Institution: King's College London		
Unit of Assessment: 10 Mathematical Sciences		
Title of case study: Implementation of Deep Neural Networks for Option Pricing in the finance sector		
Period when the underpinning research was undertaken: October 2018 – December 2020		
Details of staff conducting the underpinning research from the submitting unit:		
Name(s): Dr Blanka Horvath	**Role(s) (e.g. job title):** Lecturer in Financial Mathematics	**Period(s) employed by submitting HEI:** From 01/09/2018
Period when the claimed impact occurred: 2019 - 2020		
Is this case study continued from a case study submitted in 2014? N		

1. **Summary of the impact** (indicative maximum 100 words)

Research in financial mathematics at King's College London has led to the development of new, deep-learning-based algorithms that can efficiently and accurately calculate derivative prices for modern market models that are challenging to handle computationally with traditional techniques. These algorithms can be used for both standard models that are already part of financial firms' quantitative libraries and for the modern class of 'rough volatility' models. King's College London's research has enabled market models of financial institutions to run up to 30,000 times faster than before, allowing them to be used as practical tools in the finance industry for the first time.

These algorithms have been implemented by a range of financial sector institutions, including banks (Morgan Stanley, Danske Bank, Natixis, DEKA), fund management companies (Synergis Capital Management, Atlantic House), and financial education and consulting firms (WBS, Thalesians). The speed and accuracy of the algorithms have already provided significant competitive advantage to their users, improved risk management, reduced computing costs and enabled firms to remain at the cutting edge of financial technology.

This research is regarded as a game changer within the financial sector, facilitating new, previously impossible, work, and making existing processes more efficient and accurate.

2. **Underpinning research** (indicative maximum 500 words)

In quantitative modelling, there is typically a trade-off between speed and accuracy. Widely used traditional stochastic financial models have the benefit of rapid calibration, while newer, more precise frameworks – such as 'rough volatility' – have proven too slow to be put into production.

Rough volatility models have been shown to capture many important statistical characteristics of financial time series better than the models that are currently used in the industry. Despite their advantages, rough volatility models had not been embraced by the financial industry (prior to King's College London's research), due to the prohibitively long computational time required to calibrate and run them with traditional techniques. King's College London research has shown how deep neural networks can turbocharge pricing and calibration for existing volatility frameworks but also transform rough volatility modelling from a theoretical pipe dream to a practicable reality.

Dr Blanka Horvath, of the King's College London Financial Mathematics group, has led the development of a new pricing and calibration approach, based on deep learning techniques, called Deep Pricing. A fundamental breakthrough was made in [R1, R2] which identified efficient machine learning approaches to speed up pricing and calibration. In this approach, the option prices corresponding to a stochastic financial model are learned by a deep neural network (DNN). With that, the stochasticity is removed from pricing and calibration algorithms, which speeds up the calibration process and saves ongoing computation costs: all possible parameter combinations of the underlying stochastic model are mapped directly to a specific shape of the

corresponding option price surface (or equivalently, the implied volatility surface). The surface is compressed to a representative grid, which helps further increase numerical efficiency, achieving a speed-up by a factor of up to 30,000.

Although the primary goal was to develop algorithms for rough volatility models, the results have unveiled a universal method that applies more generally to most of the common stochastic financial models. In particular, the techniques can be applied to the standard models that already exist in quantitative libraries of financial institutions. This allows financial institutions to run their models more frequently, to closely monitor the evolution of risks in real time, and to swiftly react as market conditions change. In collaboration with researchers from JP Morgan [R3] and Standard Chartered Bank [R4] Horvath has trained DNNs directly on data of stock prices and indices, which allow the removal of dependence on any particular stochastic model. This permits the simulation of price path dynamics that match market observations even more precisely than stochastic models, in a model-free way that is purely dictated by the data.

3. **References to the research** (indicative maximum of six references)

[R1] Deep learning volatility: a deep neural network perspective on pricing and calibration in (rough) volatility models, Horvath, B., Muguruza, A. & Tomas, M., 26 Oct 2020, In: Quantitative Finance, 21, 1, p. 11-27; DOI: 10.1080/14697688.2020.1817974

[R2] On deep calibration of (rough) stochastic volatility models, Bayer, C., Horvath, B., Muguruza, A., Stemper, B., Tomas, M.; 22 Aug 2019, In: https://arxiv.org/pdf/1908.08806.pdf

[R3] A Data-driven Market Simulator of Financial Time Series for Small Data Environments, Buehler, H., Horvath, B., Lyons, T., Perez Arribas, I., Wood, B., 21 Jun 2020, In: https://arxiv.org/pdf/2006.14498.pdf

[R4] Data Anonymisation, Outlier Detection and Fighting Overfitting with Restricted Boltzmann Machines, Kondratyev, A., Schwarz, C., Horvath, B., 27 Jan 2020, In: SSRN Electronic Journal, 2020, p. 1-27; DOI: 10.2139/ssrn.3526436

4. **Details of the impact** (indicative maximum 750 words)

Horvath's research on deep pricing and calibration [R1, R2] has been widely recognised in the financial sector as a significant advancement, as demonstrated by the author's receipt of the 2020 Rising Star in Quantitative Finance award from Risk Magazine, the leading international risk management, derivatives and regulation related publication for the sector. A Director of Quantitative Analytics at a major European bank, stated at that time, *"It normally takes 30 minutes to calibrate the XVA [risk management] model. With this neural network approach, it might take seconds – and it opens [us] up to other models that, before, were not even considered."* [S9]

Its emergence sparked considerable research and development, and an ongoing change in risk management and trading practices across the industry. As one leading European bank's (Danske) representative put it, *"This kind of investment technology for investment banks is like life or death: you either onboard it soon or you are out of business in the long run."* [S1]

Organisations which have implemented Horvath's research and changed their practices include banks, fund management companies, and financial education and consulting firms. We explore some examples of each of these below.

Impact on banks

Morgan Stanley is a multinational investment bank and financial services company with a revenue of USD46,700,000,000 in 2020 (the second largest investment bank globally in 2020). The bank's fixed income business saw particular value in Horvath's work, describing her paper on Deep Learning Volatility (DLV) [R1] as *"*an outstanding paper which is one of the most famous in the field*", with respect to their continual efforts to improve model performance and speed. According to the Global Head of Macro Data Strategies at Morgan Stanley, "Having faster models allows us to react faster to market events which is extremely important especially when it comes to automated trading… it is also important in terms of reducing compute costs".* [S3].

They first encountered Horvath's research at a 2018 conference sponsored by Morgan Stanley (*Machine Learning Methods for Assessing Financial Risk, Budapest, October 2018*), which led to follow-up discussions and an in-house presentation at one of Morgan Stanley's London offices. They implemented a calibration engine based on the DLV algorithms in 2020, devoting 1 FTE for an equivalent of 6 months for its implementation and experimentation. This has enabled them to *have "better risk management for 'exotic' products by using more complicated realistic models"* (such as rough volatility), while for their proprietary in-house models the advantages have been speed and stability of the calibration [S3].

Furthermore, it has enabled the firm to do live automated pricing and combine it with other pricing methods dynamically. *"This essentially means automated pricing at the trader's desk as information streams in, for which a standard integration pricing method would be too slow"* [S3]. An algorithm needs to react in microseconds to be able to do this, which is only possible with the new technology. It has provided them a thousand-fold increase in speed: the time to carry out typical calculations of 100 swaptions has been reduced from several seconds to a matter of milliseconds. Essentially, it also allows the firm to make more efficient use of computing resources and reduces computing costs which are substantial factors in any options desk of the business. [text removed for publication] [S3]

Danske Bank is one of the largest financial institutions in Northern Europe with over 22,000 employees in 2020. The bank has deployed a general-purpose pricing and risk engine based on neural networks in 2020, inspired by Horvath's DLV research [R1] on deep pricing and calibration. Horvath presented the research to Danske Bank's staff in Copenhagen in 2019. Two of the bank's most senior quantitative analysts have spent a year working on the implementation, refinement, and integration of deep pricing with the Bank's key technologies. [S1]

According to a Chief Quantitative Analyst at Danske Bank [S1], DLV stands out among pioneering deep pricing literature due to:

1. DLV is the first paper to ever offer a tractable solution for rough volatility models, with the potential to finally deploy this superior class of models on trading desks.

2. DLV successfully applied neural network technology to resolve a problem that could not be effectively resolved by traditional means for several years. DLV demonstrated to the industry that modern Machine Learning was indeed capable of resolving some of our hardest problems, and so it enabled a vast amount of further research into deep pricing, including a substantial investment from Danske Bank.

3. DLV also showed that successful results were not obtained by a blind application of a technology developed in a different field, but a careful and principled integration with more traditional financial modelling.

Danske Bank's main interest is in the speed and accuracy of the algorithm: being able to reproduce prices and risks computed with nested simulations thousands of times faster, which gives the Bank an immense competitive advantage, ultimately increasing its market share. In particular, the speed of the algorithms enables large risk calculations on traders' desks, providing real-time risk assessment critical to pricing and risk management. It also greatly reduces the time spent on calculations for regulatory compliance, improves risk management and enables the bank to offer lower prices to customers. [S1]

The Chief Quantitative Analyst at Danske Bank explains, *"The addition of deep pricing enables us to produce comprehensive risk reports and shields us from mispricing or unexpected losses, giving us confidence to increase exotic business while many others are pulling out. At the same time, a fast, reliable computation of capital regulations liberates considerable opportunities that we may have had to turn down otherwise. This ability, for now unique in the industry, […] gives us a considerable edge over the competition in the Derivatives business. All of this also results in societal benefits, because we can do business more comfortably and offer better prices to our customers. Deep pricing and its many extensions and generalizations (like Danske Bank's differential ML) are becoming a key part of the industrial pricing and risk toolkit […] and has the potential to massively improve risk management practices in the industry in the short term."* [S1]

Natixis is a large French investment bank employing approximately 80,000 people, implemented a calibration engine in 2019, based on Horvath's research and algorithms. They had employed Aitor Muguruza (co-author on the DLV paper [R1]) on a part-time basis through 2018-2019: as a result, Natixis was one of the first banks where DLV was implemented and tested. The benefit has been an immediate performance improvement in in-house pricing models, specifically faster computational times, which gives a competitive advantage. Furthermore, the engine provides flexibility in the models, and facilitates conducting R&D on the models: Natixis currently has an ongoing R&D programme based on the insights of this research. A Natixis representative has stated that *"this is shaping the way quant finance and calibration will be going forward in the future."* [S4]

DEKA is a German-based bank with around 4,600 employees, implemented neural network prototypes for simple models in 2020, following a discussion with Horvath at a practitioners' conference (The 15th Quantitative Finance Conference, Rome, October 2019). This implementation was inspired by the research and algorithms presented in Deep Learning Volatility. [text removed for publication] The model is aimed at enabling their traders to take and hedge positions in complex interest rate derivatives. The fast and accurate calibration to market data that Horvath's algorithm provides is the crucial component in these models. [S2]

Impact on fund management companies

Synergis Capital Management is a medium sized hedge fund employing around 50 people, has implemented a calibration engine with a rough volatility framework, based on Horvath's research, after hiring Aitor Muguruza as their Head Quant based on his achievements and profile as one of the recipients of the Rising Star Award [S9]. This implementation allows them to include rough volatility models that work across markets, enabling them to act more quickly in the market. It has also improved their understanding of rough volatility models for option pricing. Synergis has had two full-time interns working on this (5 person months in total) plus 1 full time employee dedicating 50% of his time. Interns are regularly hired to fine-tune this process and explore further opportunities that are based on DLV. A testimonial from Synergis states that, *"This implementation shows clients that we are on the forefront of research, which enhances our appeal to clients, ultimately leading to increasing [our] market share."* [S5]

Atlantic House Investments, part of the Catley Lakeman May Group, is an asset management firm managing over GBP1,500,000,000 in funds (in July 2020). They have taken a significant interest in further developing Horvath's research. They have invested a total of [text removed for publication] in a research project which began in 2020, funding a PhD student to further develop the pricing and market simulation frameworks to more realistic market scenarios. [S6]

Impact on financial education and consulting firms

World Business Strategies (WBS Training Ltd.) is a company which organizes industry-leading workshops and conferences for the capital markets and treasury divisions of investment companies worldwide, has implemented a prototype of DLV provided by Horvath, following a referral from Thalesians Ltd. (see below). This prototype is used to teach DLV to the company's training clients. Horvath's research has enabled WBS to increase their client base and has led to an improvement in customer satisfaction. The research presented in Horvath's papers is taught to WBS's consulting and training clients and is regularly discussed at the company's conferences. It is disseminated to conference attendees, attracting further consulting and training clients. The inclusion of Horvath's research has boosted the attendance of conferences and trainings by up to 20%. WBS conference attendees and students on WBS trainings regularly express their interest in Horvath's work on DLV and Market Generators since they are at the forefront of the current market practice. [S7]

Thalesians Ltd. is a neocybernetics and quantitative finance company, have implemented in 2019 a version of Horvath's DLV algorithm in their proprietary software libraries, taking on a new member of staff to carry out the work. Their introduction to DLV was via a presentation given by Horvath at a Thalesians' Seminar (London, April 2019). They testify that Horvath's research has significantly increased their standing with existing clients and has enabled them to attract two

major new institutional clients. Their existing clients report productivity improvements of 30-70%, resulting from the implementation of Horvath's algorithms.

The benefits that have resulted from Thalesians onboarding this research into their business include:

- an improvement of the performance of the existing derivative pricing software libraries by 50% (speed and precision of calculations);
- an increase in the community size by 20%;
- onboarding two new institutional clients;
- improved software development processes;
- a more modern approach to modelling incorporating machine learning (ML);
- improved understanding of ML and rough volatility models;
- more efficient use of developers' time;
- a significant contribution to their marketing efforts. [S8]

Thalesians Ltd. have stated that, *"Horvath's research is leading to an improvement of the speed and quality of derivatives pricing throughout the industry. This enables better evaluation of existing derivative products and the inclusion of new and more complex derivative products in portfolios. The use of these algorithms leads to a major competitive advantage over traditional pricing methods and ultimately an increase in market share. On our part, we see these advances as game changers in quantitative finance, both facilitating new, previously impossible, work, and making existing processes more efficient."* [S8]

5. **Sources to corroborate the impact** (indicative maximum of 10 references)

[S1] Testimonial from Chief Quantitative Analyst at Danske Bank

[S2] Testimonial from Quantitative Analyst at DEKA

[S3] Testimonial from Global Head of FID Data Strats at Morgan Stanley

[S4] Testimonial from Head of Quantitative Research at Natixis

[S5] Testimonial from Head of Quantitative Modelling and Data Analytics at Synergis

[S6] KCL and Catley Lakeman May Group's research studentship contract

[S7] Testimonial from Managing Director at World Business Strategies (WBS Training Ltd)

[S8] Testimonial from Director at Thalesians Ltd

[S9] Risk Awards 2020: New machine learning techniques bring 'rough volatility' models to life