Executive Summary

This thesis represents a comprehensive overview of the current state of affairs in the class of rough stochastic volatility models. The analysis spans from the theoretical derivation of two prominent instances of the class, the rough Bergomi (rBergomi, Bayer et al. (2016)) and rough Heston (rHeston, Jaisson and Rosenbaum (2016)) models, to rigorous empirical calibration experiments carried out on the SPX and VIX financial markets. Inspired by the celebrated joint SPX & VIX calibration problem, the aim of the study is to consolidate and extend previous findings as well as shine a light on critical volatility modelling details that might have been overlooked in the past.

Starting at the origins of fractional stochastic volatility modelling laid by Comte and Renault (1998), the progression towards rough fractional stochastic volatility models is explained through the timely evolution of empirical findings prompted by the need and desire to parsimoniously recover key stylised facts observed in volatility time series. The paradigm was formally established in the seminal paper of Gatheral et al. (2018) with the introduction of the Rough Fraction Stochastic Volatility model (RFSV), a natural mathematical representation of the microstructural foundations identified in the irregularities of high-frequency volatility paths. The main results of the paper are replicated on more recent data as to explain the logistics behind the precursor to all rough volatility models. A special focus is given to the construction and evolution of the Hurst index, which defines their fractional Brownian Motion (fBM) volatility drivers. As an element of novelty, attention is drawn to the comparative evolution of Hurst index estimates against VIX levels.

The two classic option pricing models of Bergomi (Bergomi (2005)) and Heston (Heston (1993)) are revisited under the rough volatility paradigm. The re-phrasing is the subject of multiple applications of standard stochastic analysis results that are elaborated on so to highlight derivation assumptions that may hinder or explain the models' practical usage. The efficient pricing of both models is described, namely an adaption of the Hybrid Monte Carlo simulation scheme for Brownian semistationary processes (Bennedsen et al. (2017)) to rBergomi and a close to instantaneous Fourier inversion technique via rational Padé approximation theory (Gatheral and Radoičić (2019)) adopted for rHeston. The fact that both models are defined by sets of only 3 time-homogeneous theoretical parameters alongside a forward variance curve (external market information) parameter, facilitates the manual calibration to SPX and VIX volatility smiles and surfaces. The rHeston model is manually calibrated to the VIX option market potentially for the first time in known literature. Although the results of Bayer et al. (2016) and El Euch et al. (2019) serve as benchmark, the manual calibration procedure has been modified in order to handle the forward variance curve as a purely theoretical quantity rather than approximating it from variance swap data. The discussion which ensues is of great relevance for the forthcoming deep learning implementations and importantly brings into question what a better- or worse-off performance than the one attained with the original method might mean in terms of the models' well-definedness.

Traditional pricing methods and manual calibration are memory- and time-intensive computational procedures that render most stochastic volatility models unsuited to broad industry adoption. Recent developments in the emerging field of deep learning technologies offer a potential resolve to the shortcoming by replacing costly pricing map evaluations from model parameters to afferent Black-Scholes implied volatilities (IVs), or vice-versa, with almost instantaneous feed-forward runs of an artificial neural network approximation, the so-called deep IV (pricing) map. Different calibration approaches distinguish themselves based on how the map is defined. The one-step approach, also known as the prices-to-model parameters (PtM) method, has re-kindled researcher's active interest in deep pricing techniques with the breakthrough paper of Hernandez (2016). It seeks to solve the

so-called inverse pricing problem by learning the inverse pricing map from market IVs to model parameters. Feeding market option data into the deep pricer amounts to a complete calibration routine as it directly returns calibrated model parameters. Despite the great promise of the concept, both Hernandez and later on the extensive research of Bayer et al. (2019) reveal a subpar performance on unseen data. For this reason, the alternative two-step approach, i.e. the model parameters-to-prices method (MtP) has been elected towards deep calibration. The procedure combines traditional pricing methods with a standard calibration routine such as Levenberg-Marquardt (Levenberg (1944), Marquardt (1963)), the latter of which may benefit from the automatic differentiability of the network that promptly renders accurate Jacobian approximations (Bayer and Stemper (2018)). The only downside to the two-step approach is the resource-intensive synthetic data generation process which, however, is a one-off procedure.

Two-step calibration may be formulated either point- or grid-wise, depending on how the pricing map is parameterised. Respectively, option information such as strikes and maturities may either be cast as input parameters of the pricing map alongside model parameters (Bayer and Stemper (2018)), or be pre-set across a fixed strike-maturity grid on which the full model-induced IV surface is learnt "pixel by pixel" (Horvath et al. (2021)). Both of those pioneering methods are hereby considered for evaluation against their benchmark results. The findings of Bayer and Stemper (2018) are replicated for rBergomi and extended to what is potentially the first implementation of the point-based approach for rHeston. The predictive performance of the learnt IV maps on previously unseen synthetic data is oftentimes superior to what has been reported in the source material. However, deep point-based calibration to the SPX and VIX options proves rather unsatisfactory. The results highlight a major limitation of the approach in that there is insufficient expert knowledge of the model parameters' prior distributions and joint behaviour to effectively sample a representative training set point-wise. Grid-based calibration on the other hand turns out to be an extremely powerful tool. Under the network design proposed in the empirical research work of Rømer (2022), a main source of inspiration and ambitious benchmark for this report, extremely accurate replications of SPX volatility dynamics are attainable with either model. What is more, it turns out that rHeston can in fact be calibrated to the VIX, albeit not jointly with the SPX. A marginal yet guaranteed, consistent improvement over the calibration procedure proposed in Rømer (2022) has been identified. The idea is to calibrate the deep model to an IV surface using a piecewise constant parameterisation of the forward variance curve to begin with, then, keeping the optimised time-homogeneous parameters constant, to perform a second calibration, this time specifially of the forward variance curve to each individual IV slice of the full surface in turn. The procedure might be dubbed "grid-based deep calibration with a free forward variance curve".

In closing, rough volatility models are a truly captivating research area in quantitative finance. They provide the most accurate descriptions of volatility dynamics known to date. Some of the greatest derivatives pricing and hedging advancements in recent years are attributed to early representatives of the rough volatility class, leaving ample scope for the paradigm's untapped potential to be exploited further. By deploying the latest deep learning advancements to leverage the full mathematical properties of two prominent models, there is little room for interpretation that their aptness to capture the dynamics of the SPX and VIX European option markets is a nearly exhausted research question. Although neither model is capable of jointly calibrating the indices, the present comprehensive empirical study reinforces their ability to fit the SPX volatility surface remarkably well, and reveals the previously unrecognised capacity of the latter do also so for the VIX. Another pivotal finding concerns the high regularity of VIX volatility paths, hereby identified as close to that of a standard Brownian Motion and profoundly antithetic to SPX volatility drivers. It might of course be the case that the cryptic interplay between either model's parameters generates precise volatility fits at the expense of the parameters - such as the Hurst exponent indicative of volatility path roughness - losing their original significance. The empirical results cannot completely rule out that this is the case, meaning that the true roughness of volatility paths may after all be shared between the SPX and VIX. On the other hand, it is equally unclear whether the most recent development, the quadratic rough Heston extension (Gatheral et al. (2020)) which was shown to solve the joint calibration problem on several handpicked daily option chains, is not afflicted by the same fundamental change in model parameters' physical interpretation through calibration. Which, if either rough volatility model captures the true financial market behaviour, is a question left to future research.

Bibliography

C. Bayer and B. Stemper. Deep calibration of rough stochastic volatility models, 2018.

C. Bayer, P. Friz, and J. Gatheral. Pricing under rough volatility. *Quantitative Finance*, 16(6): 887–904, 2016.

C. Bayer, B. Horvath, A. Muguruza, B. Stemper, and M. Tomas. On deep calibration of (rough) stochastic volatility models, 2019.

M. Bennedsen, A. Lunde, and M. S. Pakkanen. Hybrid scheme for brownian semistationary processes. *Finance and Stochastics*, 21:931–965, 2017.

L. Bergomi. Smile dynamics ii. *Risk October*, pages 67–73, 2005.

F. Comte and E. Renault. Long memory in continuous-time stochastic volatility models. *Mathematical finance*, 8(4):291–323, 1998.

O. El Euch, J. Gatheral, and M. Rosenbaum. Roughening heston. *Risk*, pages 84–89, 2019.

J. Gatheral and R. Radoičić. Rational approximation of the rough heston solution. *International Journal of Theoretical and Applied Finance*, 22(03):1950010, 2019.

J. Gatheral, T. Jaisson, and M. Rosenbaum. Volatility is rough. *Quantitative finance*, 18(6):933–949, 2018.

J. Gatheral, P. Jusselin, and M. Rosenbaum. The quadratic rough heston model and the joint s&p 500/vix smile calibration problem. *arXiv preprint arXiv:2001.01789*, 2020.

A. Hernandez. Model calibration with neural networks. *Available at SSRN 2812140*, 2016.

S. L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies*, 6(2):327–343, 1993.

B. Horvath, A. Muguruza, and M. Tomas. Deep learning volatility: a deep neural network perspective on pricing and calibration in (rough) volatility models. *Quantitative Finance*, 21(1):11–27, 2021.

T. Jaisson and M. Rosenbaum. Rough fractional diffusions as scaling limits of nearly unstable heavy tailed hawkes processes. *The Annals of Applied Probability*, 26(5):2860–2882, 2016.

K. Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, 2(2):164–168, 1944.

D. W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.

S. E. Rømer. Empirical analysis of rough and classical stochastic volatility models to the spx and vix markets. *Quantitative Finance*, 22(10):1805–1838, 2022.