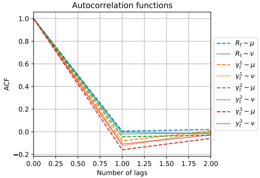

Figure: Autocorrelation function of $(R_t, \Delta\gamma_t)$

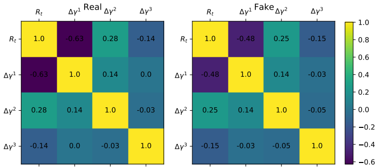

Figure: Cross-correls of $(R_t, \Delta\gamma_t)$

Disclaimer

Opinions and estimates constitute our judgement as of the date of this Material, are for informational purposes only and are subject to change without notice. It is not a research report and is not intended as such. Past performance is not indicative of future results. This Material is not the product of J.P. Morgans Research Department and therefore, has not been prepared in accordance with legal requirements to promote the independence of research, including but not limited to, the prohibition on the dealing ahead of the dissemination of investment research. This Material is not intended as research, a recommendation, advice, offer or solicitation for the purchase or sale of any financial product or service, or to be used in any way for evaluating the merits of participating in any transaction. Please consult your own advisors regarding legal, tax, accounting or any other aspects including suitability implications for your particular circumstances. J.P. Morgan disclaims any responsibility or liability whatsoever for the quality, accuracy or completeness of the information herein, and for any reliance on, or use of this material in any way. Important disclosures at: www.jpmorgan.com/disclosures.