



Arora, Manish Rajkumar (2025) Deep learning calibration framework for detecting asset price bubbles from option prices. PhD thesis.

<https://theses.gla.ac.uk/85088/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>  
[research-enlighten@glasgow.ac.uk](mailto:research-enlighten@glasgow.ac.uk)



University  
of Glasgow

# Deep Learning Calibration Framework for Detecting Asset Price Bubbles from Option Prices.

Thesis by:

Manish Rajkumar Arora

Submitted in fulfilment of the requirements of the **Degree of Doctor  
of Philosophy** in Accounting and Finance.

Adam Smith Business School, College of Social Sciences  
University of Glasgow

March, 2025



# Abstract

Asset price bubbles are associated with exuberant trading and unsustainable price increases, eventually culminating in abrupt collapses, causing widespread socioeconomic and financial devastation. Such phenomena have become increasingly frequent across various asset classes in recent times, and given the interconnectedness of global markets, the potential damage from bubbles bursts has significantly amplified. This is evidenced by the enduring repercussions of the *Global Financial Crisis* from nearly two decades ago.

This research is motivated to enhancing efficiency of bubble detection from option prices by employing neural networks, such that exuberance in the underlying asset is examined more comprehensively. Following an exhaustive overview of historical occurrences and detection methods, a three-step approach under the framework of the local martingale theory of bubbles was preferred. It captures forward looking expectations of market participants, and overcomes joint-hypothesis related issues, as opposed to traditional methods. The three-step approach relies on calibrating a sophisticated stochastic volatility jump diffusion model to market put observations, prior to testing for exuberance in call option prices. However, computational inefficiencies during calibration, limit its ability in extracting crucial information from option prices.

To overcome this roadblock, a deep calibration framework is constructed such that an optimally trained neural network is employed as a numerical solver for the desired stochastic process. This framework boosts computational efficiency by orders of magnitude, without sacrificing accuracy. It enables the extraction of crucial information regarding the formation of bubbles from the entire surface of option prices, without any compromises. Such calibrations further allow for exploration of bubbles within different maturity groups, and even across the lifetime of call options. At first, the deep calibration framework is applied to observe bubbles in the S&P 500 index, and then for a more recent case study on selected technology stocks. Finally, factors influencing the formation exuberance are examined, which also doubles down as a robustness test for the methodology.

Construction of the deep calibration framework, improves tractability of the three-step approach, making it more attractive for practitioners. Typically, despite possessing greater sophistication, inefficient stochastic processes are overlooked due to their computationally cumbersome nature, creating a trade-off between accuracy and efficiency. This trade-off is overcome with the application of neural networks, allowing for a deeper exploration and subsequently, greater comprehension of call option and underlying price bubbles. The significant boost in efficiency ensures that practitioners are able to extract real-time information regarding bubble formations, repeatedly, at superior speeds. The contributions of this research are important for policymakers and institutions seeking to properly manage risk and adjust positions in abidance with the ever-fluctuating essence of financial markets.

# Table of Contents

Abstract .....	3
List of Tables .....	9
List of Figures .....	11
List of Abbreviations .....	13
Dedication .....	15
Acknowledgements .....	17
Author's Declaration .....	19
Chapter 1: Introduction .....	21
Chapter 2: Literature Review .....	29
2.1    Historical Overview of Bubbles .....	29
2.2    Traditional Bubble Detection Methods .....	34
2.2.1    SADF and GSADF Tests.....	35
2.2.2    LPPLS Model.....	37
2.2.3    Criticism of the GSADF Test and LPPLS Model .....	38
2.3    Local Martingale Theory of Bubbles .....	39
2.3.1    Theoretical Background .....	40
2.3.2    Bubble Detection using Option Prices .....	52
2.3.3    Option Price Surface for Bubble Detection.....	63
2.4    Neural Networks: Option Pricing and Calibration.....	65
2.4.1    Artificial Neural Networks.....	66
2.4.2    Calibrating Stochastic Volatility Models with Neural Networks .....	69
2.4.3    One-Step Calibration Approach .....	72
2.4.4    Two-Step Calibration Approach.....	75
2.5    Summary and Discussion .....	81
Chapter 3: Methodology.....	87
3.1    Three-Step Approach.....	88
3.1.1    Step I: Model Selection .....	88
3.1.2    Step II: Model Calibration and Estimation .....	90
3.1.3    Step III: Estimation and Testing of Call Option Bubbles.....	92
3.2    Deep Learning Calibration Framework.....	94

3.2.1	Neural Networks: Multi-Layer Perceptron (MLPs).....	95
3.2.2	Step I: Forward Pass .....	104
3.2.3	Step II: Backward Pass .....	120
3.3	Summary and Discussion.....	125
Chapter 4: Model Estimation .....		129
4.1	Research Hypotheses .....	130
4.2	Market Data: S&P 500 Index.....	132
4.2.1	Filtration and Cleaning .....	133
4.2.2	Summary Statistics.....	135
4.3	Comparing Calibration of Stochastic Processes .....	137
4.4	Deep GSVJD Calibration.....	141
4.4.1	Monte Carlo Simulations vs. Neural Networks .....	142
4.4.2	GSVJD Calibration to the Entire Surface .....	145
4.4.3	GSVJD Calibration for Bubble Detection .....	148
4.5	Summary and Discussion.....	151
Chapter 5: Empirical Analysis .....		155
5.1	Call Options Data: S&P 500 Index .....	157
5.2	Call Option Bubbles: Most Liquid Cross Section.....	163
5.2.1	Key Statistics .....	164
5.2.2	Market Analysis .....	167
5.2.3	Model Performance.....	176
5.3	Call Option Bubbles: Maturity Groups.....	178
5.3.1	Key Statistics .....	179
5.3.2	Market Analysis .....	182
5.3.3	Model Performance.....	188
5.4	Call Option Bubbles: Option Lifetime.....	189
5.4.1	Key Statistics .....	190
5.4.2	Market Analysis .....	193
5.4.3	Model Performance.....	205
5.5	Summary and Discussion.....	207
Chapter 6: Case Study.....		211
6.1	Market Data .....	213
6.2	GSVJD Calibration for Bubble Detection .....	215
6.3	Call Option Bubbles: Most Liquid Cross Section.....	218

6.3.1	Key Statistics.....	218
6.3.2	Market Analysis.....	221
6.3.3	Model Performance .....	230
6.4	Factors Influencing the Formation of Bubbles.....	232
6.5	Summary and Discussion .....	239
Chapter 7: Conclusion.....		243
7.1	Summary, Contributions and Future Implications .....	243
References .....		251
Appendix .....		267
A.	Deep Calibration Framework: Training and Testing.....	267
A.1	Forward Pass .....	268
A.2	Backward Pass.....	283





# List of Tables

<i>Table 3.1: Hyperparameter Random Search Range.....</i>	<i>108</i>
<i>Table 3.2: Inputs and Outputs of the Forward Pass.....</i>	<i>115</i>
<i>Table 4.1: Options Data Filtration and Cleaning Procedures.....</i>	<i>133</i>
<i>Table 4.2: Summary Statistics of S&amp;P 500 Put options from HCV and Entire Surface datasets.....</i>	<i>136</i>
<i>Table 4.3: Calibration Performances of Stochastic Processes. ....</i>	<i>138</i>
<i>Table 4.4: Performance of Deep Calibration to HCV put options.....</i>	<i>143</i>
<i>Table 4.5: Performance of Deep Calibration to Entire Surface put options.....</i>	<i>146</i>
<i>Table 4.6: Summary Statistics of Deep Calibration to HCV put options. ....</i>	<i>149</i>
<i>Table 4.7: Summary Statistics of Deep Calibration to Entire Surface put options. ....</i>	<i>150</i>
<i>Table 5.1: Summary Statistics of Deep Calibration to HCV call options. ....</i>	<i>158</i>
<i>Table 5.2: Summary Statistics of daily most liquid call options from each Maturity Group. ....</i>	<i>159</i>
<i>Table 5.3: Summary Statistics of characteristics over the lifetime of Call Option contracts. ....</i>	<i>160</i>
<i>Table 5.4: GSVJD Performances on days with significant HCV call bubbles.....</i>	<i>165</i>
<i>Table 5.5: Model Performance.....</i>	<i>177</i>
<i>Table 5.6: Calibration performances on days with significant bubbles across maturity groups.....</i>	<i>180</i>
<i>Table 5.7: Model Calibration Performances over Maturity Groups.....</i>	<i>188</i>
<i>Table 5.8: Surface Calibration Performance over Call Option Lifetimes. ....</i>	<i>191</i>
<i>Table 5.9: Model Calibration Performance over Call Option Lifetime.....</i>	<i>205</i>
<i>Table 6.1: Summary Statistics of Options relevant to the Case Study.....</i>	<i>214</i>
<i>Table 6.2: Summary Statistics of GSVJD parameters relevant to the Case Study. ....</i>	<i>216</i>
<i>Table 6.3: GSVJD Performances relevant to the Case Study.....</i>	<i>219</i>
<i>Table 6.4: Model Performance relevant to the Case Study.....</i>	<i>231</i>
<i>Table 6.5: Correlation between Quadratic Variation and Call Option Bubbles.....</i>	<i>233</i>
<i>Table 6.6: Results for Regression (51). ....</i>	<i>234</i>
<i>Table 6.7: Results for Regression (52). ....</i>	<i>236</i>

<i>Table 6.8: Results for Regression (53) and (54). .....</i>	<i>238</i>
<i>Table A.1: Summary Statistics of the Random Search and Optimal Training Datasets. ....</i>	<i>272</i>
<i>Table A.2: Top Performing Architectures from the random search. ....</i>	<i>274</i>
<i>Table A.3: Optimal Training Performance. ....</i>	<i>281</i>
<i>Table A.4: Calibration Performances during Backward Pass.....</i>	<i>288</i>

# List of Figures

<i>Figure 3.1: Mechanisms of a Perceptron.</i>	97
<i>Figure 3.2: Multi-Layer Perceptron (MLP) Framework.</i>	98
<i>Figure 3.3: Grid vs. Random Search.</i>	105
<i>Figure 3.4: Illustration of a 3-fold cross validation.</i>	112
<i>Figure 3.5: Illustration of the Forward Pass.</i>	118
<i>Figure 3.6: Illustration of the Forward and Backward Pass.</i>	121
<i>Figure 4.1: Calibration Performances of Stochastic Processes.</i>	139
<i>Figure 4.2: GSVJD Calibration Performances: S&amp;P 500 Index (2019 – 2022).</i>	140
<i>Figure 4.3: GSVJD Calibration Performances: Benchmark vs. Deep Calibration Framework.</i>	144
<i>Figure 4.4: Deep Calibration Framework to the Entire Surface.</i>	147
<i>Figure 5.1: S&amp;P 500 Call Option Bubbles (HCV Calibration).</i>	169
<i>Figure 5.2: S&amp;P 500 Call Option Bubbles (Surface Calibration).</i>	173
<i>Figure 5.3: S&amp;P 500 Call Option Bubbles (Surface Calibration).</i>	183
<i>Figure 5.4: S&amp;P 500 Bubbles in the Lifetime of Call Options expiring in 2020.</i>	195
<i>Figure 5.5: S&amp;P 500 Bubbles in the Lifetime of Call Options expiring in 2021.</i>	199
<i>Figure 5.6: S&amp;P 500 Bubbles in the Lifetime of Call Options expiring in 2022.</i>	203
<i>Figure 6.1: NVDA Bubbles in HCV Call Options.</i>	221
<i>Figure 6.2: META Bubbles in HCV Call Options.</i>	223
<i>Figure 6.3: MSFT Bubbles in HCV Call Options.</i>	226
<i>Figure 6.4: AMZN Bubbles in HCV Call Options.</i>	227
<i>Figure 6.5: AMD Bubbles in HCV Call Options.</i>	229
<i>Figure A.1: Creation of Synthetic Dataset.</i>	270
<i>Figure A.2: Hyperparameter Range of Architectures in the top 10<sup>th</sup> percentile.</i>	275
<i>Figure A.3: Performances of top 10<sup>th</sup> percentile architectures across various Hyperparameters.</i>	277
<i>Figure A.4: History of Loss during Optimal Training.</i>	282
<i>Figure A.5: Backward Pass Calibrations using Optimal Architectures.</i>	285



# List of Abbreviations

ADF	Augmented Dickey-Fuller (Test)
AI	Artificial Intelligence
ANN	Artificial Neural Network
ATM	At the Money
BGD	Batch Gradient Descent
BSADF	Backward Supremum Augmented Dickey-Fuller
CaNN	Calibrated Artificial Neural Network
CEV	Constant Elasticity of Variance (Model)
CPI	Consumer Price Index
CPU	Central Processing Unit
DF	Dickey-Fuller (Test)
ELMM	Equivalent Local Martingale Measure
ELU	Exponential Linear Unit
ENUIMM	Equivalent Not Uniformly Integrable Martingale Measure
EUIMM	Equivalent Uniformly Integrable Martingale Measure
FDA	Food and Drug Administration
FOMC	Federal Open Market Meeting
GPU	Graphic Processing Unit
GSADF	Genealised Supremum Augmented Dickey-Fuller (Test)
GSVJD	Generalised Stochastic Volatility Jump Diffusion (Model)
HCV	Highest Cumulative Volume
IPO	Initial Public Offering
ITM	In the Money
LLR	Log-Likelihood Ratio
LPPLS	Log Periodic Power Law Singularity (Model)
LSTM	Long Short Term Memory (Neural Network)
MAE	Mean Absolute Error
MBD	Mini Batch Gradient Descent
MLP	Multi-Layer Perceptron (Neural Network)

MMA	Money Market Account
MSE	Mean Square Error
ND	No Dominance
NFLVR	No Free Lunch Vanishing Risk
OLS	Ordinary Least Square
OTM	Out of the Money
RAM	Random Access Memory
RIC	Refinitiv Identification Code
RMSE	Root Mean Square Error
RMSP	Root Mean Square Error Propagation
SABR	Stochastic Alpha Beta Rho (Model)
SADF	Supremum Augmented Dickey-Fuller (Test)
SGD	Stochastic Gradient Descent
TPU	Tensor Processing Unit
USMCA	United States-Mexico-Canada Agreement
WHO	World Health Organisation

# Dedication

*This thesis is dedicated to my grandfather, Banarsidas Arora.  
I never got to meet you, but I was blessed with the opportunity to honour and continue  
your legacy.*





# Acknowledgements

I express my gratitude to Prof. Charalampos Stasinakis and Dr. Panagiotis Karavitis for their supervision throughout the years. My sincerest of appreciation to the Adam Smith Business School Administration members, especially Angela Foster, Lorna Baillie, Sheena Phillips, Lorna Wilson, Christine Athrone, Ania Doswell, and Lukas Nespor, for their dedication towards always assisting me. This project was heavily reliant on the High Performance Compute Clusters, and I am deeply grateful to the allocated team's patience and determination in supporting my research needs.

Throughout this journey, I have been fortunate to work as a Graduate Teaching Assistant to Prof. Christian Ewald, Dr. Ankush Agarwal, Prof. Bart Taub, Dr. Chi-Hsiou Daniel Hung, Dr. Hormoz Ramian, Dr. Miguel Herculano, and Dr. Nnanna Oledibe; and as a research assistant to Dr. Yu-Lin Hsu and Prof. John Finch. This has been the most memorable and enjoyable part of my voyage, and I am very grateful for the opportunities to work with such esteemed company. Furthermore, I am sincerely appreciative of the advice and support provided by Dr. Ankush Agarwal, Dr. Hormoz Ramian, Dr. Miguel Herculano, Dr. Betty Wu and Jasmin Schar, at critical stages of my journey.

I am grateful to the support shown by my friends Michael and Abdul; it is always fun to talk about anything under the sun. I have been blessed with meeting several amazing people over the years. My aunts, Dr. Poonam Kurungati and Dr. Usha Manchanda, it is your achievements, that inspired me from childhood to pursue a doctorate degree. To a special teacher, Mr. Pratap Joshi, thank you teaching me how to learn. My friends back home, our ambitions have scattered us all over the globe, but those memories always bring back a smile. A special mention to the best girl in the world, with the most contagious ear-to-ear smile, Cairo.

I moved to Glasgow to pursue a doctorate degree, but little did I know that I would find you, Sneha. Your love, and support is the light that helped me navigate through the darkest of tunnels. Through all the hardships, your warm embrace was all that was needed,

who knew, it would be such a simple thing. Thank you for always being there and believing in me.

None of this would be possible without my parents, Rajkumar Arora and Vimal Arora. Thank you for always believing and pushing me to my utmost limits. To my father, Rajkumar Arora, you have always been my role model and taught me to dig as deep as required to overcome. Thank you for helping me understand that I have not only been pursuing a Ph.D. in Finance, but also in life. To my mother, all my childhood, you have passed on knowledge and values to me, which have been the most important education I have and will ever receive. I am forever grateful to my parents, for all the support, knowledge and strength you have provided me. You have been my role models, supervisors, protectors, providers, mentors, supporters, and best friends.

Lastly, to Hanuman Ji, I have called upon you through the toughest of moments, and you showed me the way. Thank you for your being my mentor, protector and friend.

# Author's Declaration

“I declare that, except where explicit reference is made to the contribution of others, that this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.”

Printed Name: Manish Rajkumar Arora

Signature:



# Chapter 1:

# Introduction

Asset price bubbles, often fuelled by investor speculation, occur when the market price deviates from its fundamental value. Historical examples range from *Tulipmania* in the 17<sup>th</sup> century; to the relatively more recent U.S. housing bubble, which caused the *Global Financial Crisis*, 2008. Initially, bubbles are associated with a joyous period of investors revelling in profitable trade, till prices reach an unsustainable level. At this stage, even the minutest of notions can trigger widespread panic sales, causing prices to rapidly collapse, leaving behind a wake of socioeconomic and financial devastation. Despite such consequences, the frequency of bubbles over the past few decades, across various asset classes, has significantly increased.

Given the constantly growing interconnectedness of global financial markets, potential damages from bubble bursts, could be beyond fathomable. To this date, several economies are yet to entirely recover from the detrimental effects of the *Global Financial Crisis*, which occurred nearly two decades ago. Hence, with great emphasis, there is an urgent need for developing a robust and accurate early warning detection system, such that timely measures can be implemented for damage control. This has piqued the interest of several academics and practitioners, including a certain famous victim of the *South Sea* bubble, Sir Isaac

Newton, who declared in defeat, “*I can calculate the motions of heavenly bodies, but not the madness of people.*” Several scholars have debated the rationality of investor behaviour during the initial phases of such events. Regardless, the identification of bubbles is crucial, considering the damage they are capable of inflicting.

Traditional methods aim to identify and date the presence of asset price bubbles by attempting to model fundamental values. The fundamental value of an asset is largely unknown, and therefore difficult to compute/validate. Consequently, traditional methods, including the popular regression techniques (Phillips et al., 2011, 2015) and LPPLS models (Johansen et al., 1999, 2000), suffer from a joint-hypothesis issue. The inconclusiveness of findings from these methods creates adverse implications for the implementation of appropriate and timely risk management measures. Furthermore, traditional methods require a large timeseries of historical data and focus on deriving information from spot markets. The former results in vulnerabilities towards structural breaks, whereas the latter prevents capturing the forward-looking nature of bubbles (Fusari et al., 2024).

The local martingale theory of bubbles is designed to identify exuberance over short time horizons, without computing the asset’s fundamental value. It overcomes joint hypothesis related concerns and eliminates vulnerabilities to structural breaks by not requiring large timeseries of data. Moreover, the theory provides a mathematical and economic framework for exploring the existence of bubbles in the underlying, by extracting information from option prices. In contrast to spot instruments, derivatives reveal information about the forward looking expectations of market participants. Bubbles are forward looking in nature, as they reflect preferences for purchasing an asset for the sole purpose of reselling it at a higher price, as opposed to holding it till liquidation. Therefore, capturing forward looking expectations from derivative markets is crucial for enhancing the ability to detect the formation of asset price bubbles.

In consistency with the local martingale theory of bubbles, Jarrow et al. (2010) provide a framework for detecting bubbles from option prices. The foundations of the framework are based on put options not revealing bubbles due to their bounded payoffs, in contrast to call options having unbounded payoffs and being capable of exhibiting exuberance. Fusari et al.

(2024) took advantage of these properties and proposed a three-step approach for detecting asset price bubbles. The method involves calibrating a sophisticated stochastic process to market put options for successfully capturing relevant forward looking information regarding the formation of call option bubbles, without suffering from a joint-hypothesis issue. Moreover, the three-step approach only requires real-time information, making it immune to structural breaks. However, given the need for calibrating a sophisticated stochastic process, the method suffers from large computational inefficiencies, which force its application to a single volatility smile, to strike a balance between efficiency and robustness. Furthermore, coupled with poor tractability, the associated inefficiencies make practitioners averse towards implementing the method for real-time bubble detection.

This research favours the three-step approach over other bubble detection methods, given its ability to overcome the joint-hypothesis issue, capture forward looking expectations of market participants, and reveal immunity to structural breaks. The method allows for accurate and robust bubble detection, but displays weaknesses with respect to efficiency and tractability, with the former restricting application to a single volatility smile. The computational burden arises during calibration of a stochastic volatility jump diffusion process. Nevertheless, this process is highly sophisticated and plays a crucial role in capturing relevant market information, given its ability to address price jumps, volatility mean-reversion and strict local martingale tendencies. Therefore, motivated by the continuously growing need for an accurate and robust approach, which is also efficient and tractable enough to meet practical requirements, this study proposes the application of neural networks. A deep calibration framework is suggested in attempts to answer the following research question: *How to improve the efficiency and tractability of the three-step approach and subsequently extend its application for bubble detection?*

The deep calibration framework is the main contribution of the thesis. Its importance is cemented during this research, through the comparison of various stochastic volatility processes during market turmoil periods. This comparison is the second contribution of the thesis, as it reveals processes capable of capturing price jumps and strict local martingale tendencies to be superior for bubble detection, despite the inefficiencies associated with their calibration. The sophistication of a stochastic process must not be sacrificed for efficiency,



as this would create vulnerabilities to the joint hypothesis issue. Therefore, the application of neural networks is recommended. According to the universal approximation theorem (Cybenko, 1989; Hornik et al., 1989; Hornik, 1991) neural networks are capable of estimating any function, with a given level of accuracy. They mimic the functions of an animal's central nervous system to understand complex non-linear relationships between input and output variables. Furthermore, given modern computational advancements, neural networks are capable of pricing thousands of option contracts, over thousands of underlying instruments, at high frequencies (De Spiegeleer et al., 2018).

The motivation behind proposing a deep calibration framework is to employ neural networks as numerical solvers for the stochastic process. The intention is to boost efficiency and tractability, without sacrificing accuracy, such that the three-step approach becomes more favourable amongst practitioners. The deep calibration framework is revealed to bolster efficiency by several large orders of magnitude and approximately double the accuracy of fitting market put options. This provides further certainty over the method not suffering from the joint-hypothesis issue. Moreover, improvements in both areas, allow for an additional contribution by expanding the application of the three-step approach to extract information from all traded options. Given calibrations to the entire surface of daily options, this research proposes the exploration of asset price bubble formations across various call option maturities, and the lifetime of certain call contracts. The former is inspired from Jarrow & Kwok (2021, 2024) but original with respect to the application of the three-step approach, whereas the latter is an entirely novel approach for examining the presence of exuberance. Both avenues enhance the comprehension of bubbles and improves the ability to implement timely and appropriate risk management measures during periods of exuberance. The remainder of this thesis is structured in the following manner, with relevant contributions from each chapter being highlighted.

Bubbles in literature are explored in *Chapter 2*, by detailing their existence, evaluating traditional methods, and arguing the superiority of the three-step approach from Fusari et al. (2024), given its ability to overcome the joint-hypothesis issue, reveal immunity to structural breaks and capture the forward-looking expectations of market participants from option prices. However, computational costs associated with calibrating a highly sophisticated

stochastic volatility process, limit the extraction of information to a single volatility smile. The abundant information present in the various smiles across the option price surface is overlooked. To address these computational costs, neural networks are proposed. Literature regarding their application to option pricing is discussed, which is followed by evaluating two approaches of calibrating stochastic volatility models. First, the one-step approach, which directly outputs model parameters, is computationally efficient but lacks output validity. Next, the two-step approach, involving an optimizer to match output with market observations, calibrates parameters from the input layer, allowing for real-time validation. In favour of the latter and inspired by Liu et al. (2019) and Horvath et al. (2021), a deep calibration framework is proposed to enhance the three-step approach for bubble detection.

*Chapter 0* begins by introducing the three-step approach, with each stage being thoroughly explained. The first step reveals preference for a generalised stochastic volatility jump diffusion (GSVJD) process, that is capable of detecting strict local martingale tendencies. The GSVJD model, in the second step, is calibrated to market put options, prior to statistically testing for the presence of exuberance in call options. Given the lack of a closed form solution, estimations of the GSVJD model must rely on Monte Carlo simulations. It is a highly sophisticated process, nevertheless, given its computational cumbersome nature, the GSVJD is likely to be overlooked, in favour of tractability. Hence, in order to allow for efficient, accurate and robust real-time bubble detection, inspired by Liu et al. (2019) and Horvath et al. (2021), a deep calibration framework is proposed for the GSVJD model, as the first and main contribution of this thesis. The construction of this framework, and integration into the three-step approach, along with a detailed explanation regarding the mechanisms of neural networks, are also provided.

The application of the deep calibration framework is empirically tested in *Chapter 4*. This chapter compares performances of various stochastic processes, with each being calibrated to the S&P 500 options data using Monte Carlo simulations, over a period comprising of the COVID-19 induced market crash. This comparison is the second contribution of the thesis, as it reveals the GSVJD model to be the most accurate but least efficient. It contributes to relevant literature by demonstrating the superiority of stochastic volatility jump diffusions models, during period of market crashes, when identifying bubbles. The ability to capture

jumps and strict local martingale tendencies in the underlying price improves accuracy, and therefore, sophistication of the GSVJD model must not be sacrificed for efficiency, especially during a market crash period. Hence, an optimal neural network architecture for learning the dynamics of the process is identified through 3-fold cross-validation across a broad range of hyperparameters, using an extensive random search. Performances of the deep calibration framework are compared against those from Monte Carlo simulations. It nearly doubles accuracy and provides a 254-fold improvement in efficiency. The deep calibration framework for estimating the GSVJD model forms the second contribution of this chapter. Furthermore, due its flexibility and computational speed, the framework was applied towards calibrating from the entire daily option price surface. This allowed for bubble detection to be conducted by capturing the forward-looking expectations of market participants to a greater extent. This broader calibration approach, across a wide range of daily option maturities, for bubble detection marks the third contribution of the chapter.

*Chapter 5* is dedicated to the detection of bubbles in the S&P 500 index. Parameters of the GSVJD model are calibrated from the daily most liquid cross sections, and surfaces of put options, using the deep calibration framework. Its first contributions focus on identifying exuberance within the most liquid call options, with respect to each parametric set. Next, parameters obtained from daily surfaces are used to observe bubbles in call options belonging to three maturity groups, marking the third contribution. This chapter further contributes by employing the same parameters towards revealing exuberance over the lifetime of call options. The expansion of the three-step approach to GSVJD calibrations from the daily surface of put options, acknowledges exuberance during the COVID-19 induced crash, and immediate recovery period. Identified bubbles in regard to the latter are captured better when examining various call option maturities and option lifetimes. However, the findings associated with the former provide a novel contribution, as exuberance is detected in association with circuit breakers, displaying consistency with the local martingale theory of bubbles.

A case study is conducted in *Chapter 6*, applying the deep calibration framework to selected technology stocks. It contributes by detecting bubbles, from calibrating GSVJD parameters to the most liquid maturity and entire surface of put options. In addition, following Fusari et

al. (2024), factors influencing the formation of bubbles are explored, along with the robustness of the methodology. A negative relationship between bubble magnitudes, and call option volume is discovered. This contribution suggests that the presence of bubbles may signal an upcoming bearish period, given the drop in call option volume. Furthermore, the proximity to earnings announcements and high underlying volumes, positively impact the occurrence of exuberance. Finally, in *Chapter 7*, the thesis is summarised and concluded, with *Appendix A* providing additional information on the training and testing of the deep calibration framework. It is on this note, that the research proceeds to providing a comprehensive overview of relevant literature.



# Chapter 2:

# Literature Review

Asset price bubbles reflect deviations in market prices from the fundamental value. A large section of economic and financial literature is dedicated towards comprehending and timestamping such phenomena. The identification of bubbles is crucial, as their existence is often followed by a crash, which more often than not causes severe crises. Initially, asset prices rise to unsustainable levels, and upon realizing this, market participants take corrective measures, causing prices to plunge, subsequently bursting the bubble. Depending on the severity of the price plunge, a financial crisis could be triggered. There are several well-documented historical bubble episodes, which are carefully examined next, to stress the importance of requiring a robust detection methodology.

## 2.1 Historical Overview of Bubbles

The earliest occurrence of an asset price bubble, as documented in Mackay (1841), can be traced back to ‘*Tulipmania*’, during the 17<sup>th</sup> century Dutch Golden age. A massive spike in demand for a rare species of Tulips was experienced by 1634, as they became symbolic of

luxury and status. Tulips typically bloom in April and May, however, producers realized that their bulbs could be uprooted and sold. In response to the growing demand, speculators entered the market, offering forward contracts on the bulbs. A bubble formed, contracts were purchased, with expectations of selling the bulbs at higher prices in the future. At the peak of the bubble, these contracts had exchanged hands multiple times. Eventually, the bubble collapsed in 1637, since buyers were unable to pay the exuberant prices initially agreed upon when entering into the contract. The ‘*Tulipmania*’ episode did not have any significant long term impacts on the Dutch economy, barring the breaking of trust between participants. However, in contrast to current times, financial markets across nations, and various asset classes, were not as heavily integrated.

Consider two basic theories regarding asset pricing, from Malkiel (1985). First, the ‘*firm-foundation*’ theory, which attempts to define an intrinsic value based on an analysis of fundamentals. Second, the ‘*greater fool*’ theory, that states the worth of the asset is determined by whatever another investor is willing to pay. The latter was evident during ‘*Tulipmania*’, as individuals purchased contracts, at high prices, solely under the perception of selling at even greater valuations to a bigger fool. The *South Sea* and *Mississippi* bubbles from the early 18<sup>th</sup> century document similar investor behaviour, which eventually resulted in significant economic loss (Scherbina & Schlusche, 2014). Additionally, both reveal the growing economic risks of financial bubbles within integrated markets<sup>1</sup>.

Upon taking a government IOU of nearly £10 million in 1711, an English Firm, *South Sea Company* was granted a monopoly over trade to the Southern Seas. Initially, profits were made from providing shares to the public in exchange of government securities, that were undertaken by the company. There was a strong demand for the stock, and despite failed quests, deceptive public appearances to reveal prosperity kept prices stable. In 1720, the company offered to fund the entire national debt, which drove the public into a buying frenzy. Nevertheless, as the true value of the company was realised, a widespread panic caused the share price to collapse. Given the promise to finance national debt, indirectly, the

---

<sup>1</sup> See Malkiel (1985) and Scherbina and Schlusche (2014) for more details on the respective events.

company was funding the war against France. Its collapse put the entire national economy under major threat.

Similarly, an effective marketing programme combined with exaggeration of wealth in *Louisiana*, the French based *Mississippi Company*, experienced a massive hike in its share price. In 1719, it obtained rights to collect most of France's taxes, and to trade outside Europe, in addition to absorbing the entire nation's debt. Stock sales were used to fund, company operations and national bonds. Approaching January 1720, share prices started to fall, as investors started to sell their equity. However, sellers were made to accept paper notes, instead of gold coins. As the bubble collapsed, more notes were printed, and the excess money supply left the economy in a state of hyperinflation.

Evidence of such behaviour continued to plague the economic wellbeing of nations overtime, and given the growing integration of financial markets, even that of the globe was not spared during more recent events. The '*Roaring Twenties*' witnessed by the U.S. economy, experienced significant growth, following an expansion in the construction and consumer goods sectors. Stock market growth was supported by the introduction of purchasing shares on margin, which fuelled speculative behaviour. Eventually, in October 1929, the boom period suddenly halted, with the collapse of stock and real estate markets. The crash severely drained household wealth, kick-starting the *Great Depression* period. Likewise, Japanese stock and real estate markets experienced massive growth at the end of the 1980s. Prior to the bubble bursting, all of the land in Japan, was valued at 4 times greater than that in the U.S. (Scherbina & Schlusche, 2014). The collapse was devastating, resulting in Japan's *lost decade*.

Not shortly after, in the final decade of the 20<sup>th</sup> century, U.S. stock markets experienced a major boom, following rapid growth in the internet sector, and closely linked fields. Given the excessive demand for internet stocks and IPOs, the *Dot-com* bubble was fueled by banks reducing standards for taking such companies public. A large number of firms were yet to record profits, whereas several revealed identical business models whilst competing in the same market. The firms were destined to fail, and when the eventual crash occurred, the U.S. economy was forced into a recessionary period, for approximately 8 months (Quinn &



Turner, 2020). The dot-com bubble did not have any large scale global impacts; however, it changed the perceptions of academics about the formation of bubbles .

Previously, academia rarely entertained the notion of asset price bubbles, with some rare instances arguing the ‘rationality’ of such phenomenon. Bubbles are deemed rational, given investors willingly purchase overvalued assets to resell at higher prices. Alternatively, when such actions are driven by herd behaviour or psychological factors, bubbles are considered irrational. The irrational behaviour of agents prevents new information from being entirely incorporated into prices. However, it is difficult to determine the rationality, as the concept is not clearly defined. For instance, some participants could favour non-financial gains and invest due to social or environmental benefits provided by a certain company. Furthermore, if the definition of rationality is extended to include misinformed and ideological investments, then no behaviour can be considered irrational. The reader is referred to Quinn & Turner (2023) for a more detailed debate regarding rational and irrational bubbles. Nevertheless, this debate is futile for the focus of the current study, aligning with relevant recent literature. Since the *Global Financial Crisis*, 2008, the existence of bubbles was no longer questioned, with disputes over rationality taking a backseat. Moreover, scholarly efforts urgently shift towards developing an early detection system.

The *Subprime Mortgage Crisis*, and subsequent *Global Financial Crisis* in 2008 were triggered by the collapse of a U.S. housing bubble. Exuberance in housing prices was created by homeowners gaining easy access by collateralizing their properties. The U.S. housing sector experienced massive growth, and to meet demands, banks imposed loose standards for mortgage recipients. Furthermore, innovative debt instruments attracted heavy international investments. The collapse of the housing market forced several borrowers to miss mortgage payments, with approximately \$7 trillion in home equity being erased (Quinn & Turner, 2020). The subsequent global crisis left long-lasting global effects, with several European nations entering a deep recession. For instance, between 2007 and 2009, the GDP per capita for Ireland, Spain and United Kingdom fell by 11.2%, 4.9% and 6.1%, with youth unemployment rates peaking at 30.8%, 55.5%, and 21.3%, respectively (Quinn & Turner, 2020). Furthermore, Spain’s GDP per capita continued to fall till 2013. This raised large concerns about the risk of tightly integrated global financial markets. Consequently, central

banks increased their interest in risk management. The crisis revealed the capability of bubbles causing enormous and long-lasting harm to global economies. Additionally, it raised fears over their more frequent occurrences than previously perceived. After the *Dot-com* bubble and *Global Financial Crisis*, Quinn & Turner (2023) highlight more recent occurrences in the Chinese Stock market and Bitcoin prices.

The aftermath of bubble bursts is commonly associated with financial crisis periods, along with fractured trust within banking and regulatory systems. Alternatively, Quinn & Turner (2023) highlight that bubbles can result in the development of transformative technology, and financial innovation, which stimulate a foundation for future technological developments and economic growth. Consider the *Dot-com* bubble, a period during which internet companies received tremendous capital investments. It is well established that majority of the firms failed, however, those that survived, paved that path for modern internet applications. Such developments would not have been possible without the capital received during the bubble period. Furthermore, financial innovations from such periods, contribute significantly to future economic growth prospects. Venture capitalism, which fuelled the *Dot-com* bubble, continues to be the beating heart of Silicon Valley. However, when regulatory standards drop, such innovations are abused to cater to the purely speculative demand of investors. Subsequently, asset prices reach unsustainable levels, prior collapsing and triggering socioeconomic and financial devastation. For example, companies being issued IPOs without proper vetting during the *Dot-com* bubble, or construction and trading of Mortgage-Backed Securities in the build up to the *Global Financial Crisis*.

Naturally, timely and strict regulatory measures, along with proper risk management practices are capable of halting such unethical practices and controlling the formation of bubbles. However, timely detection has proven to be great challenge. It is crucial to distinguish between growth in asset prices that are driven by fundamentals, as opposed to those purely fuelled by speculation. The requirement of an accurate early detection warning system has motivated academics and practitioners. Given the frequent occurrence of asset price bubbles, which possess a major threat to the wellbeing of closely integrated global financial markets, it is vital to evaluate the approaches designed for their identification.

## 2.2 Traditional Bubble Detection Methods

The existence of bubbles is attributed to the willingness of investors to pay in excess of the fundamental value, believing that the asset can be sold at an even greater price. Hence, the traded asset price comprises of two components: the fundamental value, and bubble. Several methods have been developed for capturing the latter by computing the former. However, these methods attribute failure of validating the model to the existence of bubbles (Gürkaynak, 2008), making findings inconclusive, given vulnerability to joint-hypothesis related issues (Jarrow, 2015; Fusari et al., 2024).

For further comprehension, abiding by Gürkaynak (2008), the following methods are evaluated *Variance Bounds Test* (Shiller, 1981; LeRoy & Porter, 1981), *West's two-step Test* (West, 1987), *Integration/Cointegration Tests* (Diba & Grossman, 1988a, 1988b), and *Intrinsic Bubbles* (Froot & Obstfeld, 1991; Driffill & Sola, 1998). Each approach tests the null hypothesis of the asset price equating its fundamental value, implying the absence of bubbles. The *Variance Bounds Test* acknowledges bubbles, when the variance of the actual price exceeds its bound computed from that of the fundamental value. Whereas *West's two-step Test* rules out model misspecification in favor of revealing an existing bubble. Both tests seek to reject the null hypothesis by a process of elimination, with respect to low price volatility of model misspecification. However, they fail to consider any theoretical properties of bubbles.

The *Integration/Cointegration Test* incorporates non-observant fundamentals, and a theoretical property, that rational bubbles cannot start, and if observed, they must have always existed. A Dickey-Fuller (DF) test is implemented to assess stationarity of price and dividends, along with a cointegration test. Bubbles are non-existent in the presence of stationarity and cointegration. However, the non-stationarity of bubbles holds, only because they are unable to restart after a collapse. Evans (1991) reveal that bubbles are capable of growing, even after collapsing to a small non-zero positive value. This discloses the weaknesses of implementing standard unit root tests for identifying exuberance. Finally, the *Intrinsic Bubbles* approach considers the potential of bubbles not being correlated with fundamentals and growing exogenously. Exuberance is non-existent in the presence of a

linear relationship between price and dividends. Alternatively, a non-linear relationship signals an explosive divergence between the two factors. Froot & Obstfeld (1991) modelled dividends using a random walk, whereas Driffill & Sola (1998) considered a regime-switching approach. The results were contradictory, as the former observes exuberance, whereas the latter reasons non-linearity to regime shifts. Regardless of the justification, both Froot & Obstfeld (1991) and Driffill & Sola (1998) find some form of non-linearity. This reinforces the claim that traditional methods suffer from a joint-hypothesis related issue. The inability to distinguish between misspecified fundamentals and bubbles makes findings about exuberance inconclusive. All that is revealed is the present value model fails to fit market data, given the exclusion of stylized factors (Gürkaynak, 2008). This leaves the interpretation of the result to the perception of practitioners. Naturally, having a less restrictive fundamental pricing model, incorporating either stylized factors or sophistication of the process would improve fitting. However, it could also result in a high rate of false negative detection. Hence, the root of the problem is identified as the estimation of the fundamental value.

Traditional methods face several challenges with the accurate identification and timestamping of bubbles. Nevertheless, over time, the quality of bubble detection techniques has improved, and the existence of such a phenomena started to receive more acceptance amongst academics and practitioners. A major contributing factor was the increased global frequency of asset price bubbles and crash events since the penultimate decades of the 20<sup>th</sup> century. Previously, economic and financial literature debated the existence of bubbles, and if so, whether they could be classified as rational or irrational (Quinn & Turner, 2023). However, since their existence is questioned less, over the last two decades literature has been streamlined towards developing early warning detection systems that allow timely intervention for effective risk management. Next, two popular methodologies, employed by academics and practitioners for empirical analysis, are discussed.

### 2.2.1 SADF and GSADF Tests

The difficulties faced by standard unit root tests have been well documented in Evans (1991). Phillips et al. (2011) suggest using recursive regression techniques on right-tailed unit root

tests for detecting mild explosive behavior in asset prices. The forward recursive regressions involve repeated estimations of an autoregressive model. The ADF test coefficient is computed over a subset of the sample period, which is incremented by an observation at each pass. If the supremum of the ADF (SADF) coefficient exceeds respective critical values, the null hypothesis of a unit root process is rejected in favour of a mildly explosive one. The SADF test can reveal presence of explosiveness but cannot timestamp bubbles. For timestamping, ADF statistics computed over the sample are considered. Chronologically, if the ADF statistic exceeds its critical value, the origination date of the bubble is marked. Whereas termination is acknowledged, when for the first time, following origination, the critical values are greater.

In the presence of multiple bubbles, the Phillips et al. (2011) approach was found to be inconsistent, and with reduced power. The inconsistency stems from recursive estimations of the autoregressive coefficient for identifying the second bubble, being dominated by presence/lack of explosiveness from prior periods. Specifically, if the duration of the first bubble is greater (smaller) than that of the second, the latter is overlooked (detected with a delay). The aspirations of developing a robust warning system for identifying financial bubbles, inspired Phillips et al. (2015) to propose a Generalised SADF (GSADF) test for detecting explosiveness, and a recursive backward regression technique for dating bubbles. In contrast to the SADF, the GSADF test records supremum of ADF statistics, across windows with varying start and end points. Additionally, Phillips et al. (2015) timestamp bubbles using a backward SADF (BSADF) test. The ADF regression is estimated over backward expanding windows, with their supremum being compared to critical values, for acknowledging the start/end of bubbles, similar to Phillips et al. (2011). The BSADF approach improves flexibility and is consistent with the timestamping of bubbles, as it considers various properties from even within the subsample. It is revealed superior to methods from Phillips et al. (2011), and Homm & Breitung (2012), in addition to a sequential version of the former.

The methodology is popular amongst scholars for detecting bubbles across a wide range of asset classes. Su et al. (2017), Perifanis (2019) and Khan et al. (2021) applied the approach to crude oil prices, whereas Corbet et al. (2018), Park & Yang (2024) and Osman et al.

(2024), used it to explore the formation of bubbles across multiple cryptocurrencies. Zhao et al. (2021) implemented the test on Chinese stock indices, whereas Nguyen & Waters (2022) concentrated on the S&P 500 index. The latter acknowledged the *dot-com* and *housing* bubble periods, along with more recent occurrences in December 2020 and January 2021, but also provided evidence of the method being sensitive to sample and time series selection. Kwong & Wong (2022) paired the dating of bubbles with Value-at-Risk (VaR) models, and Demmler & Fernández (2024) tested for bubbles in the NASDAQ index. Furthermore, Acharya (2024) considered individual constituents of the S&P 500 index that recorded returns in excess of a certain threshold. Interestingly, Zhang & Yao (2016) examined S&P 500 index data over two centuries, and revealed the LPPLS method to be more consistent with fewer false signals, in comparison to the BSADF method.

### 2.2.2 LPPLS Model

The Log-Periodic Power Law Singularity (LPPLS) model was established by Johansen et al. (1999, 2000). It identifies bubbles and crashes by capturing the faster-than-exponential growth and accelerating log-periodic oscillations in asset price trajectories. The former characteristic addresses unsustainable growth (with a finite crash time), stemming from positive feedback, and is represented by the power law singularity component. When singularity is reached, the bubble bursts/crashes, and the regime shifts away from super-exponential growth. The latter acknowledges competition between traders exhibiting herding behavior and fundamental valuation investors, which respectively align with positive and negative mean-reverting feedback (Shu & Song, 2024). It is depicted by the log-periodic function, which regards the hierarchical flow of accelerated panic that eventually terminates the bubbles.

The asset price is assumed to follow a standard diffusion process with a drift and discontinuous jumps. The coefficient of the latter quantifies the amplitude of the crash, with jumps being governed by the hazard rate. The hazard rate mimics the behavior of noise traders and measures the probability of a crash occurring in the next instance, given that it is yet to be experienced. Essentially, providing the probability of a large group of agents placing sell orders at the same time. The drift component increases with hazard rate to

compensate investors for growing crash risk. A two-step calibration approach is presented in Filimonov & Sornette (2013) that stabilizes, and improves efficiency of fitting LPPLS parameters to market observations for bubble detection. Furthermore, Sornette et al. (2015) introduce confidence indicators that measure sensitivity of bubble patterns to the fitting window interval time.

The application of the LPPLS model for bubble detection over a wide range of assets, is frequently seen across literature. Zhang & Yao (2016) identify bubbles in crude oil prices, whereas Shu & Zhu (2020) discovered multiple clusters in the CSI 300 index. The Russell 2000, Wilshire 5000, S&P 400 and S&P 500 indices are analyzed by Shu et al. (2021) to determine if the 2020 market crash was caused by exogenous factors. It is important to indicate that the LPPLS method is notorious for overlooking the impact of exogenous factors. Furthermore, Song et al. (2022) extended this analysis to 10 global indices, whereas Shu & Song (2024) examined the S&P 500 index after March 2020.

### 2.2.3 Criticism of the GSADF Test and LPPLS Model

It is well documented that the LPPLS model is unable to detect bubbles caused by exogenous events such as monetary, fiscal, geopolitical shocks, or those occurring due to a global pandemic (Shu et al., 2021; Song et al., 2022; Shu & Song, 2024). Park & Yang (2024) reveal the over sensitive nature of the method to initial conditions, with even the smallest of changes reducing detection accuracy. Additionally, it has a requirement for the expected asset price to be increasing throughout a bubble period, which is not always the case (Brée & Joseph, 2013; Karimov, 2017).

A bubble can persist, even though expectations are decreasing, as traders fail to synchronize their positions (Abreu & Brunnermeier, 2003). A similar interpretation can be made for the methods from Phillips et al. (2011, 2015), which consider the bubble component to be a submartingale, with explosive expectations. Both, the recursive regression tests and LPPLS models, make assumptions over the existence of bubbles, with respect to the drift of the price process (Protter, 2013). Furthermore, the approaches attempt to model the fundamental

value, and despite improving on traditional methods, are still vulnerable to the joint-hypothesis issue.

In the former, Protter (2013) reveal that Phillips et al. (2011, 2015) do not take into account the terminal payoff when computing the fundamental value. Whereas there is a lack of empirical importance towards the LPPLS fit to market observations, in comparison to the attention given to its parametric values. Calibrating parameters to market prices, which potentially exhibit bubbles, for determining the fundamental value can be erroneous, and adversely impact the conclusiveness of results. It is therefore suggested to explore the local martingale theory of bubbles, which does not require estimation of the fundamental value for detecting exuberance in asset prices.

## 2.3 Local Martingale Theory of Bubbles

A bubble is observed when the market price of the asset, exceeds its fundamental value, making the latter a vital component for identification. However, the fundamental value of an asset is vague, and ultimately unknown (Protter, 2016), making it difficult to spot bubbles prior to bursting (Cox & Hobson, 2005). Furthermore, the estimation of this component requires forecasting of future cash flows and risk-free rates, which as indicated by Chaim & Laurini (2019), can be impacted by exacerbated (disappointed) expectations during bubble (crash) periods. Traditional methods attempt to model the fundamental value and suffer from joint-hypothesis related issues.

The local martingale theory of bubbles overcomes this issue, making computation of the fundamental value redundant. Recall, the recursive regression and LPPLS methods, focus on the drift of the price process and model bubbles as submartingales, with explosive characteristics. In contrast, the local martingale theory concentrates on the diffusion component, and its growth with respect to the asset price. In the presence of bubbles, the asset price process would display strict local martingale tendencies (Delbaen & Shirakawa, 2002; Kotani, 2006; Mijatović & Urusov, 2012). The behavior of a strict local martingale processes reveals typical traits of bubbles, as they spike to a high value prior to collapsing and settling at a lower one. Economically, it can be interpreted as an asset price deviating



from its fundamental value, highlighting the willingness of investors to overpay for a potential upside (Stahl & Blauth, 2024). Hence, the aim is to detect bubbles by identifying whether or not asset prices reveal strict local martingale tendencies, as opposed to being a true martingale.

Bubbles are categorised by Jarrow et al. (2007) as Type I, II and III, in accordance with the liquidation time of the asset. Traditional methods aim to capture type II bubbles, which occur in assets with finite lives over infinite horizons. Given the infinite horizon, a large time series of data is required, which is problematic when financial markets experience structural breaks (Choi & Jarrow, 2022; Fusari et al., 2024). Additionally, Fusari et al. (2024) acknowledge the reliance of these methods to concentrate on extracting backward looking information from historical prices. Alternatively, type III bubbles prevail within assets with bounded lives in a finite horizon and are not impacted by structural breaks. They arise when the price process is a strict local martingale and can be estimated over shorter time periods. Such bubbles capture the willingness of investors to purchase an asset with the sole purpose of reselling it at a higher price in the future. The local martingale theory of bubbles is engineered towards detecting type III instances.

### 2.3.1 Theoretical Background

The foundations of the theory were built in the works of Loewenstein & Willard (2000a, 2000b), Cox & Hobson (2005), Heston et al. (2007), and Jarrow et al. (2007, 2010). In abidance to these papers, an overview of the theoretical background, is provided. The reader can be referred to Protter (2013) and Jarrow (2015) for a more detailed synopsis. Furthermore, for simplicity and convenience, notations are consistent with Jarrow (2015) and Fusari et al. (2024).

### 2.3.1.1 Economy: Fundamental Value and Bubbles

A continuous or discrete time model may be considered with interval  $[0, T]$ . For simplicity and consistency with the related literature, a finite horizon continuous time model<sup>2</sup> is presented. The randomness in the economy is characterized by  $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ , with filtration  $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ , satisfying the usual hypotheses<sup>3</sup>. The sigma-algebras,  $\mathcal{F}_t$  are subsets of the sample space,  $\Omega$  with probability measure,  $\mathbb{P}$ , determining the occurrence of each. The market comprises two traded assets: a risky asset (stock) and locally riskless<sup>4</sup> money market account (MMA). Markets are assumed to be competitive and frictionless, implying that traders are price takers<sup>5</sup>, with trading costs and constraints being non-existent. At time  $t$ , given a nonnegative default-free rate of interest,  $r_t$ , value of the MMA can be denoted by  $B_t = e^{\int_0^t r_s ds}$ .

$$G_t = S_t + B_t \int_0^t \frac{1}{B_s} dD_s \quad (1)$$

Considering liquidation time  $u \leq T$ , let  $D = (D_t)_{0 \leq t \leq u}$ , represent a nonnegative càdlàg semimartingale process of cumulative cash flows obtained from holding the risky asset. The market price of the risky asset is defined as a nonnegative càdlàg semimartingale,  $S$ , with  $S_u = X$ , at  $u$ . There are no further cash flows, and the value of the risky asset is invested in the MMA, such that  $S_t = B_t \frac{X}{B_u}$ , and cumulative gains are defined by  $G_t$ ,  $G_t \geq 0$ . The asset does not exist beyond the liquidation date, hence each process is stopped prior to  $u$ , shifting focus to the  $[0, u]$  horizon (Jarrow et al., 2010). Intuitively, considering the risky asset to be a stock, it could potentially cease to exist on a future date, given bankruptcy, a buyout/merger, or the company being broken up by antitrust laws. Hence, the lifetime of an asset is a stopping time (Protter, 2013).

---

<sup>2</sup> A discrete model can be procured by setting the price process to be constant between consecutive time periods (Jarrow, 2015; Fusari et al., 2024).

<sup>3</sup> The usual hypothesis is met when  $\mathbb{F}$  is right-continuous, and the probability space is complete (P. E. Protter, 2005).

<sup>4</sup> An asset locally riskless when its returns over the next time step are known with probability = 1.

<sup>5</sup> Traders believe their actions do not have a quantity impact on the market price.

### 2.3.1.2 No Free Lunch Vanishing Risk

The semimartingale property of the stock price process, is a necessary, however not sufficient condition for eliminating arbitrage (Protter, 2013). Therefore, the economy must be assumed to satisfy No Free Lunch Vanishing Risk<sup>6</sup> (NFLVR) conditions. It is a technical extension of the minimal structure provided by the standard no arbitrage definition (Jarrow, 2015). Both exclude self-financing trade strategies with nonnegative liquidation values that require no investment, whereas NFLVR additionally discards limiting arbitrage opportunities<sup>7</sup> (Jarrow et al., 2007, 2010; Jarrow, 2015). Hence, only admissible trade strategies are considered. A self-financing trade strategy bound from below, is admissible, presenting as a lower bound on the wealth process of an individual (Jarrow et al., 2007, 2010). Individuals are prevented from borrowing, if their debt becomes excessive, resulting from constraints on negative wealth (Loewenstein & Willard, 2000a, 2000b), limiting the replication of asset payoffs, and ruling out arbitrage opportunities in the form of doubling strategies<sup>8</sup>.

Shorting the bubble payoff, and purchasing a cheaper replicating trade strategy would imply presence of an arbitrage opportunity (Heston et al., 2007). However, given the unbounded nature of losses that could potentially stem from a short position, and constraints on negative wealth, arbitrage is ruled out. Short selling an asset with a bubble is not admissible, given that losses could violate the lower bound (Cox & Hobson, 2005; Heston et al., 2007). The lower bound for admissible trade strategies can be interpreted as constraints enforced by regulators/traders, given rational price expectations. Most importantly, it is the reason behind the existence of bubbles (Jarrow et al., 2007, 2010). Consider short constraints being non-existent, then if a bubble were to occur, the dramatic increase in prices, will be entirely offset by short positions. The formation of bubbles would be prevented given that traders possess

---

<sup>6</sup> The NFLVR condition is assumed to hold, not directly applied during pricing (Protter, 2013).

<sup>7</sup> Limits of sequences of self-financing trade strategies, requiring zero investment, which have a small loss probability.

<sup>8</sup> Strategies that generate positive returns on borrowed funds, with a probability of 1, during a finite time period (Jarrow et al., 2007).

unbounded wealth process with no borrowing limitations<sup>9</sup>. Shorting constraints<sup>10</sup> make taking such positions inadmissible as the trade would be terminated if losses cross a certain threshold. This relieves the downward selling pressure on the asset, allowing for the existence of asset price bubbles (Cox & Hobson, 2005; Heston et al., 2007; Jarrow et al., 2007).

### 2.3.1.3 Fundamental Value and Bubble Detection

NFLVR is crucial condition for the existence of price bubbles. In accordance with the first fundamental theorem of asset pricing (Delbaen & Schachermayer, 1994), if and only if a probability measure,  $\mathbb{Q}$  equivalent to  $\mathbb{P}$ , exists, such that  $G_t/B_t$  is a local martingale, NFLVR holds. Hence,  $\mathbb{Q}$  is referred<sup>11</sup> to as an equivalent local martingale measure (ELMM). A stochastic process,  $X_t$ , is considered a local martingale if and only if, when stopped on a sequence of stoppage times,  $\tau_n$ , for  $\lim_{n \rightarrow \infty} \tau_n = \infty$ , is a martingale. Such a process can either be a uniformly integrable martingale, a true martingale or strict local martingale (Jarrow, 2015). In the latter,  $G_t$  would not be a true martingale, rather supermartingale, as it is expected to decrease over time (Laurini & Chaim, 2021).

$$FV_t = E_{\mathbb{Q}} \left( \frac{S_u}{B_u} + \int_0^u \frac{1}{B_s} dD_s \middle| \mathcal{F}_t \right) B_t \quad (2)$$

$$\beta_t = S_t - FV_t \quad (3)$$

The fundamental value of the risky asset  $FV_t$ , is the price a trader is willing to pay to purchase the asset, such that it is held till liquidation (Jarrow et al., 2007, 2010). It incorporates the expectation, under  $\mathbb{Q}$ , of the discounted liquidation value and sum of cash flows. A bubble,  $\beta_t$  reflects a trader's preference to sell the risky asset at a higher price in the future, rather than holding it till liquidation. It reveals the difference between the market price and fundamental value and occurs when the latter is exceeded. The existence of a

---

<sup>9</sup> In such scenarios traders would not face issues meeting the requirement of maintaining the unbounded marked-to-market losses.

<sup>10</sup> There two types of shorting constraints relevant to the argument; Structural limitations (limited ability and/or costs of borrowing the asset), and those associated risks with short positions (Jarrow et al., 2007).

<sup>11</sup> It is also commonly referred to as the risk neutral measure.

bubble emphasizes that the resale price of the risky asset is superior to that paid for holding it till liquidation. It captures the bigger fool belief (Jarrow, 2015), as the investor (the fool) purchases the asset, believing in the possibility of reselling in the future at a greater price to another (the bigger fool).

The second fundamental theorem of asset pricing reveals that if markets are complete, there exists a unique ELMM,  $\mathbb{Q}$ , in contrast to the infinite possibilities prevalent in incomplete markets (Jarrow et al., 1999). In either case, the fundamental value is determined under the  $\mathbb{Q}$  measure, assuming a unique one is selected by the market, when incomplete (Jarrow et al., 2010). The theorem can be mathematically represented with respect to  $G_t$ , in abidance with Jarrow et al. (2010). Let  $\mathcal{M}_{loc}(G_t)$  denote a set of ELMMs, then  $|\mathcal{M}_{loc}(G_t)| = 1$  and  $|\mathcal{M}_{loc}(G_t)| \geq 2$ , for complete and incomplete markets, respectively, with  $|\cdot|$  representing cardinality. Furthermore, consider  $\mathcal{M}_{UI}(G_t)$  and  $\mathcal{M}_{NUI}(G_t)$  to be mutually exclusive non-empty subsets of  $\mathcal{M}_{loc}(G_t)$ , in the given order, comprising of Equivalent Uniformly Integrable Martingale Measure (EUIMMs) and Equivalent Not Uniformly Integrable Martingale Measures (ENUIMMs). Therefore, if markets are complete,  $\mathcal{M}_{loc}(G_t)$  will contain a single element, either from  $\mathcal{M}_{UI}(G_t)$  or  $\mathcal{M}_{NUI}(G_t)$ .

### 2.3.1.4 Asset Price Bubbles: Complete Markets

The second fundamental theorem reveals that  $|\mathcal{M}_{loc}(G_t)| = 1$ , comprising of either an EUIMM or ENUIMM. The presence of bubbles in the prices of risky assets, and derivatives written on them, have been documented in (Cox & Hobson, 2005; Heston et al., 2007). However, it is Jarrow et al. (2007) that characterise a theorem for asset price bubbles under the NFLVR condition in complete markets. Given the characteristics of the asset's liquidation values, if  $\beta_t \neq 0$ , then there are three and only three possibilities.

1.  $\beta_t$  is a local martingale, possibly a uniformly integrable martingale, under  $\mathbb{Q}$ , if  $P(u = \infty) > 0$ .
2.  $\beta_t$  is a local martingale but not a uniformly integrable martingale when unbounded, if  $P(u < \infty) = 1$ .

3.  $\beta_t$  is a strict local martingale under  $\mathbb{Q}$ , if  $u$  is a bounded stopping time.

The final possibility is the most economically interesting, as  $\beta_t$  must be a strict local martingale for its existence (Protter, 2013). Given the asset price comprises of two components; the fundamental value, which is a true martingale, and the bubble; if  $\beta_t$  is a strict local martingale, by extension so would the asset price. Hence, if and only if the asset price process is a strict local martingale a bubble exists. Considering the three types of bubbles,  $S_t$  reveals the following decomposition.

$$S_t = FV_t + \beta_t = FV_t + (\beta_t^1 + \beta_t^2 + \beta_t^3) \quad (4)$$

The bubble component is further split to depict type I, II, and III occurrences, such that  $\beta_t = \beta_t^1 + \beta_t^2 + \beta_t^3$ . Type I bubbles ( $\beta_t^1$ ) are a càdlàg non-negative uniformly integrable martingale, occurring when the asset has an infinite life, with a payoff at  $u = \infty$ . Jarrow et al. (2007) consider the example of fiat money and state such bubbles to be economically uninteresting. Given the asset does not expire, nor pay cashflows, its fundamental value is equivalent to 0, then  $\beta_t^1$  converges to  $S_{u=\infty}$ , almost surely. Hence, the entire value of the asset comes from a bubble, making it trivial from an economic perspective. Contrastingly, type II and III bubbles have strong economic relevance.

Type II bubbles ( $\beta_t^2$ ) are witnessed as càdlàg non-negative non-uniformly integrable martingales in assets with unbounded finite lives, with  $\beta_u^2 = 0$  ( $\beta_u^2$  converges to 0, almost surely). Non-uniformly integrability is a critical threshold for such bubbles. A bubble is observed when the market price exceeds the asset's fundamental value, therefore an arbitrage opportunity can exist when shorting the former and longing the latter of the price components. Jarrow et al. (2007) highlight that such trading strategies must be terminated prior to  $u$ , at some finite point. There is a possibility that the bubble can exist after the termination date, making the strategy 'risky' and therefore, not an arbitrage. Whereas, for type III bubbles ( $\beta_t^3$ ), it is the admissibility condition that prevents arbitrage. Type III bubbles, classified as càdlàg non-negative supermartingales (and strict local martingales), with  $\beta_t^3$  converging to 0, almost surely, materialise in asset with bounded lives. At maturity,  $\beta_u^3 = 0$ , for an asset with a bounded life. Taking a short position in the asset is inadmissible, therefore given NFLVR, type III bubbles are allowed to exist.

The time horizon  $[0, u]$  considered is finite, with  $P(u < \infty) = 1$ , as the asset does not exist beyond its liquidation date. Therefore, only type II and III bubbles exist in the current framework.

i)  $\beta \geq 0$ .

ii) *If the bubble has a finite maturity ( $u < \infty$ ), then it must burst on or before  $u$ , hence*

$$\beta_u 1_{\{u < \infty\}} = 0.$$

iii) *If  $\beta_t = 0$ , then  $\beta_v = 0$ , for  $v \geq t$ , stating that bubbles cannot restart after bursting, and must be present in the start of the model to be existent.*

The non-negativity<sup>12</sup> of the bubble component is revealed in i), indicating that the fundamental value of the asset is the lower bound to its market price, since traders have the option to retrade (Jarrow, 2021). If not traded, the fundamental value will equate to the market price since the asset is held till liquidation. However, if retraded, then the market price will exceed the fundamental value, with the bubble component capturing willingness of traders to sell at a higher price, instead of holding till liquidation. The final condition poses as a weakness to this model, preventing the birth/restart of bubbles, since there is only one ELMM. Additionally, given the presence of bubbles in the stock price, the put-call parity, which is almost never empirically violated, does not hold (Cox & Hobson, 2005; Heston et al., 2007). Hence, Jarrow et al. (2007) introduced the No Dominance (ND) condition from (Merton, 1973).

### 2.3.1.5 No Dominance

The ND condition<sup>13</sup> is satisfied, when there are no dominated assets, implying that given all things equal, financial agents have preference for more over less. An asset is dominated if  $a$

---

<sup>12</sup> Protter (2013) reveal bubbles can be negative for foreign exchange. For instance, assume the USD is in a bubble with respect to EUR, then the EUR is in a negative bubble specific to the USD.

<sup>13</sup> View Jarrow et al. (2007) for the extended definition.

*trade strategy with lower cost exists, such that its cashflows and liquidation value are always greater or equal (and strictly greater with a positive probability) to that of the asset* (Jarrow, 2015). In contrast to NFLVR, which is a mispricing opportunity that can be exploited<sup>14</sup>, ND compares two investments, ensuring equivalence of demand and supply (Jarrow, 2015, 2021). For instance, if an asset is dominated, it will not be held by traders and therefore be in excess supply. Moreover, it rules out the possibility of suicide strategies (Jarrow et al., 2010).

The satisfaction of ND ensures that the put-call parity holds by definition (Merton, 1973; Jarrow et al., 2007, 2010). Given NFLVR and ND are satisfied, according to the third fundamental theorem of asset pricing (Jarrow, 2021), there exists an equivalent probability measure, such the  $G_t$  would be a martingale. In complete markets, since  $|\mathcal{M}_{loc}(G_t)| = 1$ ,  $G_t$  will be a martingale and not a strict local martingale under  $\mathbb{Q}$ , ruling out the possibility of bubbles. Initially this conclusion was arrived at by (Jarrow et al., 2007), by considering two approaches for obtaining cash flows. The first involves purchasing and holding the asset, whereas in the second, an admissible trade strategy is entered to replicate the fundamental value. If a bubble (type II or III) were to exist, then the fundamental value would be exceeded by the market price, making the cost of the trade strategy cheaper. Since both approaches provide similar payoffs, the trade strategy would be dominant. Therefore, the existence of type II and III bubbles would be a direct violation of the ND condition. Hence, in a complete market, satisfying NFLVR and ND asset price bubbles are non-existent<sup>15</sup>.

### **2.3.1.6 Asset Price Bubbles: Incomplete Markets**

Given infinitely many ELMMs within incomplete markets, it is not feasible to rule out the presence of bubbles across all possibilities. The selection of ELMMs determines the fundamental value, subsequently the existence of bubbles. Considering NFLVR and ND conditions are satisfied, the third fundamental theorem states that an equivalent martingale measure must exist (Jarrow, 2021). Complete markets comprise of a single measure, within

---

<sup>14</sup> Eventually prices adjust, and the opportunity disappears.

<sup>15</sup> In complete markets, satisfying NFLVR and ND, type I bubbles exist, however they are not economically important (Jarrow et al., 2007), and not visible in the finite horizon framework.



which the asset price process must be a martingale. However, incomplete markets reveal infinite possibilities, with just one required to be an equivalent martingale measure. Therefore, the existence of bubbles distinguishes between martingale and local martingale measures cannot be ruled out.

In contrast to complete markets, different ELMMs are exhibited across time, resembling regime shifts (or alternatively, changing perceptions about the fundamental value), in the economy (Jarrow et al., 2010; Protter, 2013; Biagini et al., 2014). The fundamental value is determined by a selected local martingale measure, resembling the market chosen regime (given sufficient number of traded derivatives). Jarrow et al. (2010) investigated such occurrence by considering random stopping times (representing a regime switch), and stochastic processes to respectively signify the number of shifts and characteristics of regimes. Recall  $\mathcal{M}_{UI}(G_t)$  and  $\mathcal{M}_{NUI}(G_t)$  to be mutually exclusive non-empty subsets of  $\mathcal{M}_{loc}(G_t)$ , comprising of EUIMMs and ENUIMMs, respectively. Consider a dynamic market to have selected the  $Q_{UI} \in \mathcal{M}_{UI}(G_t)$  measure, under a particular regime. If a switch to  $R_{UI} \in \mathcal{M}_{UI}(G_t)$  would occur, the fundamental price would be independent<sup>16</sup> of the change, as  $G_t$  is a uniformly integrable martingale in both measures. However, if the measure were to shift to  $R_{NUI} \in \mathcal{M}_{NUI}(G_t)$ , then the fundamental value of the asset would decrease. When bubbles are absent, the fundamental value of the asset is equivalent to the arbitrage-free market price. However, under  $R_{NUI}$  that need not be the case, as the fundamental value would be less than or equal to the market price, resulting in the birth of a bubble.

The ELMM is selected by the market, depending on the state of the economy, and number of regime shifts (Jarrow et al., 2010). Any shift in the ELMM would capture random structural shifts in the economy that correspond to “*changing beliefs, preferences, endowments, institutional structure, market clearing mechanisms, or political/regulatory considerations*”, and potentially give birth to bubbles (Jarrow et al., 2010; Jarrow, 2015). Protter (2013) and Biagini et al. (2014) provide alternative economic interpretations for the shifting of measures. The former comments on bubbles being born due to ‘easy money’,

---

<sup>16</sup> Fundamental value of other assets could change.

with the changing measure capturing increased liquidity due to speculators gaining access to large pools of funds. Whereas in the latter, the switch resembles a change in investor perceptions, from optimism to pessimism, reflecting that the market price is exceeding the fundamental value of the asset. Even in the case of static incomplete markets, which do not have regime shifts (in contrast to dynamic markets), Jarrow et al. (2010), reveal the possibility of bubbles. In complete markets, ND rules out such occurrences, given the possibility to create dominant admissible replicating trade strategy. Static markets are the first natural generalization of complete markets but are not entirely complete. Hence, such strategies may not exist, therefore, the presence of bubbles cannot be excluded.

Despite the satisfaction of NFLVR and ND conditions, bubbles are existent within incomplete markets, due to shifting ELMMs corresponding to regime shifts in economy. Similar to complete markets, if  $\beta_t \neq 0$ , then there are three and only three possibilities: Type I ( $\beta_t^1$ ), II ( $\beta_t^2$ ), and III ( $\beta_t^3$ ) bubbles. However, since the time horizon considered is finite, only the latter two are prevalent with the following conditions (Jarrow et al., 2010).

$$i) \quad \beta_u = 0.$$

$$ii) \quad \beta_u 1_{\{u < \infty\}} = 0.$$

$$iii) \quad \text{If } \beta_t = 0, \text{ then } \beta_v = 0, \text{ for } v \geq t.$$

$$iv) \quad \text{In the scenario of no cashflows, } S_t = E_{\mathbb{Q}}\left(\frac{S_T}{B_T} \middle| \mathcal{F}_t\right) B_t + \beta_t^3 - E_{\mathbb{Q}}\left(\frac{\beta_t^3}{B_T} \middle| \mathcal{F}_t\right) B_t, \text{ for } t \leq T \leq u.$$

The first three condition are similar to those witnessed when markets are complete. Bubbles are non-negative and must burst before/on maturity,  $u$ , and cannot restart indicating that they should be present at the start of the model to prevail. Recall, condition iii) was a major weakness from Jarrow et al. (2007), and is only valid when the market is static, with no regime shifts. In dynamic markets, since there are infinite ELMMs across time, along with the likelihood of regime shifts, the birth/restart of bubbles is possible, provided a decrease in the fundamental value of the asset, due to a switch from  $\mathcal{Q}_{UI} \in \mathcal{M}_{UI}(G_t)$  to  $\mathcal{R}_{NUI} \in \mathcal{M}_{NUI}(G_t)$ . The final condition reveals market price,  $S_t$  to vary from its conditional

expectation, due to a type III bubble component, since the fundamental value,  $\beta_t^1$ , and  $\beta_t^2$  are martingales. Hence,  $S_t$  is not a martingale, consistent with bubbles existing if and only if it is a strict local martingale. The NFLVR conditions rules out bubbles as arbitrage opportunities, and given ND, the put-call parity holds (Jarrow, 2021). The characteristics regarding the presence of bubbles in a single asset can be extended to a vector containing multiple (Jarrow et al., 2010). Therefore, bubbles can exist in market indexes, along the same theorem and conditions<sup>17</sup>.

### 2.3.1.7 Empirical Application

According to the local martingale theory, a bubble exists if and only if the asset price process is a strict local martingale (Jarrow et al., 2007, 2010). The difference between a true martingale and strict local martingale are observed by analysing volatility behaviour. If the volatility integral from Delbaen & Shirakawa (2002), Kotani (2006), and Mijatović & Urusov (2012) is finite, the price process reveal strict local martingale tendencies, indicating the presence of type III bubbles. A bubble exists if the volatility grows faster than the stock price. The employment of non-parametric estimators for computing volatility over observable stock prices was considered by Jarrow et al., (2011). However, estimates only for visible price points can be gather, making it difficult to examine the volatility tail. Hence, the authors proposed implementation of parametric and non-parametric extrapolation techniques. The former comprises of selecting a parametric function form for the diffusion coefficient and calibrating it to volatility estimations prior to extrapolating. The non-parametric method implements a reproducing kernel function in a Hilbert space to extrapolate, by fitting over the estimated volatility values from observed stock prices.

Chaim & Laurini (2019) and Laurini & Chaim (2021) explore the formation of bubbles in across various asset classes, using non-parametric and parametric methods. The non-parametric approach is similar to Jarrow et al., (2011), whereas the parametric method calibrates a stochastic volatility process from Andersen & Piterbarg (2007) to market observations. Its parameter space can be partitioned into two mutually exclusive sets, which

---

<sup>17</sup> Just because a constituent reveals a bubble, that does not imply the same for the index (Jarrow et al., 2010).

respectively indicate the absence or presence of strict local martingale properties. Building on Jarrow et al., (2011), Choi & Jarrow (2022) propose a technique based on a modified convex hull of estimated volatilities. Following the estimation of volatility from observed price intervals, a power function is selected to fit lower and upper convex hulls. Next, parameters from the respective hulls are examined to reveal the rate at which volatility is rising with respect to the price. Choi & Jarrow (2024) improve on the method by incorporating cash flows, unequal price observation times and intervals during variance computations, along with corrections for small sample bias in the volatility estimator. Moreover, a strong robustness check was implemented, standard errors during the estimation of volatility bounds were adjusted for heteroskedasticity and autocorrelation, and probabilities of bubbles during an inconclusive outcome were provided.

Aforementioned methods identify bubbles within a selected period; but shape of the asset's volatility curve can change (Protter, 2013). Therefore, a bubble could exist during one period, burst and disappear, and even be reborn. The issue with Jarrow et al., (2011) and related works, is the inability to date specific instances of bubbles. Whilst a predefined period is a must for non-parametric methods, it is not required when calibrating parameters from market observations. Daily calibrations of stochastic processes allow for timestamping short-term bubbles. However, Chaim & Laurini (2019) and Laurini & Chaim (2021) conduct annual calibrations. Recall, strict local martingale tendencies capture short-term trading profits investors seek to make by purchasing for reselling at a higher price. Several works of literature (Piiroinen et al., 2018; Stahl & Blauth, 2024; Biagini et al., 2024; Fusari et al., 2024) in this area, document the existence short term bubbles to be more frequent than perceived. In addition to the inability of capturing such occurrences, the calibration over a long-time duration creates vulnerabilities to structural breaks.

Given stock prices evolve globally, Obayashi et al. (2017) highlight the need to consider some form of time dependence. If not, upon entering a bubble, the stock will remain in one till all of time. The authors expand on Jarrow et al. (2011) by proposing daily volatility estimation within a rolling window. The bubble signals are smoothened out by employing a Hidden Markov model, aiming to correct false positive occurrences corresponding to flash instances of high volatility. Furthermore, to deal with the occasional minor signals indicating

isolated bubbles, the start (death) of exuberance is marked by a certain threshold for rising (falling) prices. Empirically, there might be greater instances of short bubbles, when accounting for the bias created by the latter filter. A long short term memory (LSTM) neural network is trained by Bashchenko & Marchal, (2020) to comprehend the approach from Obayashi et al. (2017), such that it can output classifications of either true martingale or strict local martingale, after receiving price inputs.

Despite addressing the issue of dating bubbles, the approach, similar to those discussed, utilises information from the spot market for detection. The selection of instrument is vital for the kind of information the practitioner is attempting to extract. Spot markets comprise of backward-looking information, as opposed to the forward-looking type, prevalent in derivative markets (Han & Kuo, 2017). The forward-looking nature of options, allow prices to reflect future expectations of market participants, regarding several economic factors (Funahashi, 2023), making them ideal for bubble detection (Fusari et al., 2024).

### 2.3.2 Bubble Detection using Option Prices

Bubbles are forward-looking in nature and reveal the expectations of traders that purchase only to resell at a higher price. Such information cannot be captured in spot markets, as opposed to those for options. The growth of trading activity within options over a wide range of strike prices and maturity durations has increased their employment for signaling towards the existence of bubbles in the underlying prices. Options are abundant in information, on the forward-looking expectations of market participants, and even more, across varying contract durations. Incorporating the foundations of local martingale theory of bubbles, methods utilising option prices, are focused on capturing forward-looking expectations, instead of historical occurrences/trends, and require only current market observations. Such approaches are less vulnerable to hindsight bias and subsequently improve real time bubble detection. Furthermore, as acknowledged by Jarrow (2015), Jarrow (2016) and Fusari et al. (2024), within a specific framework, bubble detection using option prices can entirely overcome the joint-hypothesis related issue.

### 2.3.2.1 Theoretical Background

Option contracts comprise of bubbles, when the underlying price process reveals strict local martingale tendencies (Cox & Hobson, 2005; Heston et al., 2007; Jarrow et al., 2007). This resembles that market price of the option exceeding its fundamental value, or cheapest initial cost of replication trade strategy, obtained from market prices of the underlying risky asset. Therefore, exuberance in the underlying, impacts options prices, and could potentially result in the contingent claim exhibiting a bubble (Jarrow et al., 2007). Options have bounded lives and display only type III bubbles. Nevertheless, even the presence of type I and II bubbles in the underlying price can impact on the option price (Jarrow et al., 2007).

Focusing on complete markets, Cox & Hobson (2005) explore the impact of bubbles on option prices. When the price process is a true martingale, put-call parity holds, as opposed to when exuberance is exhibited. Put options have bounded payoffs, hence the static arbitrage conditions are not violated. Alternatively, call options have unbounded payoffs, which could result in large fundamental values during bubble periods. In such scenarios, traders are required to post collateral for shorting strategies. These collateral requirements ensure the put-call parity holds, as arbitrage through short positions would violate the admissibility condition. The fair price of the call option is equivalent to its fundamental value, plus the collateral requirement. The latter is interpreted as the bubble component. If there is no collateral requirement, then the option price is equivalent to its fundamental value. Therefore, when the underlying displays a bubble, standard option pricing methods are only applicable to put options, not call contracts. Given that the put-call parity holds, and the asset price process is a strict local martingale, then bubbles in the underlying would signal towards the exuberance in the call option price (Heston et al. 2007).

The NFLVR condition considers only admissible trade strategies, ruling out type III bubbles as arbitrage. The admissibility condition makes strategies that short options revealing bubbles (e.g. Covered Calls) unfeasible, given limitations on marked-to-market losses faced by investors. Under NFVLR, Cox & Hobson (2005) and Heston et al. (2007) reveal option price bubbles in complete markets. However, arbitrage can still exist in the form of suicide strategies, only if the trader is willing to take up the losses (Heston et al., 2007). The presence

of bubbles in the underlying or option prices would result in the violation of the put-call parity. Hence, Jarrow et al. (2007) introduced the ND condition, to ensure that the put-call parity holds, and suicide strategies are prevented. Suicide strategies let agents to throw wealth away by engaging in a reverse doubling strategy and therefore allow for the existence of bubbles in riskless assets, subsequently within prices of those that have bounded payoffs. When NFLVR and ND conditions are met, such assets cannot display bubbles, which is critical for empirical testing (Jarrow, 2015; Fusari et al., 2024).

Given the satisfaction of both conditions, Jarrow et al. (2007) show bubbles to be absent in complete markets. Hence, motivating Jarrow et al. (2010) to reveal their existence with incomplete markets. The framework from Jarrow et al. (2007, 2010) enables using option prices to overcome the joint-hypothesis issue during bubble detection. It takes advantage of put options not exhibiting bubbles and recommends capturing market features by accurately estimating traded prices, since market and fundamental values must align. Next, using this information from put option estimations, the fundamental value of call options can be computed and compared to market observations for bubble detection. Under the ND condition, call option price bubbles signal towards exuberance in the underlying, with the magnitude of the former acting as the lower bound to that of the latter. There are two benefits from fitting the model to put option prices, instead of those associated with the underlying. First, the option market, comprises of forward-looking information on expectations of market participants. Second, underlying and call option prices are capable of revealing bubbles, therefore, directly fitting a model to their market observations could provide inconclusive results.

### 2.3.2.2 European Option Bubbles

Options have bounded lives and therefore are capable of revealing only type III bubbles, when the market is incomplete, and satisfying NFLVR and ND conditions. Put options cannot display bubbles as their payoffs are bounded from above, by strike prices. In contrast, call options have unbounded payoffs, and are capable of reveal bubbles. Consider European-style call and put options, with strike price,  $\mathbb{K}$ , and maturity  $\tau$ , at  $\tau < T$ . Let the payoffs at  $\tau$ , be represented by  $C_\tau(\mathbb{K}, \tau) = \max\{S_\tau - \mathbb{K}, 0\}$  and  $P_\tau(\mathbb{K}, \tau) = \max\{\mathbb{K} - S_\tau, 0\}$ , along

with prices of the respective options,  $C_t(\mathbb{K}, \tau)$ , and  $P_t(\mathbb{K}, \tau)$ , assuming that the risky asset has no cash flows<sup>18</sup>.

$$P_t(\mathbb{K}, \tau) = E_{\mathbb{Q}} \left( \frac{\max\{\mathbb{K} - S_{\tau}, 0\}}{B_{\tau}} \middle| \mathcal{F}_t \right) B_t \quad (5)$$

$$C_t(\mathbb{K}, \tau) = E_{\mathbb{Q}} \left( \frac{\max\{S_{\tau} - \mathbb{K}, 0\}}{B_{\tau}} \middle| \mathcal{F}_t \right) B_t \quad (6)$$

Prior to exploring bubbles in put and call options, one must examine the instance of zero-coupon bonds, with fundamental value,  $p^*(t, \tau) = E_t \left[ \frac{1}{B_{\tau}} \right] B_t$ . Given nonnegative interest rates, the market value  $p(t, \tau) \leq 1$ , for all  $0 \leq t \leq \tau \leq T$ , is bounded from above. Under NFLVR,  $p(t, \tau)/B_{\tau}$ , is a local  $\mathbb{Q}$ -martingale. Given  $B_{\tau} \geq 0$ ,  $p(t, \tau) \leq 1$ ,  $p(t, \tau)/B_{\tau}$  is bounded, and hence a martingale. Bubbles exist when rational traders have the incentive to purchase an asset, for the sole purpose of reselling at a greater price in the future. This incentive is lost for assets with bounded prices, as through backward induction, traders will realise that it cannot be purchased only for reselling at a higher price, resulting in  $p^*(t, \tau) = p(t, \tau)$  (Jarrow et al., 2010; Jarrow, 2015; Fusari et al., 2024). Similarly, for  $p(t, \tau) \leq 1$ , put options have payoffs that are bounded from above, and hence, their prices do not exhibit bubbles. Put option prices are uniformly integrable  $\mathbb{Q}$ -martingales, with fundamental value,  $P_t^*(\mathbb{K}, \tau)$  equivalent to market price,  $P_t(\mathbb{K}, \tau)$ . Importantly, when the underlying asset displays a bubble, put options can be priced using standard option pricing methodologies (Jarrow et al., 2010; Protter, 2013; Jarrow, 2015). Contrastingly, call options have unbounded payoffs, and can exhibit price bubbles.

Consider  $V_t^f$  as the market price of a forward contract, with maturity,  $\tau$  and payoff,  $\{S_{\tau} - \mathbb{K}\}$ , at  $t$ . Defining  $V_t^{f*}$ ,  $P_t^*(\mathbb{K}, \tau)$  and  $C_t^*(\mathbb{K}, \tau)$  to be the fundamental values of the forward contract, put option and call option respectively, it can be shown that the put-call parity holds;  $V_t^{f*} = C_t^*(\mathbb{K}, \tau) - P_t^*(\mathbb{K}, \tau)$ . However, under the presence of bubbles, this need not be necessarily true for market prices. The ND condition by definition allows for the put-call parity to hold, even in the presence of bubbles. Therefore, in an incomplete market,

---

<sup>18</sup> This assumption has been placed for simplicity, and consistency with previous literature. It will be relaxed during the empirical testing of bubbles.



satisfying NFLVR and ND, call options inherit the underlying risky asset's bubble. At maturity,  $\tau$ , for  $t \leq \tau \leq T$ ,

$$S_\tau - K = \max\{S_\tau - K, 0\} - \max\{K - S_\tau, 0\} \quad (7)$$

taking conditional expectations under  $\mathbb{Q}$ , where  $p(t, \tau) = E_{\mathbb{Q}}\left(\frac{1}{B_\tau} \middle| \mathcal{F}_t\right) B_t$ , given  $r_t \geq 0$ :

$$\begin{aligned} E_{\mathbb{Q}}\left(\frac{S_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t - E_{\mathbb{Q}}\left(\frac{K}{B_\tau} \middle| \mathcal{F}_t\right) B_t \\ = E_{\mathbb{Q}}\left(\frac{\max\{S_\tau - K, 0\}}{B_\tau} \middle| \mathcal{F}_t\right) B_t - E_{\mathbb{Q}}\left(\frac{\max\{K - S_\tau, 0\}}{B_\tau} \middle| \mathcal{F}_t\right) B_t \end{aligned}$$

$$E_{\mathbb{Q}}\left(\frac{S_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t = E_{\mathbb{Q}}\left(\frac{\max\{S_\tau - K, 0\}}{B_\tau} \middle| \mathcal{F}_t\right) B_t - P_t(K, \tau) + p(t, \tau)K$$

from condition iv) from 1.3.1.6,  $S_t = E_{\mathbb{Q}}\left(\frac{S_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t + \beta_t - E_{\mathbb{Q}}\left(\frac{\beta_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t$ , hence

$$\begin{aligned} S_t - \beta_t + E_{\mathbb{Q}}\left(\frac{\beta_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t &= E_{\mathbb{Q}}\left(\frac{\max\{S_\tau - K, 0\}}{B_\tau} \middle| \mathcal{F}_t\right) B_t - P_t(K, \tau) + p(t, \tau)K \\ S_t - p(t, \tau)K + P_t(K, \tau) &= E_{\mathbb{Q}}\left(\frac{\max\{S_\tau - K, 0\}}{B_\tau} \middle| \mathcal{F}_t\right) B_t + \beta_t - E_{\mathbb{Q}}\left(\frac{\beta_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t \end{aligned}$$

using the definition of the put-call parity,  $S_t - p(t, \tau)K + P_t(K, \tau) = C_t(K, \tau)$ :

$$C_t(K, \tau) = E_{\mathbb{Q}}\left(\frac{\max\{S_\tau - K, 0\}}{B_\tau} \middle| \mathcal{F}_t\right) B_t + \beta_t - E_{\mathbb{Q}}\left(\frac{\beta_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t \quad (8)$$

$$C_t(K, \tau) = C_t^*(K, \tau) + \beta_t - E_{\mathbb{Q}}\left(\frac{\beta_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t \quad (9)$$

$$B_t(K, \tau) = \beta_t - E_{\mathbb{Q}}\left(\frac{\beta_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t = [C_t(K, \tau) - C_t^*(K, \tau)] \times N \quad (10)$$

It can be witnessed in (9) that European call options depict bubbles, when its market price,  $C_t(K, \tau)$  exceeds its fundamental option value,  $C_t^*(K, \tau)$ , implying  $\delta_t(\tau) \geq 0$ , where  $\delta_t(\tau) = \beta_t - E_{\mathbb{Q}}\left(\frac{\beta_\tau}{B_\tau} \middle| \mathcal{F}_t\right) B_t$ , is a supermartingale with  $\delta_\tau(\tau) = 0$ . Under the NFLVR

and ND conditions, bubbles in the underlying and call options are linearly related<sup>19</sup>. The bubble magnitude,  $\mathcal{B}_t(\mathbb{K}, \tau) = [C_t(\mathbb{K}, \tau) - C_t^*(\mathbb{K}, \tau)] \times \mathbb{N}$  ( $\mathbb{N}$  resembles the lot size of the option), is independent of  $\mathbb{K}$ , and vanishes as  $\tau$  approaches 0, with  $E_{\mathbb{Q}}\left(\frac{\beta_{\tau}}{B_{\tau}} \middle| \mathcal{F}_t\right) B_{\tau} = \beta_{\tau}$ . Hence, the bubble can be estimated by employing an appropriate pricing model to compute  $C_t^*(\mathbb{K}, \tau)$ , given  $C_t(\mathbb{K}, \tau)$  observed in the market. However, in the presence of bubbles, standard option pricing techniques are not valid for computing the call option price (Cox & Hobson, 2005; Heston et al., 2007; Jarrow et al., 2007, 2010).

### 2.3.2.3 American Option Bubbles

American-styled option contracts, in contrast to their European counterparts, allow for early exercise, any time prior to maturity. Given this possibility, they do not exhibit bubbles, even if witnessed in the underlying price. Considering  $S_t$  as the price of the underlying risky asset denoted in units of the MMA and stopping time  $\zeta$ , the fundamental value of an American call can be defined by  $C_t^{A*}(\mathbb{K}, \tau)$ .

$$C_t^{A*}(\mathbb{K}, \tau) = \sup_{\zeta \in [t, \tau]} E_{\mathbb{Q}}\left(\max\left\{S_{\zeta} - \frac{\mathbb{K}}{B_{\zeta}}, 0\right\} \middle| \mathcal{F}_t\right) \quad (11)$$

The supremum would incorporate price bubbles, with  $\zeta$  being stopped early, with a strictly positive probability (Protter, 2013). Assume the existence of a bubble, and the ability to purchase a European call option. It would be advantageous to enter into a short maturity contract, and enjoy rewards earned from the explosive price hike. However, when entering into a longer contract, if the price collapses, the investor is subjected to plummeting payoffs, and a subsequent drop in option prices. It would be beneficial to have the choice of an early exercise. Therefore, American call options, will be exercised early in the presence of an underlying bubble, which adds value to the stopping time and creates a difference between the price of American and European call options, as revealed in (12). If absent, then the stopping time would have no value, resulting in  $C_t^A(\mathbb{K}, \tau) = C_t^{A*}(\mathbb{K}, \tau) = C_t^E(\mathbb{K}, \tau) =$

---

<sup>19</sup> If the ND condition, is violated, the linear relation between call option and underlying price bubbles does not hold, with the possibility of each displaying separate bubbles. Furthermore, suicide strategies would exist, allowing for bubbles in put options, and riskless assets.

$C_t^{E*}(\mathbb{K}, \tau)$ , with  $C_t^A(\mathbb{K}, \tau)$  and  $C_t^E(\mathbb{K}, \tau)$  being the respective market prices of American and equivalent European call options.

$$C_t^{A*}(\mathbb{K}, \tau) - C_t^{E*}(\mathbb{K}, \tau) = \beta_t - E_{\mathbb{Q}}\left(\frac{\beta_{\tau}}{B_{\tau}} \middle| \mathcal{F}_t\right) B_t > 0 \quad (12)$$

$$C_t^A(\mathbb{K}, \tau) = C_t^{A*}(\mathbb{K}, \tau) = C_t^{E*}(\mathbb{K}, \tau) + \beta_t - E_{\mathbb{Q}}\left(\frac{\beta_{\tau}}{B_{\tau}} \middle| \mathcal{F}_t\right) B_t \quad (13)$$

$$\mathcal{B}_t(\mathbb{K}, \tau) = \beta_t - E_{\mathbb{Q}}\left(\frac{\beta_{\tau}}{B_{\tau}} \middle| \mathcal{F}_t\right) B_t = [C_t^A(\mathbb{K}, \tau) - C_t^{E*}(\mathbb{K}, \tau)] \times \mathbb{N} \quad (14)$$

American call options, as opposed to their European correspondents are unable to reveal bubbles, hence  $C_t^A(\mathbb{K}, \tau) = C_t^{A*}(\mathbb{K}, \tau)$ . The early exercise value reflects the bubble component, which along with the  $C_t^{E*}(\mathbb{K}, \tau)$  make up the constituents of  $C_t^A(\mathbb{K}, \tau)$ . The bubble component is a nonnegative supermartingale that is expected to decline. It has a similar impact on the price of the underlying stock, and the subsequent decision of an early exercise, as a continuous dividend payout<sup>20</sup>. Therefore, in accordance with (14) for American call options, the fundamental value of a corresponding European call can be utilised detect bubbles in the underlying price (Jarrow et al., 2007, 2010). Furthermore, it must be highlighted that American put options, similar to their European counterparts, due to bounded payoffs, do not display bubbles.

### 2.3.2.4 Empirical Testing of Bubbles with Option Prices

A bubble is observed when the fundamental value of an asset is exceeded by its market price. Under NFLVR and ND conditions, if a bubble exists in the underlying, there is no ELMM, such that  $C_t(\mathbb{K}, \tau) = C_t^*(\mathbb{K}, \tau)$ . Hence, standard risk-neutral measures of pricing may not be favorable (Jarrow et al., 2010). Computing the fundamental value of call options for bubble detection can result in inconclusive results, similar to traditional methods. A joint hypothesis issue would arise from attempts to model asset prices that display bubbles.

For  $S_u = 0$ , Jarrow (2015) highlight this problem using the null hypothesis in (15), which reveals the nonexistence of type II bubbles. Market prices are observable, whereas the

---

<sup>20</sup> If the underlying does not pay a dividend, the value and incentive of an early exercise would be non-existent (Merton, 1973).

fundamental value must be determined by a particular model for  $\{r, D, \mathbb{Q}\}$ . This gives rise to the joint hypothesis issue, given difficulty in testing for fundamental values, independent from evaluating bubbles. Similarly, despite migration to a continuous-time framework, for the null hypothesis in (16) the joint hypothesis problem persists, when a stochastic model for  $\{r, D, \mathbb{S}_u, u, \mathbb{Q}\}$  is employed. However, two other approaches for bubble detection within this framework that overcome the joint-hypothesis issue, exist (Jarrow et al., 2010; Jarrow, 2015; Jarrow, 2016).

$$\beta_t = \mathbb{S}_t - \sum_{v=t}^{\infty} E_{\mathbb{Q}}\left(\frac{\Delta D_v}{B_v} \middle| \mathcal{F}_t\right) B_t = 0 \quad (15)$$

$$\beta_t = \mathbb{S}_t - E_{\mathbb{Q}}\left(\frac{\mathbb{S}_u}{B_u} + \int_0^u \frac{1}{B_s} dD_s \middle| \mathcal{F}_t\right) B_t = 0 \quad (16)$$

First, a stochastic process for  $\mathbb{S}_t$  can be defined by segregating parameters into two mutually exclusive and exhaustive sets, that respectively do and do not exhibit bubbles. Given that it is significantly less challenging to independently test for the hypothesis of  $\mathbb{S}$  over a historic time series, in comparison to that of  $\{r, D, \mathbb{S}_u, u, \mathbb{Q}\}^{21}$ , the problem can be overcome as witnessed in Chaim & Laurini (2019) and Laurini & Chaim (2021). However, both works fail to capture forward looking expectations, given the practice of fitting to underlying prices. Alternatively, option prices can be investigated to test the presence of bubbles in the underlying.

Recall, call options, in contrast to their put counterparts, have unbounded payoffs, and can reveal price bubbles. Furthermore, given the ND condition, bubbles in the option price are linearly related to those in the underlying. Such an approach would require the estimation of  $\{D_t, \mathbb{S}_t, P_t, C_t\}$ . Since put options do not exhibit price bubbles, they can be priced using standard option-pricing methods, with precision to market quotes, used to validate models. The validated model would automatically specify  $\{D_t, \mathbb{S}_t, \mathbb{Q}\}$ , while computing  $P_t^{E*}$ , such that  $C_t^{E*}$  can be priced and compared with either  $C_t^E$  or  $C_t^A$  to test for bubbles in the call option and underlying prices, in accordance with (10) or (14), respectively. This approach provides for a stronger, and importantly, a real-time simultaneous validation for identifying

---

<sup>21</sup>  $\{\mathbb{S}_u, u, \mathbb{Q}\}$  are almost not observable for all asset price processes.

bubbles. Given the forward-looking characteristics of bubbles, options of similar nature, are best equipped to test for their existence (Fusari et al., 2024).

### 2.3.2.5 Empirical Application

The key for detecting bubbles from option prices is selecting a stochastic process that reveals strict local martingale tendencies in the underlying price<sup>22</sup>. Sophisticated stochastic volatility processes, revealing strict local martingale tendencies (see Sin, 1998; Andersen & Piterbarg, 2007; Lions & Musiela, 2007) can be calibrated to daily market observations. Protter (2016) make the case for the inclusion of price jumps, especially when considering illiquid instruments, as the price process of such assets would reveal changes only in the form of jumps. The operation of fitting model parameters of such stochastic processes to a time series of historical asset prices makes detection reliant on backward-looking information. This has adverse impacts on the quality of real-time identification.

The Stochastic Alpha Beta Rho (SABR) model is a popular choice amongst practitioners for robust estimation of implied volatility surfaces. It has a closed form solution and can successfully capture the forward variance curve. Furthermore, by fixing the lognormality of forward prices, the correlation parameter is adept of revealing strict local martingale tendencies in the underlying price process. Piironen et al. (2018) takes advantage of this property and introduce a martingale defect indicator<sup>23</sup> for bubble detection. The indicator is computed from calibrating SABR parameters from market volatility smiles by acknowledging risk preferences of market participants. Stahl & Blauth (2024) adopted the SABR model and martingale defect indicator, to conduct a two-fold analysis for bubble detection. First, the model was calibrated to 1-month volatility smiles, followed by evaluation of the entire surface. The monitoring of surfaces for martingale defects allows investors to adjust their positions by identifying irrationally priced stocks. Furthermore,

---

<sup>22</sup> The reader is referred to Sin (1998), Andersen and Piterbarg (2007) and Lions and Musiela (2007) for more information on these processes.

<sup>23</sup> The indicator signals for bubbles when the correlation parameter is positive, implying that stock prices and volatility are positively correlated. In other words, options become more expensive with increasing underlying prices, revealing optimistic risk-seeking preferences of market participants.

regulators are better equipped for improving market monitoring and protecting retail investors. Calibrating to a full surface generates more frequent bubble instances, with greater persistence. The larger number of datapoints reduces the likelihood of calibration errors and prevents breaks in the continuity of the event. However, it also smoothens anomalies within the lower maturity smile components of the surface, resulting in greater false negative signals.

Call options across different strike prices and maturities are capable of revealing full information about the underlying price distribution. Biagini et al. (2024) propose a model independent deep learning methodology for estimating the probability of a bubble, from examining all available daily call options prices. The network is trained to classify dynamics of the underlying as either strict local martingale, or true martingale. It was successful in learning the Sin (1998), displaced-CEV, and SABR processes, individually and collectively. Although, neural networks can be trained to learn more sophisticated stochastic process, such as those in Andersen & Piterbarg (2007) and Fusari et al. (2024) for bubble detection. The feed-forward nature of the approach makes it vulnerable to the joint-hypothesis issue. Parametric calibration of the selected process, as witnessed in Piiroinen et al. (2018) and Stahl & Blauth (2024), assist with this problem. However, both consider out-of-the-money (OTM) call options during calibrations, which results in the same problem.

In accordance with Jarrow (2015), the joint-hypothesis issue is overcome by calibrating to market put options, as their prices do not reveal bubbles. The fundamental value of put options is equivalent to market price. Hence, such calibrations reveal the fullest extent to which the model captures underlying dynamics, without suffering from the joint-hypothesis issue. The three-step approach from Fusari et al. (2024) is designed to specifically cater towards this requirement. It uses a stochastic volatility jump diffusion model, capable of capturing strict local martingale tendencies in the underlying process. Parameters of the model are calibrated to put options, prior to pricing call contracts. The final step involves a statistical test designed to identify significant call option bubbles. Based on the linear relationship between call and underlying price bubbles, the statistical test reduces the reliance on observing bubbles through the parametric conditions of the stochastic model.

Alternatively, Jarrow & Kwok (2021) propose a non-parametric approach to infer bubbles in the underlying from options data, without imposing the ND condition. The method involves constructing bounds for the bubble magnitude, based on the strike range and maturity of daily traded options. The bounds are expressed as weighted linear combination of put and call options. However, the bubble is only identified up to an interval, as the state price distribution is not entirely revealed beyond the available strike range. It is important to consider the less frequently traded extreme moneyness options, as they comprise the most information on the tails of the distribution. A tail-truncated state price distribution is introduced to compute fundamental values. The fundamental value comprises of a truncation bias, which vanishes with the expansion of the strike intervals. Jarrow & Kwok (2024) implement the approach for differing between quadratic variation or autoregressive related explosiveness during bubbles, within the S&P 500 index, and several selected stocks. As opposed to the non-parametric method, the three-step approach (Fusari et al., 2024) is built on the foundations of the NFLVR and ND assumptions, ensuring that the put-call parity is not violated. Furthermore, it does not suffer from the truncation bias and also allows for an independent real-time validation of the put option's fundamental value.

The three-step approach (Fusari et al., 2024) overcomes three major issues with respect to bubble detection. First, traditional methods focus on type II bubbles and require a large sample, creating a trade-off between model asymptotic and possible structural breaks. Next, such methods lack consensus for estimating the fundamental value, giving rise to a joint-hypothesis issue. Even under the local martingale theory, those from Piiroinen et al. (2018), Stahl & Blauth (2024), and Biagini et al. (2024) are vulnerable. Methods relying on estimating finiteness of the volatility integral, are immune to this issue, but struggle with timestamping bubbles. Third, the forward-looking nature of price bubbles are captured by options, in contrast to the backward-looking information from spot markets. All three issues are overcome by calibrating a stochastic volatility jump diffusion process to daily market observations of put options. Therefore, the three-step approach has been selected for bubble detection in this study. However, the stochastic process relies on Monte Carlo simulations, making daily calibrations a computationally cumbersome task. Therefore, Fusari et al. (2024) compromised by using on the daily most liquid put option smile, to strike a balance

between robustness and efficiency. Nevertheless, considering the entire surface is important, as seen in Stahl & Blauth (2024), and Biagini et al. (2024).

### 2.3.3 Option Price Surface for Bubble Detection

The three-step approach has been established as the method of choice for bubble detection in this study, due to its superior ability in overcoming the joint-hypothesis issue, whilst capturing the forward-looking nature of bubbles. However, daily parameter calibrations of the desired stochastic process, are computationally cumbersome, forcing Fusari et al. (2024) to strike a balance between robustness and efficiency by considering only the daily most liquid volatility smiles. Even though this is common industrial practice, Stahl & Blauth (2024) and Biagini et al. (2024), highlight the benefits of considering the entire surface during bubble detection.

Selection of instruments depends heavily on the type of information required by the practitioner. When attempting to capture future expectations regarding economic factors, the forward-looking nature of options, make them an ideal candidate (Funahashi, 2023). This further reveals the benefits of examining option prices for identifying bubbles in the underlying. An implied volatility surface is constructed from all traded options on a given day and plays a crucial role in obtaining such information. Its shape and level are conventionally stable, however, large movements in either reveal important changes in market conditions (Ackerer et al., 2020). The surfaces are heavily impacted by option demand and supply, which are alter in abidance with market expectations. Therefore, calibrating pricing models to the surface is detrimental for extracting forward-looking information (Funahashi, 2023). Hence, implied volatility surfaces play a key role in observing exuberance, and investment and risk-management related decision-making.

The surfaces are subjected to stochastic alterations over time, and a non-flat instantaneous profile, which is depicted by the presence of a ‘smile’, and ‘skew’ (Cont et al., 2002). They reveal individual components to comprise of high positive autocorrelation and mean-reverting behaviour, with the latter occurring more frequently over shorter maturities. Given dependence on market demand and supply, and the stochastic nature of surfaces, volatilities



could possibly exhibit jumps and/or heavy-tail increments. A negative relationship exists between implied volatilities and the returns of the underlying asset, with respect to the strike price (Romo, 2014). This is referred to as a volatility skew, and if the OTM put and call options have greater implied volatilities than the ATM counterpart, a ‘smile’ occurs. The shape is important for understanding underlying returns (Ackerer et al., 2020). If the smile, is U-shaped, then the tail of underlying returns distribution is thicker. Alternatively, a skew indicates to a distribution being thicker on one-side. If that side is the left, underlying prices are more likely to experience larger losses than gains. Recall, Jarrow et al. (2011) emphasised on the extrapolation of the tail, in attempts to capture the presence of such important information.

The negative relationship between returns from the underlying, and implied volatilities are well documented in literature. Cont & da Fonseca (2002) display the two to be negatively correlated, whereas Poulsen et al. (2009) show a stronger negative correlation between ATM implied volatilities and returns. Cao et al. (2020) provide evidence supporting volatility feedback, where increased volatility has a positive effect on the required rate of return causing prices to fall. Furthermore, change in equity return does not cause all components of the implied volatility surface to alter with the same magnitude. In relevance to bubble detection, each smile of the surface corresponds to a different maturity, hence participants with different investor horizons have varying perceptions about current equity returns. This hints towards the potential variety in the formation of bubbles, with respect to each individual smile. Implied volatilities and underlying returns are not perfectly correlated (Cont et al., 2002). Shifts in global implied volatilities are negatively correlated with underlying returns, whereas relative movements have a small correlation.

An abundance of information, regarding the underlying asset, and future expectations on market and economic conditions, can be extracted by calibrating to the entire surface. The smiles and skewness displayed by implied volatilities over varying maturity, can be explained by stochastic volatility models (Funahashi, 2023). The parameters of such models reveal large amounts of information regarding market expectations and option metrics such as; risk-neutral variance, skewness, and variance risk premium, when calibrated from implied volatility surfaces (Ulrich & Walther, 2020). Despite the effectiveness of stochastic

volatility models, Hoshisashi et al. (2023) state that estimation of volatility is complex due to limited input data, low liquidity, and the prevalence of noise. Furthermore, the most liquid option cross-section does not provide enough information on the underlying asset for the calibration of local volatility, jump-diffusion, and stochastic volatility (with/without jumps) models. In such instances even if the models are calibrated, they cannot guarantee realistic future scenarios (Cont et al., 2002). Calibration from implied volatility surfaces allows to capture information on predictive signals and expectations regarding the underlying asset (Guidolin & Wang, 2023). However, such information can differ substantially across the surface (Ulrich & Walther, 2020), over the various levels of maturity. Therefore, the entirety of the information obtained from the surface is only partially acknowledged when utilising options with the highest cumulative volume.

It is precisely due to this reason that this study emphasises on applying the three-step approach to option contracts over several maturities. The three-step approach employs a stochastic volatility jump diffusion model, with the ability to capture strict local martingale tendencies in the underlying price. It is a highly sophisticated process, engineered to extract information regarding the underlying dynamics, to a maximum extent. Its high-dimensional parametric vector excels at doing so but reliance on Monte Carlo simulations, given a lack of a closed form solution, creates a massive computational burden. Hence, Fusari et al. (2024) had opted for considering the most liquid option cross-section. However, provided preference of capturing forward-looking expectations of market participants from the entire available surface, a necessity to tackle the issue of computational efficiency, without sacrificing robustness, is created. Therefore, this study seeks the assistance of artificial neural networks.

## 2.4 Neural Networks: Option Pricing and Calibration

The three-step approach (Fusari et al., 2024) exploits the forward looking nature of option markets, and overcomes the joint-hypothesis issue during the real-time bubble detection. It implements a stochastic volatility jump diffusion model for pricing options. However, there exists a major bottleneck in its practical application. Given the unavailability of an analytical solution, Monte Carlo simulations are employed as a numerical solver for pricing. Such

computational heavy inefficiencies would force practitioners to overlook the obvious qualities of the three-step approach.

As an alternative, one could consider fitting a stochastic volatility model with fewer parameters, as in Andersen & Piterbarg (2007). However, once again, Monte Carlo methods must be relied upon, due to an unavailability of closed form solution. Additionally, the process is unable to capture the possible instances of jumps in the underlying price. Other stochastic volatility models with or without jumps, that possess analytical solutions, such as those in Merton (1976), Heston, (1993), and (Bates 1996), are unable to detect strict local martingale properties in the underlying price process. Alternatively, the SABR reveals this trait (Piiroinen et al., 2018; Stahl & Blauth, 2024), but fails to address the mean-reverting nature of volatility, making it only suitable for pricing short-term options (Gatheral 2012). Solely from a quality-based stance, the process from Fusari et al. (2024) is superior, however given the computational burden during calibrations, practitioners would be reluctant towards its application. Generally, common industrial practice prefers tractability over quality and accuracy, hence, despite its flaws, the Black & Scholes (1973) model is remains popular (Horvath et al., 2021). Given the computational inefficiency of highly sophisticated models, it is common to settle for striking a balance between robustness and efficiency. It is with the motivation of preventing practitioners from being forced to sacrifice quality over tractability, that this research turns to artificial neural networks, for efficient real time bubble detection using option prices.

### 2.4.1 Artificial Neural Networks

Artificial Neural Networks (ANNs) are sets of interconnected nodes, that determine given output(s) for a set of input(s) by mathematically imitating the functions of the human nervous system. The nodes, also known as neurons, like their biological namesake, are the building blocks of the system. They are grouped together, in stacked layers, to allow for the flow of information, such that complex relationships between inputs and outputs can be learned. The analogy from Culkin & Das (2017) of a '*a child touching a hot plate, and learning not to go near one*' can be resorted to for explaining the basic structural mechanisms of neural networks, via understanding those of the nervous system. Decisions in the nervous system

are made by considering a wide range of stimuli and passing them through layers of neurons. Narrowing down to the stimuli of pain from touching a hot plate; upon contact, information is passed through neurons to the brain and, back such that the child instantly pulls away. Additionally, that information is stored, and the child learns not to repeat the action.

The extension of this mechanism for decision making by learning from financial data, is the foundation to the contributions within this research. McCulloch & Pitts (1943) were the first to reveal the universal computational capabilities of ANNs formed from simple mathematical representations of a neuron. The Universal Approximation theorem (Cybenko, 1989; Hornik et al., 1989; Hornik, 1991) confirmed this by proving neural networks with a single hidden layer and sigmoid activation can approximate any function, within a given level of accuracy. For ANNs the wide range of stimuli that need to be understood, exist in the form of datasets with several input and output variables. The relationship between these variables is learned by passing information through the layer(s) of neurons and stored as weights and biases associated with each node. Given sufficient data, the network can be trained to deliver successful outputs, repeatedly (Culkin & Das, 2017).

Nearly a boundless number of calculations are required at high frequent intervals in the financial derivatives market, giving rise to a desperate need for speed (Culkin & Das, 2017; De Spiegeleer et al., 2018). There exists a continuous demand for calculating and updating model parameters, hedge positions, and instrument prices for risk management and profitability purposes. These metrics are relevant only over a short period of time due to the fluctuating nature of market dynamics. Given the availability of big datasets, and superior computational power, practitioners can take advantage of deep learning networks, for improved speed and pattern recognition (Culkin & Das, 2017). The latter stems from neural networks, unlike traditional econometric models being able to capture nonlinearities within inputs and outputs. In line with the universal approximation theorem, it is possible for neural networks, to replicate and at times, better the actions of analysts, risk-mangers, and traders.

In the context of the current research, options are the most crucial derivative for bubble detection, and subsequently for risk management. Since the emergence of the Black & Scholes (1973) model, the popularity of option contracts amongst traders has significantly

grown. The drawbacks of the model are well documented given its assumptions of constant volatility and lognormality of the underlying price. Stochastic volatility models (with/without jumps) overcome these issues, however, are computationally cumbersome, given reliance on iterative optimization techniques for parametric calibrations. It was towards the end of the 20<sup>th</sup> century, during which the contributions of Malliaris & Salchenberger (1993, 1996) and Hutchinson et al. (1994), pioneered the path of using ANNs for estimating options prices/implied volatilities. The authors reveal the pricing effectiveness of a simple single layer feedforward networks in modelling nonlinear dependencies. Since then, given the level of computational innovations experienced over time, several scholars have experimented with complex network architectures to bolster option pricing<sup>24</sup>.

The stochastic nature of asset pricing requires nonlinear and multivariate functions, which makes option pricing a complicated task (Ivaşcu, 2021). Hutchinson et al. (1994) highlight three characteristics possessed by neural networks, in contrast to traditional parametric approaches, that make them advantageous for option pricing. First, they do not rely on parametric restrictions or assumption but rather determine the relationship between inputs and outputs solely from data. Second, they are adaptive and responsive to structural change, as a consequence of directly learning from data. Finally, neural networks are simple to implement, and capable of entailing a wide range of derivatives and asset price dynamics. This final trait is vital for obtaining an estimation speed that is greater by orders of magnitude. De Spiegeleer et al. (2018) paint a vivid picture for the relevance of the previous statement with regards to options. The calculations for pricing such contracts on a given day, are very similar. A synopsis of such instruments can be revealed by a few parameters such as strike price, and maturity. The same applies to the market, which can be represented by interest rates, or even the parameters of a stochastic volatility model. This is ideal for machine learning tasks as the low dimensional input parameters can be summarized for the model to learn the actual function for pricing even thousands of instruments in seconds, with a strong level of accuracy.

---

<sup>24</sup> Refer to Ruf & Wang (2020), for a comprehensive overview of this literature.

It is these remarkable advancements and capabilities, that inspired the exploitation of neural networks for improving the efficiency of the three-step approach. The model calibration to daily put options is a tedious task. Furthermore, given the frequent nature of market fluctuations, calibrated parameters rapidly become obsolete, demonstrating the need for quick and repetitive estimations, to capture frequently varying market features. Monte Carlo simulations are unfit to meet such demands and reduce favourability of the approach. Neural networks, when employed as numerical solvers, are capable to calibrating stochastic volatility models, in seconds, with a high degree of accuracy (Liu et al., 2019; Horvath et al., 2021). Therefore, a deep calibration framework for the stochastic volatility jump diffusion model, is explored to improve the efficiency of the three-step approach for real-time bubble detection.

## 2.4.2 Calibrating Stochastic Volatility Models with Neural Networks

An option pricing model is calibrated when the value of set of benchmark instruments, correspond to their market counterparts (Cont, 2010). In the absence of analytical solutions, this procedure is heavily reliant on iterative computationally expensive optimisation techniques. In general, calibration of asset pricing models to option prices, across different levels of moneyness and maturities, is crucial for the approximating the risk neutral measure (Dimitroff et al., 2018). It is vital operation for maintaining the value of trading books, in correspondence to updated market prices. Innumerable calculations are required for determining the risk neutral measure at very frequent intervals for real-time risk management (portfolio simulations, stress tests, and hedge position adjustments), pricing of exotic derivatives (De Spiegeleer et al., 2018; Dimitroff et al., 2018; Liu et al., 2019), or real-time bubble detection.

Despite supreme quality of stochastic volatility models, the computational burden associated with implementation, makes practitioners favour the inferior Black & Scholes (1973) model on the grounds of tractability (Bayer & Stemper, 2018; Horvath et al., 2021). An asset pricing model is selected by balancing efficiency and accuracy (Büchel et al., 2022). The efficient runtime of stochastic volatility models is hampered by the number of parameters needed to be calibrated, along with reliance on Monte Carlo simulations. Hernandez (2016)

highlight that from a computational point of view, such models can be accepted, by making certain compromises. As with the three-step approach, poor tractability restricted Fusari et al. (2024) to most liquid options. This is consistent with common industry practices, however vulnerable to not capturing all relevant information regarding forward looking expectations of market participants.

It is difficult to attain analytical solutions for multi-dimensional high quality stochastic volatility models (Liu et al., 2019), which are forced to rely on Monte Carlo simulations. A large body of literature is dedicated towards improving this facet, most notably by implementing Fourier transformations (Carr & Madan, 1999; Fang & Oosterlee, 2009). Significant success has been witnessed in attempts to boost efficiency. However, the sheer volume of the task makes it computationally inefficient to use such models, even if they possess analytical solutions. There is need to update calculations at high-frequencies due to their tendencies of becoming obsolete due to changing market dynamics (Bayer et al., 2019; Karatas et al., 2019). These calculations need to be conducted thousands of times, for thousands of derivatives, which are written on thousands of underlying assets, at high frequency of intervals (De Spiegeleer et al., 2018). Therefore, the key to computational feasibility of stochastic volatility models, relies on the ability of a numerical solver to compute values of a large range of derivatives, simultaneously. This can be achieved by implementing neural networks as numerical solvers during option pricing operations.

Neural networks are successful at extracting features and detecting patterns from large datasets, given their strong ability to approximate non-linear functions. Hence, they can be implemented as universal function approximators, without assuming a mathematical relationship between inputs and outputs (Liu et al., 2019). When applying neural networks as high-dimensional approximators, computational efficiency can be obtained with simultaneous estimation of the desired output, within seconds, across several moneyness and maturity levels (Liu et al., 2019; Horvath et al., 2021). The ease in parallel processing provides for superior computational efficiency, by orders of magnitude, over a wide range of inputs, in comparison to Monte Carlo simulations (Hernandez, 2016; Liu et al., 2019). This allows neural networks to meet the frequent calibration requirements in derivative markets and additionally ensure improved efficiency of the three-step approach.

Calibration with the assistance of neural networks is centred around taking advantage of the universal approximation theorem by replacing expensive Monte Carlo simulations with cheap pricing runs. The expedition of option price/implied volatility estimations allows for rapid parametric calibration. A fully connected neural network is capable of producing model parameters to map market implied volatilities (Bayer & Stemper, 2018). Once trained, the network can be optimized to provide minimal output losses at super-fast processing speeds (Karatas et al., 2019). For a given set of parameters, implied volatility surfaces can be directly generated, without defining the stochastic dynamics of the underlying asset (Horvath et al., 2021). Neural networks can improve calibration, regardless of the asset pricing model (Hernandez, 2016).

Neural networks are shown to map market data to the existing market regime (Dimitroff et al., 2018). Additionally, though the choice of architecture might vary, using the machine learning approach as numerical solvers during calibrations is robust to the choice of model (Hernandez, 2016; Horvath et al., 2021). A vital component in the application of neural networks for calibration is the training phase, during which dynamics of the selected asset pricing model are learned. It is a highly computationally cumbersome task, especially with regards to creating the training dataset (market/synthetic), and seeking an optimal network architecture. However, this phase is a one-off procedure conducted offline. Once trained, calibrations are performed online, within a matter of seconds, using options spanning over several moneyness and maturity levels.

There are two kinds of approaches considered across literature, for calibrating asset pricing models with neural networks (Bayer et al., 2019). First, the one-step approach, which trains the network on market inputs to obtain parameters as an output of a feedforward network. Second, the two-step approach, which implements the neural network as a numerical solver, to learn the mapping of parameters onto prices/implied volatilities. An optimizer is utilised to calibrate parameters from market quotes. Both approaches follow the ideology of a one-time computationally intensive training period, succeeded by rapid calibration from market data.



### 2.4.3 One-Step Calibration Approach

The one-step calibration approach is motivated by the ability of neural networks to be a universal functional approximator. It aims to estimate parameters of an asset price model, as outputs of a fully trained network. Hence, an additional optimisation step is made redundant, providing significantly improved computational efficiency, in comparison to traditional calibration methods.

Traditional parametric models lack analytical solutions, therefore a strong reliance on optimisation techniques is observed for solving high-dimensional calibration problems. This hurdle can be overcome at substantial boosted levels of computational efficiency, by exploiting the expertise of neural networks in determining mathematical relationships between input and output variables. A well-trained network develops the ability to accurately map inputs to parameters of the desired stochastic volatility model. Training is a computationally heavy procedure as it requires the creation of a sufficiently large training dataset to avoid overfitting, along with the determination of the optimal neural network architecture, via a search over a wide range of hyperparameters. On the bright side, training can be conducted off-line, and once completed, real-time calibrations are carried out simultaneously, at lightening speeds, for options across several maturities and moneyness levels, without requiring an additional optimisation step. The estimation of parameters is simplified to the basic task of matrix-multiplication within a trained network, that are feasible even on personal computers.

The concept of calibrating asset price models using neural networks was pioneered by Hernandez (2016), when considering the Hull and White model for GBP swaptions. A feedforward network, with input nodes corresponding to ATM volatilities and discretely sampled yield curve components, along with output nodes representing the respective parameters of the model, was constructed. The training dataset was constructed over historical data and hence required the Hull and White model to be calibrated to market quotes to obtain parameters. During testing on out-of-sample data, performance was strong only up to a period of 6 to 12 months beyond the training period. This is a direct consequence of training being conducted on historical data, as the network only learns market scenarios from

the selected period. Hence, robustness is adversely affected, requiring regular updates to the training dataset, and network, to learn the frequently fluctuating dynamics of the market. Therefore, the computationally intensive training procedure would no longer be a one-time ordeal.

It was hinted that convolutional neural networks could be beneficial when dealing with equity implied volatility surfaces (Hernandez, 2016). This was explored by Dimitroff et al. (2018) in attempts to calibrate the Heston stochastic volatility model. The motivation stemmed from the 2D shape of implied volatility surfaces, and corresponding inputs such as moneyness, strike prices, and forward values. When stacked into a 3D tensor, in resemblance with an image, each input represents a channel<sup>25</sup>. The network transforms input data into the most relevant features required for parametric estimation. Once trained, a single forward pass would produce fitted Heston parameters, at high levels of efficiency. Training data was constructed from daily calibration of the pricing model from historical stock data. Despite considering a large equity dataset for calibration, the proposed method suffers from overfitting, providing weak performances on out-of-sample (unseen) data. This signals for caution during forecasting, amongst real-time applications. The overfitting is caused by the poor quality of training data, highlighting the need to possess the characteristics representing various market scenarios.

A case for eliminating the additional calibration step was further continued by Itkin (2019), with the suggestion of inverse mapping for training neural networks to estimate the Black and Scholes model. First, a map of European call options prices to model inputs (including implied volatilities) was created. Next, variables were rearranged, to create an inverse map, for generating implied volatilities, with the call options prices as inputs. In this order, the network was trained, such that implied volatilities from market variables could be obtained. As witnessed with the earlier approaches, a high degree of computational efficiency was obtained. However, there are two concerns with regards to this approach, when considering stochastic volatility models. First, when assessing robustness, an additional validation step would be required. Second, and most importantly, the inverse mapping procedure was

---

<sup>25</sup> Continuing with the image analogy, each channel would represent colours along the red, green, and blue (RGB) spectrum.

conducted for calibrating a single parameter. It is not flexible enough to calibrate multiple parameters. This can be confirmed by Liu et al. (2021), who used inverse mapping for training neural networks to obtain American option implied volatilities. However, they prefer a two-step approach from Liu et al. (2019), for determining implied volatility and dividend yields.

In related work, Stone (2020) propose convolutional neural networks for calibrating the Hölder regularity of the rough Bergomi model from market volatility observations. Simulated sample paths of the asset pricing model were used for training the network, instead of calibrating parameters from historical data. It is faster to create a synthetic training set. Moreover, the practitioner, via choice of parametric combinations, can control the inclusion and exclusion of extreme market scenarios. Roeder & Dimitroff (2020) made direct evaluation with the two-step approach presented in Horvath et al. (2021), using similar synthetic training datasets. A feedforward network was trained to calibrate rough Bergomi and Heston model parameters from implied volatilities. It was found to be superior with regards to the former, but struggled during the latter, due to the presence of various parameter sets corresponding to extremely similar implied volatility surfaces<sup>26</sup>. Once again, in both cases, computational efficiency is successfully obtained, but at the cost of sacrificing the ability to perform real-time validation. Implied volatility surfaces are required to be generated separately to assess the accuracy of calibrations.

The one-step approach does provide substantial efficiency in comparison to traditional parametric calibration methods, by eliminating the optimisation step. After the one-time off-line, computationally expensive training of the network is complete, a simple forward pass delivers parametric estimates within seconds. However, the removal of the optimisation step creates more issues that it is intended to resolve. Bayer (2019) reveals that the ‘black box’ nature of the one-step approach does not sit well with regulatory bodies, given its application for risk management. The ambiguity of the network design leaves parametric calibration unexplained, as the approach concentrates on reducing estimation errors, rather than accurately computing observed market features, such as implied volatilities or options

---

<sup>26</sup> This issue was not faced by Horvath et al. (2021).

prices. The issue with generalization is a direct consequence of eliminating the optimisation step, forcing a separate action for validation. There is a lack of control during calibration, and hence, mapping to unseen market data cannot be entirely trusted.

Real-time validation is crucial for bubble detection, as it provides an instantaneous insight on overcoming the joint-hypothesis issue. Therefore, the one-step approach is not ideal for such applications. However, not all is lost, as the method reveals the superior ability of neural networks to be universal functional approximators, at high levels of efficiency. It is only a matter of control that is required to take advantage of such capabilities. The two-step approach adds an additional optimisation step, which provides control by ensuring real-time validation with respect to observed market features.

#### 2.4.4 Two-Step Calibration Approach

The three-step approach for bubble detection overcomes the joint-hypothesis issue by calibrating a stochastic volatility jump diffusion model to market quotes of put options. However, it faces a major bottleneck for industrial applications. The one-step method for using neural networks to calibrate parameters of stochastic processes, improves computational efficiency but does not provide for real-time validation. The desire for improved efficiency could force the practitioner to sacrifice accuracy, which would be catastrophic for hedging and risk management applications. Therefore, a strong need arises for controlled computational efficiency during calibration.

Conceptually, the one-step and two-step approaches are similar, with the neural network undergoing a one-time offline computationally cumbersome training procedure, followed by rapid online calibrations. The two-step approach trains the network with model parameters as inputs, such that dynamics of the asset pricing process are well learned to accurately estimate implied volatilities/option prices. Once trained, optimal weights and biases at each node are frozen and used during calibration. An optimiser is employed to obtain parameters by using the trained network as a numerical solver to minimise the loss function corresponding to market observations. Since the network is trained, the second step comprises of iterative, but rapid matrix multiplications between optimal weights and biases,

and parametric inputs. This approach is built on the foundation of neural network capabilities in accurately estimating option prices and implied volatilities. Furthermore, when trained on synthetic dataset, it can deliver strong performances even during extreme scenarios.

#### 2.4.4.1 Option Pricing with Neural Networks

Applying neural networks for solving problems related to option markets grew significantly over the past three decades. Malliaris & Salchenberger (1993, 1996) and Hutchinson et al. (1994), inspired by the universal approximation theorem used a single small hidden layer network. Malliaris & Salchenberger (1993) challenged the Black and Scholes model, using a network of similar inputs, in addition to lags of closed option and underlying prices. Implied volatilities were estimated without making assumption on the underlying model, distribution and correlation between variables, in Malliaris & Salchenberger (1996). In Hutchinson et al. (1994) neural networks were trained to learn Monte Carlo simulations of the Black and Scholes model.

Despite success of neural networks, Malliaris & Salchenberger (1993) and Hutchinson et al. (1994), highlighted certain flaws. The former reveal issues regarding computational power, and a missing theorem for developing optimal architectures, given the sensitivity of performances towards hyperparameter combinations. Over time, computationally efficiency has improved, allowing for testing hyperparameter search theorems, by experimenting with deeper (more hidden layers) and wider (more nodes) networks. In the latter, neural networks were found to be highly data-intensive, and vulnerable to shortage of historical data for training. Furthermore, it was revealed that if the underlying price model is well understood, it will be preferred over the neural network.

The data-intensive issue can be tackled by improving training with the consideration of deeper/wider approach. Additionally, it is common practice to utilise synthetic datasets, which can be constructed without an upper size limit. The second issue highlighted by Hutchinson et al. (1994) is relevant to the modern application of neural networks. Jang & Lee (2019) tested various machine learning models, including neural networks, against econometric jump models, across different option maturities. The latter were found to be

superior, given their ability to address volatility smiles, fat tails, and volatility clustering with jumps, which allow for accurately capturing the risk-neutral measure in the market. Ivaşcu (2021) compare various machine learning models against parametric approaches while pricing crude oil call options, highlighting preference for XGBoost and LightGBM models, followed by neural networks.

The poor performance of neural networks in these instances arise from two sources. First, Ivaşcu (2021) focused on using market data for training the network. A synthetic training dataset can overcome this issue by providing the practitioner with control over what scenarios the network should learn. Second, neural networks are tested against benchmark models, which have already been established to have superior quality. The aim should not be to compete, but rather train the network to learn the dynamics of high-quality stochastic volatility models. Ivaşcu (2021) reveal that neural networks learn the relationship between input and output variables from market data and lack economic interpretation. If the network is trained on a synthetic dataset created by a selected pricing approach, in addition to learning the model, it will receive an economic and mathematical foundation. The superiority of such asset pricing models should not be considered as a flaw, but rather an inspiration to exploit the universal approximation theorem by using neural networks as numerical solvers. This would improve the computational efficiency of such models and make them attractive for industrial practice. The universal approximation theorem, according to Telgarsky (2016) would hold, if the hidden layer is exponentially large. Given the computational advancements over the past couple of decades, scholars began experimenting with both deeper and wider networks. Culkin & Das (2017) expanded on work of Hutchinson et al. (1994) by training a deeper feedforward neural network to learn the dynamics of the Black and Scholes model. The increased depth allowed for the capturing of subtle non-linearities in the relationship between inputs and outputs. Karatas et al. (2019) took this a step further by synthetically training deep networks to learn diffusion and jump processes, with and without stochastic volatility models for pricing European, American and Barrier options.

In McGhee (2018), a single layer neural network was able to learn the dynamics of the SABR model, and estimate implied volatilities, when fed model parameters, maturity and strike prices as inputs. Computation was 10,000 (65) times faster than the finite-difference

(numerical integration) schemes. This application of neural networks can be extended to any asset pricing model, as witnessed in Liu et al. (2019). The intention was to replace the models with a computationally efficient numerical solver. Implied volatility surfaces corresponding to stochastic volatility models, were generated at rapid pace, with high levels of accuracy. Neural networks were trained to learn the Heston and Black and Scholes models, respectively. Price outputs from the former, were further used as inputs in another network to generate implied volatility surfaces. Several other works in literature have documented the capabilities of neural networks in estimating option prices/implied volatilities. The reader is recommended to see Ruf & Wang (2020) for a comprehensive overview of literature on the topic. The ability of neural networks to learn dynamics of high-dimensional top-quality asset pricing model, and provide outputs with superior computational efficiency, is the inspiration behind the two-step calibration approach.

#### 2.4.4.2 Two-Step Calibration

The two-step calibration approach begins with training a neural network to learn the dynamics of a given asset pricing model. The first step comprises of training, which is computationally heavy and conducted offline. The network learns to estimate option price/implied volatilities from inputs. In contrast to the one-step approach, there is a preference for synthetically generated training datasets. Hence, the practitioner has the freedom to select sampling distributions (Bayer & Stemper, 2018). Given that prior to calibration, little is known about model parameters, synthetic datasets allow for experimentation over a wide range of combinations, without size restrictions. Additionally, synthetic datasets are faster to create, as parameters do not require estimation. In the second step, the trained network is saved and paired with an optimizer for rapid online calibration from market observations. The major advantage of this approach is that calibration takes place by minimising the error between market observations, rather than parameters. It overcomes concerns over generalization and the ‘black-box’ nature of neural networks, by providing a calibration approach that can be controlled and validated in real-time (Bayer et al., 2019). A combination of the Levenberg-Marquardt optimizer with neural networks was proposed by Bayer & Stemper (2018) for calibrating the Heston and rough Bergomi asset

pricing models. Training datasets comprised of options priced using the Fourier method and Monte Carlo simulations, respectively, before computing implied volatilities.

There are two-different approaches to training neural networks for parametric calibrations: pointwise and grid-based (Bayer et al., 2019). In the former, as witnessed in Bayer & Stemper (2018), strike prices (or moneyness) and maturity are fed as inputs, with a single output node, corresponding to respective implied volatilities/prices. Alternatively, within the latter, the output layer comprises of multiple nodes, resembling a grid. Each node represents a unique strike-maturity pair, therefore only model parameters are required as inputs. Bayer et al. (2019) use a grid-based network and compare it to an inverse-mapping convolutional neural network one-step approach, for calibrating the rough Bergomi model. Calibration from implied volatility surfaces with both approaches was rapid, with the latter being slightly faster. However, the two-step approach provided real-time validation, and robust parameters. Grid-based approaches reduce the input dimensions and allow for exploiting the data structure by incorporating updates from neighbouring outputs into the learning process (Horvath et al., 2021). Nevertheless, they are not flexible and heavily reliant on interpolation and extrapolation techniques for estimating missing implied volatilities.

The flexibility of the pointwise approach allows for considering all traded options during daily calibrations. Even though the output layer consists of a single node, Liu et al. (2019) reveals that it is possible to simultaneously price an array of options during parametric calibration. The Calibration artificial Neural Network (CaNN) framework is popular amongst practitioners, being applied to calibrate Heston and Bates models (Liu et al., 2019), implied volatilities and dividend yields from American Options (Liu et al., 2021), and the Trolle and Schwartz model from market European swaptions (Büchel et al., 2022). In the latter, it ensures simultaneous pricing across 3 dimensions, strike prices, option expiry and swap tenors. On the other hand, a similar deep grid-based approach was applied by Horvath et al. (2021), for calibrating parameters of the Heston and Rough Bergomi models from daily implied volatility surfaces. It reveals that a network, trained over significantly fewer epochs, with much lesser nodes, can obtain accurate calibrations at rapid speeds, even on a standard personal computer. Both respective methods from Liu et al. (2019) and Horvath et al. (2021) act as a source of inspiration for the main contribution of this research.



In both works, networks are first trained on synthetic datasets, to learn the selected asset pricing process for estimating implied volatilities across a wide range moneyness and maturity terms. Liu et al. (2019) use the COS-method, whereas Horvath et al. (2021) employ Monte Carlo simulations to price options for training datasets. In the latter instance, neural network estimation makes it redundant to calculate payoffs of the option during calibration, despite the model lacking a closed-form solution. The first step also comprises of a validation phase, during which performance of the trained network is tested on its ability to fit a synthetic test dataset. Büchel et al. (2022) take this a step further, by using market calibrated parameters of the benchmark while testing. The eradication of overfitting concerns with the implementation of such validation phases is vital in determining the optimal architecture, and subsequently the weights and biases, within the hidden layers, that are crucial for the calibration phase.

For the second stage, Differential Evolution is heavily favoured as the optimiser in the aforementioned works. Given a well-trained neural network from the previous stage, the hidden layers are frozen, market data is fed into the output layer, and nodes corresponding to parameters in the input layer become learnable. The speed of calibration is heavily dependent on the rate at which implied volatilities/option prices are estimated. The iterative nature of optimisers, along with slow evaluations conducted by the Monte Carlo simulations, created a major computational burden, when calibrating stochastic volatility models. Liu et al. (2019) revealed benefits of combining the differential evolution optimiser with neural networks. This allowed the entire population during the selection stage to be addressed simultaneously, taking advantage of the ability possessed by a neural network to estimate wide ranges of market observations. Büchel et al. (2022) suggested the use of parametric estimation from the previous day, as initial values for the optimiser. Both techniques reduced the task of calibration to a matter of seconds. However, the approach from Horvath et al. (2021) shows that such levels of efficiency can be attained by simply considering a smaller (fewer nodes) network.

The two-step approach is a deterministic calibration framework, that relies on the existing knowledge and experience of the practitioner, with respect to traditional methods (Horvath

et al., 2021). Its performances are not dependent on the market period, but rather on the selected asset pricing model (Büchel et al., 2022). The method benefits from cheap and instantaneous forward runs of the network during the estimation of market observations, significantly boosting the iterative optimisation step. Even though it is marginally slower than its one-step counterpart, the two-step approach is favourable amongst regulators (Bayer et al., 2019; Ruf & Wang, 2020; Büchel et al., 2022), so much so that it can be treated on par with stochastic volatility models. The additional optimisation ensures real-time testing and interpretation of outputs, with respect to benchmark, making it simpler to explain results or any form of deviations.

Neural networks are solely employed as tools for improving computational efficiency of high-quality asset pricing models. Hence, regulators are not concerned about the ‘black box’ nature, notoriously associated with them. Furthermore, synthetic datasets ensure more robust training and testing stages, reducing the probability of over-fitting. The practitioner is flexible to decide size and distribution of variables in the dataset, possessing the ability to include several scenarios that are not necessarily captured by historic data. Hence, performances of the two-step approach are not dependent on the market regime, but instead on the selection of the model, in contrast to the one-step approach. Finally, improved robustness, and reduced overfitting tendencies, can ensure that the network does not require frequent/any retraining. It is with these benefits in mind that a pointwise two-step approach inspired by the works of Liu et al. (2019) and Horvath et al. (2021) is proposed for calibrating the sophisticated stochastic volatility jump diffusion model, within the three-step approach (Fusari et al., 2024) for real-time bubble detection.

## 2.5 Summary and Discussion

Bubbles are revealed when the price deviates from its fundamental value. This phenomenon captures willingness of investors to purchase the asset, even at exaggerated valuation, for the sole purpose of reselling it at an even higher price. There are several well documented historical price bubbles, across various asset types, dating back to *Tulipmania*, in the 17<sup>th</sup> century, which eventually collapse, causing mass socioeconomic and financial losses. Alarming, in recent times, the frequency at which these events occur, has significantly

increased. Moreover, the growing integration of global financial markets, and interlinkage across different asset classes, has vastly scaled the potential of devastation incurred from bubbles bursting. This was witnessed during the *Global Financial Crisis, 2008*, which imprinted long-lasting effects economies worldwide, to the extent that some are yet to recover. Hence, with great emphasis, there is an urgent need for developing a robust and accurate early warning detection system, such that timely measures can be implemented for damage control. This has piqued the interest of several academics and practitioners, contributing to the growing literature of bubble detection methods.

Traditional methods identify bubbles by modelling the fundamental value of the asset. However, these approaches suffered from a joint-hypothesis related issue, which made findings inconclusive. Over the past decade, popularity of implementing recursive regression techniques (Phillips et al., 2011, 2015) and the LPPLS model (Johansen et al., 1999, 2000) grew, given their improved ability to detect and timestamp bubbles. Nevertheless, similar to traditional methods, they are plagued by joint hypothesis issues, given focus on modelling the fundamental value. Additionally, the LPPLS model is unable to consider the impact of exogenous factors on bubble formation. Akin traditional methods, both techniques seek out type II bubbles, which exist on an infinite horizon. Large timeseries of data are required, which create vulnerability to potential structural breaks within the sample. Therefore, preferences shifted towards exploring the local martingale theory of bubbles.

The local martingale theory is designed to detect type III bubbles by focusing on the diffusion component of the asset price process. Such bubbles exist within assets with bounded lives, capturing short-term trade strategies developed by investors, for making a quick profit. In accordance with the theory, bubbles are revealed when the asset price is a strict local martingale. This property captures the deviation from fundamentals, and is observed when volatility grows at a rate, in excess of that experienced by the price. Therefore, exuberance is only detected by modelling the asset's volatility, making the need to estimate fundamental values, entirely redundant, and eradicating the notion of a joint-hypothesis related issue. Furthermore, since type III bubbles exist within short horizons, large timeseries of dataset are not required, eliminating vulnerability towards structure

breaks, and reducing heavy reliance on extracting backward-looking information from historical data. The estimation of the volatility function for bubble detection was witnessed in Jarrow et al. (2011), inspiring several future works. However, it was only till the application of a Hidden Markov Model (Obayashi et al., 2017), that allowed for the date-stamping bubbles.

An alternative approach for detecting bubbles, as observed at annual frequencies in Chaim & Laurini (2019) and Laurini & Chaim (2021), is the calibration of stochastic volatility models. These large timesteps increased vulnerability to structural breaks. Moreover, the model was calibrated to spot markets, which are not capable of revealing forward-looking information. In contrast option prices excel at capturing these characteristics. The literature exploring option prices for detecting bubbles in the underlying price has been gaining substantial attention from scholars. Put option prices are unable to reveal bubbles, given bounded payoffs, however, call options have unbounded payoffs, allowing their prices to reflect exuberance. Additionally, call option bubbles signal towards exuberance in the underlying price, with the magnitude of the former, acting as the lower bound for that of the latter.

The SABR model was calibrated to daily volatility smiles in Piironen et al. (2018) and Stahl & Blauth (2024), with the latter extending its application to the entire surface. A deep neural network was trained in Biagini et al. (2024) to learn various stochastic volatility models for identifying bubbles in the underlying, from call option prices. Whereas, Jarrow & Kwok (2021, 2024) employed a nonparametric method to infer bubbles from option data. However, the three-step approach (Fusari et al., 2024) stands out in comparison to its counterparts. It provides a solid foundation for overcoming the joint-hypothesis issue by calibrating a sophisticated stochastic volatility jump diffusion model to market put option prices. Also, reliance on parametric conditions for detecting strict local martingale tendencies is reduced, due a statistical test designed to observe call options bubbles, signally towards the presence of similar phenomena in the underlying price. Hence, this study proposes the implementation of the three-step approach for detecting asset price bubbles, on the grounds of its superior ability to overcome the joint-hypothesis issue and capture forward-looking expectations of market participants.

The three-step approach calibrates parameters for a stochastic volatility jump diffusion model, by fitting it to market put observations. The lack of a closed form solution forced reliance on Monte Carlo simulations for estimation. This created a large computational burden, and in order to strike balance between robustness and efficiency, Fusari et al. (2024) were forced to calibrate parameters from, only the daily most liquid option smile. Despite being common practice, such calibrations fail to capture the full extent of information present across the entire surface. Biagini et al. (2024) and Stahl & Blauth (2024) shows benefits of utilising the entire option price surfaces for bubble detection. The abundance of information present across various option maturities improves comprehension of forward-looking expectations, especially with regards to changes in the underlying price, and subsequently, betters the quality of bubble detection. Biagini et al. (2024) directly examine bubbles from call prices, whereas Stahl & Blauth (2024) use the SABR model, for which a closed form solution exists. Nevertheless, the three-step approach is preferred.

In comparison to the method from Biagini et al. (2024), it provides an additional validation step by calibrating to put options. Furthermore, as opposed to the SABR, the stochastic process employed in the three-step approach, captures price jumps and mean-reverting characteristics of volatility. Therefore, the focus shifts towards improving the computational efficiency of the three-step approach, especially during the calibration phase. The industrial application of stochastic processes reveal strong preference for tractability over accuracy (Horvath et al., 2021). However, compromising on accuracy during bubble detection can result in devastating consequences due to inadequate and untimely risk management. Hence, to improve efficiency of the three-step approach, without sacrificing accuracy, this study proposes the application of neural networks.

The universal approximation theorem (Cybenko, 1989; Hornik et al., 1989; Hornik, 1991) reveals that neural networks can be exploited for estimating any function, with a given level of accuracy. This pursued several works to tackle pricing problems in the derivative markets, given the highly frequent need for thousands of calculations, across thousands of contracts, written over thousands of underlying instruments (De Spiegeleer et al., 2018). Two forms of neural network operations have been employed to overcome the trade-off between accuracy

and tractability during calibrations. First, the one-step approach, directly estimate parameters as outputs, from market observations. Second, the two-step approach trains the network to estimate market observations and employs an optimizer for calibration.

The additional step provides real-time validation to the calibrated parameters, which is satisfactory for regulatory bodies. In context to bubble detection, it ensures real-time testing of potential joint-hypothesis issues. The two-step calibration also allows training and testing on synthetic datasets, giving the practitioner control over its size, and distribution of input parameters. This improves the robustness, reducing or even eliminating the need for retraining. Furthermore, it ensures performances of the two-step approach to be solely dependent on the selected model. Therefore, inspired by the works of Liu et al. (2019) and Horvath et al. (2021) a two-step neural network calibration framework is proposed for improving the computational efficiency of the three-step approach, and subsequently expanding its application for asset price bubble detection.

The neural network calibration framework is employed towards estimating the stochastic volatility jump diffusion process from market put options data. Given, the massive boost is computational efficiency and flexibility, the stochastic process can be calibrated from daily put option surfaces, instead of being restricted to the most liquid cross section. Furthermore, the study proposes the employment of parameters from surface calibrations to explore bubble formations within different call option maturity groups, and over the lifetime of option contracts, respectively. In brief, this research proposes a neural network calibration framework to enhance the computational efficiency of the three-step approach, such that it can be applied over a vast range of option maturities. In doing so, the ability to capture forward looking expectations of market participants for bubble detection, is improved. Additionally, the computational boost provided to a highly sophisticated stochastic process, makes the three-step approach more attractive for real-time bubble detection. Hence, improving capabilities towards implementing timely and appropriate risk management measures, for reducing socioeconomic and financial damages, consequent of a crash.



# Chapter 3:

# Methodology

In the previous chapter, following an in-depth examination of asset price bubble detection methods, preference for the three-step approach (Fusari et al., 2024) was revealed. The method identifies bubbles in the underlying from option prices, in accordance with principles from the local martingale theory of bubbles. Recall, in contrast to traditional methods, it overcomes the trade-off between model asymptotics and structural breaks; eliminates the joint-hypothesis issue and captures the forward-looking nature of bubbles. Nevertheless, given the computational burden associated with its applications, a deep neural network calibration framework is required. The current chapter delves into discussing the three-step approach and construction of the deep calibration framework, in abundance with the theoretical background of the local martingale theory from (2.3), bubble detection using option prices (2.3.2).



## 3.1 Three-Step Approach

A three-step approach was designed by Fusari et al. (2024) under the framework of the local martingale theory, for bubble detection using option prices. Options display type III bubbles, which capture short-term trade strategies, intended to perceive the future resale value of the underlying. Therefore, such bubbles can be identified and quantified using real-time prices, with reliance on historical data and backward looking time windows made redundant. Furthermore, the three-step approach provides a real-time validation of the modelling process and overcomes the notorious joint-hypothesis issue that has plagued traditional methods.

Under the assumption that markets are incomplete and satisfying NFLVR and ND, options are priced using a stochastic volatility jump diffusion model, which relies on Monte Carlo simulations for computation. It is important that the selected model is capable of detecting strict local martingale tendencies in the underlying price process. Therefore, the first step is focused on selecting the appropriate model for option pricing. Next, the chosen process is calibrated to daily market put option observations, such that estimated parameters are used for pricing call options. It is in this step that the joint hypothesis is overcome in real-time, without requiring independent testing. Finally, the bubble is estimated and tested for statistical significance by examining the difference between the prevailing market and model-estimated (fundamental) call option prices.

### 3.1.1 Step I: Model Selection

The specification of the correct parametric model for the underlying price process is vital, requiring it to admit martingale and strict local martingale tendencies. Therefore, popular processes such as those from Merton (1976), Heston (1993) and Bates (1996), are ruled out. The SABR model (P. S. Hagan et al., 2002) is a viable candidate, as revealed in Piironen et al. (2018) and Stahl & Blauth (2024). However, it is unable to capture the mean-reverting

nature of volatility, hence making it unsuitable for pricing long-term options (Gatheral, 2012).

$$\frac{dS_t}{S_t} = \sqrt{V_t}dW_t + dJ_t - \lambda Mdt \quad (17)$$

$$dV_t = \kappa(\bar{v} - V_t)dt + \sigma_v V_t^p dZ_t \quad (18)$$

Stochastic volatility processes, that capture strict local martingale tendencies, such as those from Sin (1998) and Andersen & Piterbarg, (2007), have been previously witnessed in the works of Biagini et al. (2024), and, Chaim & Laurini (2019) & Laurini & Chaim (2021), respectively. In contrast to the SABR, these models are capable of addressing the mean-reverting nature of volatility. However, Fusari et al. (2024) suggest an improvement by taking into account price jumps, and recommend the Generalised Stochastic Volatility Jump Diffusion (GSVJD) model, as revealed in (17) and (18).

The model is represented under the ELMM  $\mathbb{Q}$ , which according to the assumptions of NFLVR and incomplete markets, is allocated by the market to determine the fundamental value of the asset. The price of the underlying is normalised with respect to the MMA. Hence, through integration by parts, (17) and (18) represent a process for discounted prices by eliminating the riskless drift component. The GSVJD captures strict local martingale tendencies in the underlying price, conditional to values of parameters  $\rho$  and  $p$ . For  $p \leq 0.5$  or  $p \geq 1.5$ , the underlying process is a martingale. When  $0.5 < p < 1.5$ , the underlying price process is a martingale, if  $\rho \leq 0$ , but is a strict local martingale for  $\rho > 0$ . In the latter scenario, explosive behaviour is present in the volatility is greater, relative to the underlying asset's price hike. It abides by the same intuition revealed in Delbaen & Shirakawa (2002), Kotani (2006), and Mijatović & Urusov (2012), where bubbles to existing if and only if,  $\int_{\epsilon}^{\infty} \frac{s}{\sigma(s)} < \infty$ , for  $\epsilon > 0$ .

The parameter vector  $\theta = \{V_0, \bar{v}, \kappa, \sigma_v, \rho, p, \mu_y, \sigma_y, \lambda\}$ , has 9 constituents<sup>27</sup>. The martingale jump term is denoted by  $(dJ_t - \lambda Mdt)$ , where  $dJ_t = d(\sum_{j=1}^{N_t}(Y_j - 1)) = (Y_t -$

---

<sup>27</sup> Parameters are explained in the order of representation in  $\theta$ ; initial volatility, long-term mean of the variance, mean-reversion speed of variance, volatility of volatility, correlation coefficient between standard Brownian

1) $dN_t$ ; with jump size  $Y_i \sim \mathcal{LN}(\mu_y, \sigma_y)$ , counting process represented by  $N_t \sim \text{Poisson}(\lambda t)$ ,  $\lambda$  being the intensity of total number of jumps till time  $t$ , and  $J_t - \lambda M dt$  is a martingale, with jump-compensator  $M = E[Y_j - 1] = e^{(\mu_y + 0.5\sigma_y^2)} - 1$ . In addition, the GSVJD model nests the stochastic processes from Heston (1993), Bates (1996), and Andersen & Piterbarg (2007). When  $\lambda = 0$ , it resembles the process from Andersen & Piterbarg (2007). Whereas, with  $p = 0.5$ , for  $\lambda = 0$  and  $\lambda > 0$ , it represents the Heston (1993) and Bates (1996) model, respectively. Notice, in the latter instance, the given the fixed condition on parameter  $p$ , the processes are unable to capture strict local martingale tendencies.

### 3.1.2 Step II: Model Calibration and Estimation

The selection of the sophisticated GSVJD model is justified by its ability to reflect strict local martingale tendencies in the underlying prices, via specific combinations of  $p$  and  $\rho$  parameters. In addition, it includes jumps in the stochastic processes, which have been found to improve option pricing as documented in the well renowned works of (Merton, 1976; Bates, 1996). A major drawback is that the GSVJD lacks a closed form solution, and therefore, computation is reliant on Monte Carlo simulations.

$$V_{i,t+\Delta t} = V_{i,t} + \kappa(\bar{v} - V_{i,t})\Delta t + \sigma_v V_{i,t}^p \sqrt{\Delta t} B_{i,1} \quad (19)$$

$$\tilde{S}_{i,t+\Delta t} = \tilde{S}_{i,t} \exp \left[ -(\lambda M + 0.5 \cdot V_{i,t})\Delta t + \sqrt{V_{i,t}\Delta t} (\rho B_{i,1} + \sqrt{1 - \rho^2} B_{i,2}) + J \right] \quad (20)$$

For  $i = 1, \dots, MC$ , and  $\Delta t = \frac{1}{365 \times 5}$ , the value of the stock ( $\tilde{S}$ ) with accumulated dividends ( $\delta$ ) from  $t$ , till maturity ( $t + \tau$ ), is computed in accordance to the discretised versions of the GSVJD model, in (19) and (20). Estimation is conducted in abundance with Broadie & Kaya (2006), with  $B_{i,1}$  and  $B_{i,2}$  representing independent standard normal variables. Given jump size,  $Y_i \sim \mathcal{LN}(\mu_y, \sigma_y)$ , at first (20) can be computed, by ignoring the jump component,  $J$ , to obtain  $\hat{S}_{i,t+\tau}$ . Next, number of jumps occurring during the time horizon,

---

motions  $W_t$  and  $Z_t$ , parameter  $p$  (impacts the volatility of volatility), the moments of  $Y_i$ , and intensity of total number of jumps till time  $t$ .

$N_{i,\tau} \sim \text{Poisson}(\lambda\tau)$  are simulated along with independent jump sizes. Finally, by multiplying  $\hat{S}_{i,t+\tau}$  with the jump sizes,  $\tilde{S}_{i,t+\tau}$  can be determined, where  $\tilde{S}_{i,t+\tau} = \hat{S}_{i,t+\tau} \prod_{n=1}^{N_{i,\tau}} Y_{i,n}$ . Since option holders do not receive dividends, the fundamental values the European-styled put and call options from both models, are calculated as per (21) and (22), respectively, where  $K = \mathbb{K}/B_{t+\tau}$  represents the strike price denoted in terms of the MMA, and  $S_{i,t+\tau} = \tilde{S}_{i,t+\tau} e^{-\delta(t+\tau-t)}$ .

$$P_t^*(\theta, K, \tau) = \frac{1}{MC} \sum_{i=1}^{MC} (K - S_{i,t+\tau})^+ \quad (21)$$

$$C_t^*(\theta, K, \tau) = \frac{1}{MC} \sum_{i=1}^{MC} (S_{i,t+\tau} - K)^+ \quad (22)$$

Since put option prices do not reveal bubbles, the objective is to calibrate  $\theta$ , such that  $P_t^*(\theta, K, \tau)$  is aligned with its American-/European-styled market counterpart,  $P_t^{A/E}(K, \tau)$ . This ensures that the three-step approach does not suffer from the joint-hypothesis related issues. Prices are converted to implied volatilities for  $j$  strikes,  $\sigma_t(\theta_t, K_{t,j}, \tau_t)$ , and  $\sigma_t(K_{t,j}, \tau_t)$ , respectively, such that the root mean squared error,  $RMSE_{\sigma,t}$ , is minimised. Implied volatilities are preferred over prices, as the objective function is normalised, and errors can be interpreted in terms percentage (Fusari et al., 2024).

$$RMSE_{\sigma,t} = \min_{\theta_t} \sqrt{\frac{1}{N_t^p} [\sigma_t(K_{t,j}, \tau_t) - \sigma_t(\theta_t, K_{t,j}, \tau_t)]^2} \quad (23)$$

The objective function is minimised over  $N_t^p$  put options to obtain the GSVJD parametric vector, corresponding to the selected model, where  $\theta_t = \theta_{GSVJD,t}$ . Recall, as shown in (14), even for American options, the fundamental value of a European call counterpart, is required for detecting and quantifying bubbles. Therefore, option fundamental values required for bubble detection in this study, will be computed in accordance with (21), (22), and (23),

using Differential Evolution<sup>28</sup> (Storn & Price, 1997). Even though local optimisers are faster, they require an initial value, which makes their performance dependent on quality of the guess, and vulnerable to insufficient knowledge from the practitioner's behalf. Upon completion of calibration,  $\theta_t$  will be used to compute the fundamental value of call options, as per (22).

### 3.1.3 Step III: Estimation and Testing of Call Option Bubbles

Model-implied call option prices,  $C_t^{A/E}(K_{t,j}, \tau_t)$ , across  $N_t^c$  available strikes corresponding to maturity  $\tau_t$ , are computed using  $\theta_t$  obtained from calibrating to  $\sigma_t(K_{t,j}, \tau_t)$ . The estimation and testing of bubbles, in both the underlying and call option prices, is robust and free of the joint-hypothesis problem. Call options display bubbles when their fundamental value is exceeded by that quoted in the market. For a given strike price,  $\{K_{t,j}\}_{j=1, \dots, N_t^c}$ , the bubble in an option is denoted by  $\mathcal{B}_t(K_{t,j}, \tau_t)$ , with  $\mathbb{N}$  representing the option lot size.

$$\mathcal{B}_t(K_{t,j}, \tau_t) = [C_t^{A/E}(K_{t,j}, \tau_t) - C_t^*(K_{t,j}, \tau_t)] \times \mathbb{N} \quad (24)$$

Options trade within the bid-ask spread, rather than precisely at the mid-point, and given that the determination of the *true* market price is vague, Fusari et al. (2024) suggest setting  $\mathcal{B}_t(K_{t,j}, \tau_t) = 0$ , when  $C_t^{Bid}(K_{t,j}, \tau_t) \leq C_t^*(\theta_t, K_{t,j}, \tau_t) \leq C_t^{Ask}(K_{t,j}, \tau_t)$ . If  $C_t^{Bid}(K_{t,j}, \tau_t) > C_t^*(\theta_t, K_{t,j}, \tau_t)$ , or  $C_t^*(\theta_t, K_{t,j}, \tau_t) < C_t^{Ask}(K_{t,j}, \tau_t)$ ,  $\mathcal{B}_t(K_{t,j}, \tau_t)$  is computed according to (24). Since there are  $N_t^c$  call options, average of  $\mathcal{B}_t(K_{t,j}, \tau_t)$  across the various strike prices is considered, prior to normalising with respect to  $\mathbb{S}$ .

$$\widehat{\mathbb{B}}_t = \frac{1}{\mathbb{S}_t} \sum_n^{N_t^c} \frac{\widehat{\mathcal{B}}_t(K_{t,j}, \tau_t)}{N_t^c} \quad (25)$$

The average magnitude of the call option bubble  $\widehat{\mathbb{B}}_t$ , can be interpreted as a percentage of  $\mathbb{S}_t$ , acting as a lower bound to the size of that in the underlying price. After the estimation

---

<sup>28</sup> Differential Evolution was implemented as per the `scipy.optimize` library in Python ([https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.differential\\_evolution.html](https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.differential_evolution.html)), with default functional values being utilised.

stages, only call options, that adhering to the ND<sup>29</sup> condition (Merton, 1973) are considered. This ensures that the put-call parity, and hence a linear relationship between bubbles in the option and underlying prices is established. If so, the presence of a bubble in call option prices would signal exuberance in the underlying, and vice versa. This acts as the key motivation behind the development of the conditional test by Fusari et al. (2024).

$$\widehat{\mathbb{B}}_t = \mathbb{B}_t + \varepsilon_t; \varepsilon_t \sim \mathcal{N}(0, \sigma_t) \quad (26)$$

The conditional test is built on the foundations that negative bubbles ( $\widehat{\mathbb{B}}_t < 0$ ) occur due to potential model misspecifications or observation errors. According to the local martingale theory, the fundamental value of an asset acts as the lower bound for its market price. The null hypothesis,  $\mathbb{B}_t = 0$ , reveals a non-existent bubble,  $\varepsilon_t$  represents the model error and noise, with  $\sigma_t$  being utilised to construct a time-varying threshold. The construction of the threshold relies on  $\widehat{\mathbb{B}}_t < 0$  observations. Negative bubbles exist due to two reasons. First, observation error in market prices, especially in context to OTM options with large bid-ask spreads, along with asynchronous trading, and potential trading constraints. Second, measurement error due to model misspecification, given that call option bubble measurements rely on the difference between the market and fundamental price.

$$\widehat{\sigma}_t^2 = \sum_{i \in Y_-} \frac{\widehat{\mathbb{B}}_t^2}{N_-} \quad (27)$$

The negative bubbles are utilised to construct the threshold for testing the null hypothesis. From (26) it can be deduced that if  $\widehat{\mathbb{B}}_t < 0$ , then  $\mathbb{B}_t + \varepsilon_t < 0$ , and given that bubbles are nonnegative,  $\mathbb{B}_t > 0$ ,  $\varepsilon_t$  will be negative and larger in magnitude,  $\varepsilon_t < -\mathbb{B}_t \leq 0$ . Under the null, given  $\mathbb{B}_t = 0$ ,  $\widehat{\mathbb{B}}_t = \varepsilon_t < 0$ , implies that the observed bubble is an error. On these foundations, an estimator for the unbiased variance is created, by considering a small window, before  $t$ . Daily bubble estimates, at observation times  $i = t - k, \dots, t$ , where  $\widehat{\mathbb{B}}_t < 0$ , are collected and stored in  $Y_-$ , which consist of  $N_-$  elements, such that the unbiased estimator,  $\widehat{\sigma}_t^2$  is obtained from (27). Under the null, even instances of  $\widehat{\mathbb{B}}_t > 0$ , are considered as an error, implying that when  $\mathbb{B}_t > 0$ , it must be larger than  $\varepsilon_t$ , which are negative. Hence,

---

<sup>29</sup> European Call options must adhere to  $P_t^E(K, \tau) + S_t - KB_{t,\tau} \geq C_t^E(K, \tau)$ , whereas for American styled contracts, the condition implied  $C_t^A(K, \tau) \leq S_t$ .

including such observations in the construction of  $\sigma_t^2$ , would generate a bias<sup>30</sup>. The null hypothesis is rejected, if  $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ , with the latter being the time-varying threshold, comprising of level of significance ( $\alpha$ ), and the unbiased estimator ( $\widehat{\sigma}_t$ ). The latter permits using all available information on a given day and hence allows for real-time bubble detection. In contrast to Fusari et al. (2024), who consider the 10% level of significance and a window size of 180 days for constructing the unbiased estimator, each of the two variables are explored over a range of values, in this study.

## 3.2 Deep Learning Calibration Framework

The three-step approach (Fusari et al., 2024) exploits the forward looking nature of option markets for bubble detection. Additionally, in contrast to traditional methods, it overcomes the joint hypothesis issue in real time to provide relevance and robustness to findings. The approach implements the GSVJD model for pricing options by considering the stochastic nature of volatility, along with jumps in the underlying price process.

The GSVJD process possesses high quality capabilities in pricing and capturing information across various option maturities. However, there exists a major bottleneck in industrial application of the model. Due to the unavailability of an analytical solution, Monte Carlo simulations are employed as a numerical solver. Hence, Fusari et al. (2024) were forced to strike a balance between robustness and efficiency, by calibrating the GSVJD parameters from only the daily most liquid option maturities. The benefits of considering the entire surface of prices for bubble detection have been highlighted in Stahl & Blauth (2024) and Biagini et al. (2024). Calibrating the GSVJD to the entire surface, allows bubble identification to take place using greater quality of information regarding the forward-

---

<sup>30</sup> Fusari et al. (2024) highlight that pricing American option without considering early exercise, adds a model misspecification in the form of inflated put option prices. This results in higher put implied volatilities, and subsequently, those associated with the fundamental value of call options, to suffer a similar faith. Therefore, the magnitudes of bubbles are underestimated, hampering the likelihood of identifying significant instances of exuberance in call options. If so, this phenomenon would resemble a similar impact as the presence of stock lending fees, which raise the cost of shorting, and therefore, inflate put option prices.

looking expectations of market participants. Therefore, this study employs neural networks to boost computational efficiency, without adversely affecting accuracy.

Inspired by the two-step methodology from Liu et al. (2019) and Horvath et al. (2021), a deep calibration of the three-step approach for real-time bubble detection is proposed. In this approach, neural networks are applied as numerical solvers with the sole purpose of expediting calculations. First, the optimal architecture of the network must be determined, and trained on a synthetically generated dataset. The training and validation process is referred to as the *forward pass*, similar to CaNN framework (Liu et al. 2019). It ensures that the dynamics of the GSVJD model are well learned by the network. In line with Liu et al. (2019), Horvath et al. (2021), and Büchel et al. (2022), a Multi-Layer Perceptron (MLP) model will be employed. The *forward pass* is computationally cumbersome but fortunately, a one-time offline operation. During the second step, *backward pass*, the trained network is frozen and paired with an optimiser for rapid online calibrations. It is due to such efficiency, that the calibration of the GSVJD model can be extended to the entire available option price surface. The goal is to employ neural networks towards calibrating the GSVJD model, such that the need for Monte Carlo simulations is made redundant. For the remainder of the research, the two-stages of the calibration approach will be referred to as the *forward* and *backward pass*, respectively.

### 3.2.1 Neural Networks: Multi-Layer Perceptron (MLPs)

Artificial neural networks (ANNs) lie in the realm of machine learning, between the engineering and artificial intelligence disciplines (Lek & Park, 2008). Corresponding to the former, error functions are utilized for optimisation, whereas for the latter, they rely on heuristic methods that lack theoretical background to support decision making. This provides ANNs with the ability to comprehend complex non-linear dependencies between financial variables, by mimicking the biological nervous system. The connection of neurons, allows flow of information, similar to the power of synaptic links within the biological structure (Pothina & Nagaraja, 2023). The simplest of architectures prevail in the form of feedforward networks, characterized by the direction of information flow. The addition of a hidden layer



with an appropriate activation function makes them capable of general modelling (Lavine & Blank, 2009), abiding by the universal approximation theorem.

A Multi-Layer Perceptron (MLP) is a form of feedforward neural network, comprising of at least a single hidden layer, sandwiched between input and output layers. An MLP is the most popular neural network deployed in practice, given its ability to extract nonlinear input-output relationships without requiring prior knowledge on specific functional forms, and simplistic nature of execution. It comprises of two steps: forward-propagation, and backward-propagation. First, inputs enter the network, and pass through hidden layers, to generate output(s). All nodes, barring those within the same layer, are interconnected, allowing for such unidirectional flow of information. Hidden layers play a crucial role, assigning weights, and biases, and transforming inputs by using nonlinear activation functions. The provision of output(s) by relying on the heuristic of an activation function, aligns with the artificial intelligence discipline. Next, estimated output(s) is compared to the target(s), via a loss function. The iterative calculations of output by updating weights and biases at each neuron to minimize the loss function is known as backward-propagation, linking to the engineering discipline. A combination of elements from both disciplines ensures the neural network is capable of making decisions based on a particular training dataset. Obviously, the workings of MLP networks are not as straightforward, comprising of several mathematical operations. Hence a detailed explanation of each step is provided below. Furthermore, given preference for a pointwise two-step calibration approach, the neural network will comprise of a single output node.

### *3.2.1.1 Forward Propagation*

The forward propagation step is classified by unidirectional flow of information, from feeding the neural network inputs, to the estimation of outputs. All nodes in consecutive layers are interconnected, building a pathway, via constituents of hidden layers for determining the relationship between input and output variables. A basic artificial node that receives input to produce output via transformative operations is known as a perceptron. It is the mathematical representation of a biological neuron, capable of making decisions by analysing data.


$$z = \phi \left( \sum_{i=1}^n w_i \cdot x_i + b \right) \quad (28)$$

$$y(x) = F(x|\varphi), \quad \varphi = (W_1, B_1, W_2, B_2, \dots, W_h, B_h) \quad (29)$$

$$z_{m,1} = \phi \left( \sum_{i=1}^n w_{i,m,1} \cdot x_i + b_{m,1} \right) \quad (30)$$

The mechanisms of a perceptron receiving  $n$  inputs ( $x$ ), and producing output,  $z$ , is illustrated in *Figure 3.1*. Inputs entering the neuron are assigned an individual weight,  $w_i$  for  $i = 1, 2, \dots, n$ . A constant term,  $b$ , known as the bias, is added to the weighted sum of inputs, prior to being passed through an activation function  $\phi(\cdot)$  for computing output. Activation functions play a crucial role in comprehending the complex input-output dependencies. They are activated for performing non-linear transformations on input values (or not), in relevance to the received data (Jindal et al., 2022). The adjustment of weights only permits movement of output along the activation curve, whereas a bias acts as an anchor, allowing for upward/downward shifts to improve representation of the input space (Hagan et al., 2014). The generation of output, by feeding a perceptron,  $n$  inputs, is represented in (28).

**Figure 3.1: Mechanisms of a Perceptron.**




Perceptrons mimic biological neurons, whereas the MLP network imitates the entire nervous system. A neural network has multiple layers (at least one hidden), each consisting of several stacked perceptrons. The nodes across successive layers are interconnected, creating a link between associated input and output patterns. Furthermore, the output node in the final layer, is connected to all neurons within the network. In the hidden layers, neurons are stacked and

arranged in parallel (Lavine & Blank, 2009), allowing for decision making based on multiple tests on input parametric combinations. Therefore, the activation function at each perceptron aims to produce output, as inputs for nodes in the next layer.

The forward propagation stages of an MLP network comprising of  $h$ -hidden layers, is illustrated in *Figure 3.2*. An input layer, with  $n$  nodes is considered, such that mechanisms of all neurons from the first hidden layer, are consistent with Figure 3.1. A minor adjustment of adding two subscript indices is required, to ensure the output from a particular node, within a certain hidden layer can be identified. For instance,  $z_{1,1}$  would denote output of the initial node, in the first hidden layer. The extension to subsequent layers is straightforward, which consider outputs from the previous, as inputs. Overall, there are  $j$  neurons within each hidden layer, that can develop an understanding of relationship between input and output variables, by assigning weights and biases. This forward flow of information can be mathematically rendered by considering the output ( $y$ ), to be a function of inputs ( $x$ ), and MLP parameters ( $\varphi$ ). The latter comprises of weight matrices ( $W_l$ ) and bias vectors ( $b_l$ ) associated with the  $l^{th}$  hidden layer, for  $l = 1, 2, \dots, h$ .

**Figure 3.2: Multi-Layer Perceptron (MLP) Framework.**



*Note: Representation of the forward-flow of information from  $n$  inputs to the output node, via  $h$  hidden layers, each consisting of  $j$  neurons.*

It is worth revealing the process of obtaining outputs from each layer, to illustrate the determination of the final output ( $y$ ) value. Consider the first hidden layer ( $l = 1$ ), where inputs of the network are directly received and processed at each neuron, such that the layer produces a vector of outputs,  $Z_1$ , where  $z_{m,1} \in Z_1$ , for  $m = 1, 2, \dots, j$ . Upon closer inspection, output at the  $m^{th}$  node is obtained by allocating weights,  $w_{i,m,1} \in W_1$ , corresponding inputs  $x_i \in X$ , for  $i = 1, 2, \dots, n$ , and adding a bias  $b_{m,1} \in B_1$ , prior to passing it through an activation function,  $\phi(\cdot)$ .

$$\begin{bmatrix} Z_{1,1} \\ Z_{2,1} \\ Z_{3,1} \\ \vdots \\ Z_{j,1} \end{bmatrix} = \phi \left( \begin{bmatrix} w_{1,1,1} & w_{2,1,1} & w_{3,1,1} & \dots & w_{n,1,1} \\ w_{1,2,1} & w_{2,2,1} & w_{3,2,1} & \dots & w_{n,2,1} \\ w_{1,3,1} & w_{2,3,1} & w_{3,3,1} & \dots & w_{n,3,1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ w_{1,m,1} & w_{2,m,1} & w_{3,m,1} & \dots & w_{n,m,1} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{bmatrix} + \begin{bmatrix} b_{1,1} \\ b_{2,1} \\ b_{3,1} \\ \vdots \\ b_{j,1} \end{bmatrix} \right) \quad (31)$$

$$Z_1 = \phi(W_1 \cdot X + B_1) \quad (32)$$

$$\begin{aligned} Z_2 &= \phi(W_2 \cdot Z_1 + B_2) \\ &\vdots \end{aligned} \quad (33)$$

$$y = \phi(W_{(h+1)} \cdot Z_h + B_{(h+1)})$$

The entire mathematical operation of output generation from the first hidden layer, is revealed in (31) and simplified in (32), with respect to vector and matrix notations. Neurons of the next hidden layer, receive  $Z_1$  as inputs, such that outputs,  $Z_2$  are obtained as depicted in (33). This process of using outputs from the previous layer as inputs, continues to till the final layer<sup>31</sup>. In the output layer, results from all the hidden layers are combined and evaluated (Lavine & Blank, 2009). The knowledge accumulated in weights and biases associated with each neuron, is utilised to produce the final output,  $y$ , so it can be tested against the actual target value. Testing is conducted by considering the distance between the estimated and actual values, with the assistance of a loss function. The network aims to minimise the loss function by iteratively calculating output from adjusting weights and biases, at every neuron, during the backward propagation step.

---

<sup>31</sup> Given  $h$  hidden layers, since the output layer is in succession of the last, weights and biases allocated are indexed as  $(h + 1)$ .

### 3.2.1.2 Backward Propagation

Weights (biases) emphasise (deemphasise) the influence of individual input variables (Lavine & Blank, 2009). They resemble the most important determinants of the MLP network for comprehending the relationship between input and output variables. Once established, the knowledge of the MLP can be tested by assessing distance, between estimated output,  $y(x|\varphi)$ , and targeted/actual output,  $y(x)$ , via the loss function,  $L(\varphi)$ . The most popular variant of  $L(\varphi)$  is the mean squared error (MSE) function, seen in (34). In the backpropagation stage, gradient decent techniques are utilised to re-adjust parameters,  $\varphi$ , such that  $L(\varphi)$  is minimised.

$$\arg \min_{\varphi} L(\varphi) = \frac{1}{N} \sum_{k=1}^N (y_k(x) - y_k(x|\varphi))^2 \quad (34)$$

Gradient descent methods work towards seeking the steepest gradient to find an optimum point, for a given function (Christensen & Bastien, 2016). The method is performed on the entire error surface by modifying  $\varphi$ , in proportion to the gradient at a given location (Lek & Park, 2008). Since the goal is to minimise  $L(\varphi)$ ;  $\varphi$  is altered such that the search moves in the direction of negative gradients, for obtaining the minimum value (Pothina & Nagaraja, 2023). The gradients are computed by taking the partial differential of  $L(\varphi)$ , with respect to each parametric component in  $\varphi$ . To comprehend the mechanisms of the approach, consider the analogy of an individual stuck at the top of a mountainous terrain (close resemblance to the input surface) on a foggy night, searching for a way down (descent). The heavy fog restricts vision, preventing the individual from seeing beyond the next step. By assessing all available information at the current position, a step towards the steepest decline is taken, with the aim of moving closer to being downhill. This process is repeated at each successive step in attempts to reach the bottom of the mountains (global minimum). Naturally, there is a strong preference for reaching ground-level with the fewest steps.

There are three-variants of the optimisation technique: Batch-Gradient Descent (BGD), Stochastic Gradient Descent, and Mini-Batch Gradient Descent (MBD). The BGD computes gradients with respect to network parameters, by considering the entire dataset, in contrast to SGD, which uses each pair of actual and estimated outputs (Ruder, 2017; Tiwari & Young

Chong, 2020). BGD ensures stable convergence for convex and non-convex surfaces, to global minimum and local minima, respectively. However, using an entire dataset for each parametric update, makes the approach slow and intractable, especially for large datasets that do not fit memory requirements. The SGD is faster and capable of performing frequent updates, resulting in improved convergence. It eradicates repeated calculations across similar targets, that cause computational inefficiencies in the BGD (Ruder, 2017). Alternatively, the frequency of updates can induce noise, and potentially result in poor quality loss function approximations (Bonetto & Latzko, 2020; Tiwari & Young Chong, 2020). Such noise can cause jumps to new and potentially better local minima, but overshooting could prevent the global convergence (Ruder, 2017)<sup>32</sup>.

The selection between SGD and BGD optimisation methods, comprises of a trade-off between accuracy and efficiency. The MBD incorporates the best of both and overcome the aforementioned issues. It splits the training dataset into fixed-size mini-batches, and performs parametric updates on each (Ruder, 2017; Tiwari & Young Chong, 2020; Izadkhah, 2022). Typically, MBD is preferred, and commonly referred to as SGD<sup>33</sup>, given the effectiveness of their combination. The term ‘stochastic’ refers to the behaviour of the learning process, where the actual gradient is replaced by that obtained from random samples during each iteration (Izadkhah, 2022; Zhou et al., 2022). A neural network is usually trained on large dataset, which may not be computationally feasible with BGD (Zhou et al., 2022). Instead, the SGD with mini batches can be employed for updating weights and biases. Variance during updates is reduced, providing stable convergence, without facing major computational problems (Bonetto & Latzko, 2020).

$$W_{l,q+1} = W_{l,q} - \eta \frac{\partial L(\varphi)}{\partial W_l} \quad (35)$$

$$B_{l,q+1} = B_{l,q} - \eta \frac{\partial L(\varphi)}{\partial B_l} \quad (36)$$

---

<sup>32</sup> Ruder (2017) also highlight that a slow decreasing learning rate, can make the SGD provide similar convergence as the BGD.

<sup>33</sup> From here onwards, MBD will be referred to as SGD, in abidance with common convention.

Once output is estimated and compared to the actual value, the iterative process of minimising the loss function, starts. The backpropagation step is based on the assumption that the parameters contribute to some portion of the error (Lavine & Blank, 2009). Consider the  $q^{th}$  iteration (epoch) of the network, during which  $W_{l,q}$  and  $B_{l,q}$ , representing the current weights and bias from the  $l^{th}$  hidden layer, are utilised for generating an output, and subsequently approximating the loss function. Weight and biases for the next epoch ( $q + 1$ ) are updated by the product of the learning rate ( $\eta$ ) and partial derivative of the loss function (gradient), as revealed in (35) and (36), respectively. In context to the SGD, the process is repeated across all batch samples of the training dataset, taking into account the average gradient (Izadkhah, 2022). Such iterations continue till a limit, or an acceptable value of the loss function has been reached.

The issue with using gradient descent optimisation is that good convergence is not guaranteed (Ruder, 2017). The learning rate plays a big role in the updating of parameters, as it explains how many units should be moved from the current position to improve the pace at which the loss function is minimised (Pothina & Nagaraja, 2023). Adhering to the mountain analogy, it would be equivalent to the range of visibility, which impacts the size of the step taken, in the selected direction. It is rather difficult to settle on an optimal learning rate value (Lek & Park, 2008; Lavine & Blank, 2009; Ruder, 2017). Smaller values are beneficial for narrowing down a minimum, however such convergence can be extremely computationally cumbersome. Larger learning rates, on the other hand, improve efficiency, but could cause the loss function to fluctuate around the minimum, or even diverge, resulting in unsatisfactory and instable learning. Scheduled learning rates (e.g. declining learning rates after certain number epochs) can improve performance, however, they have to be defined prior to training, and may fail to adapt to the properties of the dataset (Darken et al., 1992). Regardless, the same learning rate at a given epoch would apply for all updates, which might not be preferable, as those occurring rarely should receive larger gradients.

In addition to the issues regarding selection of the learning rate, the network is vulnerable to getting trapped in a suboptimal local minima (Ruder, 2017). This is a consequence of the presence of saddle points, where the dimensions move in opposing directions (Dauphin et al., 2014). In such instances, training can be excruciating, with no improvements, given the

small gradient. Recall the analogy of the mountain, where the fog limits vision to only the next step. If the individual were to get stuck in a saddle point, based on current information, it would be extremely difficult to find the next step with the steepest descent, to improve the current position. In such instances, the steepness of the successive step will slowly disappear, and the individual will be stuck at the same position. Returning to the training of neural networks, this is referred to as the vanishing gradient problem, as the size of updates, tends to 0, with each passing iteration. It is very difficult to escape such scenarios, hence large learning rates should be considered (Lavine & Blank, 2009), however, setting a desired schedule prior to training is not possible.

Adaptive learning rate algorithms, such as Adagrad (Duchi et al., 2011) and RMSProp (Tieleman & Hinton, 2012) improve the robustness of SGD optimisation by computing adaptive learning rates for each parameter. The former is favoured when gradients are sparse, whereas the latter is preferred in on-line and non-stationary settings. The Adaptive Moment Estimation (Adam) optimiser (Kingma & Ba, 2017) combines benefits from both methods, and eliminates the selection trade-off. Bias-correction and momentum are added to the RMSProp such that performances improve during the presence of sparse gradients (Ruder, 2017). It is favourable amongst practitioners in the domain of option pricing (Liu et al., 2019; Horvath et al., 2021). For conciseness in its mathematical representation, let parameters from the  $q^{th}$  epoch and  $l^{th}$  hidden layer be represented by  $\varphi_{l,q}$ , where  $\varphi_{l,q} = \{W_{l,q}, B_{l,q}\}$ .

$$\varphi_{l,q+1} = \varphi_{l,q} - \frac{\eta}{\sqrt{\hat{v}_q} + \epsilon} \hat{m}_q \quad (37)$$

$$\hat{m}_q = \frac{m_q}{1 - \beta_{1,q}}, \hat{v}_q = \frac{v_q}{1 - \beta_{2,q}} \quad (38)$$

$$m_q = \beta_1 m_{q-1} + (1 - \beta_1) \frac{\partial L(\varphi)}{\partial \varphi_{l,q}}, v_q = \beta_2 v_{q-1} + (1 - \beta_2) \left( \frac{\partial L(\varphi)}{\partial \varphi_{l,q}} \right)^2 \quad (39)$$

The gradients at each epoch,  $\varphi_{l,q+1}$  are obtained from computing adaptive learning rates with respect to each parameter. The estimates of first (mean) and second (variance) moments of gradients at each epoch, with exponentially decaying rates of  $\beta_1$  and  $\beta_2$ , are represented by  $m_q$  and  $v_q$ , respectively. Similar to RMSProp,  $v_q$  resembles an exponentially decaying average of squared gradients from previous iterations. Furthermore, Adam also considers  $m_q$ , which alike momentum, denotes the exponentially decaying average updates from



former epochs. The moments  $m_q$  and  $v_q$  are biased towards 0 during the initial time-steps, when  $\beta_1$  and  $\beta_2$ , were close to 1 (small decay rates). This is corrected by estimating of  $\hat{m}_q$  and  $\hat{v}_q$ , providing an adaptive learning rate, for gradient estimation<sup>34</sup>.

The Adam optimizer reduces chances of the network being stuck at saddle points and subsequently improves convergences. Nevertheless, the goal of the MLP network is not necessarily to reach the global minimum, as this would result in overfitting. The network would be overtrained, by only learning features that are specific to the training dataset (Lek & Park, 2008), instead of general characteristics regarding the relationship between input and output variables. Overfitting occurs due to excessive weight growth in overly flexible networks, or training on noisy data (Lavine & Blank, 2009). The MLP must obtain a solution for weights and biases, such that it can fit training and test data, without a loss of generalization. This is possible by paying careful attention during the determination of network architecture via hyperparameter tuning, and the creation of the training dataset with a wide distributional range of inputs and outputs. Hence, the *forward pass* stage of the network-based calibration process is crucial. It prevents overfitting by ensuring the employed optimal MLP architecture vividly understands the dynamics of the GSVJD model.

### 3.2.2 Step I: Forward Pass

The first stage of the two-step calibration process is the *forward pass*. It is a one-time computationally intensive procedure, during which the MLP network is trained offline. The aim of this step is to seek an MLP architecture, such that the dynamics of the GSVJD model are well understood by the neural network. This is crucial, as weights and biases from the hidden layer, are employed in the *backward pass* stage for calibrating parameters of the GSVJD from market data. In order to prevent overfitting, a detailed random search for hyperparameter tuning is conducted, with each architecture undergoing a k-fold validation.

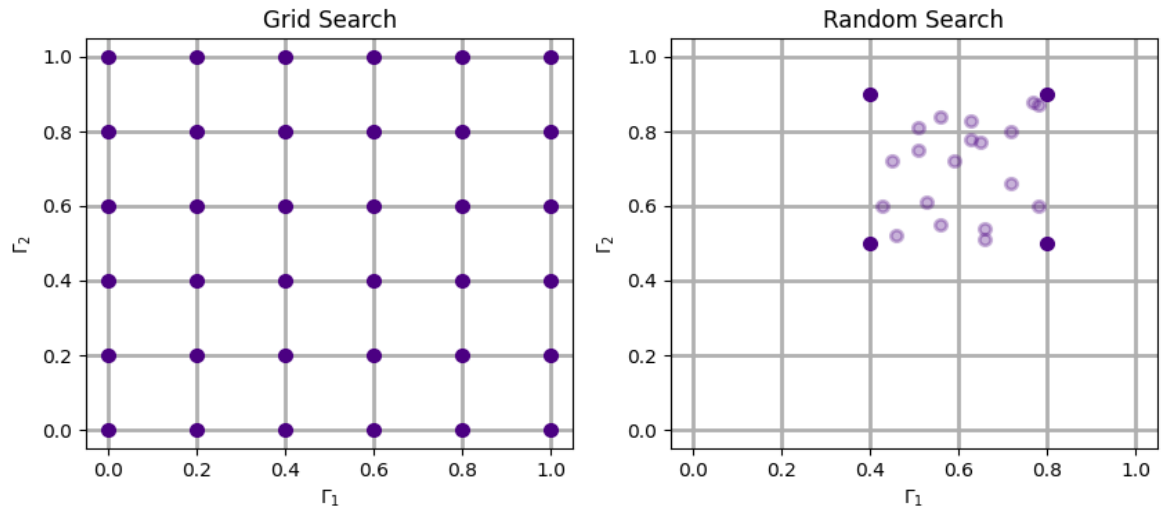
---

<sup>34</sup> The proposed default values of  $\beta_1$ ,  $\beta_2$ , and  $\epsilon$  for best performances by Kingma & Ba (2017) are met by the Keras library (<https://keras.io/api/optimizers/adam/>).

### 3.2.2.1 Random Search: Hyperparameter Optimisation

In their pioneering work within the domain of applying neural networks for estimating option prices, Malliaris & Salchenberger (1993) revealed the sensitivity of performances to the selection of hyperparameters and compositions of training data. Concerns over the latter, are diminished by constructing a synthetic training dataset. At the time of the groundbreaking work, computational powers were at their infancy, making it tedious to explore with varying ranges of hyperparametric combinations. Moreover, there was no formal theory or developed methodology for making such practices efficient. However, times have changed, the decades since have witnessed tremendous growth in computational capabilities, which enhanced research in the field of machine learning. Furthermore, a large quantity of literature is dedicated towards developing and improving hyperparameter optimisation to better application of machine learning models.

**Figure 3.3: Grid vs. Random Search**



*Note: The initial search space for configurations of hypothetical hyperparameters,  $\Gamma_1$  and  $\Gamma_2$ , with respect to the Grid (Left) and Random (Right) search methods.*

A comprehensive evaluation of various methods dedicated to this purpose was undertaken by Yang & Shami (2020). The most basic hyperparameter optimisation approach, Babysitting, also commonly referred to as the ‘*Trial and Error*’ or ‘*Grad Student Descent*’, is implemented based on guesses, and prior knowledge/experience. A model is created, and

the hyperparameter search process is only stopped when a satisfactory result is obtained, or when the ‘grad student’ approaches a deadline. The challenge is to possess sufficient knowledge such that the optimal combination is recognised, otherwise the process can be extremely long and exhaustive, with only a ‘deadline’ coming to the mercy of the practitioner. The desperate need for structure is met by the traditional Grid Search method (LaValle et al., 2004). It is a brute-force method that assesses all combinations within a specific grid. At the start, a large search space and step size is considered, which are narrowed based on previous results, till the optimal configuration is found. This approach is simple to implement, and easy to parallelize for reducing computational costs. However, the method can be extremely inefficient, over high-dimensional parametric spaces, due to the curse of dimensionality (Lorenzo et al., 2017).

The random search (Bergstra & Bengio, 2012) is similar to the grid search, but instead of testing across all hyperparameter combinations, it considers configurations, selected at random. A range of parametric vectors are determined, within which all candidates are trained (Yang & Shami, 2020). The exploitation of large parametric spaces with straightforward parallelization is possible, as a pre-determined range improves efficiency by ensuring that resources are not wasted on poorly performing combinations. However, even with expertise in determining a range for parameters, there might exist several unnecessary independent explorations over weak architectures. As with the case of the grid search method, the random search cannot independently direct itself to well-performing regions. An iterative algorithm for their hyperparameter optimisation, Bayesian Optimisation (Snoek et al., 2012) allows for such direction by using a surrogate model and acquisition function. The former aims to fit all observation points into the objective function for producing a predictive distribution. Various versions of surrogate models have been observed, using Gaussian Process, Random Forest, and tree-Parzen estimators (Yang & Shami, 2020). The latter selects points on the hyperparameter space by assessing those that have not been sampled, and those that are likely to obtain global minimum based on the posterior distribution. Similar benefits are observed with the application of Genetic Algorithms, which are favourable over the grid and random search methods, for large parametric spaces (Liashchynskyi & Liashchynskyi, 2019).

The Bayesian Optimisation and Genetic Algorithms approaches significantly improve computational efficiency when exploring large parametric spaces. However, they are not favourable for training under the *forward pass*. Considering the dependence on previous results for determining the direction to the global minimum, parallelization is prevented, translating to be a burden, rather than a boon for the practitioner. The various network architectures would be trained in parallel, with the search space of hyperparameter configuration being constructed in accordance with expert knowledge obtained from Liu et al. (2019), Horvath et al. (2021) and Büchel et al. (2022). The union of expert knowledge and the random search method, has been found favourable by Liu et al. (2019) for hyperparameter tuning, and would therefore be used to seek the optimal forward pass architecture. Selecting the random search is further justified by strong parallel processing abilities, and the availability of expert knowledge from past literature.

The optimal MLP architecture for the *forward pass*, is determined by an extensive search over 720 unique hyperparametric combinations. The aim is to discover an architecture that provides a balance between both computational efficiency and accuracy (Büchel et al., 2022). The output layer activation and optimiser for the loss function, are fixed to ‘*linear*’ and ‘*Adam*’, respectively. Furthermore, number of epochs, which resembles the number of iterations during the training process, are set to 200. This contradicts the preference for 8,000 and 5,000 iterations in Liu et al. (2019) and Büchel et al. (2022), respectively. The choice of 200 epochs is motivated by not wanting to overtrain the network, and reducing computational burden, with the scaling of inputs and outputs providing a further performance boost (Horvath et al., 2021). The random search will be conducted over six varying parameters: neurons, hidden layers, batch size, learning rate, (hidden layer) activation function, and  $L_2$  regularisation, with the range for each displayed in *Table 3.1*.

The primary architecture of a network is determined by its width and depth. The former refers to the number of neurons, whereas the latter alludes to the amount of layers. Both dimensions play a vital role in influencing the capacity and performance of the model. A neuron holds crucial information, regarding the understanding of a given relationship between input and output variables. Hence, increasing the width of the model would improve its ability to comprehend complex data patterns. On the other hand, a deep network is capable

of better understating the intricate structure (LeCun et al., 2015) or hierarchical representation (Tomasini & Wyart, 2024) in high-dimensional data.

**Table 3.1: Hyperparameter Random Search Range.**

*Note: Hyperparameter range comprising of 720 unique architecture combinations for the random search to determine the optimal neural network for the forward pass.*

Hyperparameter	Range
Neurons	[10, 20, 30, 40]
Hidden Layers	[2, 3, 4]
Batch Size	[2048, 4096, 8192]
Learning Rate ( $\eta$ )	$[1.0 \times 10^{-6}, 1.0 \times 10^{-2}, 10^{-1}]$
Hidden Layer Activation	ReLU, ELU
Output Layer Activation	Linear
Optimiser	Adam
$L_2$ regularization ( $\lambda$ )	$[0, 1.0 \times 10^{-3}]$
Epochs	200

Theoretically, according to the universal approximation theorem, neural networks with a single hidden layer can learn any function to a certain level of accuracy. However, as witnessed in McGhee (2018), a large number of neurons might be required, risking the ability to learn and generalize the function (Goodfellow et al., 2016). Increasing the width of a network, also worsens computational inefficiency, given the exponential rise in the number of parameters to be estimated. Shallow networks are notoriously subjected towards the curse of dimensionality, (Poggio & Liao, 2018; Wojtowysch & Weinan, 2020). Deep networks are capable of not only improving prediction accuracy, but also computational efficiency (Goodfellow et al., 2016). However, Zagoruyko & Komodakis (2017) highlight the problem of diminishing feature reuse arises from making the network too deep. It was discovered that a wider network with 16 layers outperformed one with 1000, and doubling of layers, only improved performances slightly, hence not justifying the increased computational burden.

Naturally, balance between the network depth and width is required. Inclusion of an additional layer, beyond the already existing four, does not better performances (Bayer et al., 2019). Therefore, only networks with 2, 3, or 4 hidden layers are considered for the random search. Regarding the width, to avoid computational inefficiencies, the range of neurons experimented upon vary between 10 and 40, at increments of 10. This selection is consistent with the number of nodes found for a two-step calibration process in Horvath et al. (2021). Another source of motivation stems from the operations of the *backward pass*, during which the network is merely a numerical solver for calibrating the GSVJD parameters. This task boils down to matrix operations, and if the network were to suffer from the curse of dimensionality, given the high-dimensional nature of GSVJD parameters, calibration would be slow and deter the practitioner from preferring the suggested approach.

Given the general architecture of the network has been established, focus can shift to a the remaining hyperparameters. A key stage of information processing within each neuron is the passage through a nonlinear activation function. It allows the network to understand complex data patterns, by transforming the output at each neuron within a certain range (Tharsanee et al., 2021; Subah & Deb, 2023; Ravikumar et al., 2024). Activation functions are the main reason behind successful application of neural networks to financial variables. As witnessed within the early works of Malliaris & Salchenberger (1993, 1996) and Hutchinson et al. (1994), the popular choice of activation function for neurons in the hidden layers was Sigmoid. Given its restrictive range between 0 and 1, when the neuron provides outputs at each tail, the function saturates, and gradient during backpropagation tends to 0 (Subah & Deb, 2023). This phenomenon is also known as the vanishing gradient issue and slightly improves by using the hyperbolic tangent (Tanh) function, given the extended -1 to 1 output range. However, at a slightly increased computational cost, the outputs when close to both tails make the network vulnerable to vanishing gradients (Tharsanee et al., 2021).

$$\text{ReLU: } \phi(x) = \max(x, 0) \quad (40)$$

$$\text{ELU: } \phi(x) = \begin{cases} x, & x \geq 0 \\ \alpha(e^x - 1), & x \leq 0 \end{cases} \quad (41)$$

The rectified linear unit (ReLU) improves computational efficiency and convergence by overcoming the vanishing gradient issue (Glorot et al., 2011). However, given the deactivation of neurons that provide negative outputs, the employment of the ReLU function,

can result in a *dead ReLU* situation, during which weights are updated without any activated neurons (Subah & Deb, 2023). Despite this issue and the proposal of several alternatives, it remains a popular choice amongst practitioners, even during the application of the two-step calibration approach (Liu et al., 2019; Büchel et al., 2022). The inclusion of negative activation output is required to improve upon the drawbacks of the ReLU function. Therefore, in abidance with Horvath et al. (2021), the exponential linear unit (ELU) is also explored during the random search. ELU considers negative outputs, by taking the exponential, and scaling them with a scaled parameter,  $\alpha > 0$ . It has been found to further bolster efficiency and accuracy (Clevert et al., 2016).

The selection of ReLU and ELU activation functions, assist in protecting against the vanishing gradient issue. To ensure smoother optimisation during training, it is vital to have a closer look at the learning rate ( $\eta$ ), batch sizes and  $L_2$  regularisation hyperparameter ranges. As discussed earlier, the learning rate explains the amount of units moved along the error surface, from the current position, such that the neural network function is optimized. There exists a trade-off, whilst assigning a value to  $\eta$ , as smaller rates slow convergence, whereas larger ones instable learning. Hence, a wide range of learning rates, between  $1.0 \times 10^{-6}$  and  $1.0 \times 10^{-2}$  are considered, at increments of  $10^{-1}$ . The Adam optimiser provides further improvements to training, by adapting learning rates, with respect to weights and biases during each epoch. This feature is a galvanization of the SGD<sup>35</sup> optimiser, which performs parametric updates by considering mini batches of training data.

A similar trade-off exists when selecting the batch size hyperparameter. The best form of generalization is achieved from considering a batch size of 1, however, given large dimensions of the training dataset, it would be computationally unfeasible to proceed in such a manner. Smaller batch sizes boost convergence towards flat sections of the error surface, allowing for improved generalization (Keskar et al., 2017), but worsen training time, and are prone to have high variances during gradient estimation (Goodfellow et al., 2016). Larger batch sizes, on the other hand, improve computational efficiency, and provide more accurate

---

<sup>35</sup> In this context, the stochastic gradient descent (SGD) is combined with the Mini-Batch Gradient Descent (MBD) approach, such that batches of training data are considered for parametric updates, rather than individual components.

gradients. They ensure information extraction features of a neural network are not underutilized. Although, at the cost of a loss of generalisation, given tendencies to converge to local minima (Keskar et al., 2017). Overall, selection of batch sizes boils down to the size of the training dataset, and available processing power. In regard to the latter, following a general rule of thumb<sup>36</sup>, it is recommended to opt for power of 2 batch size values (Goodfellow et al., 2016). Keeping in mind a requirement for a large training dataset, due to memory related issues, training would be unfeasible or highly inefficient on smaller values. Therefore, batch sizes of 2048, 4096 and 8192 were selected, to ensure a balance between gradient estimations, generalisation, and computational efficiency.

The final hyperparameter is the regularization term ( $\lambda$ ), during  $L_2$  regularisation<sup>37</sup>. The goal of regularisation is to prevent overfitting by reducing the generalisation error (Goodfellow et al., 2016). It modifies the loss function, by adding a penalty term, with respect to either the weights or biases values. To stay clear of any issues related to underfitting, the penalty term is obtained from the neural network weights. The concept focuses on taking the network from a potential overfitting regime, towards matching the true data-generating process. The penalty term comprises of a parameter,  $\lambda$ , which controls the degree to regularisation, and is tested for two values, 0 (no regularisation), and  $1.0 \times 10^{-3}$ . Given the implementation of  $L_2$  regularisation, with respect to weights, at each hidden layer, batch normalization and dropping rates were not considered for tackling overfitting<sup>38</sup>.

### 3.2.2.2 K-Fold Cross Validation

During the random search, each of the 720 unique neural network architectures are trained to determine the most optimal combination of hyperparameters for the *forward pass*. To further reduce overfitting, inspired by Liu et al. (2019), a k-fold cross validation training approach, with  $k = 3$  is considered, as illustrated in *Figure 3.4*. The mechanisms are

---

<sup>36</sup> The rule stems from the provision of GPUs, even on personal computers.

<sup>37</sup> In comparison with  $L_1$  regularisation,  $L_2$  regularisers use squared values within penalty terms, allowing them to be smooth and continuous, resulting in smoother, stable, and robust optimisation (Hastie et al., 2009).

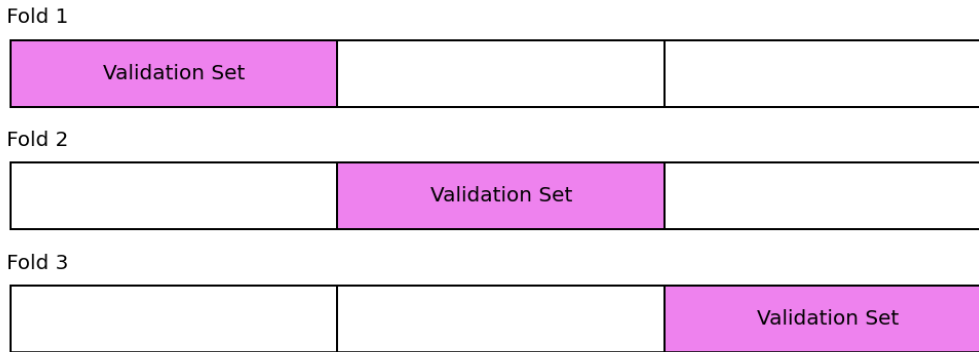
<sup>38</sup> Upon experimenting with Batch normalization and dropping rates, these hyperparameters were not considered during the training of the optimal *forward pass* in Liu et al. (2019).



straightforward, beginning with splitting the training dataset into  $k$  subsets. During each fold (experiment), a subset will be allocated for validation, whereas the remaining  $k - 1$  sets are used to train the selected network architecture. The final metrics for the given architecture are computed by averaging the errors over the  $k$ -folds. This procedure is repeated across all hyperparameter combinations, with each architecture being ranked based on the averaged performances.

A major advantage of the validation approach is its simplicity, however, given the size of the training dataset, and number of hyperparametric combinations, concerns regarding computational inefficiency arise. Additionally, error metrics obtained from each fold, can have a high degree of variability, especially when the dimensions of the training subsets are small (Lavine & Blank, 2009; Refaeilzadeh et al., 2009; Shobha & Rangaswamy, 2018). Small datasets, especially those constructed on historical data, could suffer from overlapping of training data. Hence, the training dataset should be large enough, such that when split into 3 equivalent subsections, and there is sufficient data for each fold.

**Figure 3.4: Illustration of a 3-fold cross validation.**



*Note: The training dataset is split into 3 components, with each acting as the validation set, in correspondence to the respective fold.*

Since there is a preference for synthetically generated training datasets, there are no concerns over the overlapping of training data. Finally, computational costs arising from training each of the 720 network architectures, over 3 folds, can be daunting. Nevertheless, a significant proportion of the costs are reduced by opting for 200 epochs. Moreover, the simplicity of

parallelizing the random search, along with access to GPU hardware and High-Performance Computing clusters, assists in overcoming this issue.

$$MSE = \frac{1}{N} \sum_{k=1}^N (y_k(x) - y_k(x|\varphi))^2 \quad (42)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{k=1}^N (y_k(x) - y_k(x|\varphi))^2} \quad (43)$$

$$MAE = \frac{1}{N} \sum_{k=1}^N |y_k(x) - y_k(x|\varphi)| \quad (44)$$

The aim of the training procedure is to determine optimal weights and biases, such that the Mean Squared Error (MSE) metric between estimated output,  $y_k(x|\varphi)$  and actual output,  $y_k(x)$  is minimised. Additionally, the Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) metrics are considered for assessing the performance of the each hyperparametric combination. The MSE metric averages the squared difference between  $y_k(x|\varphi)$  and  $y_k(x)$ , and places greater emphasis on larger errors. The RMSE penalizes larger errors to an even greater extent, with respect to smaller deviations. In addition, it is obtained within the same units, as the target variable, hence allowing for simpler interpretation. Both these metrics are suitable, when a greater emphasis is required on large errors. On the other hand, MAE measures the average absolute deviation between  $y_k(x|\varphi)$  and  $y_k(x)$ . In contrast to the MSE and RMSE, penalization of large errors is no different from that of smaller ones, making the metric robust to outliers. The average value of all three metrics will be taken, across the 3-fold of validation, for each architecture explored in the random search. While assessing the performances based on all these metrics, more emphasis will be placed on the RMSE, given its stellar ability to deal with outliers.

### 3.2.2.3 Training and Validation: Optimal Architecture

The optimal neural network hyperparametric combination is determined by a random search, conducted over 720 architectures. Each unique combination undergoes a 3-fold cross-validation to improve robustness and reduce the likelihood of overfitting. The main goal of

the training procedure is to formulate weights and biases through iterative efforts at minimising the loss function between actual and estimated output values. The optimal computation of these parameters is crucial, as they are responsible for storing the extent to which a network understands the relationship between input and output variables. It is this level of understanding that will be stored and exploited in the backward pass during rapid online GSVJD parametric calibrations, for real time bubble detection.

Simply working with a training dataset for hyperparametric optimisation is not sufficient to prevent overfitting. Recall, it is not necessarily favourable to reach the global minimum of the error function surface during training. If so, the network would have learned to estimate the outputs from the training data, rather than understanding the general dynamics/characteristics of the relationship between input and output variables. In such situations, when tested on unseen market data, the neural network will suffer from overfitting, which resembles joint-hypothesis concerns during bubble detection. Therefore, it is advisable to assess the performance of the network, at each iteration during training, on an unseen dataset.

The computations of weights and biases by minimising the loss functions with respect to both datasets, provides validation. These datasets are known as '*Validation Datasets*' and have been utilised to prevent machine learning models from overfitting. They are constructed by splitting and separating a proportion of the training data, at random. By doing so, the validation dataset would comprise of unseen data, with similar characteristics to that used for training. It is also conventional among scholars to split the training dataset into three components for training, validation, and testing. As noted, the first two are employed during the training procedure, whereas the latter is utilised to test the performance of the trained network. Intuitively, it conducts a similar practice (on different data) as validation on the final epoch. Given the implementation of the K-fold cross validation, and additional assessments conducted during the backward pass for selecting the optimal measure, testing datasets are not considered. Performances of all neural network architectures in the random search during training (and validation), are evaluated on average RMSE across the 3-fold cross validation.

### 3.2.2.4 Input and Output Variables

The primary focus is to improve computational efficiency of calibrating the GSVJD model parameters from market observations. The neural network is proposed as a numerical solver in place of the computationally cumbersome Monte Carlo simulations in the three-step approach (Fusari et al., 2024). In the three-step approach, the GSVJD model was calibrated to daily market put option implied volatilities to overcome the joint-hypothesis issue. Recall, put option have bounded payoffs, due to which their prices cannot reveal bubbles. The accurate calibration to market put implied volatilities would ensure that the fundamental and market values of the option are aligned, consistent with the scenario during which asset price bubbles are absent. The obtained parameters would correctly resemble the current market regime (ELMM) and can be utilised for pricing call options for bubble detection.

**Table 3.2: Inputs and Outputs of the Forward Pass.**

Inputs ( $x$ )	Output ( $y$ )
$\theta = \{V_0, \bar{v}, \kappa, \sigma_v, \rho, p, \mu_J, \sigma_J, \lambda_J\}$	$\sigma_P(\theta, \Omega)$
$\Omega = \{mon = (\mathbb{K}/F), \tau\}$	

For a moment, consider only the computation of implied volatilities from Monte Carlo simulations of the GSVJD. The model is fed inputs in the form of 9 parametric values, the underlying and strike price, risk-free rate, dividend yield and the option's time to maturity, to compute price. The price is then passed through an iterative solver to generate the corresponding implied volatility. The entire purpose behind being meticulous with each aspect of the training stage was to ensure that the neural network learns such dynamics of the GSVJD model, for computing implied volatilities. The network directly estimates the output from inputs, making the need of two solvers for the separate generation of prices and implied volatilities redundant. This provides a massive computational boost during the optimisation process, considering random parameters are inserted into the model for estimation, at each iteration, in attempts to minimise the objective function, with respect to market observations. The input and output characteristics of the *forward pass* are revealed in Table 3.2.

Universally, during calibration procedures, parameters are considered as inputs, whereas the variable needing estimation for comparison to actual observations, the output. Similarly, the network is fed all components of the GSVJD parametric vector,  $\theta$  and  $\Omega = \{mon \text{ and } \tau\}$ , as inputs, such that put implied volatilities,  $\sigma_p(\theta, \Omega)$ , are generated as the output. Training is conducted to ensure that the nonlinear relationship between these variables is well understood by the neural network. The selection of  $\sigma_p(\theta, \Omega)$  as the output variable ensures that calibration during the backward pass would be conducted, to overcome the joint-hypothesis issue during bubble detection.

Recall, the calibration of stochastic volatility models can be conducted using two structural forms: *pointwise* or *grid-based*. In the latter, implied volatilities from options over a wide range of strike prices and maturities are considered in a fixed-grid structure. The output layer comprises of multiple neurons, each resembling a location on the grid, corresponding to a unique combination of strike prices and maturity. This allows for simultaneous computations of all implied volatilities that have been considered by the grid. Given the economic resemblance of each output node, the need to include strike prices and maturities as inputs is made redundant. However, the output nodes are fixed, and those options with differing characteristics cannot be examined. This is a major drawback, especially during the application of the three-step approach for bubble detection.

When calibrating GSVJD parameters, Fusari et al. (2024) use daily options corresponding to maturities with the highest cumulative volume. The maturities and number of options selected, vary over each day. In order to accommodate these changes, the output layer structure would require altering, which translates to re-training the network. The computational costs of such a practice significantly outweigh the benefits obtained from reducing the dimensions of the input layer, and simultaneous estimation of implied volatilities. Alternatively, given the computational boost from the deep calibration framework, it would be possible to include all daily put options during calibrations, as recommended by the current study. A grid can be formed over the components of daily implied volatility surfaces, constructed over a fixed range of strike prices and maturities see Horvath et al. (2021). However, these surfaces would require using interpolation and extrapolation techniques, which would induce an additional error during the calibration

procedure. Moreover, calibrating parameters to all traded options on a given day would also not be possible, given the varying values of strike prices, maturities, and number of options across time.


In the *pointwise* structure, the output layer comprises of a single node, allowing the network to only estimate the implied volatility of one option. This adds to the calibration time, when applied to several contracts during an iterative optimisation process. Furthermore, the single output node does not hold any economic resemblance with respect to the strike price and maturity of the corresponding option. Hence, these two characteristics must be fed as inputs, increasing the dimension of the first layer. Ironically, the single output node is the main advantage of this structure, as it allows for a great amount of flexibility, which is crucial for bubble detection, when considering daily options over a large sample period. Since the strike price and maturity are inputs, minor alterations in their values allow for the consideration of any desirable option contracts. Therefore, the *pointwise* structure can simply be utilised to accommodate contracts from, both the daily highest cumulative volume maturities, and entire price surface. This superior flexibility creates a preference for this structural form, when utilising the two-step calibration approach on the GSVJD model. Hence, the neural network will comprise of a single output node, with two additional input nodes, corresponding to moneyness ( $mon$ ) and maturity ( $\tau$ ).

Computing implied volatility requires the following inputs: underlying, option and strike prices, along with the risk-free rate, maturity, and dividend yields. Therefore, in addition to the GSVJD parameters, the *pointwise* neural network would need inputs that can capture impact of these variables for the accurate generation of implied volatilities. Ruf & Wang (2020) reveal that the underlying and strike price can be considered as either individual inputs, or collectively in the form of moneyness. Scholars have been found to prefer the latter, as it reduces the likelihood of overfitting, in addition to decreasing the dimensions of the input layer. Overall, the generalisation ability of the neural network is improved. Therefore, moneyness,  $mon = (K/F)$ , consistent with the requirements of the *pointwise* approach, is included as an additional input, but with a slight alteration. Instead of computing it with respect to the underlying price, the forward price,  $F = Se^{(r-q)\tau}$ , is utilised. This is common industrial practice, given the greater consistency with the Black and Scholes

formula. Additionally, the forward price accommodates the impact of the risk-free rate and dividend yields, over the remaining life of the option.

An option contract is defined by its strike price and time to maturity ( $\tau$ ). The *mon* input variable takes the former characteristic under consideration, which leave the requirement of defining the latter as the final network input. In contrast to the risk-free rate and dividend yield, maturity is specific to each contract and has a significant impact on the price of the option. Hence, despite its involvement in computing moneyness, maturity must be accounted for as a separate input. This ensures that implied volatilities are well defined and estimated, in accordance to their unique moneyness and maturity combinations. Therefore, the neural network will be trained, such that it understands the dynamics of the GSVJD model by learning the complex relationships between inputs,  $x = \{\theta, mon, \tau\}$  and put option implied volatility output,  $y = \sigma_p(\theta, \Omega)$ .

*Figure 3.5: Illustration of the Forward Pass.*



*Note: The neural network learns dynamics of the GSVJD model to estimate normalised output,  $\sigma_p(\theta, \Omega)'$ , from scaled inputs:  $\theta'$ ,  $mon'$ , and  $\tau'$ .*

The *forward pass* step seeks to best learns the features of the GSVJD model. As illustrated in *Figure 3.5*, this is achieved by feeding it the following inputs,  $x = \{\theta, mon, \tau\}$  to compute implied volatilities,  $y = \sigma_p(\theta, \Omega)$ . Observe, the input variables are scaled,  $x' = \{\theta', mon', \tau'\}$ , prior to entering network for the computation of normalised output,  $y' = \sigma_p(\theta, \Omega)'$ . This operation boosts overall estimation and performance of the network and prevents it from being stuck at local minima. The selected inputs add strong economic

interpretation to the estimation of implied volatilities during the forward pass. However, it must not be forgotten that a neural network relies heavily on data for comprehension of the relationship between input and output variables. In such instances, Büchel et al. (2022) warn the practitioner about sensitivity of networks towards diverging dimensions of input variables, especially with respect to parameters of stochastic volatility models. If not carefully dealt with, the neural network would face overfitting issues. This would hamper both, the accuracy of estimating implied volatilities in the *forward pass*, and calibration of parameters in the *backward pass*.

$$x' = \frac{2x - (x_{max} + x_{min})}{(x_{max} - x_{min})} \quad (45)$$

$$y' = \frac{y - y_{mean}}{y_{SD}} \quad (46)$$

The solution lays with scaling input and output variables, based on a pre-defined range. Inspired by Horvath et al. (2021) and Büchel et al. (2022), all input variables are scaled as per (45), ensuring  $x' \in [-1, 1]$ . Whereas output is normalised in (46) by subtracting the mean ( $y_{mean}$ ) and dividing by standard deviation ( $y_{SD}$ ). These operations make the loss function more symmetrical, allowing for easier and faster convergence, such that the neural network overcomes the vanishing gradient problem. Additionally, Büchel et al. (2022) reveal the scaling of model parameters to be beneficial during the calibration step. The guess range of the optimizer is restricted between -1 and 1, allowing for fast and more accurate calibrations.

It must be documented that all scaling operations of input and output variables are conducted with respect to the training dataset. Therefore, during inverse-scaling to obtain calibrated parameters and implied volatilities, all estimates will be within the range on which the network was trained. This further appreciates using a synthetic dataset during training. In historical datasets, all parameters are scaled according to observations from past market scenarios. If the backward pass were to stumble upon an extreme regime, the calibrated parameters would not be reliable, given restrictions from scaling bounds. When employing a synthetic dataset, the practitioner controls the range of scaling bounds, by choice of including various scenarios, ensuring the benefits of such operations are experienced to the



fullest. On this note, the random search for hyperparametric optimisation can be conducted with each combination undergoing a 3-fold validation. Upon completion the *forward pass* step would present an optimally trained network, the hidden layers of which would be frozen, and implemented for calibrating GSVJD parameters during the *backward pass* step.

### 3.2.3 Step II: Backward Pass

Consider the one-step approach, which calibrates parameters of an asset pricing model, by estimating them as outputs. This would allow for the simultaneous calibrations from market observations over a vast time period, in a matter of seconds. However, given that parameters are estimated as outputs, and not calibrated to fit market data using an optimizer, the approach is highly vulnerable to overfitting and robustness related issues. The one-step method trains the network to minimise the error function between actual and estimated parameters, by taking market observations as inputs. Hence, the dynamics of the chosen model are not understood, but rather the network learns to fit the training dataset.

Additionally, there is a need to utilise historical data for training, which results in the network performing poorly when applied to unseen time periods, especially those with extreme market scenarios. The performances of the network are not captured by the parametric output, but instead by the ability of such estimations to compute option prices/IMPLIED volatilities. A separate validation step is required, which reduces efficiency. If skipped, regulatory concerns regarding the black box' nature of neural network approximations would arise. In regards, to bubble detection, the method is vulnerable to the joint-hypothesis issue.


The desire for real-time validation to overcome the joint hypothesis issue during bubble detection is satisfied by the two-step approach. The *backward pass* is dedicated to pairing an optimizer with the trained neural network (from the *forward pass*) for calibrating parameters to market put option IMPLIED volatilities. The additional optimisation step, though inefficient, adds extra validation to retrieved parameters, and relaxes regulatory concerns over the 'black box' nature of neural network approximations. A well-trained neural network from the *forward pass* step is crucial for assuring a strong and reliable backward pass. It

ensures that the neural network has developed a concrete understanding of the GSVJD model dynamics.

### 3.2.3.1 Backward Pass: Framework

During the *forward pass*, networks were trained to estimate put option implied volatilities,  $\sigma_p(\theta, \Omega)$  from receiving the GSVJD model parameters ( $\theta$ ), moneyness ( $mon$ ) and option maturity ( $\tau$ ), as inputs. The trained hidden layers comprise of weights and bias parameters, at each node, that store crucial information regarding the comprehension of the GSVJD model. This information builds the core foundation of the *backward pass* stage, allowing the network to transition from an estimator for implied volatilities, into a simple, yet extremely rapid numerical solver for parametric calibrations from market observations.

**Figure 3.6: Illustration of the Forward and Backward Pass.**



*Note: During the forward pass (top), the neural network is trained to estimate normalised output,  $\sigma_p(\theta, \Omega)'$ , from scaled inputs:  $\theta'$ ,  $mon'$ , and  $\tau'$ . Once trained, the hidden layers are frozen and carried to backward pass (bottom), for calibrating  $\theta'$ , from market put option implied volatilities,  $\sigma_p(mon, \tau)'$ . The remaining inputs,  $mon'$ , and  $\tau'$ , are obtained from the market, set in correspondence to  $\sigma_p(mon, \tau)'$ .*

The trained hidden layers from the *forward pass* are frozen and carried on to the *backward pass*, for parametric calibrations, as illustrated in Figure 3.6. In this phase, the neural network is paired with an optimizer, such that the inverse problem of calibration is tackled

with respect to the given market output. The network is already trained, so weights and biases within its hidden layers are known. Hence, the *backward pass* becomes deterministic, with only simple matrix multiplications required to map parameters to market observations (Büchel et al., 2022). Furthermore, it can be conducted online, at high levels of computational efficiency, such that the time for calibrations is reduced to a matter of seconds. This ensures that the frequent recalibration requirements of the high-dimensional GSVJD parametric vectors, cater to the consistently changing nature of markets. Therefore, improving the tractability and overall preference of the practitioner, for utilising the three-step approach for real-time bubble detection.

The calibration operations are conducted in the following manner. At first, market put option implied volatilities are introduced as targeted outputs in the loss function. Alternatively, they can be interpreted as fixed components of the output layer, that are connected to the loss function. The optimizer works towards minimising the loss function, by feeding parameter combinations into the network, for estimating implied volatilities. These combinations are received in the input layer, at the 9 nodes corresponding to each component of the scaled GSVJD parametric vector,  $\theta'$ . Hence, these 9 input nodes, that were previously fixed during the *forward pass* training, become learnable. The remaining 2, designated to  $mon'$ , and  $\tau'$ , are fixed to relevant market information, corresponding to the observed implied volatility at the output node.

At each optimising iteration, the parameter guess enter the input layer, and along with the 2 observable inputs, get transferred to the hidden layer. Recall, all weights and biases have already been determined during the *forward pass*, hence, a simple series of matrix multiplications, within each layer provides an output for the next. At the final layer, the network estimates an output, which is then fed into the loss function for computing the error, with respect to its market counterpart. The iterative optimisation continues, with fresh parametric guesses, until the loss function is minimised or tolerance with respect to the error/number of trials, is reached. Essentially, the forward propagation phase of the neural network is exploited for rapid computations to estimate implied volatilities, to allow for superior efficiency during optimisation of the loss function.

The parameters calibrated from the *backward pass* would be obtained in the scaled form, and require an inversing operation, to be utilised for pricing and subsequent bubble detection. The scaling of the parameters is advantageous, as it improves the efficiency and convergence quality of the optimizer. Notice, the structure of the network comprises of a single output node, given preference for the flexibility provided by the *pointwise* approach. This indicates that during calibrations, implied volatilities are estimated one at a time, rather than simultaneously as witnessed during the selection of the *grid-based* approach. However, given its lack of flexibility, the latter is not an ideal candidate for real-time bubble detection. Though the approach is marginally more inefficient, the selection of the differential evolution optimizer, with superior parallel computing abilities when paired with neural networks, allow to overcome this issue.

### 3.2.3.2 Backward Pass: Optimisation

The selection of the optimisation algorithm is crucial for achieving accurate calibrations of the GSVJD model. The Levenburg and Marquardt algorithm is a popular choice of gradient-based optimisation methods, amongst Hernandez (2016), Bayer & Stemper (2018); and Bayer et al. (2019), for calibrating stochastic volatility models. Whereas Differential Evolution is a gradient-free optimiser. It has also been employed by Liu et al. (2019) during the calibration of the Heston and Bates option pricing models. The abilities of both techniques, when paired with a deep two-step approach for calibrating implied volatility surfaces, were compared by Horvath et al., (2021). Though the Levenburg and Marquardt approach was found to be more balanced in regard to accuracy and computational efficiency, it was inferior to the differential evolution optimiser. In general, gradient-based methods are quicker at converging, however, provide sufficient, rather than necessary conditions for doing so. Hence, they struggle when applied to non-convex functions, requiring the assistance of superior initial guess. Alternatively, gradient-free methods are relatively slower, but make up for the computational inefficiency by being capable of finding a global solution, regardless of the characteristics of the objective function.

Gradient-based approaches struggle with non-convex functions, due to the presence of multiple local-minima, and performances could be limited to the knowledge of the

practitioner. It is common to utilise parameters from the previous time period, as an initial guess, however, this practice might not be suitable for volatile phases. The differential evolution optimiser does require an initial guess, however due to its stochastic nature, it is not reliant on the quality of initial values provided to find the global minimum. The algorithm generates initial values at random, hence, it has no issues dealing with non-convex functions. Furthermore, Liu et al. (2019) reveal a two-fold speed achieved from pairing neural networks with differential evolution.

These benefits of the differential evolution algorithm are explained with respect to GSVJD calibrations. First, implied volatilities, corresponding to a set of parameters, are computed at once. Even though there is one output node, all input sets, each comprising of the parameters that prevail in the rest, and a unique moneyness-maturity combination, are fed into the network simultaneously to estimate implied volatilities. Second, during the *backward pass*, at each iteration of the optimiser, the generated parameter candidates enter the network at once, resulting in the estimation of all analogous output simultaneously. In contrast, Monte Carlo simulations estimate output for each candidate, individually. Büchel et al. (2022) suggest using parameters from the previous day to further enhance efficiency, however, this is avoided given their unsuitability to volatile periods/sudden market crashes.

$$RMSE_{\sigma,t} = \sqrt{\frac{1}{N_p} \sum_{k=1}^{N_p} (\sigma_P(\theta, \Omega)' - \sigma_P(mon, \tau)')^2} \quad (47)$$

$$SSE_{\sigma,t} = \sum_{k=1}^{N_p} (\sigma_P(\theta, \Omega)' - \sigma_P(mon, \tau)')^2 + \lambda \quad (48)$$

Differential evolution is selected for optimizing the loss function, given its superior parallel computation abilities, when paired with a neural network, and independence from the quality of initial guess, regardless of the loss function. There are two loss functions considered, separately, during the *backward pass*;  $RMSE_{\sigma,t}$  and inspired by Liu et al. (2019), sum of squared error ( $SSE_{\sigma,t}$ ) with a penalty term,  $\lambda = 1 \times 10^{-6}$ . The *backward pass* not only to calibrates the GSVJD model at high speeds but will also be used to select the best architecture for bubble detection. The latter alludes to selecting top performing architectures from the *forward pass* and employing them to calibrate GSVJD parameters from market data. The parameters from the best performer will be selected for computing the fundamental

value of call options, and subsequently detecting bubbles. Therefore, the evaluation of these architectures, using two loss functions, adds robustness to the selection process. As observed in (47) and (48), the Differential Evolution optimiser works towards minimising the both loss functions, with respect to normalised implied volatilities. Upon reversing of the scale, the performances of each architecture are assessed using  $MSE_{\sigma,t}$ ,  $RMSE_{\sigma,t}$ , and  $MAE_{\sigma,t}$  metrics, see (42), (43), and (44).

### 3.3 Summary and Discussion

The foundation of the three-step approach is based on the local martingale theory of bubbles, which reveals that a risky asset can exhibit bubbles, when an incomplete market satisfies the NFLVR and ND conditions. Due to their bounded maturity, options exhibit type III bubbles, which under the ND condition, share a linear relationship with exuberance in the underlying asset. Put options have bounded payoffs, and cannot display a bubble, and hence their market and fundamental prices must align. However, call options have unbounded payoffs, and are capable of revealing bubbles, as investors are willing to purchase them for the sole purpose of reselling in the future, at a higher price. The three-step approach examines option prices and exploits the linear relationship between call option and underlying price bubbles. Options reveal expectations regarding the future value traders are willing to pay for the underlying asset, and hence apprehend the forward looking nature of bubbles. Furthermore, type III bubble address short term trade strategies, intended for this purpose, can be identified using real-time data, without relying on historical prices or backward-looking windows. It is crucial for the option pricing methodology to highlight strict local martingale tendencies in the underlying. Hence, the GSVJD model is employed for pricing options.

The joint hypothesis problem is overcome by calibrating the GSVJD model to market put option implied volatilities and using those parameters to estimate the fundamental value of call options. Accurate calibration ensures that the fundamental and market values align, such that put option prices do not display bubbles. Finally, a statistical test is employed for detecting call options bubbles, which under the ND condition would signal the presence of exuberance in the underlying price. The GSVJD process is impressive, as it is capable of capturing strict local martingale tendencies, mean-reverting volatility, and price jumps.

However, it lacks a closed form solution, hence forcing the practitioner to rely on Monte Carlo simulations. The computational cumbersome nature of such numerical solvers make it unfeasible to consider the entire daily option price surface, restricting the application to only a single volatility smile. In order to overcome this issue, the application of neural networks was suggested.

Neural networks are a set of interconnected nodes, that mathematically mimic the biological nervous system. According to the universal approximation theorem, they can estimate any function, to a certain degree of accuracy. The advancements in computational software and hardware over the past couple of decades have allowed the machine learning approach to price thousands of derivative contracts, at rapid speeds. It is this property, that is exploited during proposed calibration of the GSVJD model. There are two methods for calibrating such asset pricing models. First, the one-step approach, which estimates parameters as outputs of the networks. Second, the two-step approach that trains the neural network to estimate market observations and then pairs the model with an optimiser for calibration. The optimisation step allows for real-time validation, which is not only satisfactory for regulators, but also beneficial for practitioners, when attempting to overcome any potential instance of a joint-hypothesis issue during bubble detection. Therefore, the two-step approach is favoured for the motives of this research.

A two-step neural network calibration framework was proposed. It comprises of an MLP, which estimates implied volatilities from taking parameters of the GSVJD, moneyness and option maturities, as inputs. The first stage, *forward pass*, focuses on training the network to learn the dynamics of the GSVJD model for computing implied volatilities. An extensive random search, with a 3-fold cross validation over 720 hyperparameter combinations is initiated to determine the optimal neural network architecture for the *forward pass*. The second stage, *backward pass*, is engineered to calibrate the parameters from market observations.

The *backward pass* pairs the pre-trained network architecture from the *forward pass*, with an optimizer for calibrating GSVJD model parameters from market data. The trained hidden layers of the network store vital information regarding the understanding the GSVJD model

dynamics. The hidden and output layers are frozen, with the latter being allocated market data on put option implied volatilities. Additionally, input nodes designated to moneyness, and maturity values that correspond to implied volatilities in the output node, receive market observations. The only neurons in the network that are learnable are those from the input layer, assigned to each GSVJD parameter. Since hidden layers are already trained, the network solely operates as a numerical solver. The superior parallel processing abilities of the Differential Evolution in combination with neural networks, are heavily exploited. Furthermore, in contrast to the gradient-based techniques, there is no reliance on the quality of initial parameter guesses during optimisation.

The obvious benefit of the neural network framework is the massive computational boost received from making Monte Carlo simulations redundant during the calibration phase, allowing for rapid real-time bubble detection. Furthermore, the *pointwise* structure comprises of a single output node, in contrast to the fixed grid-based alternative. It provides the framework with flexibility to consider a varying number of options during daily calibrations, rather than being restricted to a fixed maturity-moneyness combination. Additionally, in contrast to the one-step approach, the additional optimization step enables instantaneous validation, which is crucial for overcoming the joint-hypothesis issue. Finally, training using synthetic data makes the network robust to the ever-changing fickle nature of markets. In the next chapter, calibration performances of the GSVJD model, using Monte Carlo simulations, and the optimal neural network architecture are compared on S&P 500 index data. Though a boost in efficiency by a large scale of magnitude is appreciated, it is also vital to ensure calibration accuracy is not sacrificed.





# Chapter 4:

# Model Estimation

The following chapter focuses on calibrating the GSVJD model, using both Monte Carlo simulations, and the proposed deep neural network calibration framework. The calibrations are conducted and compared on daily S&P 500 index data, between January 2019 and December 2022. The benefits of detecting bubbles from option prices, using the three-step approach (Fusari et al., 2024) have been well documented in the previous two chapters. The method incorporates a generalised stochastic volatility jump diffusion model (GSVJD) for option pricing, which captures strict local martingale tendencies in the underlying price process. The sophistication of the model allows it to address characteristics of asset prices to a greater extent but at a major computational cost. Hence, the application of neural networks was suggested. However, prior to indulging in the estimation of the deep calibration framework, the GSVJD model is compared to other stochastic volatility processes, to examine the tradeoff between efficiency and accuracy. The comparison reveals the superiority of the GSVJD model. Hence, justifying its consideration as the benchmark for training the neural network. Next, after intensive training and testing of the two-step deep

calibration framework<sup>39</sup>, it is revealed to outperform the benchmark, with regard to both accuracy and efficiency. Upon providing a boost in computational estimation and accuracy, the deep neural network calibration framework is applied for estimating the daily GSVJD parameters from the entire option price surface.

## 4.1 Research Hypotheses

The three-step approach, despite overcoming issues associated with traditional methods, can be unfavourable amongst practitioners given the poor efficiency and tractability associated with GSVJD calibration. The GSVJD excels at capturing forward looking expectations of market participants, however, to strike a balance between efficiency and robustness, Fusari et al. (2024) restricted calibration to the most liquid volatility smile. As highlighted in 2.3.3, using the entire surface of options enhances the ability to capture market characteristics that are crucial for bubble detection. This is a direct consequence of calibrating GSVJD parameters from various option maturities, rather than just the one. Therefore, a deep calibration framework is proposed to answer the following research question: *How to improve the efficiency and tractability of the three-step approach and subsequently extend its application for bubble detection?*

The motivation behind the research question stems from the need for an accurate and robust bubble detection method, which is also efficient and tractable enough to meet practical requirements. Issues regarding efficiency and tractability arise from the required GSVJD calibrations; hence the proposal for employing neural networks. However, prior to proceeding with the construction of the deep calibration framework, this research explores a potential solution in the form of an alternative stochastic process. It is crucial that the selected process must capture strict local martingale tendencies. Therefore, the GSVJD model is compared to the stochastic volatility process from Andersen & Piterbarg (2007). Both processes possess the ability to address strict local martingale tendencies but only the former considers price jumps. Furthermore, the Heston (1993) and Bates (1996) models, which do not capture strict local martingale characteristics, are considered to reveal the importance of this property within stochastic processes. It is hypothesised that the GSVJD

---

<sup>39</sup> Refer to *Section A*, in the Appendix for more details.

model will outperform its counterparts, however, since it requires the greatest number of parameters to be calibrated, the process will also be the least efficient. Consequently, this would create a demand for the deep calibration framework in order to ensure substandard stochastic processes are not utilised for bubble detection, for the mere reason of efficiency.

The deep calibration framework is trained to learn the dynamics of the GSVJD model and compared to the benchmark calibrations that require Monte Carlo simulations. It is hypothesised that the former would significantly boost efficiency, without sacrificing accuracy. The latter component is vital for overcoming the joint-hypothesis issue during bubble detection. Moreover, the improved efficiency and tractability of GSVJD calibrations, enable parametric estimations from all daily traded put options. Consequently, the application of the three-step approach can be extended to detecting underlying price bubbles from examining various call option maturity groups, and even from exploring the lifetime of selected options. This is possible since parameters obtained from calibrating to daily option surfaces are applicable for pricing call options with various maturities. The extended application of the three-step approach, using the deep calibration framework, is tested over daily S&P 500 options data from 2019 – 2022.

The selection of the time period plays a crucial role in the examination of exuberance. During the market crash in March 2020, the S&P 500 experienced 4 circuit breakers. In accordance with the NFLVR assumptions, the shorting constraints, artificially ensured that market prices remained above the fundamental value. Despite the dire state of the economy, the S&P 500 erased its losses and reached record highs by August 2020, with strong support from the technology sector. The recovery was fuelled by large monetary and fiscal provisions, which created an imbalance, resulting in severe inflationary pressure over the next year. However, the S&P 500 continued to surge, even throughout 2021. Therefore, this research hypothesises the index to reveal bubbles between 2020 and 2021, especially during the crash in March 2020, and subsequent recovery period. Furthermore, the difference between detecting bubbles from calibrating to the most liquid smile, and parametric estimation from the entire surface, are explored. Initially, daily parameters obtained from the most liquid put options, and entire surface are used to price the most liquid call options for bubble detection. Next, parameters estimated from the entire surface are used explore bubbles across various call

option maturities, and the lifetime of certain selected contracts. This extended application of the three-step approach is examined over the S&P 500 index, in accordance with the hypothesised presence of exuberance. Nevertheless, prior to testing the various hypotheses presented in this section, this chapter proceeds with a discussion of the relevant market data, and applied cleaning and filtering techniques, required for calibrations.

## 4.2 Market Data: S&P 500 Index

In the three-step approach, Fusari et al. (2024) highlight the importance of calibrating to market put option observations, in order to overcome the joint-hypothesis issue associated with bubble detection. Hence, calibration comparisons conducted in this chapter are carried out on market put options data. The following section details various cleaning and filtering procedures utilised to create the required dataset. Given computational inefficiencies associated with Monte Carlo simulations, all comparisons of the different stochastic processes, are conducted on the most liquid daily volatility smile. This approach provides a balance between efficiency and robustness, making it well suited for examination over a large time period. The calibration comparison, along with empirical analysis of bubble detection is conducted over daily S&P 500 Index European-styled options, between January 2, 2019, and December 30, 2022. The study period spans over 4 years, comprising of 1,000 trading days, with all market-related data being collected from Refinitiv Eikon.

The selection of the study period allows for examining the existence of bubbles prior to, during and after the COVID-19 induced market crash in March 2020. As a direct consequence of the economic impacts of the global pandemic, the S&P 500 experienced massive volatility, with fear fuelling massive selloffs. Consequently, circuit breakers were triggered on March 9, 12, 16, and 18, 2020. These events are relevant for the empirical analysis of detected bubbles, as given the implementation of such shorting constraints, the market price is prevented from collapsing to the fundamental value. Furthermore, the study period allows for analysing the robustness of various stochastic process, across different market scenarios. It is crucial to utilise a stochastic process, that can accurately calibrate to various market scenarios, including those with extreme fluctuations. This would be

beneficial for overcoming joint hypothesis related issues and subsequently improve the effectiveness of timely risk management practices.

### 4.2.1 Filtration and Cleaning

The following section is dedicated to comprehending the filtration and cleaning procedures implemented on options data. The steps of the procedure are described in *Table 4.1* and further explained below. Moreover, these steps are also applied to the call options utilised for bubble detection (see 5.1). Option prices were collected over a wide range of strikes, at increments of \$5.00. For each contract, the mid-price is considered, by computing the average of bid and ask values. Given the 15 minutes delay in the closing of option markets, with respect to the cash market for equity indices, despite best efforts of credible data sources, such as Refinitiv Eikon, an issue of price asynchrony could exist (Almeida et al., 2023; Fusari et al., 2024). Therefore, the first step in *Table 4.1* focuses on computing spot prices for the S&P 500 index, from forward prices of ATM options, via the put-call parity.

**Table 4.1: Options Data Filtration and Cleaning Procedures.**

Filtration and Cleaning Steps	Procedure
<i>Step 1</i>	Estimating Implied Spot Prices ( $S_t$ ).
<i>Step 2</i>	Applying Time to Maturity ( $\tau$ ) Conditions.
<i>Step 3</i>	Filtrating Illiquid Options using Moneyness, Implied Volatility and Trading Volume Conditions.

On a given day,  $t$ , using prices from put and call option pairs, corresponding to the same maturity,  $\tau$ ; forward prices,  $F_{t,\tau} = \mathbb{K} + e^{r\tau}(C_t(\mathbb{K}, \tau) - P_t(\mathbb{K}, \tau))$ , are computed. Next, across the five most at-the-money (ATM) options, the median of  $F_{t,\tau}$  is considered, for determining the implied spot price,  $S_t = F_{t,\tau}e^{-(r-q)\tau}$ , where  $r$  and  $q$ , respectively represent the risk-free rates and continuously compounded dividend yields. ATM options tend to be the most liquid, providing an increased likelihood of the put-call parity being satisfied, and subsequently, consistency with the NFLVR and ND assumptions. Risk-free rates are

determined by linearly interpolating the Zero-Coupon Yield Curve provided by Eikon, following Hagan and West (2006), as denoted in (49). In the equation,  $r(\tau)$  resembles the rate matched to  $\tau$ , and  $r_i$  ( $r_{i+1}$ ) corresponds to the yield curve component with the closest maturity  $t_i$  ( $t_{i+1}$ ), that is less (greater) than  $\tau$ .

$$r(\tau) = \frac{\tau - t_i}{t_{i+1} - t_i} r_{i+1} + \frac{t_{i+1} - \tau}{t_{i+1} - t_i} r_i \quad (49)$$

$$C_{t,bid}(\mathbb{K}, \tau) - P_{t,ask}(\mathbb{K}, \tau) = b_0 + b_1\$ + b_2\$ \tau + b_3\mathbb{K} + b_4\mathbb{K} \tau + b_5 D_{B,A} \quad (50)$$

The constituents of the S&P 500 index pay dividends, and hence a daily yield must be calculated. It is common practice to compute implied dividend yields from the put-call parity forward prices (Hull, 2014; Ulrich & Walther, 2020; Almeida et al., 2023). However, this approach is vulnerable to facing price asynchrony, given that the implied yield would be obtained from matching forward prices to the close index price. Although Refinitiv Eikon provides a timeseries of daily yields, option maturities are ignored during computation. Hence, the methodology provided by OptionMetrics<sup>40</sup>, in (50) is implemented, under the assumption of the put-call parity holding. The regression considers difference between bid-call and ask-put prices,  $C_{t,bid}(\mathbb{K}, \tau) - P_{t,ask}(\mathbb{K}, \tau)$ , as the dependent variable, with  $D_{B,A}$ <sup>41</sup> taking the value of 1. The remaining independent variables comprise of the underlying price ( $\$$ ), strike price ( $\mathbb{K}$ ), and time to maturity ( $\tau$ ). Dividend yields, on a given day, are estimated by utilizing data on all traded options over the past three months, excluding those with maturities of less than 15 days. In accordance with the put-call parity, the continuously compounded dividend yield is the negative of computed  $b_2$  parameter value.

After computing the implied spot prices, the next step from *Table 4.1* focuses on the maturity condition. Bubbles within short maturity options reveal insight on larger occurrences in the underlying over longer horizons (Fusari et al., 2024). Recall, call option bubble magnitudes provide a lower bound for that in the underlying. For short option contracts, potential bubbles would disappear, as the magnitude reduces with tenor. Therefore, options with maturities below 7 days were discarded. Fusari et al. (2024) also drop options with maturities greater

<sup>40</sup> [View IvyDB\\_US\\_v5.4\\_Reference\\_Manual.pdf](#).

<sup>41</sup>  $D_{B,A} = 1$ , if  $C_{t,bid}(\mathbb{K}, \tau)$  is used, and 0, otherwise.

than 50 days, due to the concentration of liquidity over shorter contracts. In contrast, given the dependence of the bubble magnitude on option maturity, to enhance detection, contracts with durations up to 365 days are included. Furthermore, despite having higher trading volumes, due to their relatively shorter lives, weekly options are overlooked by preference for monthly contracts<sup>42</sup>.

Highly illiquid contracts impose challenges during calibration; hence they were filtered out in the final step, on the ground moneyness, implied volatilities and trading volume (Piiroinen et al., 2018; Stahl and Blauth, 2024; Fusari et al., 2024). Options with no trading volumes, and those corresponding to market implied volatilities greater than 1, were discarded. It is common practice to calibrate parameters of asset pricing models to only OTM options, due to their highly liquid nature. Bubble detection approaches, notoriously suffer from a joint hypothesis issue, and when considering option prices, Jarrow (2015) suggested the accuracy in pricing put options. Two criteria were considered for moneyness<sup>43</sup>: standardised moneyness ( $m$ ), and log-moneyness ( $k$ ), with options meeting  $-10 < m < 5$ , and  $k < 0.5$ , being retained. The standardised moneyness and log-moneyness conditions ensure disregarding deep in-/out- of-the-money options and allow for the vital consideration of contracts over a wide range of moneyness in the study, without illiquidity concerns.

## 4.2.2 Summary Statistics

To balance the trade-off between computational efficiency and robustness, Fusari et al. (2024) use daily most liquid options corresponding to maturities with the highest cumulative volume ( $HCV$ ). Despite being highly liquid, these options cover only a single slice of the implied volatility surface. However, they provide a great environment for comparing calibration performances. The calibration of various stochastic processes, along with that of the two-step neural network approach are conducted on the  $HCV$  put options dataset. The entire option price surface comprises a wide range of maturities, and hence contains excess

---

<sup>42</sup> Preference for monthly options on the S&P 500 index is consistent with Fusari et al. (2024).

<sup>43</sup> Standardised moneyness ( $m$ ) =  $\frac{\ln(\mathbb{K}/S)}{\sigma_{ATM} \sqrt{\tau}}$ , where  $\sigma_{ATM}$ , represents the implied volatility of the ATM option, Log-moneyness ( $k$ ) =  $\ln \left| \mathbb{K}/F_t \right|$ .



information, on future expectations of the underlying price, and market conditions (Ulrich & Walther, 2020). Given the calibration boost from neural networks, it will be possible to compute daily GSVJD parameters from the entire option surface. The summary statistics of put options belonging to the *HCV* and *Entire Surface* dataset are revealed in *Table 4.2*.

**Table 4.2: Summary Statistics of S&P 500 Put options from HCV and Entire Surface datasets.**

*Note: Summary Statistics of S&P 500 put options across the HCV and Entire Surface datasets, are displayed in Panel A and B, respectively. The statistics are computed across each maturity,  $\tau$ , on a particular day. Volume and Open Interest (000s) report accumulated values across daily contracts with similar  $\tau$ . Bid-, Ask-, and Mid-prices are denoted in \$, whereas  $\sigma_{ATM}$ , represents implied volatilities of the most ATM option and  $N_P$  reveals number of put options. Standardized- and log-moneyness are depicted by  $m$  and  $k$ , respectively, along with  $mon = \mathbb{K}/\mathbb{S}$ .*

	Mean	SD	P25	P50	P75
Panel A: HCV					
Bid	56.99	99.84	5.00	20.70	67.20
Ask	58.36	101.75	5.40	21.40	68.70
Mid	57.67	100.79	5.20	21.05	68.00
Volume	90.68	42.78	63.23	81.85	105.89
Open Interest	1357.46	510.13	938.59	1337.78	1722.96
$m$	-2.13	2.41	-3.38	-1.55	-0.35
$k$	-0.11	0.13	-0.18	-0.09	-0.02
$mon$	0.90	0.11	0.83	0.92	0.98
$\tau$	42.91	26.98	23.00	37.00	59.00
$N_P$	161.08	45.47	128.00	155.00	188.00
$\sigma_{ATM}$	19.06%	7.69%	13.47%	17.55%	23.14%
Panel B: Entire Surface					
Bid	94.03	125.77	13.90	50.20	125.30
Ask	95.91	127.88	14.40	51.40	127.50
Mid	94.97	126.81	14.15	50.80	126.40
Volume	24.50	33.62	3.17	10.60	32.69
Open Interest	414.38	483.32	57.12	228.13	601.81
$m$	-1.59	1.95	-2.42	-1.09	-0.23
$k$	-0.12	0.14	-0.20	-0.10	-0.02
$mon$	0.89	0.12	0.82	0.91	0.98
$\tau$	167.60	101.67	80.00	156.00	253.00
$N_P$	72.82	58.42	29.00	51.00	105.00
$\sigma_{ATM}$	20.03%	6.09%	15.63%	19.47%	23.48%

Daily most liquid options have accumulated volumes that are approximately 4 times greater than those across maturities within the Entire Surface, as seen in *Table 4.2*. A wide range of strike prices are considered from each day, with two moneyness criteria implemented for filtering out illiquid contracts. The *mon* statistic reveals more than 75% of put options are OTM. Recall, only option contracts expiring within 7 and 365 days were considered. The mean maturity in the *HCV* dataset is approximately 43 days, with a median of 37.00 days. In contrast, across the *Entire Surface*, the average and median contract durations are 167.60 and 156.00 days, respectively. As opposed to a single maturity in the *HCV* dataset on a given day, the *Entire Surface* comprises of 9 unique maturities, on average. This shows that the *Entire Surface* dataset is significantly larger, consisting of an impressive range of moneyness spanning over multiple volatility smiles, with respect to each day. Furthermore, *Table 4.2* reveals approximately 73 daily put contracts traded across each maturity, in comparison to 161 for the most liquid maturity. Hence, on average, calibration to the entire surface, would be conducted on approximately 677 daily put options. For the GSVJD model, this would be computationally infeasible with Monte Carlo simulations, given the current resources at hand.

### 4.3 Comparing Calibration of Stochastic Processes

Prior to proceeding with the neural network framework, it is important to examine calibration performances of various stochastic processes. The following four model are analysed: Heston (Heston, 1993), Bates (Bates, 1996), Andersen and Piterbarg (Andersen & Piterbarg, 2007), and GSVJD (Fusari et al., 2024). Recall, from (17) and (18), it can be shown that the GSVJD nests, these three processes. When jumps are not considered ( $\lambda = 0$ ), the GSVJD resembles the Andersen and Piterbarg process, which can acknowledge strict local martingale tendencies within the price of the asset. However, for  $p = 0.5$ , such tendencies cannot be captured. Under this condition, for  $\lambda = 0$  and  $\lambda > 0$ , the GSVJD reveals Heston and Bates models, respectively, with the latter addressing price jumps.

The stochastic processes are computed using 1,000 Monte Carlo simulations<sup>44</sup>, with a timestep of  $\frac{1}{365 \times 5}$ . They are calibrated to daily most liquid S&P 500 put options, with the objective of minimising  $RMSE_{(\sigma,t)}$ , using Differential Evolution. The study opts against the predetermination of parametric values, despite it being common practice to reduce computational costs. A system with an Intel-i7 processor was not practical for daily calibrations, over the large study period. Hence, the High-Performance Compute Clusters, at the University of Glasgow were accessed<sup>45</sup>.

**Table 4.3: Calibration Performances of Stochastic Processes.**

*Note: Summary Statistics of calibration performance,  $RMSE_{(\sigma,t)}$ , and time (hours), across the entire sample period, for each stochastic process, are displayed in Panel A and B, respectively.*

	mean	std	P25	P50	P75	max
Panel A: $RMSE_{(\sigma,t)}$						
Heston	4.53%	3.64%	2.36%	3.40%	5.20%	31.19%
Bates	0.36%	0.34%	0.18%	0.28%	0.41%	3.96%
Andersen and Piterbarg	3.60%	3.09%	1.79%	2.64%	4.13%	30.26%
GSVJD	0.31%	0.35%	0.12%	0.18%	0.33%	3.96%
Panel B: Time						
Heston	0.66	0.36	0.40	0.59	0.82	2.92
Bates	1.42	0.68	0.89	1.27	1.80	5.40
Andersen and Piterbarg	1.38	0.39	1.15	1.35	1.60	3.32
GSVJD	3.56	1.20	2.67	3.34	4.19	9.54


The calibration performance and time metrics from *Table 4.3*, reveal a trade-off between computational efficiency and accuracy. The average  $RMSE_{(\sigma,t)}$  values are significantly larger for the Heston and Andersen and Piterbarg models, which do not take jumps into account. The calibration error from both models, peak on April 9, 2020, at 31.19% and 30.26%, respectively. Contrastingly, the worst performance from the Bates and GSVJD models, is lower than the sample period average from Heston error metrics, and only

<sup>44</sup> The Heston and Bates models, possess various closed form solutions, however, to maintain consistency with remaining models, Monte Carlo simulations were employed.

<sup>45</sup> For more information, refer to: [University of Glasgow - MyGlasgow - IT - High Performance Compute Clusters](#).

marginally higher than that from the Andersen and Piterbarg stochastic process. In the respective order, these performances<sup>46</sup> were witnessed on March 11, 2020, and January 27, 2021. It is evident that the inclusion of jumps, significantly improves calibration performances. Furthermore, the ability to capture strict local martingale tendencies, by not fixing  $p = 0.5$ , enhances calibration. The GSVJD provides an average  $RMSE_{(\sigma,t)}$  of 0.31%, across the sample period, compared to 0.36% from the Bates model. The difference is marginal, but the former excels with respect to the 25<sup>th</sup>, 50<sup>th</sup> (median) and 75<sup>th</sup> percentiles values. However, despite performing best, the GSVJD consists of the greatest number of parameters, and is therefore associated with the largest computational cost. It requires 3.56 hours, in comparison to the 0.66, 1.38, and 1.42 hours needed by the Heston, Andersen and Piterbarg, and Bates models, respectively.

**Figure 4.1: Calibration Performances of Stochastic Processes.**



*Note: Calibration performances of the GSVJD, Andersen and Piterbarg, Bates, and Heston stochastic processes, over daily S&P 500 put options between 2019 and 2022. The Box Plot (left) illustrates the distribution of  $\log RMSE_{(\sigma,t)}$ , whereas the historical performances (right) are revealed over the sample period.*


The computational inefficiencies associated with the GSVJD model justifies why practitioners might want to stay clear of it. However, it possesses superior quality in capturing crucial market price characteristics during calibration. The box plot, along with that of historical performances for the  $\log RMSE_{(\sigma,t)}$  metric, illustrated in *Figure 4.1*, reveal

<sup>46</sup> The Bates model delivered a very similar metric on January 27, 2021, whereas the GSVJD documented an  $RMSE_{(\sigma,t)}$  value of 1.84% on March 11, 2020.

strong and reliable calibration performances from the GSVJD model. Specifically, these plots highlight preference over the Bates, maintaining consistency with *Table 4.3*. Observe, the median performance of the GSVJD model is equivalent to that at the 25<sup>th</sup> percentile of the Bates process.

The performances of the GSVJD model in *Figure 4.2*, reveal approximately 95% of the error metrics to be below 1%, which further cements the pricing strength of the model. Naturally, across a sample of 1,000 trading days, there are bound to be exceptions in the form of sharp spikes, with two such instances standing out. First, the period aligning with the start of the COVID-19 induced lockdown, and subsequent market crash (March 9 – 16, 2020), experienced an average  $RMSE_{\sigma,t}$  of 1.56%. Second, the highest error metric of 3.96% was encountered on January 27, 2021, a day corresponding with the short squeeze on the GameStop stock. Fusari et al. (2024) resort to dropping days beyond the 95<sup>th</sup> percentile of  $RMSE_{\sigma,t}$ , however, this practice is ignored to provide a holistic analysis of bubble detection, especially during the COVID-19 induced market crash period.

**Figure 4.2: GSVJD Calibration Performances: S&P 500 Index (2019 – 2022).**



*Note: GSVJD Calibration performances measured by  $RMSE_{\sigma,t}$  (%), over daily S&P 500 put options during the study period, along with the mean, median and P95 (95<sup>th</sup> percentile) values.*

The superiority of the GSVJD model, in comparison to its counterparts, has been depicted within this section. The ability to capture volatility mean-reversion, jumps and strict local

martingale tendencies, make it effective accurately price options, in order to detect bubbles without suffering from a joint-hypothesis related issue. However, it is plagued with a major computational cost, which can be unattractive for practitioners preferring tractability. Therefore, to make the three-step approach more attractive for real-time bubble detection, a deep calibration neural network framework is required.

## 4.4 Deep GSVJD Calibration

The deep calibration framework is engineered to replace Monte Carlo simulations, during estimation of the GSVJD model, to boost efficiency without sacrificing accuracy. The framework comprises of two steps: *forward pass* and *backward pass*. The former is devoted to making the network comprehend dynamics of the GSVJD process, whereas the latter calibrates parameters from market data, using trained hidden layers from the first step. In 3.2, it was revealed that the optimal network architecture would be found by an extensive random search over 720 hyperparametric combinations, with each undergoing a 3-fold cross validation, over a synthetic dataset.

The intensive training and testing procedures for determining the optimal architecture for bubble detection are briefly summarized below. For more information, the reader is referred to *Section A* of the Appendix. Initially, all architectures underwent a 3-fold cross validation, when being trained on the *Random Search* dataset, such that the top performing networks were selected for optimal training. The *Random Search* dataset comprised of 1 million options, diminishing concerns regarding the availability of sufficient data during each validation fold. The latter phase included training over the significantly larger, *Optimal Training* dataset, constructed from 10 million options. The performances of all architecture during the optimal training phase are strong. A firm preference for hyperparameters resembling deeper and wider networks, with smaller batch sizes, is revealed. Given the close proximity of training performances, and the need to overcome the joint-hypothesis issue during bubble detection, the optimal architecture is selected based on calibrations to market data, during the *backward pass*.

The hidden layers from each optimally trained architecture are frozen and carried to the *backward pass*. Since the weights and biases within these layers are already determined, the network solely operates as a numerical solver during the calibration process. When applied to the *HCV* dataset, each network delivered strong and stable performances, with all but one being superior to the benchmark Monte Carlo calibrations of the GSVJD. The best performing architecture was the deepest and widest, with the smallest batch size. It consisted of 4 hidden layers, 40 nodes, with ELU activation and a batch size of 2048, improving computational efficiency and accuracy by the approximate respective magnitudes of 254 and 2, in comparison to the benchmark.

This enables application of the deep framework, for calibrating the GSVJD model to the entire surface. Hence, allowing for bubble detection, using the three-step approach, with an improved ability to capture the forward-looking expectations of market participants. Recall, the goal for the deep calibration framework is to replace Monte Carlo simulations, as a numerical solver. This section proceeds with comparing GSVJD calibrations from both methods, over the daily most liquid options. Next, deep framework is employed towards calibrating the stochastic process to the entire surface, prior to highlighting the wide-scale application of this practice during bubble detection.

#### 4.4.1 Monte Carlo Simulations vs. Neural Networks

Initially, during the benchmark calibration of the GSVJD model, given the large computational burden associated with Monte Carlo simulations, a trade-off between robustness and efficiency was experienced. In abidance with Fusari et al. (2024), to strike a balance, only the daily most liquid option maturities were considered. To maintain consistency during comparisons, the same practice was utilised when employing the deep calibration framework. Both approaches were applied to the *HCV* dataset (*Table 4.2*), with the goal of minimising  $RMSE_{\sigma,t}$  using Differential Evolution, for daily estimation of GSVJD parameters. Hence, the sole purpose of the neural network is to replace the computational cumbersome Monte Carlo simulations, as a numerical solver. The summary statistics of performance and calibration time, with respect to the deep calibration framework, are reported in *Table 4.4*.

As seen in *Table 4.3*, the GSVJD model provided superior performances in comparison to the Heston, Andersen and Piterbarg, and Bates stochastic processes, by recording an average  $RMSE_{\sigma,t}$  of 0.31%, over the entire sample period. However, reliance on Monte Carlo simulations created a large computational burden, motivating the recommendation of the deep calibration framework. On average, the neural network required 0.014 hours (nearly 50 seconds) for daily calibrations, in contrast to the 3.56 hours needed by the benchmark<sup>47</sup>. It increased efficiency by an order of magnitude of 254 and was also 48 times faster than the Heston model, despite estimating 4 additional parameters. Therefore, the need for heavy computational resources is entirely diminished, when the three-step approach for bubble detection is enhanced by the deep calibration framework. Nevertheless, the boost in efficiency would be futile if the deep calibration framework is unable to at least match the accuracy of the benchmark.

**Table 4.4: Performance of Deep Calibration to HCV put options.**

*Note: Summary Statistics of GSVJD calibration performance,  $RMSE_{(\sigma,t)}$ , and time (hours) using the Deep Calibration Framework, across the entire sample period.*

	mean	std	P25	P50	P75
$RMSE_{\sigma,t}$	0.17%	0.23%	0.06%	0.10%	0.20%
<i>Time</i>	0.014	0.005	0.010	0.013	0.017


The application of machine learning models is typically associated with a trade-off between accuracy and efficiency, in abundance with their respective benchmarks. It is common practice across literature to aim for improving efficiency, with the condition of at least matching the performances of the benchmark model. The deep calibration framework does one better, and outperforms its benchmark, on both grounds. As witnessed in *Table 4.4*, the neural network approximately doubled the calibration accuracy by reporting an average  $RMSE_{\sigma,t}$  of 0.17%. Interestingly, its median and 75<sup>th</sup> percentile values, 0.10% and 0.20%, are lesser than the respective 25<sup>th</sup> percentile and mean  $RMSE_{\sigma,t}$  metrics, 0.12% and 0.31%, from the Monte

<sup>47</sup> These calibrations were conducted on a personal computer with an i7-processor, in contrast to requiring High-Performance Computer Clusters, when using Monte Carlo simulations.



Carlo simulations. Naturally, the deep calibration framework also outperforms the metrics of the other stochastic process, as it is trained to learn the dynamics of the GSVJD model.

**Figure 4.3: GSVJD Calibration Performances: Benchmark vs. Deep Calibration Framework.**



Note: GSVJD Calibration performances measured by  $RMSE_{\sigma,t}$  (%), over daily S&P 500 put options between 2019 and 2022, using Monte Carlo Simulations (left) and the Deep Calibration Framework (right).

Large credit is due to the intensive training step dedicated towards selecting the optimal *forward pass* architecture. The random search, coupled with 3-fold cross validation training over a synthetic dataset, ensures that the best selected hyperparameter combination is robust to various market scenarios. Recall, the network is trained on a large synthetic dataset, such that it becomes aware of various market regimes. The benefits of this practice are observed when comparing the standard deviations of both approaches. The benchmark provides an  $RMSE_{\sigma,t}$  standard deviation of 0.35% (see Table 4.3), whereas Table 4.4 reveals 0.23% for its counterpart. This indicates strong capabilities of the neural network to calibrate parameters, by comprehending the various market regimes, across the large sample period. Furthermore, given the *forward pass* is a one-time requirement, the creation of the synthetic dataset used 10,000 Monte Carlo simulations, compared to the 1,000 employed by the benchmark calibration. The number of simulations in the latter instances was restricted, given the high computational cost associated with calibration.

Performances of the deep calibration framework can be further examined and compared to those of the benchmark in Figure 4.3. A quick glance reveals  $RMSE_{\sigma,t}$  metrics over the entire sample from the deep calibration framework, to be more stable, with fewer spikes.

Furthermore, remarkably, it can be observed that the calibration error across 95% of the sample period does not exceed 0.46%, compared to 1.02% when using the benchmark. However, similar to its counterpart, the neural network struggled during the period corresponding to the COVID-19 induced market crash lockdown (March 9-16, 2020) and on January 27, 2021, a day synonymous with the GameStop stock short squeeze. During the former period, the neural network recorded an average metric of 1.29%, whereas on latter day, revealed 3.94%. In both instances, the benchmark provided larger errors of 1.56%, and 3.96%, respectively. The performance of the deep calibration framework between March 9-16, 2020, are very appreciative, considering the extreme volatile nature of the market during this period.

#### 4.4.2 GSVJD Calibration to the Entire Surface

The development of the deep calibration framework for estimating GSVJD parameters was motivated by improving computational efficiency. It is engineered to replace Monte Carlo simulations as numerical solvers within the three-step approach, such that daily option price surfaces can be considered for bubble detection. Information regarding the underlying price, varies substantially across different levels of maturity, within the entire surface (Ulrich & Walther, 2020). Hence, to capture forward looking expectations of market participants to the fullest extent possible, the GSVJD model is calibrated to the entire surface. During previous calibrations using the daily most liquid option prices, such information was only partially acknowledged.

The deep calibration framework is applied to the *Entire Surface* dataset (see *Table 4.2*), and the summary statistics of the  $RMSE_{\sigma,t}$  metric and computational time, are displayed in *Table 4.5*. Not all options belonging to the fixed moneyness-maturity grids of implied volatility surfaces are necessarily traded. In such instances, the practitioner would be forced to rely on interpolating a selected model to market data, followed by extrapolating to the required moneyness-maturity combination. This practice would be ideal, when implementing the *grid-based* two-step neural network calibration approach, as witnessed in Horvath et al. (2021). However, it, would induce an additional error, which could hamper the conclusiveness of the detected bubbles. Therefore, this study opted to directly calibrate to

all traded options on a given day. Nevertheless, the motivation remains the same, as one would expect from using implied volatility surfaces.

**Table 4.5: Performance of Deep Calibration to Entire Surface put options.**

*Note: Summary Statistics of GSVJD calibration performance,  $RMSE_{(\sigma,t)}$ , and time (hours) using the Deep Calibration Framework, across the entire sample period.*

	mean	std	P25	P50	P75
$RMSE_{\sigma,t}$	0.75%	0.39%	0.53%	0.69%	0.86%
<i>Time</i>	0.013	0.004	0.010	0.013	0.016


As revealed previously, the optimal deep calibration architecture outperformed the benchmark, on the ground of efficiency and accuracy. Hence, it was applied to calibrating from the entire daily surface of S&P 500 put options, between 2019 and 2022. Similar to the *HCV* dataset calibration, the network was paired with a Differential Evolution optimizer, and all computations were conducted, with the aim of minimising  $RMSE_{\sigma,t}$ , on a personal computer, with an Intel i7 processor. Despite comprising of on average 9 maturities, daily calibrations to such an extent are no longer a computational burden, as a result of the massive efficiency boost provided by the neural network.

Observe from *Table 4.5*, the speed of calibration, is marginally faster than that witnessed during calibrations to the most liquid smiles. The additional consideration of maturities increases the computational time exponentially. When paired with neural networks, the Differential Evolution optimiser provides a two-fold speed up (Liu et al., 2019). Recall, the neural network computes all option implied volatilities simultaneously, without requiring an additional solver. Therefore, the deep calibration framework requires only approximately 0.013 hours (46.8 seconds) for daily parameter estimation. On average, in comparison to the benchmark, despite calibrating from a significantly larger number of put options, it provides a speed up with a magnitude of 273.

The average calibration  $RMSE_{\sigma,t}$  metric is 0.75%, much greater than those observed from calibrating to the *HCV* dataset. However, given that the *Entire Surface* dataset comprises of, on average 9 maturities for each day, these performances are justifiable. There are several relatively illiquid options presents in the sample, which create difficulties during model

fitting. The calibrations performances are still strong, given a median error of 0.69%. Additionally, the  $RMSE_{\sigma,t}$  does not respectively exceed 0.86% and 1.25%, for at least 75% and 95% of the days within the sample period.

**Figure 4.4: Deep Calibration Framework to the Entire Surface.**



*Note: GSVJD Calibration performances measured by  $RMSE_{(\sigma,t)}$  (%), over daily S&P 500 put option surfaces between 2019 and 2022, using the Deep Calibration Framework.*

From *Figure 4.4*, it can be observed that the error is stable across majority of the period, barring two instances. As witnessed with previous GSVJD calibrations, these instances correspond to the COVID-19 induced lockdown, and subsequent crash period, and January 27, 2021, a day that experienced the GameStop stock short squeeze. During that latter,  $RMSE_{\sigma,t}$  was measured at 3.21%, lower than that from earlier application of neural networks. Alternatively, regarding the former, poor performances were noticed between March 3, 2020, and April 6, 2020, with an average error metric of 2.30%. Within this period, March 9, 11, and 12, 2020 reveal  $RMSE_{\sigma,t}$  metrics of 4.01%, 4.41% and 4.11%. The extreme volatile nature of markets on these days, along with the first and last of them experiencing circuit breakers, induced pricing difficulties.

These difficulties were further exacerbated by considering several maturities, plagued with relatively illiquid options, during calibrations. In comparison with calibrations to the *HCV* dataset, this particular period of poor performances, is extended. However, it is consistent with the reasoning of extreme volatility, given the COVID-19 induced market crash. Nevertheless, on average, across the sample period, calibration performances are strong and

stable, enabling the use of estimated parameters for pricing call options, to proceed with bubble detection. The standard deviation of  $RMSE_{\sigma,t}$  in *Table 4.5*, is only marginally higher than that observed in *Table 4.3*, corresponding to benchmark calibrations.

### 4.4.3 GSVJD Calibration for Bubble Detection

The superiority of the deep calibration framework in comparison to Monte Carlo simulations has been acknowledged in *4.4.1*. Furthermore, the efficiency boost allowed for GSVJD parametric estimations from across the entire surface. Calibrations in *4.4.1* and *4.4.2* were conducted on the same differential evolution seed for direct comparisons. However, such a practice can hamper the robustness of bubble detection, given that optimizers are vulnerable to the problem of obtaining non-unique solutions (Wshah et al., 2020).

A common practice is to employ parameters from the previous day as initial values (Büchel et al., 2022), but this could be problematic when the market experiences extreme scenarios. Hence, it is difficult to rely on the  $\rho$  and  $p$  conditions for detecting bubbles. Naturally, this raises a question regarding the employment of specialized stochastic processes, capable of identifying strict local martingale tendencies, over conventional methods such as Heston and Bates models. Although computational efforts are reduced, as revealed in *Table 4.3*, accuracy is sacrificed. The GSVJD method is still better equipped for pricing options during potential periods of exuberance, making the three-step approach less vulnerable to a joint-hypothesis related issue. The impressive sophistication of the GSVJD process is precisely why a deep calibration framework was engineered to overcome its computational issues. Taking advantage of the efficiency boost provided by the deep calibration framework, daily GSVJD parameters are computed over 10 different seeds<sup>48</sup>. Summary statistics of average parameters estimated from the most liquid cross section of options, along with the  $RMSE_{\sigma,t}$  performance metric, are reported in *Table 4.6*.

---

<sup>48</sup> Given the need for calibrating the entire sample over multiple seeds, the High Performance Computer Clusters, at the University of Glasgow were accessed.

The average calibration performance across all seeds from *Table 4.6*, strongly resembles those reported in *Table 4.4*. Moreover, all seeds display their worse performance on January 27, 2021, with a mean  $RMSE_{\sigma,t}$  metric of 3.94%. Whereas, during the period between March 9 – 16, 2020, performance averages at 1.24%. The consistency across the seeds, reveals robustness of the deep calibration framework, when calibrating to put option implied volatilities. However, as observed from *Table 4.6*; the mean of standard errors over all seeds, associated with the estimation of each parameters are large, especially for  $p$  and  $\rho$ . This is consistent with Fusari et al. (2024), who acknowledge the high standard errors corresponding with daily parametric calibrations. It is due to such reasoning; conditional tests were integrated into the final phase of the three-step approach.

**Table 4.6: Summary Statistics of Deep Calibration to HCV put options.**

*Note: Summary Statistic for average daily calibrated GSVJD parameters computed from the most liquid cross section, error metric ( $RMSE_{\sigma,t}$ ), using the Deep Calibration Framework, across 10 seeds. The table reports the mean, standard deviation (std), 25<sup>th</sup>, 50<sup>th</sup> and 75<sup>th</sup> percentile (P25, P50, and P75) values, and standard errors (SE).*

	mean	std	P25	P50	P75	SE
$\kappa$	1.598	1.111	0.808	1.435	2.099	0.324
$\sigma_v$	1.309	0.549	0.917	1.327	1.712	0.149
$\bar{v}$	0.093	0.076	0.034	0.072	0.131	0.016
$V_0$	0.036	0.029	0.014	0.026	0.050	0.003
$\rho$	-0.628	0.249	-0.830	-0.644	-0.464	0.066
$p$	0.631	0.241	0.470	0.640	0.775	0.073
$\lambda$	0.192	0.303	0.080	0.116	0.180	0.046
$\mu_y$	-0.231	0.137	-0.316	-0.209	-0.138	0.038
$\sigma_y$	0.295	0.127	0.203	0.278	0.377	0.023
$RMSE_{\sigma,t}$	0.17%	0.23%	0.06%	0.11%	0.19%	-

The ND condition ensures that the call option and underlying price bubbles share a linear relationship, with the magnitude of the former acting as the lower bound to that of the latter. Therefore, the presence of type III bubbles in call option prices admit strict local martingale tendencies in the underlying process. Furthermore, through a robustness check, Fusari et al. (2024) reveal that strict local martingale conditions for parameters  $\rho$  and  $p$ , have a positive impact on the magnitude and probability of bubbles. Hence, in the latter stages of the study,

an identical robustness check is implemented to examine the consistency of bubbles identified from parameters estimated by the deep calibration framework, and the strict local martingale condition.

A similar trend is witnessed with respect to *Entire Surface* calibrations over the various seeds. The distribution of calibration performances, measured by  $RMSE_{\sigma,t}$  in Table 4.7 are near identical to those from Table 4.5. Moreover, during the period of March 3, 2020, and April 6, 2020, error averaged at 2.32%, whereas on the day of the GameStop short squeeze, it recorded 3.21%. Even calibrations to the *Entire Surface* reveal large standard errors for parameters  $\rho$  and  $p$ ; hence the detection of bubbles must rely on the statistical test, in the final phase of the three-step approach. As proposed for *HCV* calibrations, a robustness check is conducted to examine the impact of the strict local martingale conditions for parameters  $\rho$  and  $p$  (obtained from the entire surface) on the formation of bubbles.

**Table 4.7: Summary Statistics of Deep Calibration to Entire Surface put options.**

*Note: Summary Statistic for average daily calibrated GSVJD parameters computed from the Entire Surface, error metric ( $RMSE_{\sigma,t}$ ), using the Deep Calibration Framework, across 10 seeds. The table reports the mean, standard deviation (std), 25<sup>th</sup>, 50<sup>th</sup> and 75<sup>th</sup> percentile (P25, P50, and P75) values, and standard errors (SE).*

	mean	std	P25	P50	P75	SE
$\kappa$	2.849	1.238	1.894	2.698	3.606	0.436
$\sigma_v$	1.590	0.535	1.234	1.674	1.997	0.145
$\bar{v}$	0.065	0.033	0.040	0.061	0.089	0.006
$V_0$	0.045	0.061	0.015	0.029	0.057	0.001
$\rho$	-0.691	0.203	-0.846	-0.712	-0.561	0.048
$p$	0.725	0.258	0.567	0.687	0.830	0.074
$\lambda$	0.113	0.127	0.054	0.078	0.114	0.030
$\mu_y$	-0.354	0.217	-0.485	-0.320	-0.192	0.064
$\sigma_y$	0.327	0.170	0.191	0.292	0.433	0.038
$RMSE_{\sigma,t}$	0.74%	0.39%	0.53%	0.70%	0.86%	-

Despite revealing large standard errors, proceeding with the average GSVJD parameters improves robustness during bubble detection. Overall, application of neural networks for calibrating the GSVJD model are favourable over the alternative reliance on Monte Carlo simulations. The improved accuracy of calibrating to market put options reveals strong

capabilities of overcoming the joint hypothesis issue during bubble detection. Furthermore, highly efficient nature of practice makes the three-step approach more tractable, and therefore attractive to practitioners for real-time bubble identification. The deep calibration framework is integrated only to replace Monte Carlo simulations as a numerical solver when estimating parameters of the GSVJD model. Inspired by Biagini et al. (2024), parameters from each seed are employed for pricing call options using (19), (20) and (22), with average modelled prices being considered for bubble detection. This practice is implemented for both sets of parameters, attempting to improve the robustness of the deep calibration framework during bubble detection.

It is crucial to empirically analyse the bubbles detected from integrating the deep calibration framework into the three-step approach. This action is undertaken in the next chapter, with respect to options from the daily most liquid cross section, and entire surface. Calibrating GSVJD model for bubble detection, to the entire surface, is beneficial, as it allows to capture the forward-looking expectations of market participants to a larger extent. Since these parameters are obtained from all traded options on a given day, they are applicable towards detecting bubbles across various call option maturity groups, instead of being restricted to the most liquid one. Furthermore, the presence of a bubble, across the lifetime of an option can also be examined, given the historical calibration of parameters. In real-time such applications are difficult to implement, however the massive computational boost received from the deep calibration framework, make them an exciting avenue to explore. In all instances, average modelled call option prices, across all 10 seeds are considered.

## 4.5 Summary and Discussion

After constructing a deep framework for calibrating the GSVJD model in the previous chapter, the focus shifted to computing its parameters from market put options. Daily European-styled options, written on the S&P 500 index, between January 2, 2019, and December 30, 2022, were selected for this purpose. The sample includes 1,000 trading days, over a span of 4 years. In this chapter, two datasets are introduced, with both undergoing intensive filtering and cleaning procedures to limit issues with respect to pricing. The first,



*HCV*, consists of options that correspond to the most liquid maturity, whereas the second, *Entire Surface*, considers all traded options, on a given day.

The *HCV* dataset allows for a balance between robustness and efficiency during calibration, which is crucial when employing Monte Carlo simulations. It is utilised for comparing the accuracy and efficiency of the benchmark calibrations to other stochastic processes, in addition to those from the neural network framework. The *Entire Surface* dataset, on average considers daily options across 9 maturities, as compared to the singular inclusion within the *HCV*. This roughly translates to the daily calibrations over 677 options when employing the former, as opposed to 161 with respect to its counterpart. Calibrations to all traded options enabled capturing the forward looking expectations of market participants to a greater extent. However, such practices are also plagued with a large computational cost, motivating a need for improving efficiency. First, the GSVJD model is compared to other stochastic processes, on the grounds of calibrating to market put options, in order to seek for a more efficient process, with a similar range of accuracy.

The GSVJD process is compared against the Heston, Bates, and Andersen and Piterbarg models, all of which are nested within it, subject to some parametric adjustments. Recall, the GSVJD is capable of addressing strict local martingale tendencies and jumps within the asset price process. In regard to the others, the former characteristic is solely acknowledged by the Andersen and Piterbarg model, whereas the latter is only captured by the Bates process. The Heston process does not reflect either and has the least number of parameters required to be computed. Hence, it is the most efficient in comparison to the other processes. The comparison reveals advantages for methods that capture jumps, and address strict local martingale tendencies, with the former providing greater benefits. Overall, as hypothesised (4.1) the GSVJD process outperformed its counterparts, when calibrated to the *HCV* dataset. Nevertheless, the supreme performances came at a cost, given the need for estimating 9 parameters during daily calibrations, on average, 3.56 hours were required. The highly inefficient nature of calibrations makes the GSVJD process unattractive to practitioners, who prefer tractability. However, its quality cannot be ignored during real-time bubble detection. Selecting an alternative process would improve efficiency, but could hamper the conclusiveness of identified bubbles, given the potential existence of a joint-hypothesis

related issue. Therefore, a requirement for applying the deep calibration framework to the GSVJD process, for parametric estimations arose.

The deep calibration framework significantly boosted computational efficiency while even improving accuracy. It operates in two phases: *forward pass*, where neural networks learn the GSVJD dynamics, and *backward pass*, during which parameters are calibrated from market data. The architecture for the *forward pass* was determined from a random search, with the best performing networks being selected for optimal training. In the *backward pass*, frozen hidden layers from each optimally trained network were used as numerical solvers, achieving calibration speeds up to 542 times faster than Monte Carlo simulations, with most networks also outperforming the benchmark, in regard to accuracy<sup>49</sup>. The top architecture featured 4 hidden layers, 40 nodes, with ELU activation and batch size of 2,048, demonstrating a 254-fold increase in efficiency and approximately doubling the calibration accuracy. The significant enhancement in efficiency is consistent with the hypothesis from 4.1, but in regard to accuracy it outperforms the original requirements of preventing large scale sacrifices. A detailed comparison of GSVJD model calibrations from Monte Carlo simulations, and the optimal neural network framework, was conducted. Performances of the latter were already established as superior, but were also revealed to be stabler, consisting significantly fewer error metric spikes, across the sample period. Interestingly, both approaches performed poorly at similar instances, which corresponded to the COVID-19 induced market crash, and GameStop stock short squeeze.

The improved calibration accuracy and efficiency enabled the application of the deep calibration framework to the *Entire Surface* dataset. Remarkably, despite considering a larger number of options, the speed of calibration was not impacted. This could be credited to the benefits of parallel processing achieved from pairing the Differential Evolution optimiser with neural networks. The average error metric was much higher, in comparison to the *HCV* dataset calibrations. Additionally, calibrating to the *Entire Surface* delivered poor performances over a greater period during the COVID-19 induced market crash. However, this is justifiable given the extreme volatile nature of markets, and the potential presence of

---

<sup>49</sup> See Appendix (Section A) for more information on the performance of neural networks during the *backward pass* stage.

relatively illiquid options, across multiple maturities considered during calibrations. Overall, the error metric was stable, building a strong case for continuing with the detection of bubbles, following calibration of the GSVJD model to all daily traded options. Furthermore, it allows for examining the presence of bubbles within various option maturities, and during the lifetime of options. Such applications are difficult to implement in real-time, however the deep calibration framework, enabled their exploration in the next chapter, to test the remaining hypotheses from *4.1*.

# Chapter 5:

# Empirical Analysis

In Chapter 4, the deep calibration framework was applied to estimate GSVJD parameters from daily S&P 500 put options. Over a four-year period, calibration to the most liquid options, revealed a boost in accuracy and computational efficiency by 2 and 254 orders of magnitude, respectively. The improved accuracy demonstrated the neural network's strong comprehension of GSVJD dynamics. Whereas superior efficiency allowed for calibrating from the entire daily put option surfaces. When applied to the entire surface, despite considering a remarkably greater number of option contracts, average calibration time, remained unchanged over the sample period.

The primary purpose of the deep calibration framework is to replace Monte Carlo simulations in the three-step approach. In this approach, the GSVJD is calibrated to daily put options to ensure alignment between fundamental values and market prices. Given that such options cannot display bubbles, accurate calibrations overcome the joint-hypothesis issue, which has notoriously plagued the conclusiveness of traditional methods. Proceeding from this stage, parameters are utilised to compute the fundamental value of call options for bubble detection. If the fundamental value is exceeded by the market price, a call option bubble is witnessed, which signals the underlying process to be a strict local martingale. The

final stage of the three-step approach is dedicated to statistically testing the difference between the fundamental values and market prices.

Recall, under ND, call options with maturity  $\tau$ , can reveal bubbles in the underlying. Since  $\tau < T$ , Fusari et al. (2024) admit the underlying bubble cannot be observed to its fullest extent. It is only possible to capture the entire bubble via call options prices, when  $\tau = T$ , however, as the model horizon is unknown, this is not achievable in practice. Therefore, examining call options can only signal the presence of an underlying bubble. Specifically, for  $\tau < T$ , the magnitude of call option bubbles, act as a lower bound to those in the underlying. Bubble magnitude grows with  $\tau$ , and reaches 0 at option expiry. On a given day, the size, potential existence of a call option bubble, and the ability to address it in the underlying, varies over contract maturities. In Stahl & Blauth (2024) and Biagini et al. (2024), bubbles have been identified by making use of option surfaces. These works motivated this study to calibrate GSVJD parameters from the entire put option surface, such that information on forward-looking expectations of market participants is captured to a greater extent. However, these studies do not examine and differentiate between the magnitude (and existence) of call option bubbles, over different maturities. The current chapter contributes to existing literature by employing GSVJD parameters computed from daily put options within the *HCV* and *Entire Surface* datasets (Table 4.2), to identify, analyse and compare, the presence of S&P 500 index price bubbles, according to various call option maturity categories.

The deep calibration framework plays a massive role in detecting bubbles across various call option maturities. It is due to the computational boost provided by neural networks, that GSVJD parameters can be calibrated from the *Entire Surface* dataset. Moreover, the ability to calibrate a sophisticated stochastic volatility jump diffusion model, such as the GSVJD, within seconds; makes the application of the three-step approach for real time bubble detection more attractive to practitioners. Daily model parameters are calibrated from all traded options, and hence applicable to computing the fundamental call option value, across various maturities. Furthermore, over the sample period of 4 years, daily parameters can be applied to estimating a time series of fundamental values for specific call contracts.

Therefore, creating the possibility of examining the formation of bubbles over the lifetime of options. This is an additional contribution to the asset pricing bubble detection literature.

The following chapter begins with GSVJD parameters calibrated from the *HCV* (Table 4.6) and *Entire Surface* (Table 4.7) being utilised to observe bubbles in the daily most liquid cross section of call options. Next, the latter set of parameters are considered for identifying exuberance within *short* ( $8 \leq \tau < 60$ ), *medium* ( $60 \leq \tau < 180$ ), and *long* ( $180 \leq \tau < 364$ ) call option maturity groups. The categorisation of these call option maturities for bubble detection, is inspired by Jarrow & Kwok (2021, 2024). Finally, those most liquid call options revealing exuberance, according to GSVJD parameters from the entire surface, are selected for examining the formation of bubbles over the option's lifetime. Market analysis of detected bubbles is conducted in reference to reports published by Zacks Investment Research, Inc., unless specified. Prior to delving into the detection of daily S&P 500 bubbles, an overview of required call option datasets, is provided.

## 5.1 Call Options Data: S&P 500 Index

Bubbles in the S&P 500 index are identified by examining European-styled call option contracts, across various maturity groups. In alliance with daily put options from Table 4.2, contracts over a wide strike range at increments of \$5.00, were selected, and filtered using similar moneyness and maturity criteria, with computation of relevant risk-free rates and dividend yields using (49) and (50), respectively (see 4.2.1 for more details). The sample period spans over 4 years between January 2, 2019, and December 30, 2022, comprising of 1,000 trading days, with all market related data being collected from Refinitiv Eikon. Since, call options are utilised for bubble detection all selected contracts satisfy the ND condition from Merton (1973). It is only when this condition is met that bubbles in the underlying and call option prices share a linear relationship, with the magnitude of the latter, acting as a lower bound to that of the former.

The summary statistics of the *HCV* call option dataset is presented in Table 5.1. This dataset comprises of the daily most liquid cross section of call options over the sample period. Similar to *HCV* put options from Table 4.2, the call *HCV* dataset revealed 75% of its

contracts were due to expire within 59 days, showing relatively higher levels of liquidity amongst shorter contracts. A comparison between put and call options, showed the former as, more liquid by an approximate factor of 1.54, and OTM by a greater extent. If the maturity filtering were stricter, as in Fusari et al. (2024), and dropped options with maturity exceeding 50 days, then more than 25% of the study period would not comprise of the most liquid option maturity. The first task of this chapter is to utilise GSVJD parameters calibrated to put options from *HCV* and *Entire Surface* datasets (see Table 4.2), to identify for bubbles in the call *HCV* dataset. Next, parameters from the *Entire Surface* are employed for detecting bubbles across various maturity groups.

**Table 5.1: Summary Statistics of Deep Calibration to HCV call options.**

*Note: Summary Statistics of S&P 500 Call options across the HCV dataset. The statistics are computed across each maturity,  $\tau$ , on a particular day. Volume (000s) reports accumulated values across daily contracts with similar  $\tau$ . Bid-, Ask-, and Mid-Prices are denoted in \$, along with  $mon = \mathbb{K}/\$$ .*

	Mean	SD	P25	P50	P75
Bid	98.87	166.17	3.60	35.10	125.80
Ask	100.59	167.77	4.00	36.10	128.10
Mid	99.73	166.96	3.80	35.60	126.95
Volume	58.59	29.93	38.02	51.94	72.12
$mon$	1.03	0.09	0.99	1.03	1.08
$\tau$	42.91	26.98	23.00	37.00	59.00

The examination of bubbles across various maturity groups is conducted by dividing daily call contracts into three categories; *short*, *medium*, and *long*, with summary statistics for each reported in Table 5.2. The *short* group corresponds to options with  $8 \leq \tau < 60$  days. Since potential bubbles disappear over very small maturities, options with  $\tau \leq 7$  days were discarded. The upper bound for this group was set at 60 days to abide by the average maturity of *HCV* call options (see Table 5.1). The bounds for the *medium* ( $60 \leq \tau < 180$ ) and *long* ( $180 \leq \tau < 365$ ) maturity groups, are set in the accordance to (Jarrow & Kwok, 2021, 2024). Throughout the study, given low liquidity, option contracts with  $\tau > 365$  days were discarded. Additionally, since bubbles are dependent on the option maturity, a single cross-section must be selected for each day. Therefore, for each group, only the daily cross-sections of call options, belonging to the maturity with the highest cumulative volume were considered.

**Table 5.2: Summary Statistics of daily most liquid call options from each Maturity Group.**

Note: Summary Statistics of the daily most liquid call options from each maturity group are displayed in respective panels. Volume (000s) reports accumulated values for daily contracts, across a single  $\tau$ , on a given day. Bid-, Ask-, and Mid-Prices are denoted in \$, along with  $mon = \mathbb{K}/\$$ .

	mean	std	P25	P50	P75
Panel A: Short ( $8 \leq \tau < 60$ )					
Bid	146.10	179.38	35.70	93.30	185.60
Ask	148.65	181.02	37.00	95.40	189.20
Mid	147.38	180.19	36.35	94.40	187.45
Volume	45.43	29.04	26.15	38.84	57.39
$mon$	0.99	0.07	0.96	0.99	1.02
$\tau$	30.81	14.07	18.00	30.00	42.00
Panel B: Medium ( $60 \leq \tau < 180$ ).					
Bid	213.46	218.78	84.70	154.90	254.20
Ask	216.05	220.45	86.10	156.70	257.60
Mid	214.76	219.61	85.40	155.80	255.85
Volume	24.93	25.33	6.49	16.49	36.61
$mon$	0.99	0.09	0.96	1.00	1.02
$\tau$	93.59	24.48	74.00	87.00	108.00
Panel C: Long ( $180 \leq \tau < 365$ ).					
Bid	324.92	272.80	147.80	255.30	409.35
Ask	329.72	274.48	150.93	260.55	415.08
Mid	327.32	273.61	149.66	257.80	412.25
Volume	4.46	5.21	1.40	2.83	5.31
$mon$	0.98	0.12	0.93	0.99	1.04
$\tau$	253.50	54.74	205.00	240.00	302.00

The average maturity of contracts in the *short* group is 30.81 days, compared to 42.91, 93.50 and 253.50 days for the respective *HCV*, *medium*, and *long* call options. The accumulated volume over the daily cross section of options depicts liquidity. As expected, the *short* maturity group has the most liquid call contracts. The average accumulated volume is approximately 2 and 10 times greater than that within the *medium* and *long* groups, respectively. In comparison to the *HCV* dataset (Table 5.1), mean volume is only 1.29 times smaller, showing close resemblance. Overall, the liquidity of the option decreases with its maturity, which is prominently observed over the remaining two groups. Furthermore, the



distribution of moneyness is similar across the three maturity groups, heavily concentrated around ATM options.

**Table 5.3: Summary Statistics of characteristics over the lifetime of Call Option contracts.**

*Note: Summary statistics of daily historical time series of call options that revealed bubbles, in accordance with GSVJD parameters calibrated from the Entire Surface. Bid-, Ask-, and Mid-Prices are denoted in \$, whereas maturity,  $\tau$ , is stated in days, along with  $mon = \mathbb{K}/\$$ . Since the study only considers monthly S&P 500 contracts, all options expire on the third Friday of each month. Call options are categorized by their RICs in respective panels, which resemble the month and year of the contract's maturity.*

RIC (Maturity)	mean	std	P25	P50	P75
Panel A: C20 (March 2020)					
Bid	150.96	169.28	43.20	103.20	193.00
Ask	153.34	171.27	44.20	104.90	195.00
Mid	152.15	170.27	43.70	103.85	193.55
$mon$	0.99	0.08	0.95	0.99	1.02
$\tau$	186.12	103.35	98.50	186.00	274.50
Panel B: E20 (May 2020)					
Bid	145.91	153.72	41.05	104.00	195.00
Ask	149.81	155.95	44.25	106.90	199.25
Mid	147.86	154.82	42.65	105.60	197.27
$mon$	1.03	0.11	0.98	1.02	1.09
$\tau$	99.60	40.71	70.00	100.00	133.00
Panel C: F20 (June 2020)					
Bid	198.48	183.67	70.50	155.30	264.40
Ask	201.62	185.56	72.80	157.00	268.80
Mid	200.05	184.61	71.55	155.90	266.80
$mon$	0.99	0.11	0.94	0.99	1.03
$\tau$	185.43	103.83	94.75	185.50	275.25
Panel D: G20 (July 2020)					
Bid	178.57	180.29	66.00	128.10	224.10
Ask	181.40	181.86	68.20	129.80	228.10
Mid	179.98	181.07	67.15	129.05	226.05
$mon$	0.99	0.09	0.95	0.99	1.03
$\tau$	91.60	48.99	49.75	91.50	133.25

Table 5.3 continued...

Table 5.3 continued...

Panel E: I20 (September 2020)					
Bid	230.15	207.46	96.98	170.35	296.42
Ask	233.03	208.75	99.00	172.60	300.00
Mid	231.59	208.10	98.21	171.52	298.19
<i>mon</i>	0.98	0.11	0.93	0.99	1.01
$\tau$	182.71	102.83	94.00	182.00	270.00
Panel F: L20 (December 2020)					
Bid	263.44	238.15	106.45	190.90	334.75
Ask	267.43	239.76	109.10	194.20	340.05
Mid	265.43	238.93	108.05	192.50	337.55
<i>mon</i>	0.98	0.12	0.93	0.99	1.02
$\tau$	184.73	102.22	97.25	184.50	273.25
Panel G: C21 (March 2021)					
Bid	236.12	202.59	105.20	185.00	299.83
Ask	240.35	204.90	107.00	188.05	306.75
Mid	238.24	203.70	106.10	186.50	303.86
<i>mon</i>	0.97	0.08	0.94	0.98	1.01
$\tau$	186.11	103.25	98.50	186.00	274.50
Panel H: D21 (April 2021)					
Bid	196.15	173.77	83.05	154.50	254.83
Ask	199.41	175.78	84.60	156.35	259.50
Mid	197.78	174.73	83.76	155.40	257.25
<i>mon</i>	0.97	0.06	0.95	0.98	1.01
$\tau$	90.70	49.16	49.00	91.00	133.00
Panel I: F21 (June 2021)					
Bid	243.72	213.96	97.80	194.70	315.30
Ask	247.66	216.47	99.20	198.10	322.05
Mid	245.69	215.17	98.58	196.95	317.95
<i>mon</i>	0.97	0.07	0.94	0.98	1.01
$\tau$	185.49	103.85	94.75	185.50	275.25
Panel J: I21 (September 2021)					
Bid	226.72	220.09	88.30	168.90	283.20
Ask	228.74	221.45	89.10	170.70	286.10
Mid	227.73	220.76	88.65	169.95	284.55
<i>mon</i>	0.97	0.07	0.95	0.99	1.00
$\tau$	181.75	102.06	93.75	180.50	269.25

Table 5.3 continued...

Table 5.3 continued...

Panel K: L21 (December 2021)					
Bid	304.14	312.49	105.00	197.50	367.10
Ask	307.00	314.33	106.40	199.50	370.50
Mid	305.57	313.40	105.85	198.85	368.65
<i>mon</i>	0.96	0.09	0.94	0.98	1.01
$\tau$	184.76	102.23	97.25	184.50	273.25
Panel L: A22 (January 2022)					
Bid	238.77	244.67	95.28	170.55	294.30
Ask	241.16	246.27	96.60	172.50	297.92
Mid	239.96	245.47	95.89	171.58	295.50
<i>mon</i>	0.98	0.07	0.96	0.99	1.01
$\tau$	184.66	102.78	96.50	184.00	271.50
Panel M: C22 (March 2022)					
Bid	215.50	261.76	60.00	143.95	261.33
Ask	217.80	263.37	61.30	145.70	264.55
Mid	216.65	262.56	60.68	144.80	262.71
<i>mon</i>	0.99	0.08	0.97	1.00	1.03
$\tau$	185.65	103.14	97.25	185.50	274.25
Panel N: L22 (December 2022)					
Bid	256.06	257.45	88.80	184.70	331.60
Ask	259.47	259.69	90.40	187.30	336.20
Mid	257.76	258.56	89.60	186.00	334.15
<i>mon</i>	1.00	0.10	0.96	1.00	1.05
$\tau$	185.55	102.84	97.25	185.50	274.25

An entirely novel contribution of examining bubbles over the lifetime of call options is proposed in this chapter. The examination of such phenomena requires construction of a dataset, with historical contract prices, that revealed bubbles when applying GSVJD parameters from the entire surface to *HCV* call options (see 5.2). The summary statistics cross section of call option contracts that reveal bubbles on a given day, are displayed in Table 5.3. They are categorised according to expiry dates and classified by corresponding expiry RICs (Refinitiv Identification Code)<sup>50</sup>. For example, consider the RIC ‘C20’, where

<sup>50</sup> All S&P 500 index option contracts in this study, expire on the third Friday of the relevant month. For more information, please refer to RIC RULES – THOMSON REUTERS SPEED GUIDE. This can be accessed by typing ‘RULES2’ in the Eikon Search bar.

‘C’ and ‘20’ resemble an option expiring in the month and year of *March*, and *2020*, respectively. The time series prices of call contracts aligning with the expiry RICs, are collected whilst abiding by  $7 < \tau < 365$  days condition.

There are 14 unique RICs, each comprising of a wide range of strike prices, as noticed by the *mon* range in *Table 5.3*. The RICs categorize daily cross-sections of call options that have the same expiry date. Over time for each RIC, call contracts reveal decreasing maturities. Hence,  $\tau$  in *Table 5.3*, not only provides information on the number of days, within the sample period for each RIC. On most occasions,  $\tau$  ranges between 8 and 364 days, as implied by the imposed maturity conditions. However, there are some exceptions due to the consideration of only call options that abide by the ND condition. It is crucial for ND to be satisfied, as it ensures a linear relationship between call option and underlying price bubbles. Following the examination of relevant call option datasets, this study proceeds with the detection of bubbles.

## 5.2 Call Option Bubbles: Most Liquid Cross Section

The optimal deep calibration framework architecture was employed towards estimating the parameters of the GSVJD model, from the *HCV* and *Entire Surface* put option datasets. From now onwards, those parameters are respectively denoted as  $\theta_{HCV}$  and  $\theta_{Surface}$ . Daily calibrations were conducted within seconds, at high levels of accuracy, improving attractiveness of the three-step approach to practitioners. Furthermore, accurate calibrations to market put options, provided strong evidence of overcoming the joint hypothesis issue. In this section, fundamental values of the most liquid call options are computed for detecting daily bubbles in the S&P 500 index. A comparison is made amongst exuberance observed from capturing forward looking expectations for the most liquid smile or entire surface.

A bubble is revealed when the market price of the call option exceeds it fundamental value. Following comparing modelled and market prices, the magnitude of the call option bubble, as a percentage of the underlying price,  $\widehat{\mathbb{B}}_t$ , is computed (see 3.1.3), such that exuberance on a given day, is significant, if  $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ , where  $\alpha$  corresponds to the level of significance, and  $\widehat{\sigma}_t$  is dependent past information over a given number of days. A significant call option

bubble signals the presence of strict local martingale behavior in the underlying price. Recall, under the ND condition bubbles in call option and underlying prices share a linear relationship, with the magnitude of the former acting as the lower bound to that of the latter. The empirical analysis of daily exuberance observed in the S&P 500 index, during January 2019 and December 2022, is initiated by selecting four different window sizes: 30, 60, 90, and 180 days, over three levels of significance,  $\alpha = \{10\%, 5\%, 1\%\}$ .

The window size determines the extent to which information from historical observations are considered, whilst examining the significance of the bubble. This might make the method seem backward looking, however, past information required is with respect to call option prices, which by nature are forward-looking. A trade-off is prevalent during the selection of window sizes, which control the quantity of past information utilised for assessing the significance of exuberance. The probability of false detection increases with smaller windows, whereas larger ones improve on this issue due to their more conservative nature. However, the conservative property possessed by longer window lengths, could downplay the impact of volatile periods, raising the probability of false negative errors.

The selection of a window length can be interpreted as being analogous to the risk-appetite of investors, or alternatively the strictness in authority imposed by regulators (when demanding capital reserves in accordance with the observed period of exuberance). For instance, a more conservative detection approach would signal lower regulatory pressure on capital requirements, and a riskier outlook from investors, who would reveal looser preferences towards protecting positions from the potential presence of bubbles. Hence, it is integral to consider different window sizes, and levels of significance, when constructing the time-varying threshold,  $\alpha \hat{\sigma}_t$ , during real-time bubble detection.

## 5.2.1 Key Statistics

Traditional methods directly estimate the fundamental value and provide no consensus of assessing the validity of detected bubbles. Alternatively, the three-step approach gives validity by estimating the fundamental value of call options using GSVJD parameters, calibrated from market put option prices. Recall, due to their bounded payoffs, put options

cannot display bubbles, and therefore their fundamental and market values must align. By ensuring this, bubble detection does not suffer from the joint hypothesis issue. The  $RMSE_{\sigma,t}$  metric measures the degree of accuracy between modelled and market put option implied volatilities during calibrations. It is an important metric for assessing the presence of joint-hypothesis related issues that could plague the conclusiveness of observed bubbles.

As witnessed in *Table 4.6* and *Table 4.7*,  $RMSE_{\sigma,t}$  was used to evaluate the overall accuracy of the deep calibration framework, when applied to the respective *HCV* and *Entire Surface* datasets. Similarly, the error metric is utilised to examine the performance of the GSVJD model, on days that display exuberance ( $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ ) in the S&P 500 index. Along with the  $RMSE_{\sigma,t}$  metric, the number of significant bubbles ( $\mathbb{N}_B$ ) between January 2019 and December 2022, with respect to each combination of window sizes and levels of significance, are reported in *Table 5.4*, after considering a burn-in period.

**Table 5.4: GSVJD Performances on days with significant HCV call bubbles.**

The  $RMSE_{\sigma,t}$  (%) error metric, along with number of significant bubbles ( $\mathbb{N}_B$ ), and percentage of the sample period (excluding burn-in) revealing exuberance,  $\mathbb{N}_B(\%)$ , reported for days displaying bubbles within *HCV* call options for different levels of significance ( $\alpha$ ) and windows sizes in respective panels, with regard to  $\theta_{HCV}$  and  $\theta_{Surface}$ .

	$\alpha$	$\mathbb{N}_B$	$\mathbb{N}_B(\%)$	$RMSE_{\sigma,t}(\%)$
<i>Panel A: 30 days</i>				
<i>HCV</i>	10%	53.00	5.46%	0.14%
	5%	31.00	3.19%	0.17%
	1%	11.00	1.13%	0.31%
<i>Entire Surface</i>	10%	39.00	4.02%	0.95%
	5%	28.00	2.88%	1.05%
	1%	19.00	1.96%	1.06%
<i>Panel B: 60 days</i>				
<i>HCV</i>	10%	45.00	4.78%	0.18%
	5%	24.00	2.55%	0.21%
	1%	10.00	1.06%	0.33%

*Table 5.4 continued...*

Table 5.4 continued...

	10%	34.00	3.61%	0.98%
<i>Entire Surface</i>	5%	21.00	2.23%	1.09%
	1%	13.00	1.38%	1.29%
<i>Panel C: 90 days</i>				
	10%	36.00	3.95%	0.19%
<i>HCV</i>	5%	25.00	2.74%	0.21%
	1%	11.00	1.21%	0.31%
	10%	32.00	3.51%	1.05%
<i>Entire Surface</i>	5%	25.00	2.74%	1.08%
	1%	15.00	1.65%	1.24%
<i>Panel D: 180 days</i>				
	10%	32.00	3.90%	0.18%
<i>HCV</i>	5%	19.00	2.31%	0.23%
	1%	6.00	0.73%	0.24%
	10%	31.00	3.78%	1.00%
<i>Entire Surface</i>	5%	23.00	2.80%	1.07%
	1%	17.00	2.07%	1.19%

The frequency of bubbles observed from calibrating to *HCV* put options is greater, especially over smaller window sizes and more lenient levels of significance. As the strictness of detection increases, the difference between number of significant days of exuberance from both parameter sets reduces. Moreover, at the 1% level, *Entire Surface* calibrations capture bubbles to a greater extent. Such parameters also reveal low variability in identified occurrences, across all windows. This indicates that the bubble detected from *Entire Surface* calibrations are more robust, and not victims of the trade-off associated with the selection of window lengths. Unfortunately, the same cannot be concluded about *HCV* calibrations, and hence, a comprehensive market analysis is required to examine robustness.

Alternatively,  $RMSE_{\sigma,t}$  metrics corresponding to  $\theta_{Surface}$  seem to paint a concerning picture, despite robustness across windows. The error metric on days of exuberance, across the window size are larger than those associated with *HCV* calibrations. Furthermore, in all scenarios,  $RMSE_{\sigma,t}$  are greater than the 75<sup>th</sup> percentile of 0.86% (*Table 4.7*) from over the entire sample. The metrics from  $\theta_{HCV}$  are significantly lower but do tend to exceed their respective 75<sup>th</sup> percentile (*Table 4.6*) as well. Although, this could raise concerns about a

joint-hypothesis issue, both parameters set reveal bubbles during the COVID-19 induced crash period, when the GSVJD performances poorly. This is especially relevant to  $\theta_{Surface}$ , as a large proportion of its identified instances occur between March 13, and April 30, 2020, a period during which  $RMSE_{\sigma,t}$  averaged at 1.41%. Nevertheless, these events are economically justifiable in consistency with the local martingale theory of bubbles.

The identification of such occurrences is actually reflected as a strength capturing forward looking expectations of market participants from the entire surface (see 5.2.2). Given the varying number of options used during calibrations, the  $RMSE_{\sigma,t}$  error metrics corresponding to  $\theta_{Surface}$ , are not necessarily comparable to those obtained from  $\theta_{HCV}$ . Recall, the *Entire Surface*, on average, comprises of 9 daily maturities including several relatively illiquid contracts. Regardless, calibration performances are strong, with bubbles detected being robust over the different windows. However, for further comparison and comprehension of identified exuberance, a market analysis is required.

## 5.2.2 Market Analysis

Bubble detection in the S&P 500 index prices, between 2019 and 2022, is conducted through inspecting *HCV* call options. Under the ND condition, a significant call option bubble reveals presence of strict local martingale tendencies in the underlying price process. The linear relation between the two holds, such that the magnitude of the former acts as the lower bound to that of the latter. Call option bubbles are estimated by assessing the difference between market prices and fundamental values, with the latter being computed from GSVJD parameters. The construction of the deep calibration framework significantly reduced computational inefficiency and enabled the determination of  $\theta_{HCV}$ , and  $\theta_{Surface}$ . The market analysis of bubbles is tailored towards investigating exuberance detected by using  $\theta_{HCV}$ , followed those corresponding to  $\theta_{Surface}$ . Furthermore, the appearance of significant bubbles,  $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ , over varying levels of significance and window sizes, is analysed chronologically.

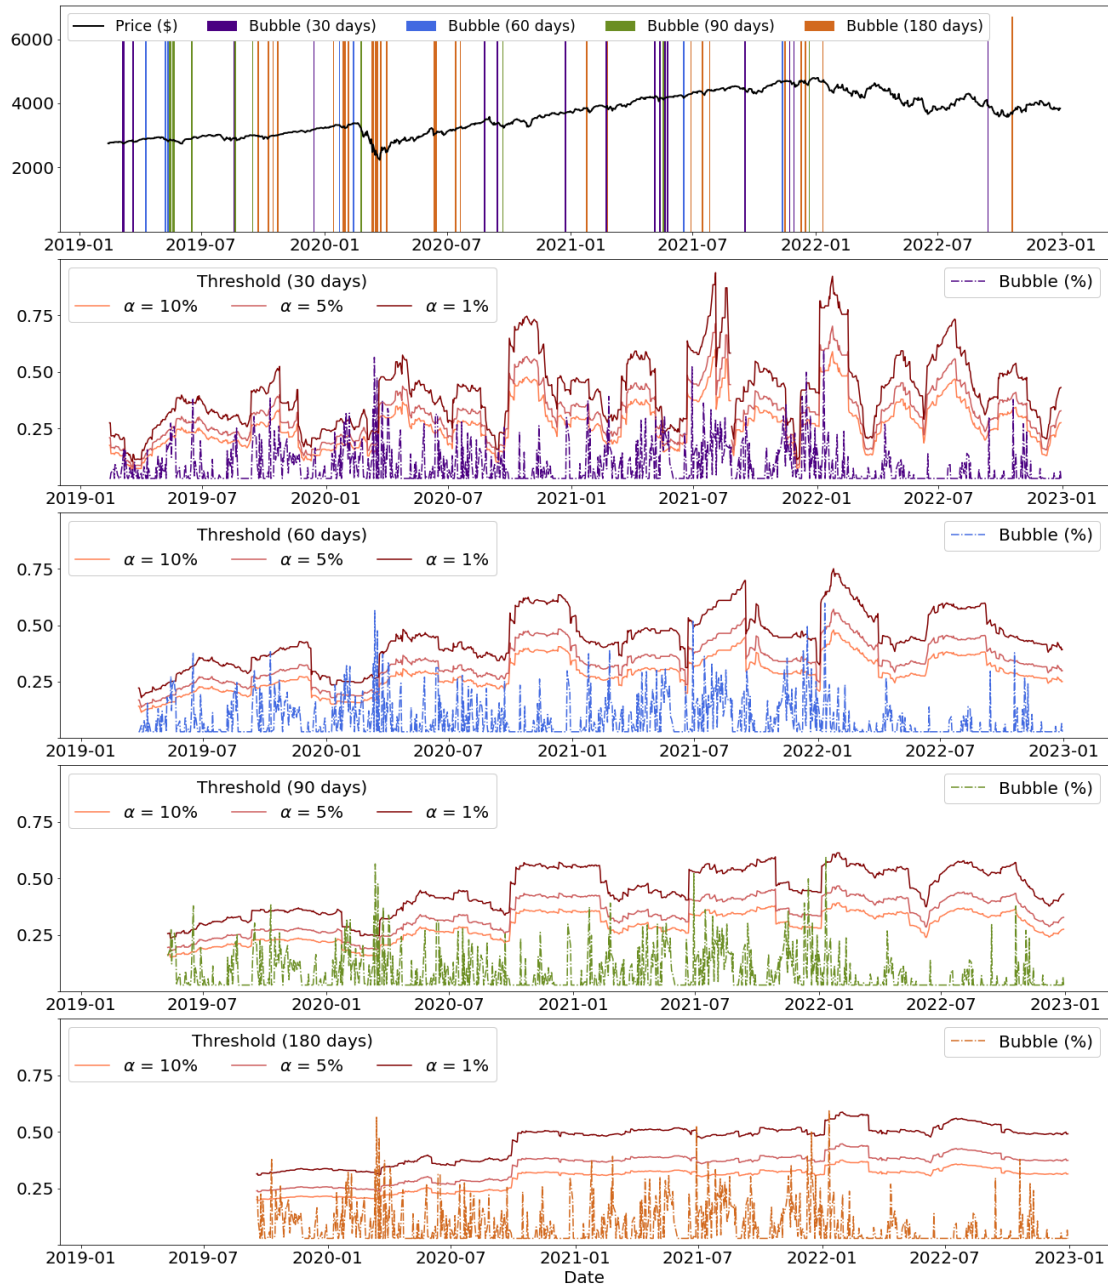


### 5.2.2.1 Bubble Detection: HCV Calibration

Significant call option bubbles ( $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ ) identified across all window sizes and levels of significance, using  $\theta_{HCV}$ , are illustrated in *Figure 5.1*. In the figure (*top*), bubbles are timestamped and displayed alongside daily S&P 500 index prices. The remaining subplots (*bottom*) demonstrate daily call bubble magnitudes ( $\widehat{\mathbb{B}}_t$ ) against the time-varying threshold ( $\alpha \widehat{\sigma}_t$ ), with respect to different window sizes and levels of significance, for each day in the sample period. Recall, the magnitude is normalised with respect to the underlying price. The accumulation of past information for estimating the time-varying threshold begins at the start of the sample. However, bubble detection initiates after a burn-in period, analogous to the selected window size. The time-varying threshold cannot be computed during this period; hence this initial component of the sample has been acknowledged as an empty space. In the timestamping plot (*top*), significant bubbles from all window sizes are displayed, therefore the initial burn-in period resembles that of the smallest (30-day) window.

The market analysis of observed bubbles in *Figure 5.1*, is conducted chronologically. The first group of events are witnessed on March 6, 8, 22 and April 10, 2019. This minor cluster is also singled out to explain the interpretation of *Figure 5.1*. Consider exuberance on March 6, which is detected at 10% significance when employing a 30-day window. Corresponding to the latter, it is timestamped (*top*) in *indigo*. Its magnitude is 0.13% of the underlying price, \$2,771.45, which approximates to \$3.60. This value resembles the lower bound for the size of the underlying price bubble, and when scaled by 100 ( $\approx \$360.29$ ), provides magnitude of the call option bubble ( $N = 100$ ). The events on March 8 and 22 are both revealed within the same window, as well. However, the bubble on April 10, is acknowledged by 30- and 60-day windows. Hence, it is timestamped (*blue*), in accordance with the largest window that reveals significance. Therefore, timestamping of exuberance in *Figure 5.1* also highlights robustness across various conservative levels of detection.

**Figure 5.1: S&P 500 Call Option Bubbles (HCV Calibration).**



*Note: Daily S&P 500 bubbles detected in HCV call options using  $\theta_{HCV}$ . Bubbles are timestamped (top) and assessed at various levels significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ), with respect to each window size (bottom).*

The aforementioned 4 bubbles coincided with growing uncertainty over U.S. and China reaching an agreement to curb the on-going trade-war. The deteriorating relations between both nations, along with a global economic slowdown, significantly hurt investor sentiments throughout the year. A cluster of 7 bubbles was observed between May 9, and 22, as strong

earning boosted markets. However, China placing retaliatory measures in response to U.S. raising tariffs on Chinese goods, along with weak economic data from the Asian nation, dampened investor confidence. The exuberance on August 19 and 21, coincided with central banks of both nations favouring interest rate cuts to tackle economic slowdown. Alternatively, the bubbles on October 9, 10, 16, and 23, aligned with improving trade relations between two countries, and a strong earnings quarter. The final bubble of 2019 was detected on December 16, marking the U.S. and China signing the first phase of a trade deal, and clearance of the USMCA.

A cluster of 8 bubbles was observed between January 14 and February 24, 2020. Investor sentiments were dominated by rising fears over travel and trade disruptions, as the spread of coronavirus cases in China accelerated. Nevertheless, strong corporate earnings, especially those from technology companies, bolstered markets. The occurrences detected on January 30, 31 and February 5 are noteworthy. The S&P 500 closed 1.77% lower on January 31, as the U.S. documented its first coronavirus case. Alternatively, the index rose by 1.50% on February 4, 2020, given increased manufacturing activity due to eased trade tensions with China. However, the exuberance on February 24, coincided with a 3.35% collapse in the S&P 500 index, as the global coronavirus pandemic worsened and demands for safe assets, such as gold and government bonds heightened. This cluster reflected the willingness of investors seeking short-term gains, especially during earnings seasons, among times of building anxiety.

Eventually, the global pandemic induced a market crash, and a cluster of 8 bubbles was identified between March 11 and April 3, 2020. A circuit breaker was implemented in the opening hours of March 12, as the S&P 500 closed 9.51% below. Exuberance was recorded on March 13, after a 9.29% rebound, but investors lost complete confidence on March 16, 2020, causing the index to crash by approximately 12%, triggering a circuit breaker. An additional circuit breaker was experienced on March 18, but bubbles were detected on March 17 and 19, with the former corresponding to a 6.00% gain. The capturing of these bubbles is partially consistent with the research hypothesis (see 4.1). During this period, the S&P 500 experienced 4 circuit breakers, which ensured that the market price did not crash to the fundamental value. The exuberance detected by  $\theta_{HCV}$  parameters occur prior to/after days

witnessing circuit breakers, indicating traders possessing the ambition to purchase only for reselling at higher prices. Nevertheless, bubbles were not detected on the days of circuit breakers. Additionally, exuberance was observed on March 24 and 25, as the S&P 500 gained 9.38%, and 1.15%, respectively. The conclusion of the crude oil price war between Russia and Saudi Arabia, coincided with bubbles on April 2 and 3, with respective returns of 2.28% and -1.51%. However, two days prior, the S&P 500 suffered consecutive losses of 1.61% and 4.41% due to the extension of lockdown measures.

The S&P 500 index was also hypothesised to reveal exuberance during its post-crash recovery phases. However,  $\theta_{HCV}$  parameters displayed very few instances till the end of the year. Most notably, bubbles were spotted on June 12 and 15, as the S&P 500 gained 1.31% and 0.83%, respectively, following a 5.89% crash on June 11, 2020. Furthermore, individual instances were captured on July 14, and 21, in addition to that on August 26. The latter is of particular interest as the S&P 500 reached an all-time high. However, consumer confidence levels fell to a 6-year low in August, due to pessimism about the short-term economic outlook. The technology sector drove market recovery, and upon receiving encouragement over COVID-19 treatments, the S&P 500 rose by 1.27% on September 14. Nevertheless, exuberance was identified amid concerns over the overvaluation of technology stocks and delayed fiscal stimulus. Similarly, despite gaining 1.05% on September 22, the index reflected a bubble, as sentiment dropped due to rising coronavirus cases in Europe. The final bubble of 2020, occurred on December 24, coinciding with rising cases of a new variant, falling consumer confidence, and even more delays in expected fiscal stimulus.

At the turn of the year, bubbles were detected on January 25, and February 23 and 25, 2021, with the latter day experiencing a large equity sell-off due to rising treasury bond yields, causing the S&P 500 to close 2.45% below. Next, a cluster of 7 bubbles was observed between May 6 and 25, a period during which, solid job market data and indications of faster-than-expected recovery from the pandemic losses, boosted market sentiments. However, concerns remained about the existing high levels of unemployment, growing inflationary pressure, and potential interest rate hikes. The escalation of such fears triggered a large selling spree of technology stocks, and subsequently exuberance on May 18. Furthermore, bubbles were exhibited over the next three consecutive days, and May 25. The technology

sector took a massive hit following China providing warnings over the issuance of cryptocurrencies, rising inflation, and the Federal Reserve hinting at tightening of the monetary policy. Moreover, the European Central Bank raised concerns over market exuberance due to the imbalance created by monetary and fiscal stimuli.

Consistent with the research hypothesis from 4.1, bubbles were observed in 2021, as the S&P 500 surged, despite growing concerns over inflation and rising coronavirus cases. The next instances occurring on June 18, July 16 and 27, and September 17, were individualistic in nature. However, towards the end of the year, a larger concentration of bubbles was observed. Strong earnings from major corporations and rallying technology stocks lifted market sentiments to reveal instances on November 12 and 16. Furthermore, bubbles were spotted on November 23, 26 and 29. The S&P 500 dropped by 2.72% on November 26, as WHO declared the Omicron variant as a source of concern. However, it recovered by 1.32%, on November 29, after assurance of the non-lethal characteristics of the variant. Finally, the exuberance on December 10, 16, 17, and 22 coincided with the expected increase in tapering of the quantitative easing programme, unexpected interest rate hikes from the Bank of England and implementation of lockdown measures in Europe.

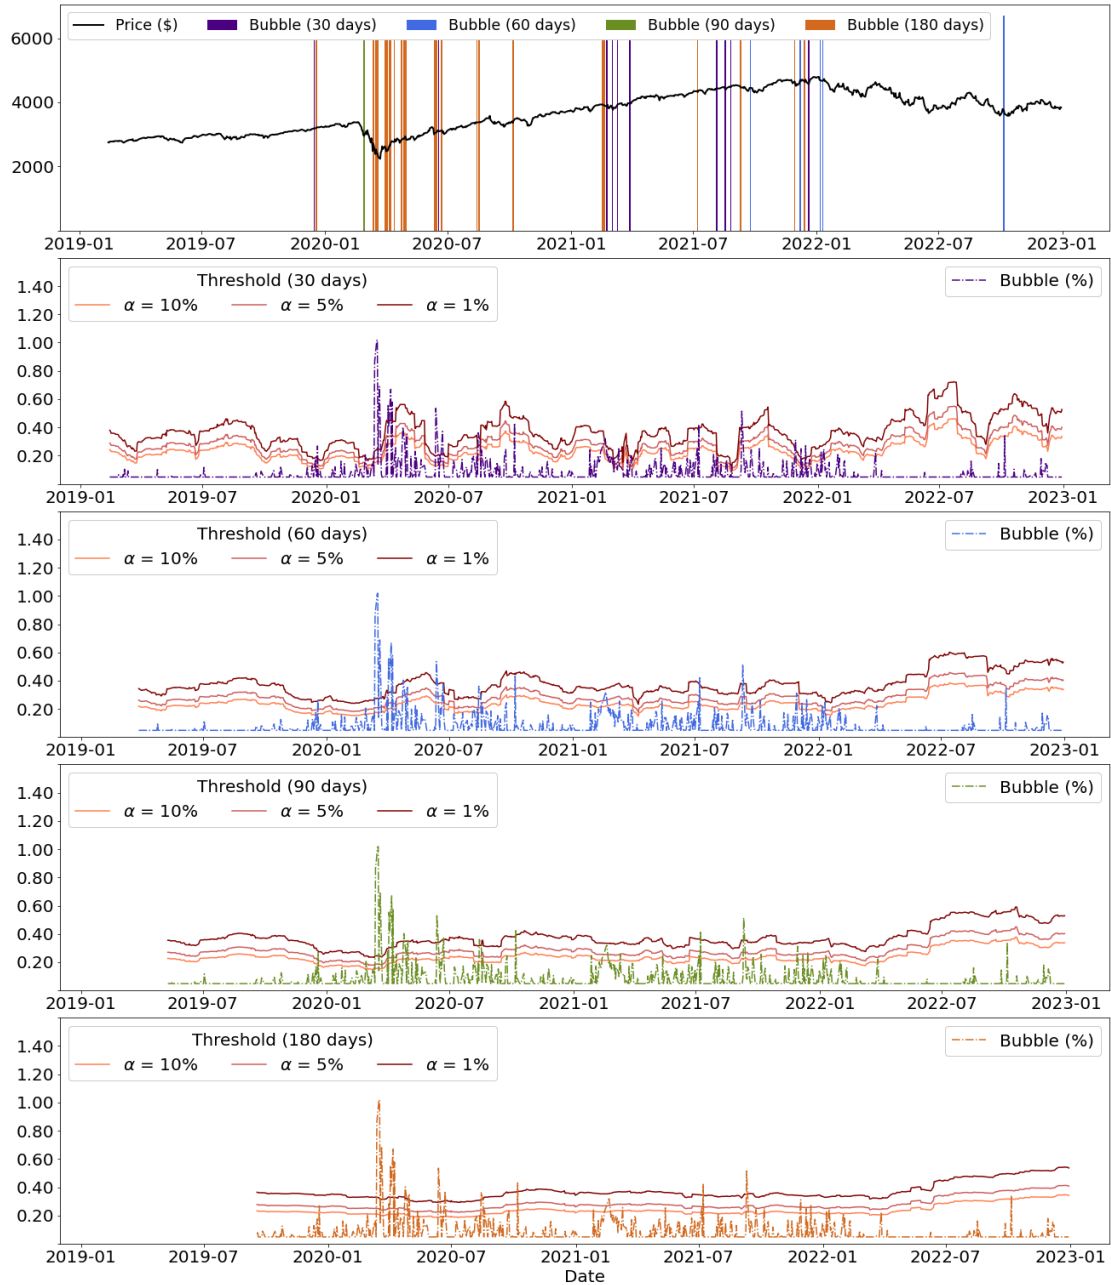
In 2022, the S&P 500 exhibited 3 bubbles on January 11, September 14 and October 20, according to  $\theta_{HCV}$  calibrations. The former corresponded to rising treasury yield hurting technology stocks, and concerns over policymakers signalling at least 3 interest rate hikes in the upcoming year. Nevertheless, following the announcement of such measures only being implemented after completion of the tapering programme, technology and growth stocks rallied. The S&P 500 collapsed by 4.32% on September 13, amid increased inflationary concerns, despite the Federal Reserve continuing to hike rates. Exuberance on the next day captured fears of an upcoming third consecutive 75 bps rate hike. The final instance coincided with the 10-year Treasury yield reaching a 14 year peak.

### 5.2.2.2 Bubble Detection: Surface Calibration

The study proceeds with identifying significant call option bubbles ( $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ ) in accordance with  $\theta_{Surface}$ . The days of exuberance are displayed in *Figure 5.2*, which follows

similar interpretation as *Figure 5.1*, with bubbles being timestamped (*top*) and their magnitudes revealed, along with time-varying thresholds across difference window lengths (*bottom*). The initial sample period corresponding to the selected window size, is burned and acknowledged as an empty space.

**Figure 5.2: S&P 500 Call Option Bubbles (Surface Calibration).**



*Note: Daily S&P 500 bubbles detected in HCV call options using  $\theta_{surface}$ . Bubbles are timestamped (top) and assessed at various levels significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ), with respect to each window size (bottom).*

Exuberance was first spotted on December 16, and 19, 2019, coinciding with U.S. and China reaching a trade-agreement, along with progression of USMCA. The former instance was captured by  $\theta_{HCV}$ , amongst the several others across 2019 that addressed concerns over a global economic slowdown and continuous failed negotiations between both nations. Moreover, in January and February 2020,  $\theta_{HCV}$  revealed exuberance, following boosted sentiments due to a strong earnings season, despite building fears over the coronavirus outbreak. During the week commencing on February 24, the S&P 500 lost 11.89%, with the largest loss of 4.41% occurring on February 27, 2020. A bubble was revealed by  $\theta_{HCV}$  on February 24, whereas  $\theta_{Surface}$  identified exuberance on February 28, acknowledging concerns of the global pandemic forcing the closure of factories in Japan, South Korea and Italy.

The coronavirus pandemic entirely crushed investor sentiment, inducing a market crash in March 2020. Bubbles were detected on March 13, 16, 17, 19 and 20, with all barring those from March 16, and 20 being acknowledged by  $\theta_{HCV}$ . Circuit breakers were implemented on March 12 and 16, as the index closed lower by 9.51% and 11.98%, respectively. The instances on March 13, and 17 coincided with respective gains of 9.29% and 11.98%, respectively. Moreover, another circuit breaker was experienced on March 18, with exuberance being observed on March 19 and 20. Once again consistency with the research hypothesis (4.1) was maintained, as immediate days following circuit breakers exhibited bubbles. This revealed strong willingness of traders to purchase for the sole purpose of reselling at a higher price, in search of a quick buck. However, in contrast to  $\theta_{HCV}$ , a bubble was detected on March 16, a day that experienced a circuit breaker, and the largest losses during this period.

Applying  $\theta_{Surface}$  revealed greater consistency with the research hypothesis (4.1), as it acknowledged exuberance during the immediate recovery of the S&P 500 index. Bubbles were observed on March 31, April 2, 3, 6, 8 and 13, 2020, with only those on April 2 and 3, being addressed by  $\theta_{HCV}$ . Following consecutive losses of 1.61% and 4.41%, the S&P 500 experienced a 2.28% rise on April 2, after the conclusion of the crude oil price war. Moreover, hikes of 7.03% and 3.41% were witnessed on April 6 and 8, respectively, as the Federal Reserve doubled down on its commitment to a dovish stance. Furthermore, bubbles

were seen on all days between April 23 and 30, barring April 28. The index experienced respective gains of 1.47% and 2.66% on April 27 and 29, given strong possibilities of reopening the economy, positive news regarding drug trials, and commitment towards a dovish monetary policy. However, rising unemployment and inflation, falling consumer spending and poor quarterly earnings from major corporations continued to weigh heavy on investor sentiment.

During the S&P 500 recovery phase, till the end of 2020,  $\theta_{Surface}$  detected few instances of exuberance, as opposed to the research hypothesis (4.1). Nevertheless, a likewise observation was made to the application of  $\theta_{HCV}$ . Bubbles were identified on June 12 and 15, after a 5.89% crash on June 11. However,  $\theta_{Surface}$  also spotted instances on June 18 and 22, with investor sentiments were dampened by warnings from the WHO over acceleration of the pandemic. Next, exuberance was witnessed on August 14 and 17, as the technology sector continued to bolster market recovery, such that all COVID-19 induced losses were erased by August 18, 2020. The last bubble of 2020 appeared on October 7, as President Trump signed off an individual stimulus package. The S&P 500 rose by 1.74% on this day, following a 1.40% loss from the previous.

The first bubble episodes in 2021 were observed on all days between February 16 and 19, and February 23. The S&P 500 underwent a 5-day losing streak which ended on February 23, as rising treasury yields and increased concerns over inflation triggered a large equity sell-off. The rising treasury yields caused a large bond sell-off, further harming technology stocks. Consequently, exuberance on March 3 coincided with a 1.31% dip. Nevertheless, following the passing of a long awaited coronavirus relief package, the S&P 500 gained 1.04% and revealed a bubble on March 11. Next, occurrences were spotted on March 29 and 30, as consumer confidence reached a post pandemic high, and technology stocks received a massive boost from the vaccination drive. However, the 10-year Treasury yield reached a 14 year high, indicating great concerns over inflation and equity overvaluation.

At the immediate turn of the year, exuberance was detected on January 6 and 10, 2022. The former coincided with expectations that the Federal Reserve might hike rates, sooner than expected. However, upon confirmation that it would occur after the tapering programme,



technology stocks rallied. The latter corresponded to rising treasury yields adversely impacting growth stocks. Recall,  $\theta_{HCV}$  captured formation of exuberance due to similar reasons, on January 11, 2022. The final bubble occurred on October 6, 2022, as the S&P 500 fell by 1.03%. It aligned with rising treasury yields, as despite aggressive interest rate hikes, inflation remained high. The Federal Reserve doubled down on its aggressive stance, even though strong job market numbers made market participants expect a tamer approach.

### 5.2.3 Model Performance

The identification of daily bubbles using  $\theta_{HCV}$  and  $\theta_{Surface}$  reflected investor sentiments over market prices exceeding the fundamental value. These bubbles were short-lived, and in most instances followed, and preceded days that experienced large negative returns. This reflected the intentions of traders to purchase shares for the sole purpose of reselling at a higher price, with sales occurring at the slightest indication of panic. The detection of bubbles relied on a statistical test (see 3.1.3), designed to exhibit significant exuberance in call option prices. Under the ND condition, bubbles in the call option and underlying prices are linearly related, such that the magnitude of the former acts as the lower bound that of the latter.

The  $RMSE_{\sigma,t}$  metrics, with respect to the test, are observed in *Table 5.5*. In abidance with Fusari et al. (2024) they are examined to assess the robustness of bubbles detected from  $\theta_{HCV}$  and  $\theta_{Surface}$ . The local martingale theorem of bubbles highlights that negative bubbles do not exist, since the price of the asset is bound from below by its fundamental value. However, they can be present in the three-step approach framework as a result of model misspecification introducing a measurement error (Fusari et al., 2024). A key component to assess validity of bubbles would be the evaluation of  $RMSE_{\sigma,t}$  on days where  $\widehat{\mathbb{B}}_t < 0$ . The negative magnitudes would occur due to noise in the model, and are expected to display larger error values, in comparison to days with  $\widehat{\mathbb{B}}_t > 0$ . This would provide evidence that the identified days of exuberance are less affected by the noise.

The  $RMSE_{\sigma,t}$  metrics corresponding to the  $\theta_{HCV}$  are greater for  $\widehat{\mathbb{B}}_t < 0$ , revealing calibrations of the GSVJD to be less impacted by noise for days with  $\widehat{\mathbb{B}}_t > 0$ . Overall, the strong calibration performances on  $HCV$  put options, provide validation for bubbles. Alternatively, the error metrics, across all three-categories are larger, when considering the *Entire Surface*. Furthermore, the  $RMSE_{\sigma,t}$  metric for such calibrations, is observed to be greater for days with  $\widehat{\mathbb{B}}_t > 0$ . However, as seen in *Table 5.5*, the difference is small, at most, not exceeding 0.11%, across all window sizes. It is important to recollect that daily calibrations on this dataset are conducted across, on average 9 volatility smiles. This approximately translates to increasing the number of options from 161 to 677, with the inclusion of several relatively illiquid contracts.

**Table 5.5: Model Performance.**

*Note:  $RMSE_{\sigma,t}$  (%) metrics with respect to  $\theta_{HCV}$  and  $\theta_{Surface}$ , corresponding to the entire sample (excluding burn-in), along days with negative ( $\widehat{\mathbb{B}}_t < 0$ ) and positive ( $\widehat{\mathbb{B}}_t > 0$ ) bubble magnitudes, across all window sizes.*

Window Size	$\theta_{HCV}$			$\theta_{Surface}$		
	<i>Entire</i>	$\widehat{\mathbb{B}}_t < 0$	$\widehat{\mathbb{B}}_t > 0$	<i>Entire</i>	$\widehat{\mathbb{B}}_t < 0$	$\widehat{\mathbb{B}}_t > 0$
30	0.167%	0.182%	0.150%	0.750%	0.713%	0.825%
60	0.168%	0.184%	0.151%	0.757%	0.722%	0.829%
90	0.171%	0.187%	0.152%	0.763%	0.728%	0.831%
180	0.171%	0.188%	0.151%	0.794%	0.769%	0.835%

The greater values of  $RMSE_{\sigma,t}$  from *Entire Surface* calibrations, are attributed to the detection of bubbles during the peak of the COVID-19 induced market crash, between March 13 and April 30, 2020. During this period, the error metric averaged at 1.41%, nearly twice the mean performance across the entire sample. Concerns over the presence of a joint-hypothesis related issue could be raised, however, the exhibition of these bubbles has comprised of an economic justification, especially since they occur close to/on days that have experienced circuit breakers. They are consistent with the NFLVR condition, which only considers admissible strategies. This provides validation to the detection of bubbles in the S&P 500 index calibrating the GSVJD model to all traded options, with the assistance of the neural network framework. Recall, even  $\theta_{HCV}$  calibrations struggled during this period, and captured almost the same events between March 13 and April 3, 2020.

Compared to  $\theta_{HCV}$ , bubble detection from *HCV* call options using  $\theta_{Surface}$  has revealed more consistency with the hypothesis from 4.1. Both exhibited bubbles on immediate days following circuit breakers in March 2020, capturing the willingness of traders to purchase for the sole purpose of reselling at a higher price to earn quick profit. Nevertheless, only the latter captured exuberance on a day that experienced a circuit breaker. Furthermore,  $\theta_{Surface}$  also addressed exuberance at the end of April 2020, as the S&P 500 index started to recover, despite large levels of unemployment. However, beyond this date, barring individualistic instances, neither set of parameters revealed bubbles when the S&P 500 surged throughout the rest of 2020 and 2021, despite the dire condition of the economy. Hence, this creates motivation to explore the application of  $\theta_{Surface}$  on various call options for bubble detection.

Options reflect the future expectations of market participants regarding the value of the underlying asset. Each cross-section of daily contracts belongs to a unique maturity and reveals forward-looking expectations over the corresponding duration. Daily calibrations to the *Entire Surface* dataset, have been beneficial for detecting bubbles, as estimated parameters capture this information to the fullest extent possible. It is well documented that call options bubble magnitudes act as the lower bound to those present in the underlying. Derivative contracts expire prior to the future liquidation date of the risky asset and hence cannot capture the entire underlying bubble. However, even in such circumstances, if bubbles are observed over small durations, they are capable of reflecting the presence of exuberance over larger periods of time (Fusari et al., 2024). Hence, as seen in Jarrow & Kwok (2021, 2024), it is worth exploring the presence of bubbles across various call option maturity groups. This is possible due to the availability of  $\theta_{Surface}$ , that are applicable to all option maturities.

### 5.3 Call Option Bubbles: Maturity Groups

The deep calibration framework improved accuracy and efficiency of estimating GSVJD model parameters from market put options data for bubble detection via the three-step approach. This allowed for the consideration of information from all traded put options on a

given day. Calibrations to several daily volatility smiles, ensured the inspection of a more complete market, enabling the capturing of forward-looking expectations to a greater extent, without relying on interpolation and extrapolation methodologies.

Calibration of the GSVJD model to the *Entire Surface* dataset has been beneficial, in reflecting information, obtained from several daily volatility smiles. For a given maturity (smile), corresponding options capture forward-looking expectations of market participants over the remaining contract life. Such call contracts expire before the future liquidation date of the underlying and hence cannot capture the full extent of exuberance in the underlying price. As highlighted by Fusari et al. (2024), bubbles in option prices are observed over small durations, and can indicate towards exuberance over larger time periods. These bubbles can be explored by analysing call option across various maturity groups. Since  $\theta_{Surface}$  are obtained from the daily surface of put options, the parameters are applicable for estimating the fundamental values of call options across various maturity cross sections. The exuberance in S&P 500 index are detected by examining three maturity groups of call options (see Table 5.2): *short* ( $8 \leq \tau < 60$ ), *medium* ( $60 \leq \tau < 180$ ), and *long* ( $180 \leq \tau < 365$ ). Only the daily most liquid cross-section of call options (within each group), that abide by the ND condition, are selected. A bubble is considered significant if  $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ , with the latter being the time-varying threshold estimated over the same window sizes and levels of significance, as in 5.2.2.

### 5.3.1 Key Statistics

The  $\theta_{Surface}$  parameters, were applied towards pricing call options for identifying exuberance, across three different maturity groups: *short*, *medium*, and *long*. Recall, the  $RMSE_{\sigma,t}$  metric measures calibration performance whilst obtaining  $\theta_{Surface}$ , and plays an important role in examining the presence of joint-hypothesis related issues. The average error metric across days exhibiting significant bubbles over various combinations of window sizes and levels of significance are reported for each maturity group, in Table 5.6.

Across all three maturity groups, the number of observations increase and decrease, with significance levels and window sizes, respectively. Detection at the 10% and 5% levels of

significance, within the *short* group provide exception of the latter. Moreover, *Table 5.6* reveals the  $RMSE_{\sigma,t}$  metric to be greater with significance in all scenarios, barring the implementation of a 180-day window on *long* maturity contracts. This maturity group identifies the least instances, maintaining consistency with type III bubbles arising due to short-term trade strategies that reveal the forward-looking nature of investor expectations. Furthermore, exuberance must be significantly larger for it to be flagged by long term investors. Alternatively, the largest number of occurrences are associated with the *medium* maturity group. Once again, maintaining consistency with the local martingale theory, given that bubble magnitude disappears as the options tends to maturity (Fusari et al., 2024).

**Table 5.6: Calibration performances on days with significant bubbles across maturity groups.**

*Note:  $RMSE_{\sigma,t}$  (%) error metric, along with number of significant bubbles ( $N_B$ ), and percentage of the sample period (excluding burn-in) revealing exuberance,  $N_B(\%)$ , reported for days displaying exuberance within most liquid cross section of call options, from short, medium and long maturity groups, in accordance with  $\theta_{Surface}$ .*

Window Size (days)	$\alpha$	$N_B$	$N_B$ (%)	$RMSE_{\sigma,t}$
Panel A: <i>Short</i> ( $8 \leq \tau < 60$ )				
30	10%	25.00	2.68%	0.79%
	5%	18.00	1.93%	0.81%
	1%	7.00	0.75%	0.95%
60	10%	27.00	2.99%	0.80%
	5%	13.00	1.44%	0.91%
	1%	7.00	0.77%	1.02%
90	10%	21.00	2.40%	0.82%
	5%	15.00	1.72%	0.84%
	1%	4.00	0.46%	1.32%
180	10%	27.00	3.44%	0.80%
	5%	16.00	2.04%	0.88%
	1%	3.00	0.38%	1.27%
Panel B: <i>Medium</i> ( $60 \leq \tau < 180$ )				
30	10%	83.00	8.55%	0.90%
	5%	65.00	6.69%	0.92%
	1%	44.00	4.53%	0.94%
60	10%	73.00	7.76%	0.94%
	5%	55.00	5.84%	0.94%
	1%	34.00	3.61%	1.03%

*Table 5.6 continued...*

Table 5.6 continued...

	10%	65.00	7.14%	0.99%
90	5%	48.00	5.27%	1.04%
	1%	28.00	3.07%	1.14%
	10%	57.00	6.94%	0.96%
180	5%	32.00	3.90%	0.99%
	1%	11.00	1.34%	1.33%
Panel C: Long ( $180 \leq \tau < 365$ )				
	10%	16.00	1.76%	1.56%
30	5%	11.00	1.21%	1.83%
	1%	5.00	0.55%	1.98%
	10%	11.00	1.25%	1.95%
60	5%	7.00	0.80%	2.33%
	1%	5.00	0.57%	1.98%
	10%	11.00	1.30%	1.94%
90	5%	7.00	0.82%	2.33%
	1%	4.00	0.47%	2.34%
	10%	8.00	1.05%	1.69%
180	5%	4.00	0.53%	2.34%
	1%	2.00	0.26%	1.76%

Since calibrations are conducted on the *Entire Surface* dataset, the error metrics can be compared those with similar correspondence, displayed in *Table 5.4*. The days exhibiting exuberance within the *short* and *medium* maturity groups, provide smaller  $RMSE_{\sigma,t}$  values in comparison to those from bubbles detected within *HCV* call options. In regard to both groups, detection at 1% significance using a 180-day window, in addition to a 90-day window on the former, act as exceptions. These exceptions are justified by majority of detected bubbles occurring during the COVID-19 induced market turmoil between March and April 2020. Alternatively, *long* maturity group comprise of larger error metrics, with respect to its counterparts, and those from *Table 5.4*. As observed in 5.2.1, calibration to the *Entire Surface* raised concerns over the presence of a potential joint-hypothesis related issue. Once again, a major contributor to this observation, is the capturing of exuberance during the COVID-19 induced market crash period, especially over higher levels of significance and greater window periods. Majority of bubbles captured using the *long* maturity group occurred between March and April 2020, explaining its alliance with high error metrics in *Table 5.6*. However, these instances are consistent with the local martingale theory of

bubbles and are closely analysed in the next subsection. Hence, asserting importance of statistical analysis being accompanied by market evaluation to assess the validity of the bubbles.

### 5.3.2 Market Analysis

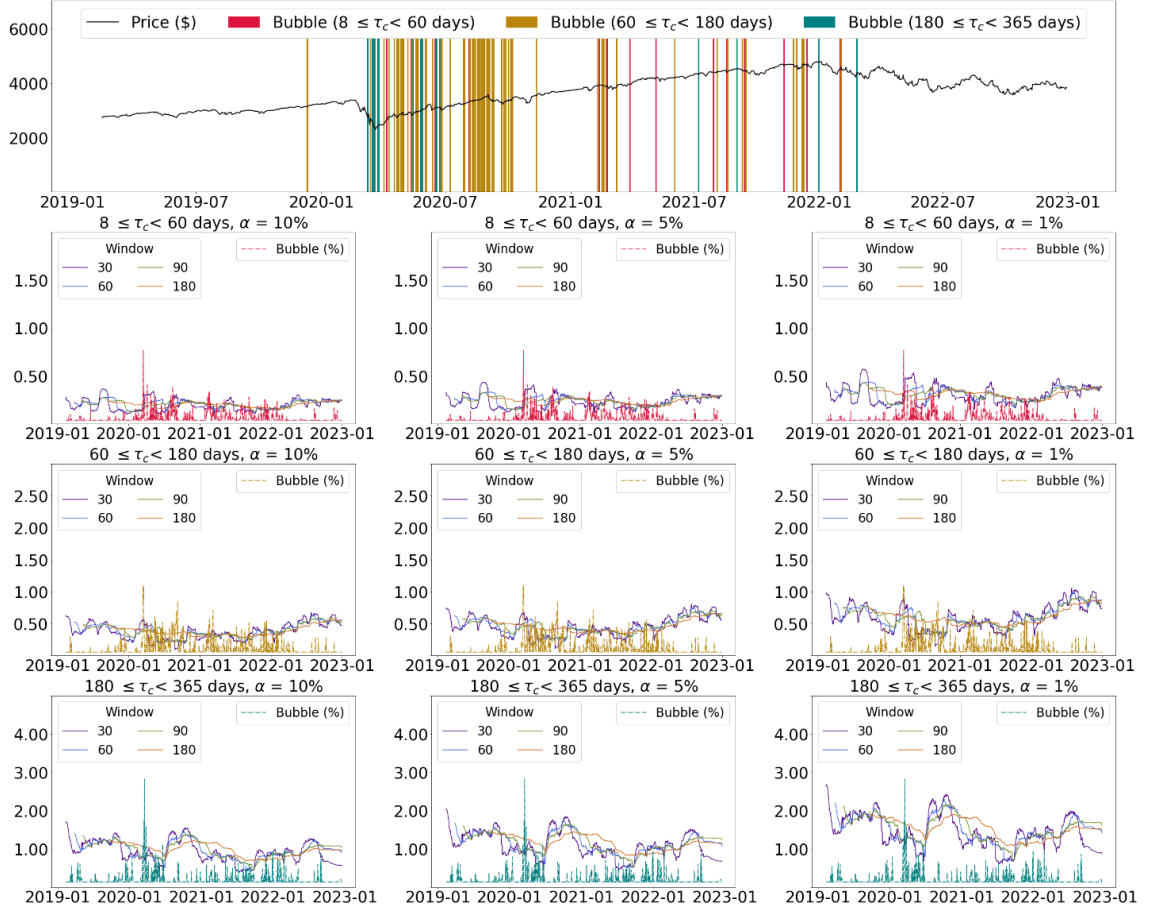
A significant call option bubble is detected when  $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ , with the latter being the time-varying threshold, estimated over different levels of significance ( $\alpha$ ) and window sizes. The detection of exuberance is conducted by inspecting the most liquid daily cross section of call options, within the three maturity groups. For each group, fundamental call values are computed using  $\theta_{Surface}$ , and compared with market observations. The identified bubbles, across  $\tau_C \in (short, medium, long)$  in the S&P 500 between January 2, 2019, and December 31, 2022, are displayed in *Figure 5.3*.

Market analysis of exuberance is conducted in chronological order, with the interpretation of *Figure 5.3* being different to those witnessed earlier. The timestamping of bubbles is revealed (*top*) along with the underlying price. Only two episodes were observed in 2019, both witnessed by  $\tau_M$  ( $60 \leq \tau < 180$  days) call options, on December 11, and 12. In *Figure 5.3*, exuberance is dated with respect to maturity groups, regardless of the window size, and levels of significance. Consider the former instance, dated in golden, since it was seen only by the medium maturity group. The subplots (*bottom*) display magnitude of call bubbles ( $\widehat{\mathbb{B}}_t$ ) against the time-varying threshold from all window sizes, at a certain level of significance, for each maturity group. Continuing with the example, the exuberance has a magnitude of 0.25% and was observed using a 30-day windows at 10% significance.

Bubbles observed by *short* and *long* maturity groups are depicted in *red* and *teal*, respectively. Consider the exuberance on March 17, 2020, which was visible in all three groups. It is timestamped in *teal*, in association with the largest maturity group that acknowledged the event. Interestingly, the bubble magnitudes, 0.74%, 0.82%, and 0.83%, were documented to increase with call maturity. This property held for daily bubbles that were detected across various maturity groups. It is consistent with the local martingale theory, as the size of exuberance increases with the option's remaining life. Hence,

examining bubbles across various maturity levels, on a given day, allows to quantify their presence in the underlying price, to a greater extent.

**Figure 5.3: S&P 500 Call Option Bubbles (Surface Calibration).**



Daily bubbles detected in the short ( $8 \leq \tau < 60$ ), medium ( $60 \leq \tau < 180$ ), and long ( $180 \leq \tau < 365$ ) call option maturity groups, written on the S&P 500 index (2019-2022), using  $\theta_{surface}$ . Bubbles are timestamped (top) and assessed at each combination of window sizes (30, 60, 90, and 180 days) and levels of significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ), with respect to each maturity group (bottom).

The detection of bubbles across various call option maturity groups shows strong consistencies with the research hypothesis (see 4.1). A cluster of 13 episodes, between March 9 and April 9, 2020, is captured across all categories. During this period,  $\tau_S$  ( $8 \leq \tau < 60$  days) call options identified bubbles only on March 17, 19, and April 6, with the latter corresponding to a 7.03% hike. Alternatively,  $\tau_M$  contracts detected bubbles on April 2 and 9. Furthermore, the same group exhibited bubbles on all days between March 13 and 30, whereas  $\tau_L$  ( $180 \leq \tau < 365$  days) contracts observed instances on March 9, 16, 17, 18, and 19. Previously, using  $\theta_{HCV}$  and  $\theta_{surface}$  on *HCV* call options, highlighted bubbles on days



succeeding those that experienced circuit breakers; March 13, 17 and 19. Moreover,  $\theta_{Surface}$  also acknowledged exuberance on March 16. Nevertheless,  $\tau_M$  and  $\tau_L$  observe episodes on days that experienced circuit breakers; March 16 and 18, in addition to March 9 being addressed by the latter. The benefits of extending the application of the three-step approach to various maturity groups, reveals more consistency with the local martingale theory of bubbles. Finally,  $\tau_L$  options flagged exuberance on March 24, 25, and 26, as the S&P 500 rallied by 9.38%, 1.15% and 6.24%, respectively.

The exploration of bubbles within the various call option maturity groups, continued to reveal consistency with the research hypothesis. During the recovery period, 31 instances were reported between April 17 and June 26, 2020. At first 10 episodes were witnessed between April 17 and May 1, 2020, with all barring that on April 29, being captured by  $\tau_M$  options. The exception, along with that on April 23, was detected by the  $\tau_S$  group. Recall, exuberance on April 23, 24, 27, and 30, were also acknowledged by *HCV* call options, when using  $\theta_{Surface}$ . All trading days between April 21 and May 1, 2020, revealed exuberance. The S&P 500 experienced gains of 2.29%, 1.39%, 1.47% and 2.66% on April 22, 24, 27 and 29, respectively. However, in the aftermath of the pandemic, investor sentiment was dominated by the massive unemployment levels, and weak corporate earnings.

The 15 events of exuberance between May 7 and June 3, 2020, were not observed in *HCV* call options. Once again, majority of the instances were acknowledged by  $\tau_M$ , with those seen on May 14; and May 15, 26 and 28, also spotted by  $\tau_S$  and  $\tau_L$  options, respectively. In addition, the  $\tau_S$  group uniquely identified exuberance on May 12 and 20. Investor confidence was recovering as the economy underwent a month and half long partial lockdown period, despite several warnings being received over the easing of restrictive measures. A hike of 3.15% was noticed on May 18, following assurance from the Federal Reserve about economic recovery, and the FDA providing approval to a vaccine. However, worsening geopolitical relations with China, and increased unemployment continued to dampen market optimism.

An additional 6 days of exuberance were spotted between June 12 and 26, 2020. The event on June 16 was only seen in  $\tau_S$  options, whereas that on June 23, was acknowledged by  $\tau_M$

and  $\tau_L$  contracts. Moreover, the bubble on June 18, was detected across all three maturity groups, while the remaining occurrence were identified by  $\tau_M$  call options. Recall, the application of  $\theta_{Surface}$  ( $\theta_{HCV}$ ) on *HCV* call options spotted exuberance on June 12, 15, 18, and 22 (June 12 and 15), as well. The episodes reflected concerns over a second coronavirus wave, weak earnings from major corporations, and restricted business activity. Furthermore, 3 instances were observed by  $\tau_M$  options, in July 2020. A massive technology stock rebound aligned with exuberance on July 8, as the coronavirus spread continued to intensify. The bubbles on July 28 and 29, corresponded to weak investor sentiment derived from disappointing earnings, falling consumer confidence, and rising cases; despite the Federal Reserve maintaining interest rates, and on-going negotiations over new fiscal stimuli.

Previously, as observed with *HCV* call options, exuberance episodes during the recovery period of the S&P 500 index in 2020, were sparse. However, that is not the case when considering call options across various maturity groups. Between August 5 and September 10, 2020; 25 bubbles were identified, with all but the first being detected by  $\tau_M$  call options. In addition to the occurrence on August 5, those of August 10, 14, 17, and 19, were also acknowledged by  $\tau_S$  contracts. The bubbles on August 14, and 17, were even revealed by *HCV* call options, using  $\theta_{Surface}$ . Initially, market sentiments were boosted by developments over the vaccine front, strong optimism regarding a fiscal stimulus agreement, and fiscal provision for airline companies. Furthermore, the S&P 500 index had erased all of its COVID-19 losses by August 18, 2020. The technology sector gained further momentum as the Federal Reserve remained committed to its expansionary policy. However, consumer confidence index dropped to a 6-year low, and weekly jobless claims began to cross 1 million. The index crashed by 3.51% on September 3, 2020, triggered by Dr. Anthony Fauci's comments from the previous day, stating that the vaccine will not be prepared till the year end. The technology sector faced a massive sell-off and lost approximately 11% of its value by September 8, 2020. The exuberance on September 9 and 10, coincided with respective returns of 2.02% and -1.76%.

The remaining bubble episodes in 2020, were captured by  $\tau_M$  call options and witnessed between September 22, and November 12. Exuberance on the September 22, and October 7, were acknowledged by applying the respective  $\theta_{HCV}$  and  $\theta_{Surface}$  on *HCV* contracts.

Bubbles were detected on September 22, 25, and 28, as the S&P 500 index grew by 1.05%, 1.60% and 1.61%, respectively. Despite rising coronavirus cases across Europe, and the possibility of implementing fresh lockdown measures, investor sentiments were lifted by a rebound of technology stocks and unveiling of the long awaited \$2.2 trillion new coronavirus aid package. The reopening of negotiations over the latter, did boost markets and aligned with exuberance on October 5, 7, and 8. The S&P 500 fell by 1.40% on October 6, as President Trump decided to postpone negotiations. However, with immediate effect, an individual stimulus package was passed. Finally, the occurrence on November 12, coincided with rising coronavirus cases, that made investor fearful over re-imposition of business restrictions.

The first episodes in 2021 were revealed between February 10 and March 29, with occurrences on February 10, 12, 24, and March 29 (February 22, March 9 and 10) being addressed solely by  $\tau_s$  ( $\tau_M$ ) call options. Both  $\tau_s$  and  $\tau_M$  captured exuberance on February 11, 16, 18 and 19, as all barring the former were detected by pricing *HCV* contracts with  $\theta_{Surface}$ . Moreover, *HCV* call options acknowledged instances of February 23 and March 11 and 29; and February 23 and 25, using  $\theta_{Surface}$  and  $\theta_{HCV}$ , respectively. The S&P 500 went on a 5-day losing streak till February 23. A large sell-off was triggered by rising treasury yields, as inflationary and equity overvaluation concerns grew, forcing technology and other growth stocks to take a dive. On March 10, the coronavirus relief package was passed by the House of Representatives. The S&P 500 benefited from improving consumer confidence, and technology sector, both supported by a strong vaccination drive. However, in the background, inflation and overvaluation fears were creeping up as the 10-year Treasury yield reached a 14 year peak.

Inflationary concerns, in addition to fears over rising unemployment and the Federal Reserve potentially raising interest rates,  $\tau_s$ , similar to  $\theta_{HCV}$ , reflected exuberance on May 6, 2021. The  $\tau_M$  call option exhibited a bubble on June 3, as economic activity and employment levels started to pick up, but inflation reached its highest since September 2008. Next, 6 episodes were spotted between July 8 and August 20. Those on July 8 and August 18 were detected by  $\tau_M$  and  $\tau_L$  groups. Furthermore, exuberance on July 30 and August 18; and August 20, were respectively identified by  $\tau_s$  and  $\tau_M$ , whereas both groups revealed bubbles on August

4 and 19. The bubble on July 8, coincided with growing concerns over global economic recovery as Asian nations experienced a resurgence in coronavirus cases. On August 18, the S&P 500 fell by 1.08%, marking a 6-consecutive day in red. Investors expressed concerns over rising coronavirus cases and the Federal Reserve prematurely pulling back the quantitative easing programme and harming recovery. However, over the next two days, technology stocks rebounded following a drop in treasury yields.

The inflationary pressure and concerns over rising coronavirus cases continued to build over the next month. The  $\tau_L$  group observed exuberance on September 2, whereas both  $\tau_S$  and  $\tau_M$  options addressed bubbles on September 10, similar to pricing  $HCV$  contracts with  $\theta_{Surface}$ . Individually, in respective order, the maturity groups identified exuberance on September 13 and 15. The S&P 500 rebounded after undergoing a 5-day losing streak that ended on September 10. Furthermore,  $\tau_S$  contracts acknowledged an event on November 10, and the  $\tau_M$  group revealed one on November 24, whereas both, along with  $\theta_{HCV}$  highlighted exuberance on November 29. Investors regained confidence over the new coronavirus variant and uncertainty over the monetary policy stance. The  $\tau_S$  and  $\tau_L$  call option detected a bubble on December 7 to coincide with a 2.07% spike. Aligning with rallying markets being held back by anxiety over inflation,  $\tau_M$  spotted an instance on December 9. As technology stocks dipped following indications of faster quantitative easing tapering, and potentially earlier than expected rate hikes,  $\tau_S$  contracts noticed a bubble on December 14. The final episode of 2021 was captured by  $\tau_L$  call options, on December 31.

The only 3 bubble episodes observed in 2022, occurred on February 1, 2 and 25. The former was detected by  $\tau_S$  options, whereas the instance on February 2 (25) were acknowledged by  $\tau_S$  and  $\tau_M$  ( $\tau_M$  and  $\tau_L$ ) options. Technology stocks suffered throughout January 2022, as market participants feared interest rates to be hiked, as early as March. Nonetheless, rising treasury yields provided bank stocks with momentum, and along with strong corporate earnings, supported the market rally. The final exuberance coincided with the S&P 500 rising by 1.50% and 2.24% on February 24 and 25, respectively, despite Russia initiating its attack on Ukraine. Investor confidence stemmed from Russia agreeing to send delegates to Ukraine, and the collective decision of the U.S., U.K. and European Union, to sanction Russia.

### 5.3.3 Model Performance

Bubble detection across different maturity groups, was possible due to the efficiency provided by the deep calibration framework in computing  $\theta_{Surface}$ . Across the entire sample,  $\tau_S$ ,  $\tau_M$  and  $\tau_L$  revealed exuberance to be short-lived, however, to last longer than that exhibited by *HCV* call options. Furthermore, bubble magnitudes increased with time to maturity, providing more insight on the potential size of exuberance in the underlying.

**Table 5.7: Model Calibration Performances over Maturity Groups.**

Note:  $RMSE_{\sigma,t}$  (%) metrics of  $\theta_{Surface}$ , corresponding to the entire sample (excluding burn-in) period, along days with negative ( $\widehat{\mathbb{B}}_t < 0$ ) and positive ( $\widehat{\mathbb{B}}_t > 0$ ) bubble magnitudes, for each maturity group (short, medium and long), across all window sizes.

Window Size	Entire	$\widehat{\mathbb{B}}_t < 0$	$\widehat{\mathbb{B}}_t > 0$
Panel A: <i>Short</i> ( $8 \leq \tau < 60$ )			
30	0.748%	0.719%	0.808%
60	0.756%	0.728%	0.812%
90	0.761%	0.734%	0.814%
180	0.793%	0.776%	0.821%
Panel B: <i>Medium</i> ( $60 \leq \tau < 180$ )			
30	0.750%	0.699%	0.842%
60	0.757%	0.707%	0.847%
90	0.763%	0.712%	0.848%
180	0.794%	0.751%	0.856%
Panel C: <i>Long</i> ( $180 \leq \tau < 365$ )			
30	0.740%	0.723%	0.784%
60	0.748%	0.732%	0.789%
90	0.753%	0.734%	0.806%
180	0.785%	0.772%	0.818%

As seen in 5.2.3, the statistical validity of the approach can be inspected by examining the  $RMSE_{\sigma,t}$  metric, presented in Table 5.7. The local martingale theory of bubbles prevents existence of negative magnitudes, given that the fundamental value acts as the lower bound to market price. Hence, Fusari et al. (2024) documented the presence of such occurrences as potential model misspecification. Overall, the  $RMSE_{\sigma,t}$  metrics highlight strong

resemblance to those from *Table 5.5*, as calibration is conducted over the *Entire Surface* dataset, and bubbles are observed around similar dates. Error on days with  $\widehat{\mathbb{B}}_t > 0$ , is marginally greater, with the maximum difference not exceeding 0.14%. This minor inferiority can be attributed to the identification of call option bubbles during the COVID-19 induced crash period, between March and April 2020. These days of exuberance are consistent with the local martingale theory. Hence, eliminating any concerns over the presence of any potential joint hypothesis related issues, and justifying the pioneering of the GSVJD calibration framework such that existence of bubbles can be examined from the perspective of various call option maturities. Given call option bubbles signal strict local martingale tendencies in the underlying price, such phenomena can be examined to a greater extent by capturing the various duration of the forward-looking expectations from the market participant's preference of option maturity.

## 5.4 Call Option Bubbles: Option Lifetime

In the previous sections, S&P 500 bubbles were examined by exploring daily call options across various maturity groups. It allowed to perceive the existence of such phenomena, with respect to the forward-looking expectations of market participants, over different time durations that were associated with the expiry of call contracts. This practice was possible due to the computational boost provided by the deep calibration framework, in regard to estimating  $\theta_{Surface}$ . Since daily calibrations were conducted across all traded options,  $\theta_{Surface}$  are capable of computing fundamental values of call options across various maturities.

Though call option bubbles, under ND, signal towards exuberance in underlying prices, they are unable to reflect the entirety of the occurrence. However, the analysis across various maturity groups revealed, to a greater extent, the size and presence of bubbles. This ability relies on the selection of call option maturities, alternatively, the duration over which the forward-looking expectations of market participants are to be examined. In the current section, it is proposed to explore bubbles within specific call option contracts. This allows for capturing forward-looking expectations regarding the formation of bubbles, over the option life, without being restricted by maturity or liquidity classifications. Once again, such

practice is possible due to the superior computational abilities of the deep calibration framework, as  $\theta_{Surface}$ , is used to estimate the required fundamental values. The daily cross section of *HCV* call options that reveal bubbles when using  $\theta_{Surface}$ , are selected and categorized by maturities, with the assistance of RICs (see 5.1).

A combined 14 unique RICs are considered, with all historical call prices of each being examined in abundance to the ND and  $7 < \tau < 365$  days conditions being satisfied. The former is to ensure that call option bubbles are linearly related to those in the underlying, whereas the latter is consistent with filtering of datasets throughout the study. Furthermore, the maturity condition also ensures that the selected call options are explored, till at most, 1 year prior to their expiry. This allows for studying the evolution of a call option bubble, and its indication towards the presence and magnitude of that in the underlying price.

### 5.4.1 Key Statistics

The  $\theta_{Surface}$  parameters are utilised for computing fundamental values of call options, within each of the 14 RIC categories. Specific to each category, daily cross-sections of call options are utilised in the construction of the time-varying threshold for identifying significant bubbles ( $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ ). Since the  $7 < \tau < 365$  days condition is implied, each RIC covers a period of approximately 1 year, within the entire sample (January 2, 2019 - December 31, 2022). Hence, to prevent loss of a large data portion, bubbles are detected using only 30- and 60-day windows, within three levels of significance. Finally, the ND condition ensures that call option bubbles signal towards strict local martingale tendencies in the underlying.

The key statistics in relation to bubbles detected over the lifetime of the various options are reported in *Table 5.8*. Prior to analysing them, it is important to acknowledge that only 11 RICs are included, as no days of exuberance were identified within the historical prices of call options expiring in May (*E20*) and July (*G20*), 2020 and April 2021 (*D21*). Furthermore, in the remaining contracts, bubble detection occurs over different time periods, making it difficult to compare error metrics across the RICs. However,  $RMSE_{\sigma,t}$  values, corresponding to exuberance displayed within contract with 2020 maturity, standout. The *C20*, *F20*, *I20*

and  $L20$  RICs reveal relatively larger average error values, in comparison to their counterparts.

**Table 5.8: Surface Calibration Performance over Call Option Lifetimes.**

*Note:  $RMSE_{\sigma,t}$  (%) error metric, along with number of significant bubbles ( $N_B$ ), and percentage of the sample period (excluding burn-in) revealing exuberance  $N_B(\%)$ , reported for days displaying  $\widehat{B}_t > \alpha \widehat{\sigma}_t$ , in respective panels for various RICs.*

Window Size (days)	$\alpha$	$N_B$	$N_B$ (%)	$RMSE_{\sigma,t}$
Panel A: C20 (March 2020)				
30	10%	4.00	1.87%	2.14%
	5%	3.00	1.40%	2.12%
60	10%	1.00	0.54%	3.09%
Panel B: F20 (June 2020)				
30	10%	9.00	4.19%	1.47%
	5%	6.00	2.79%	1.81%
	1%	5.00	2.33%	1.90%
60	10%	7.00	3.78%	1.71%
	5%	6.00	3.24%	1.81%
	1%	4.00	2.16%	1.83%
Panel C: I20 (September 2020)				
30	10%	13.00	6.13%	1.14%
	5%	8.00	3.77%	1.17%
	1%	4.00	1.89%	1.58%
60	10%	7.00	3.85%	1.50%
	5%	4.00	2.20%	1.67%
	1%	2.00	1.10%	1.95%
Panel D: L20 (December 2020)				
30	10%	36.00	16.74%	1.03%
	5%	30.00	13.95%	1.04%
	1%	17.00	7.91%	1.10%
60	10%	20.00	10.81%	1.04%
	5%	13.00	7.03%	1.15%
	1%	7.00	3.78%	1.27%

Table 5.8 continued...



Table 5.8 continued...

Panel E: C21 (March 2021)				
30	10%	5.00	2.34%	0.74%
	5%	2.00	0.93%	0.68%
	1%	1.00	0.47%	0.65%
60	10%	1.00	0.54%	0.65%
	5%	1.00	0.54%	0.65%
Panel F: F21 (June 2021)				
30	10%	13.00	6.05%	0.72%
	5%	10.00	4.65%	0.70%
	1%	2.00	0.93%	0.58%
60	10%	3.00	1.62%	0.81%
	5%	1.00	0.54%	0.87%
	1%	1.00	0.54%	0.87%
Panel G: I21(September 2021)				
30	10%	4.00	1.90%	0.78%
	5%	3.00	1.42%	0.75%
60	10%	1.00	0.55%	0.84%
	5%	1.00	0.55%	0.84%
Panel H: L21 (December 2021)				
30	10%	6.00	2.79%	0.72%
	5%	1.00	0.47%	0.84%
Panel I: A22 (January 2022)				
30	10%	3.00	1.40%	0.81%
	5%	3.00	1.40%	0.81%
Panel J: C22 (March 2022)				
30	10%	3.00	1.40%	0.76%
	5%	2.00	0.93%	0.73%
60	10%	1.00	0.54%	0.69%
Panel K: L22 (December 2022)				
30	10%	1.00	0.47%	1.12%

Similarly,  $L22$  call options displaying higher  $RMSE_{\sigma,t}$  but since only a single occurrence is detected, it cannot be categorised with 4 aforementioned RICs. The group of RICs coinciding with large  $RMSE_{\sigma,t}$  values, all expire in 2020, and since the sample period does not exceed 365 days, their error metrics are justified by the detection of exuberance within the COVID-19 induced market crash period, during which the GSVJD provides weaker calibration

performances. This is consistent with previous findings and eases concerns over the presence of potential joint-hypothesis related issues, given alliance with the local martingale theory of bubbles. The frequency of exuberance is the most within call options maturing in December 2020 (*L20*). Over the various combinations of window sizes and levels of significance, bubbles cover 3.78% to 16.74% of its specific sample period. In comparison to the rest, the proportion of the sample period depicting exuberance, ranges from 0.47% to 6.13%. The width of the range, especially at the lower end can be attributed to the lesser number of bubbles detected over stricter levels of significance and larger windows. Overall, in most instances, fewer occurrences of exuberance are detected, with average  $RMSE_{\sigma,t}$  metrics, being below 1% and close to those from  $\theta_{Surface}$  calibrations (*Table 4.7*).

## 5.4.2 Market Analysis

Recall, significant call option bubbles ( $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$ ) reveal the presence of exuberance in the underlying price. The magnitude of the former, which increases with time to maturity, acts as the lower bound to the size of the latter. The bubble within historical prices of daily cross-section of selected call options (categorized by a similar expiry date), over a period are studied. This allows for comprehending the forward-looking expectations regarding bubble formation, over the lifetime of call option contracts. Additionally, it enables the examination of factors and various market scenarios that impact investor perceptions. Such perceptions are heavily influenced by the option's remaining time to maturity. The market analysis of bubbles over the lifetime of call options is conducted chronologically for each RIC and presented in order of corresponding expiry dates, categorised by years. The RICs were selected based on the condition that they revealed bubbles when applying  $\theta_{Surface}$  to *HCV* call options.

### 5.4.2.1 Call Options Expiring in 2020

The application of  $\theta_{Surface}$  on *HCV* call options revealed bubbles in 4 call contracts, that expire in 2020. The respective *C20*, *F20*, *I20*, and *L20* RICs indicate March, June, September and December 2020 maturities. The identified bubbles, across the different levels of


significance and window sizes are seen in *Figure 5.4*. Bubbles are timestamped (*top*), and their magnitudes are examined against the time-varying thresholds (*bottom*), constructed over past information, with respect to the selected window size. For each RIC, timeseries of the cross-section of call option prices are considered, such that the  $7 < \tau < 365$  condition is satisfied. Accordingly, following a 30-day initial burn-in period, for the *C20* call options, the sample ranges between May 3, 2019, and March 12, 2020.

The *C20* call options, as seen in *Figure 5.4 (Panel A)*, reveal 4 instances of exuberance. The first on December 12, 2019, coincides with progression of the USMCA trade deal, along with U.S. and China being close to signing a trade agreement. It was observed using a 30-day window and therefore is timestamped in *indigo*. Moreover, a similar bubble was acknowledged by  $\tau_M$  call options. Next, occurrences were detected on February 27 and 28, 2020, addressing concerns over the spread of global coronavirus pandemic, and forced closure of factories across major European and Asian nations. The latter instance was identified by employing  $\theta_{Surface}$  on *HCV* call options. A circuit breaker was triggered on March 9, as the S&P 500 collapsed by 7.60%. The  $\tau_L$  contracts observed exuberance on this day, however, *C20* RIC captured a bubble on March 10, as the S&P 500 gained 4.94%. This instance was witnessed using the 30-and 60-day window and hence is timestamped in *blue*.


All three remaining RICs (*F20*, *I20*, and *L20*) detected exuberance on March 16, 17, 18 and April 8, 2020. In addition, *F20* call options revealed bubbles on March 19, 20 and April 3, whereas the latter, along with those on April 7 and 9, were identified by *I20* contracts. The application of  $\theta_{HCV}$  and  $\theta_{Surface}$ , respectively addressed occurrences on March 17, 19 and April 3; and March 16, 17, 19, 20, April 3 and 8. All three maturity groups acknowledged instances on March 17, and 19, while  $\tau_M$  and  $\tau_L$  contracts captured those on March 16, and 19. The *medium* maturity group, even witnessed the bubbles on March 20, and April 9. The detection of exuberance reveals consistency with the research hypothesis (4.1) and previous findings during this period. It is interesting to examine the magnitude of exuberance captured across the various option expiries. Consider, the bubble on March 16, which is quantified at 0.89%, 1.33% and 1.59%, respectively, in accordance with *C20*, *I20*, and *L20* RICs. This is consistent with local martingale theory of bubbles, the previous analysis of bubbles within

various maturity groups, and across the remainder of similar events corresponding to the lifetime of different options, indicating towards increasing magnitude with maturity.

**Figure 5.4: S&P 500 Bubbles in the Lifetime of Call Options expiring in 2020.**




Panel A: C20



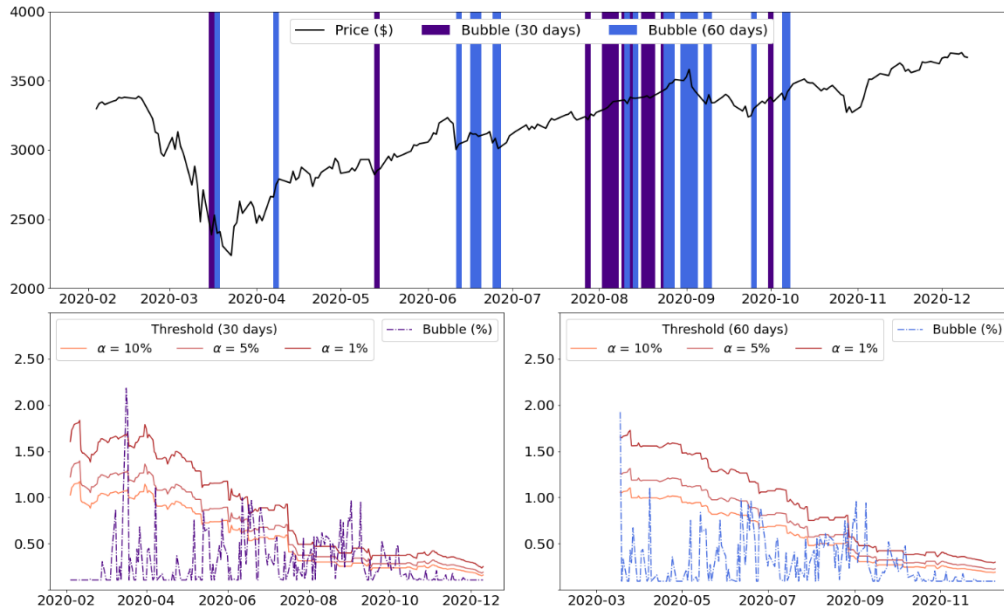
Panel B: F20

Figure 5.4 continued...

Figure 5.4 continued...



Panel C: I20



Panel D: L20

Note: Daily bubbles detected in the C20, F20, I20, and L20 call options written on the S&P 500 index, using  $\theta_{Surface}$ , displayed in Panel A, B, C, and D, respectively. Bubbles are timestamped (top) and assessed at each combination of window sizes (30, and 60 days) and levels of significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ).

The  $F20$  and  $I20$  contracts revealed exuberance on April 23, and May 14, with both being addressed by  $\tau_M$  contracts. Moreover, the former was also acknowledged by applying  $\theta_{Surface}$  to  $HCV$  and  $\tau_S$  call options. The final bubble detected by  $F20$  was witnessed on May 19, and also captured by  $\tau_M$ . The S&P 500 gained 2.29% on April 22, but continuously rising unemployment hampered investor sentiment. The bubbles on May 14 and 19, corresponded to recovering investor confidence, as the economy continued in a partial lockdown phase, with proposals for new fiscal stimuli, and developments on the coronavirus treatment front providing a further boost. Subsequently, the index hiked by 1.15% and 3.15% on May 14, and 18, respectively, prior to falling by 1.05% on May 19.

Similar to  $HCV$  and  $\tau_M$  contracts,  $L20$  options witnessed exuberance on June 12, 2020. Moreover, it identified bubbles on June 19, 25, and 26, with the latter two also being addressed by the *medium* maturity group. The  $I20$  RIC, in respective alignment with  $\tau_S$  and  $\tau_M$  call options, detected bubbles on June 16 and 23. Furthermore, both RICs captured exuberance on June 17 and 18. The latter bubble was even acknowledged by applying  $\theta_{Surface}$  to  $HCV$  and  $\tau_S$  options. Prior to the start of the cluster, on June 11, the S&P 500 plummeted by 5.89%, following panic over a second coronavirus wave. Investor sentiments received a lift from the Federal Reserve committing to its dovish stance, and the index closed in green over the next three days, recording respective wins of 1.31%, 0.83% and 1.90%. Nonetheless, the S&P 500 fell by 2.59% on June 24, before gaining 1.10%, and then dropping by 2.42%, on subsequent days. Markets were hurt from restricted business activity due to rising coronavirus cases, and weak earnings from major corporations. The  $L20$  contracts further revealed exuberance on July 28, similar to  $\tau_M$ , coinciding with falling consumer confidence, and rising cases dominating investor sentiments.

The  $L20$  contracts, next reveal 20 instances of exuberance between August 3, and September 9, 2020. During this period, the  $I20$  RIC observed bubbles on August 10, and 14. The latter, along with that on August 17, were also captured by using  $\theta_{Surface}$  on  $HCV$  and  $\tau_S$  call options. Additionally,  $\tau_S$  addressed bubbles on August 10, and 19, whereas all occurrences barring those on August 3 and 12, were acknowledged by  $\tau_M$  contracts. The discovered cluster strongly resembles that observed across the various maturity groups. It was not exhibited in  $HCV$  call options and therefore shed light on the importance of calibrating to

the entire surface and examining bubbles amongst call options with expiries. Furthermore, the capturing of these episodes, also reveals greater consistency with the research hypothesis (4.1). The initial days of the cluster witnessed a boost in market sentiments following developments on the coronavirus vaccination front, optimism over new fiscal stimulus package, and further provisions for airline companies. All COVID-19 induced losses were erased by August 18, 2020, as technology stocks that supported the market recovery, significantly benefited from the Federal Reserve's dovish outlook. Nevertheless, the S&P 500 crashed by 3.51% on September 3, following news about the highly anticipated vaccine solution not being attainable till the end of year. By September 8, the technology sector shed close to 11% of its value. The S&P 500 rose by 2.02% the next day, prior to dropping 1.76% on September 10, as a stalled labour market signalled towards slow economic recovery.


The final bubbles detected by options expiring in 2020, were observed by the *L20* RIC, on September 25, and October 1, 6, and 7. The first and last instances, were respectively identified by the application of  $\theta_{Surface}$  on  $\tau_M$  and *HCV* call options. Technology stocks rebounded towards the end of September but alternatively, a large spike in coronavirus infection was experienced across Europe. In addition to weak labour market reports, failed negotiations over the fiscal relief package, dented investor sentiments. The S&P 500 gained 1.80% on October 5 following the restart of negotiations over the fiscal stimulus package. However, it lost 1.40% the next day, as President Trump postponed such talks to after the election. Nonetheless, the index hiked by 1.74%, upon passing of an individual stimulus package with immediate effect.

#### 5.4.2.2 Call Options Expiring in 2021.


There are 4 RICs of call options expiring in 2021, *C21*, *F21*, *I21*, and *L21*, that revealed bubbles when applying  $\theta_{Surface}$  parameters to *HCV* call options. In respective order, these RICs adhere to March, June, September, and December contract terminations, and exhibited 5, 13, 4 and 6 days of exuberance over their lifetime. The bubble occurrences for each RIC are illustrated in corresponding panels, are revealed in *Figure 5.5*, which can be interpreted as its predecessor, *Figure 5.4*. Exuberance was first witnessed by *F21* options on January 29, 2021, coinciding with the S&P 500 dropping by 1.93% as market participants expressed

concerns over continued purchases of stocks similar to, and inclusive of GameStop, despite it experiencing a short squeeze two days prior.

**Figure 5.5: S&P 500 Bubbles in the Lifetime of Call Options expiring in 2021.**



Panel A: C21




Panel B: F21


Figure 5.5 continued...



Figure 5.5 continued...



Panel C: I21



Panel D: L21

Note: Daily bubbles detected in the C21, F21, I21, and L21 call options written on the S&P 500 index, using  $\theta_{\text{surface}}$ , displayed in Panel A, B, C, and D, respectively. Bubbles are timestamped (top) and assessed at each combination of window sizes (30, and 60 days) and levels of significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ).

The *F2I* RIC exhibited bubbles on February 11, 16, 17, 18, 19, 22, and 23, 2021. The instance on February 24 was captured, in addition to those on February 16, 18, 22, and 23, by *C2I* contracts. All occurrences between February 16 and 19, along with that on February 23, were acknowledged by using  $\theta_{Surface}$  on *HCV* call options, with the latter also being addressed by employing  $\theta_{HCV}$ . Additionally, each bubble between February 11 and 19 were seen within  $\tau_S$  and  $\tau_M$  contracts, barring that on February 17. The maturity groups also highlighted exuberance on February 24 and 22, respectively. Moreover, *F2I* options revealed exuberance on March 1, 2, 10, April 26 and May 3, 2021, while only that on March 10, was noticed by the  $\tau_M$  group. The S&P 500 index entered a 5-day losing streak, starting February 16, as a spike in treasury yields reflected growing concerns over inflation and equity valuation. Fears over the former were eased by the Federal Reserve maintaining an expansionary stance. Technology and other growth stocks received a boost, as the S&P 500 finished 2.37% higher on March 1. The exuberance on March 10 captured passing of the long-awaited coronavirus relief package. Furthermore, despite strong earnings from major corporations, the bubble on May 3, 2021, corresponded to growing fears over the resurgence of coronavirus infections in major emerging nations.


Next, bubbles observed on June 2, July 8 and 31, and August 4, 2021, are revealed by *I2I* contracts. The instance on July 8 was also acknowledged by employing  $\theta_{Surface}$  on *HCV*,  $\tau_M$  and  $\tau_L$  call options. Furthermore, exuberance on July 31, was detected by  $\tau_S$  and *L2I* options. The latter RIC even identified an episode on August 19, which along with the event on August 4, was captured by the  $\tau_S$  and  $\tau_M$  maturity groups. The bubble on June 2, coincided with improvements in global demand, signalled by boosted manufacturing activity, and a spike in crude oil price. Concerns over global economic recovery resurfaced, and were exhibited on July 8, after Asian nations experienced a spike in coronavirus cases. The next two days of exuberance reflected anxiety over dropping consumer sentiment, rising inflationary pressure, and increased coronavirus spread. Investor sentiments remained dampened, despite strong earnings from major corporations, and signs of economic recovery. The S&P 500 closed 1.08% below, marking a 6-day losing streak on August 18, as rising coronavirus cases and premature tapering of the quantitative easing programme threatened the pace of economic recovery. Nonetheless, technology stocks rallied over the next two days, consequent to falling treasury yields.

The final 4 bubble episodes within options expiring in 2021, were captured by  $L2I$  options, on September 14, November 10 and 29, and December 7. The former corresponded to rising inflationary pressure, along with the rapid spread of coronavirus cases hampering hiring and consumer confidence. The exuberance on November 10, also detected by  $\tau_S$  contracts, captured investors shift from technology and growth stocks, following a spike in treasury yields. Both  $\tau_S$  and  $\tau_M$ , in addition to applying  $\theta_{HCV}$  and  $\theta_{Surface}$  to  $HCV$  call options, acknowledged the occurrence on November 29, as the S&P 500 gained 1.32%. Technology stocks received a massive boost from the omicron coronavirus variant being deemed as not a source for panic. Finally, the last episode, which was even addressed by employing  $\theta_{Surface}$  for pricing  $HCV$ ,  $\tau_S$  and  $\tau_M$  call options, aligned with a 2.07% rise in the S&P 500 index. It reflected investors regaining confidence by overcoming uncertainty regarding the monetary policy and fears over the omicron variant.


#### 5.4.2.3 Call Options Expiring in 2022.

The final three  $HCV$  call options, that revealed bubbles when using  $\theta_{Surface}$  parameters, expired in January, March, and December 2022, respectively represented by  $A22$ ,  $C22$ , and  $L22$  RICs. The exuberance occurring over the lifetime of these contracts are timestamped (*top*), and quantified (*below*), within each of the corresponding panels of *Figure 5.6* which can be interpreted, in similar capacity to *Figure 5.4* and *Figure 5.5*. It can be observed that several bubbles detected by 2021 expiry contracts, were also noticed within  $A22$  and  $C22$  prices. For instance, the former addresses events on July 30, November 29, and December 7, 2021, whereas only that on November 29, was captured by the latter. The  $A22$  call options revealed exuberance on July 30, 2021, as visible in *Panel A* of *Figure 5.6*, coinciding with fears over rising inflation, increased coronavirus spreads, and dropping sentiments, even though major corporations reported strong earnings. Similar to  $\tau_L$ ,  $C22$  contracts captured a bubble on September 2, as sentiments received a boost from improvements in manufacturing and labour market activity.

**Figure 5.6: S&P 500 Bubbles in the Lifetime of Call Options expiring in 2022.**




Panel A: A22



Panel B: C22

Figure 5.6 continued...

Figure 5.6 continued...



Panel C: L22

Note: Daily bubbles detected in the A22, C22, and L22 call options written on the S&P 500 index, using  $\theta_{Surface}$ , displayed in Panel A, B, and C, respectively. Bubbles are timestamped (top) and assessed at each combination of window sizes (30 and 60 days) and levels of significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ).

Next, occurrences were observed on November 29, and December 7, in correspondence to investors regaining confidence by overcoming coronavirus and inflation related concerns. At the turn of the year, exuberance was first exhibited on February 2, 2022. It was acknowledged by C22, along with  $\tau_S$  and  $\tau_M$  options. Momentum from bank stocks and strong corporate earnings bolstered the index, despite a struggling technology sector. The final episode was detected on May 5, by the L22 RIC, reflecting interest rate hikes hindering economic growth, rather than combating supply-side inflation. Consequently, a technology sell-off triggered the S&P 500 index to fall by 3.57% lower.

### 5.4.3 Model Performance

Identification of exuberance across the lifetime of call options, allowed for capturing the forward-looking expectations of market participants over future durations. The call options were categorized in accordance with their maturity dates, and historical prices of only those associated with *HCV* call options, that revealed exuberance when applying  $\theta_{Surface}$  parameters, being selected. As observed previously, the  $RMSE_{\sigma,t}$  metric displayed in *Table 5.9*, is utilised for testing the statistical validity of the approach. Recall, the entire sample for each RIC, refers to the period over which historical call prices, adhere to the  $7 < \tau < 365$  days condition.

**Table 5.9: Model Calibration Performance over Call Option Lifetime.**

*Note:  $RMSE_{\sigma,t}$  (%) metrics from  $\theta_{Surface}$  calibration corresponding to the entire sample (excluding burn-in) period, along days with negative ( $\widehat{\mathbb{B}}_t < 0$ ) and positive ( $\widehat{\mathbb{B}}_t > 0$ ) bubble magnitudes, for each RICs, are reported in respective panels.*

Window Size (days)	Entire	$\widehat{\mathbb{B}}_t < 0$	$\widehat{\mathbb{B}}_t > 0$
Panel A: C20 (March 2020)			
30	0.640%	0.595%	0.768%
60	0.652%	0.602%	0.768%
Panel B: F20 (June 2020)			
30	0.816%	0.740%	0.957%
60	0.868%	0.807%	0.957%
Panel C: I20 (September 2020)			
30	0.925%	0.808%	1.014%
60	0.988%	0.882%	1.055%
Panel D: L20 (December 2020)			
30	0.970%	0.886%	1.033%
60	0.893%	0.800%	0.949%
Panel E: C21 (March 2021)			
30	0.745%	0.724%	0.778%
60	0.745%	0.711%	0.804%
Panel F: F21 (June 2021)			
30	0.740%	0.726%	0.782%
60	0.693%	0.654%	0.781%

*Table 5.9 continued...*

Table 5.9 continued...

Panel G: I21(September 2021)			
30	0.647%	0.630%	0.680%
60	0.677%	0.664%	0.702%
Panel H: L21 (December 2021)			
30	0.710%	0.736%	0.680%
60	0.693%	0.714%	0.672%
Panel I: A22 (January 2022)			
30	0.706%	0.723%	0.671%
60	0.691%	0.707%	0.668%
Panel J: C22 (March 2022)			
30	0.711%	0.713%	0.708%
60	0.704%	0.707%	0.700%
Panel K: L22 (December 2022)			
30	0.874%	0.853%	0.941%
60	0.882%	0.852%	0.983%

The existence of negative bubbles is prevented by the local martingale theorem by considering the fundamental value of the asset, as the lower bound to its market price. Hence, the occurrence of such phenomena is justified by the possible presence of model misspecification Fusari et al. (2024). As documented in Table 5.9, across both windows, the  $RMSE_{\sigma,t}$  metric is lower during  $\widehat{\mathbb{B}}_t > 0$  days, for *L21*, *A22* and *C22* RICs. The error associated with *I21* options is marginally higher for  $\widehat{\mathbb{B}}_t > 0$ . A likewise trend in the  $RMSE_{\sigma,t}$  values is witnessed for *C21*, *F21*, and *L21* call contracts, with a difference between average values on with positive and negative magnitudes not exceeding 0.13%. This gap reduces, with respect to both windows sizes, as the maturity of the option increases. It is a direct consequence of the  $\tau < 365$  days condition, as historical prices of later expiring RICs become clear of the COVID-19 induced crash period.

Observe, for call options expiring in 2020, not only does the gap between  $RMSE_{\sigma,t}$  metrics increase, but so does the absolute magnitude of average errors. As witnessed over previous practices of bubble detection in this study, such errors arise due to relatively weaker GSVJD calibration performances during the COVID-19 induced market crash period. The days of exuberance between March 16, and April 23, 2020, align with higher GSVJD calibration

errors. More specifically, due to greater detection of bubbles during this period, the  $F20$ ,  $I20$ , and  $L20$  RICs reveal greater  $RMSE_{\sigma,t}$  values on  $\widehat{\mathbb{B}}_t > 0$  days. However, such events coincide with the triggering of circuit breakers, which act as a lower bound to admissible trade strategies, making the existence of exuberance consistent with the local martingale theory of bubbles. This justification eliminates any concerns over the potential existence of joint-hypothesis related issues. Overall, performance of the  $\theta_{Surface}$  are consistent and robust in the detection of bubbles, throughout all RICs. They provide the ability to comprehend the signalling of exuberance in the S&P 500 index, according to forward-looking expectations of market participants, over different future durations, captured by the horizon of various option contracts.

## 5.5 Summary and Discussion

The deep calibration framework for computing GSVJD parameters was previously applied to the  $HCV$  and *Entire Surface* datasets. When compared to the benchmark Monte Carlo simulations, it provided a large boost in computational efficiency, along with a 2-fold improvement in calibration accuracy. In the current chapter,  $\theta_{HCV}$  and  $\theta_{Surface}$  were utilised for estimating the fundamental value of call options for bubble detection. Recall, under the ND condition, exuberance in the underlying and call option prices, share a linear relationship, such that the magnitude of the latter, acts as a lower bound to that of the former. The identification of bubbles was conducted over different window sizes, and various levels of significance.

First,  $HCV$  call options were examined using  $\theta_{HCV}$  and  $\theta_{Surface}$ . The former captured exuberance coinciding with growing trade tensions between U.S. and China, along with concerns over an economic slowdown, throughout 2019. Whereas, towards the end of 2019, both addressed the signing of a trade agreement between the two nations. As opposed to  $\theta_{Surface}$ , bubbles, to a greater extent, were witnessed during late January and early February, prior to the global coronavirus pandemic. The pair further acknowledged occurrences in March 2020, that aligned with the COVID-19 induced crash. Bubbles during this period, across days that experienced large losses and gains, addressed the dynamic prices



movements driven by speculative buying, with the slightest indication of panic, resulting in greater selloffs. Interestingly,  $\theta_{Surface}$  displayed exuberance on March 16, which coincided with a circuit breaker, maintaining consistency with the local martingale theorem of bubbles. Moreover, calibration to the entire surface, identified more bubbles during the immediate recovery period.

In 2021, the S&P 500 index was supported mostly by economic recovery and the technology sector, even though inflationary pressure and coronavirus cases were on the rise. These concerns remained throughout their year, with the former severely impacting technology stocks, and raising questions about their overvaluation. Large technology stock sell-offs contributed towards  $\theta_{HCV}$  revealing exuberance in February and May 2021, with  $\theta_{Surface}$  capturing a much greater extent of the former cluster. However, unchanged interest rates, continued to bolster technology stocks. In general, bubbles captured negative sentiment of investors, hinting towards market prices exceeding the fundamental value. Overall, the  $\theta_{Surface}$  was better equipped, given the use of information across several volatility smiles, to detect bubbles during dramatic price periods that were fuelled by speculative behaviours linked to clear economic triggers. A relatively larger error metric was associated with  $\theta_{Surface}$  calibrations. This was a direct consequence of detecting several bubbles between March and April 2020, a period synonymous with the coronavirus induced crash. Nevertheless, concerns over the presence of joint-hypothesis issues were eliminated, as these days of exuberance were consistent with the local martingale theory of bubbles.

Next,  $\theta_{Surface}$  were employed to examine days of exuberance, associated with call options from three maturity groups: *short* ( $\tau_S$ ), *medium* ( $\tau_M$ ), and *long* ( $\tau_L$ ). This practice allowed the detection of bubbles by comprehending forward looking expectations of market participants, over various investment horizons. This was possible as  $\theta_{Surface}$  was calibrated from the entire daily surface, allowing to compute the fundamental value of call options from various maturity groups. Collectively, across all three groups, the identified exuberance closely resembled those within *HCV* call options, when using  $\theta_{Surface}$ . However, a key difference of bubbles being spotted in larger groups was revealed. For instance, the revelations of exuberance between March and October 2020, provided more insight on investor sentiments throughout the crash and instantaneous recovery period, and revealed

greater consistency with the research hypothesis. Furthermore, bubbles were acknowledged in February, and March 2021, as the S&P 500 reached a record high, despite technology stocks suffering from rising inflation, and investors expressing concerns over the long-lasting pandemic induced effects on the economy. The overall impact of rising inflationary concerns, along with lingering impacts from the pandemic, continued to plague investor sentiments, as the S&P 500 surged with support from technology stocks. This was well captured by detected exuberance between July 2021 and February 2022.

Taking advantage of  $\theta_{surface}$  parameters being able to compute fundamental values of daily call options, across various maturities, bubbles were examined over the lifetime of certain contract cross-section. This provided an even deeper comprehension of investor perceptions regarding the formation of such phenomena. The identified bubbles were very close to those within the different maturity groups, especially with respect to the clusters in March, June, August and September 2020, aligning closely with research hypothesis. In both practices, bubbles observed on similar days, either by various call option RICs or maturity groups, were found to increase with remaining duration of the contract. This is consistent with the local martingale theory of bubbles, that reveals the magnitude of such events to nullify as the option expires. Recall, that the size of the call option price bubbles acts as the lower bound to that in the underlying. Therefore, by exhibiting a greater magnitude of the same events in call options with longer remaining lives, a larger extent of that present in the underlying can be acknowledged. Furthermore, exuberance was more often spotted towards the end of the contract, as it would have to be significantly larger for it be present during the earlier stages. This indicates towards the variation in investor perceptions that hold contracts of different durations. Such bubbles would symbolise the willingness of traders to purchase closer to maturity contracts, with aspirations of making a short-term profit.

The efficiency of the deep calibration framework, along with the improvement in accuracy, allowed for the GSVJD parameters to be computed from the entire surface of daily put options, within a matter of seconds. Firstly, this significantly bolstered the tractability of the three-step approach amongst practitioners indulging in real-time bubble detection. Next, the  $\theta_{surface}$  comprises of information regarding the forward-looking expectations of market participants extracted from various volatility smiles, in contrast to solely focusing on the

most liquid maturity. Hence, these parameters were more equipped to capture the direct impact of speculative behaviour fuelled by economic triggers. Furthermore, they allowed for examining the presence of bubbles across various call option maturity groups, and even within the lifetime of several contract cross-sections, which revealed more consistency with research hypothesis in capturing exuberance during the COVID-19 induced crash, and immediate recovery periods. This developed a deeper understanding regarding the formation of bubbles, in addition to revealing a greater extent of their magnitude in the underlying price.

# Chapter 6:

## Case Study

The widespread application of enhancing the three-step approach with neural networks, was thoroughly explored in the previous chapter. Next, this study proceeds with extending the implementation of the methodology towards identifying exuberance in various technology stocks that have American-styled option written on them. Recall, in abidance with the local martingale theory of bubbles (see 2.3.2), American calls do not exhibit bubbles. However, difference between their market prices, and the fundamental value of equivalent European-styled options, capture the bubble component, similar to reflecting the early exercise value. Therefore, the deep calibration framework can be applied to calibrate the GSVJD model to American styled market put options.

Since put options do not reveal bubbles, market prices must align with their fundamental values. Moreover, the bubble component is similar to reflecting the early exercise value, and given non-existence of the former, the fundamental value of American- and European-styled put options are considered to align. This validates the application of the deep calibration framework, without having to retrain the optimal neural network. The computed parameters are used to estimate fundamental call option values from (19), (20), and (22). The difference

between the modelled, and equivalent American-styled market call prices represents the bubble, similar to an early exercise premium. The case study provides a great opportunity to express the robustness of the neural network framework in calibrating the GSVJD model to various equities. This stellar capability is a direct consequence of the optimal hyperparameter search, and the network being trained on a synthetic dataset, rather than historical prices. The latter prepares the deep calibration framework to adjust to new markets scenarios, and different equity instruments.

Market analysis of bubbles in the S&P 500 from *Chapter 5* revealed the index, or more generally U.S. markets to recover from the devastations of the COVID-19 induced crash by receiving major support from technology stocks. The share prices of these companies continued to rally throughout the low-interest rate environments of 2021. Moreover, even after the following year, despite the Federal Reserve altering its expansionary stance, a peaked interest in the development of artificial intelligence tools for a wide variety of tasks bolstered this sector. This attracted interest in examining five particular technology stocks. Specifically, the selection of these companies is justified by the following reasons. Microsoft Corporation (*MSFT*) and Meta Platforms, Inc. (*META*) have been heavily focused on incorporating the benefits of these tools into everyday lives, by making them accessible across various domains. Machine learning and artificial intelligence algorithms require high speeds and efficiency, making them reliant on super performing graphic processors and chips. Nvidia Corporation (*NVDA*) and Advanced Micro Devices, Inc. (*AMD*), indulge in the business of manufacturing such products. In both respective fields, the aforementioned companies have experienced significantly varying magnitudes of gains in stock prices over the sample period (August 1, 2022, to April 30, 2024). They have been included in the case study, such that bubble formations in competitors experiencing different degrees of benefits from the AI boom, are examined. Additionally, Amazon.com, Inc. (*AMZN*) is included, as it is a technology stock, which does not direct beneficiary of the boom.

Alternatively, this chapter conducts a robustness check on the application of the three-step approach. Strict local martingale tendencies in underlying prices could not be revealed by relying on parameters  $p$  and  $\rho$ , given the non-unique nature of the optimiser. Therefore, signals indicating the presence of exuberance in the underlying were obtained from using a

statistical test to examine the significance of call option bubbles. Essentially, relying on the linear relationship between call option and underlying price bubbles, under ND. Given large standard errors being obtained during GSVJD calibrations, Fusari et al. (2024) revealed a robustness check for the three-step approach, via examining factors that influence the formation of bubbles, in abidance with the local martingale theory. The current chapter begins with an overview of relevant data associated with the selected technology sector stocks, prior to examining the performance of the deep calibration framework. The framework is applied to the most liquid cross section (*HCV*), and entire surface of daily put options. However, bubble detection is restricted to *HCV* call contracts, given the short-term lives of stock options. Finally, a robustness check is conducted, by examining factors that influence the formation of bubbles.

## 6.1 Market Data

The case study is geared towards assessing the robust application of the deep neural network framework for applying the three-step approach to five technology stocks: *MSFT*, *NVDA*, *AMZN*, *META*, and *AMD*. Empirical analysis is initiated over daily American-styled options written on each stock, during a period of 437 trading days. All market observations between August 1, 2022, and April 30, 2024, were collected from Refinitiv Eikon. Bubbles are detected by calibrating the GSVJD model to the daily most liquid cross section and surface of put options, for examining *HCV* call options. The summary statistics of options data for each company is provided in *Table 6.1*. Mid-prices of weekly options were computed by taking the average between bid and ask quotes. The strike increments varied with respect to the price range of the stocks. *NVDA*, *META* and *MSFT* options were selected at increments for \$2.50, whereas for *AMZN* and *AMD*, the difference between consecutive strike prices was \$1.00. Given the synchronization in the trading of the equity and option markets, the need to compute implied spot prices is made redundant. Risk-free rates were obtained from linearly interpolating the Zero-Coupon Yield Curve from Eikon, to corresponding option maturities, using (49). Moreover, in abidance with Stahl & Blauth (2024), dividend yields were assumed to be constant as per their last known payment. Finally, only options with nonzero trading volumes, satisfying  $-10 < m < 5$ , and  $k < 0.5$ , and expiring between 7 and 365 days are retained (refer to 5.1 for more information).

**Table 6.1: Summary Statistics of Options relevant to the Case Study.**

Note: Summary statistics of HCV and Entire Surface put options, and HCV call options for each stock. For a single  $\tau$ , Volume (000s) reports accumulated values for daily contracts, and  $N_p/N_c$  reveals the number of put/call contracts, on a given day. Bid-, Ask-, and Mid-Prices are denoted in \$, along with  $mon = \mathbb{K}/\$$ .

	Put						Call		
	HCV			Entire Surface			HCV		
	mean	std	P50	mean	std	P50	mean	std	P50
Panel A: MSFT									
Bid	10.44	20.15	1.26	12.58	22.95	3.70	19.77	28.73	5.70
Ask	11.11	20.92	1.48	13.30	23.61	4.05	20.72	29.64	6.05
Mid	10.78	20.53	1.35	12.94	23.28	3.85	20.24	29.19	5.88
$mon$	0.93	0.17	0.94	0.89	0.19	0.90	1.02	0.17	1.01
Volume	25.84	16.79	21.30	6.12	9.30	2.65	32.29	22.21	26.77
$N_p/N_c$	40.57	6.59	40.00	30.79	9.45	31.00	41.14	7.42	40.00
$\tau$	22.18	25.04	16.00	96.96	86.54	58.00	22.18	25.04	16.00
Panel B: NVDA									
Bid	21.33	36.76	3.75	27.04	46.87	8.35	39.88	59.51	13.30
Ask	22.02	37.64	3.95	27.82	47.80	8.75	40.89	60.52	13.80
Mid	21.68	37.20	3.85	27.43	47.33	8.55	40.38	60.01	13.55
$mon$	0.94	0.24	0.93	0.87	0.29	0.86	1.05	0.28	1.01
Volume	55.29	30.49	49.48	12.18	18.81	5.05	64.17	47.28	53.39
$N_p/N_c$	61.76	23.00	64.00	52.15	24.22	48.00	64.85	21.46	63.00
$\tau$	17.97	17.58	11.00	92.14	85.69	50.00	17.97	17.58	11.00
Panel C: AMZN									
Bid	8.07	15.25	1.22	8.65	14.87	2.86	9.61	13.68	3.05
Ask	8.31	15.49	1.30	8.96	15.17	3.05	9.91	13.97	3.20
Mid	8.19	15.36	1.26	8.81	15.01	2.95	9.76	13.83	3.12
$mon$	0.95	0.24	0.93	0.93	0.25	0.92	1.06	0.26	1.02
Volume	53.91	67.89	35.76	12.36	26.92	4.67	58.44	42.65	47.82
$N_p/N_c$	47.55	19.44	45.00	32.77	17.55	30.00	47.67	17.31	46.00
$\tau$	25.33	28.41	16.00	94.85	90.07	50.00	25.33	28.41	16.00
Panel D: META									
Bid	18.92	37.38	2.85	19.55	34.52	6.00	23.65	36.55	5.90
Ask	19.27	37.66	2.96	19.97	34.89	6.25	24.11	37.06	6.10
Mid	19.09	37.52	2.90	19.76	34.70	6.13	23.88	36.80	6.00
$mon$	0.98	0.32	0.94	0.92	0.31	0.90	1.09	0.35	1.03
Volume	30.01	30.72	20.79	6.59	13.56	2.46	31.38	23.89	25.41
$N_p/N_c$	42.67	12.38	42.00	32.61	12.57	32.00	47.82	14.08	46.00
$\tau$	25.51	30.85	14.00	93.57	89.98	49.00	25.51	30.85	14.00

Table 6.1 continued...

Table 6.1 continued...

Panel E: AMD									
Bid	6.64	11.11	1.65	7.61	11.82	2.93	8.18	11.83	2.58
Ask	6.84	11.36	1.71	7.82	12.05	3.05	8.43	12.16	2.65
Mid	6.74	11.24	1.68	7.72	11.93	3.00	8.31	11.99	2.61
<i>mon</i>	0.96	0.21	0.95	0.93	0.24	0.93	1.07	0.27	1.03
Volume	33.73	21.32	27.37	8.19	12.48	3.65	44.30	31.02	36.73
$N_P/N_C$	37.73	12.44	38.00	25.82	11.40	22.00	41.59	12.78	42.00
$\tau$	20.05	22.96	11.00	80.17	75.79	44.00	20.05	22.96	11.00

In Table 6.1, the daily average volume of *HCV* put options is approximately 4.5 times larger across all stocks. Majority of liquidity is concentrated within short-term options, given *HCV* contracts, on average expire between 17 – 25 days, compared to 80 – 96 days from the *Entire Surface* dataset. In respective order, median maturity ranges for the datasets are 11 – 16 and 44 – 58 days. Both ranges are significantly lower than those seen for S&P 500 index options (view Table 4.2 and Table 5.1). Furthermore, as opposed to S&P 500 options, call contracts written over equity reveal more liquidity. The selection of weekly options in Table 6.1, makes it is difficult to examine bubbles across various maturity groups, and over historical call prices. Hence, bubbles are only detected within daily *HCV* call options.

The options span over a wide range of moneyness, on average reflecting OTM characteristics. The inclusion of a wide range of strike prices, ensures a sufficient number of option contracts for calibration and bubble detection. Overall, stocks within the *HCV* dataset reveal, on average 37 – 61 and 41 – 64 daily put and call contracts, respectively. Whereas, for each maturity, the mean number of put options ranged between 25 and 52 in the *Entire Surface* dataset. Finally, call option prices must abide by the ND condition,  $C_t^A \leq S_t$  (Jarrow et al., 2010). This ensures that call option and underlying price bubbles share a linear relationship, with presence in the former, signalling towards strict local martingale tendencies in the latter.

## 6.2 GSVJD Calibration for Bubble Detection

The deep calibration framework is paired with the Differential Evolution optimiser to calibrate GSVJD parameters. The network was trained on the GSVJD model, for pricing European options, and its application is consistent with the local martingale theory of bubble



detection. Similar to the approach in 4.4.3, GSVJD parameters are estimated over 10 different seeds. The summary statistics of average parameters,  $\theta_{HCV}$  and  $\theta_{Surface}$ , for each stock, are presented in Table 6.2, along with respective calibration performance measures. The average  $RMSE_{\sigma,t}$  ranges between 1.37% – 2.14%, and 2.29% – 3.50%, over the constituents of the *HCV* and *Entire Surface* datasets, respectively. These metrics are higher than those seen when calibrating to the S&P 500 index. However, in consistency with Fusari et al. (2024), this is due to equity options revealing significantly larger implied volatilities. Furthermore, calibration is vulnerable to non-unique solutions, hence a statistical test is relied upon to detect call option exuberance for signaling such behaviour in the underlying.

**Table 6.2: Summary Statistics of GSVJD parameters relevant to the Case Study.**

*Note: Summary Statistic and standard errors (SE) for daily calibrated GSVJD parameters, and  $RMSE_{\sigma,t}$  (%) from HCV and Entire Surface put options of each stock, are represented in respective panels.*

	$\theta_{HCV}$						$\theta_{Surface}$					
	mean	std	P25	P50	P75	SE	mean	std	P25	P50	P75	SE
Panel A: MSFT												
$\kappa$	2.058	1.688	0.814	1.619	2.744	0.438	2.734	1.438	1.756	2.419	3.540	0.269
$\sigma_v$	1.463	0.530	1.078	1.469	1.842	0.121	1.243	0.643	0.680	1.202	1.769	0.030
$\bar{v}$	0.174	0.112	0.089	0.158	0.238	0.038	0.091	0.057	0.046	0.093	0.129	0.008
$V_0$	0.060	0.035	0.034	0.054	0.078	0.003	0.070	0.040	0.038	0.065	0.092	0.000
$\rho$	-0.569	0.385	-0.875	-0.643	-0.345	0.083	-0.586	0.326	-0.895	-0.632	-0.292	0.020
$p$	0.696	0.397	0.423	0.623	0.912	0.072	0.564	0.337	0.338	0.498	0.755	0.008
$\lambda$	0.286	0.364	0.080	0.143	0.323	0.011	0.094	0.076	0.054	0.073	0.099	0.008
$\mu_y$	0.038	0.230	-0.022	0.086	0.169	0.066	0.103	0.131	0.036	0.110	0.185	0.003
$\sigma_y$	0.529	0.220	0.359	0.516	0.681	0.061	0.668	0.135	0.588	0.666	0.759	0.018
$RMSE_{\sigma,t}$	2.14%	1.41%	1.13%	1.85%	2.76%	-	2.29%	0.86%	1.64%	2.13%	2.76%	-
Panel B: NVDA												
$\kappa$	5.658	2.429	4.088	5.628	7.268	1.196	3.733	2.240	2.001	3.708	5.269	0.168
$\sigma_v$	1.412	0.576	0.972	1.410	1.878	0.128	1.287	0.496	0.927	1.254	1.630	0.045
$\bar{v}$	0.404	0.137	0.320	0.406	0.494	0.060	0.178	0.108	0.091	0.195	0.248	0.007
$V_0$	0.191	0.123	0.103	0.150	0.251	0.014	0.251	0.117	0.162	0.234	0.312	0.001
$\rho$	-0.568	0.448	-0.960	-0.686	-0.345	0.056	-0.584	0.459	-0.975	-0.712	-0.366	0.030
$p$	0.646	0.457	0.317	0.537	0.893	0.099	0.724	0.499	0.328	0.675	1.059	0.045
$\lambda$	0.311	0.433	0.121	0.197	0.352	0.040	0.102	0.080	0.062	0.078	0.112	0.005
$\mu_y$	0.008	0.202	-0.090	0.038	0.151	0.015	-0.104	0.183	-0.192	-0.079	0.017	0.007
$\sigma_y$	0.594	0.210	0.457	0.557	0.703	0.034	0.909	0.151	0.879	0.980	1.000	0.008
$RMSE_{\sigma,t}$	2.09%	1.55%	1.22%	1.71%	2.52%	-	3.00%	1.18%	2.17%	2.64%	3.60%	-

Table 6.2 continued...

Table 6.2 continued...

Panel C: AMZN												
$\kappa$	3.090	2.380	1.252	2.522	4.522	0.740	3.221	1.795	1.788	2.823	4.888	0.100
$\sigma_v$	1.526	0.570	1.083	1.471	1.986	0.132	1.466	0.567	1.128	1.436	1.903	0.044
$\bar{v}$	0.248	0.147	0.133	0.244	0.342	0.041	0.111	0.081	0.039	0.115	0.150	0.005
$V_0$	0.111	0.071	0.058	0.095	0.150	0.008	0.141	0.088	0.069	0.136	0.195	0.001
$\rho$	-0.614	0.336	-0.903	-0.670	-0.402	0.061	-0.542	0.352	-0.860	-0.570	-0.301	0.025
$p$	0.661	0.414	0.345	0.570	0.927	0.080	0.591	0.328	0.339	0.517	0.827	0.028
$\lambda$	0.284	0.395	0.077	0.155	0.354	0.043	0.121	0.132	0.060	0.083	0.111	0.006
$\mu_y$	0.040	0.138	-0.012	0.055	0.123	0.014	0.040	0.132	-0.032	0.059	0.134	0.005
$\sigma_y$	0.587	0.220	0.423	0.583	0.736	0.035	0.727	0.181	0.607	0.757	0.861	0.011
$RMSE_{\sigma,t}$	1.42%	0.97%	0.81%	1.20%	1.69%	-	2.54%	1.21%	1.62%	2.23%	3.19%	-
Panel D: META												
$\kappa$	4.004	2.659	1.670	3.606	5.919	0.846	3.054	1.837	1.564	2.880	4.412	0.130
$\sigma_v$	1.481	0.582	1.048	1.482	1.925	0.124	1.642	0.498	1.295	1.745	2.003	0.034
$\bar{v}$	0.302	0.165	0.181	0.300	0.423	0.049	0.142	0.089	0.051	0.164	0.211	0.008
$V_0$	0.139	0.107	0.067	0.104	0.187	0.010	0.185	0.125	0.095	0.163	0.239	0.001
$\rho$	-0.579	0.362	-0.895	-0.623	-0.349	0.061	-0.562	0.375	-0.919	-0.594	-0.320	0.031
$p$	0.699	0.452	0.331	0.601	1.013	0.090	0.727	0.381	0.390	0.723	1.048	0.032
$\lambda$	0.255	0.249	0.092	0.145	0.331	0.032	0.096	0.079	0.053	0.077	0.105	0.004
$\mu_y$	-0.014	0.175	-0.106	0.024	0.105	0.014	-0.027	0.156	-0.117	-0.004	0.075	0.006
$\sigma_y$	0.589	0.257	0.388	0.560	0.804	0.029	0.856	0.157	0.791	0.897	0.988	0.009
$RMSE_{\sigma,t}$	1.37%	0.87%	0.81%	1.17%	1.72%	-	3.50%	1.43%	2.26%	3.28%	4.56%	-
Panel E: AMD												
$\kappa$	5.359	2.367	3.807	5.461	7.065	1.267	3.559	2.242	1.630	3.472	5.472	0.139
$\sigma_v$	1.450	0.580	1.037	1.450	1.908	0.140	1.174	0.511	0.756	1.183	1.528	0.030
$\bar{v}$	0.370	0.129	0.284	0.374	0.456	0.069	0.136	0.115	0.012	0.129	0.215	0.006
$V_0$	0.185	0.093	0.117	0.161	0.234	0.014	0.239	0.087	0.178	0.219	0.289	0.001
$\rho$	-0.484	0.504	-0.895	-0.600	-0.248	0.067	-0.518	0.431	-0.970	-0.541	-0.206	0.016
$p$	0.674	0.440	0.324	0.618	0.948	0.110	0.547	0.458	0.199	0.407	0.845	0.027
$\lambda$	0.378	0.324	0.160	0.266	0.496	0.048	0.144	0.114	0.079	0.106	0.165	0.008
$\mu_y$	-0.031	0.193	-0.111	0.010	0.092	0.015	-0.057	0.175	-0.156	-0.031	0.067	0.004
$\sigma_y$	0.583	0.219	0.416	0.543	0.740	0.037	0.828	0.166	0.731	0.872	0.989	0.011
$RMSE_{\sigma,t}$	1.61%	0.92%	0.97%	1.46%	2.03%	-	2.40%	0.90%	1.70%	2.24%	2.92%	-

Overall computational superiority and accuracy of the framework allow it to successfully boost the tractability of the GSVJD model and subsequently make the three-step approach attractive for real time bubble detection. Once obtained, daily parameters are inserted into (19), and (20), to price the fundamental call option values using (22). Given the network is trained on synthetic data, it can be applied to different stocks, within various market scenarios. The neural network framework simply replaces Monte Carlo simulation during the calibration of the three-step approach. The mean standard errors from Table 6.2, averaged across all seeds are relatively large. This increases reliance on the statistical test from the

three-step approach, as opposed to focusing on parametric conditions. Hence, the ND condition is vital as it ensures that bubbles in call options and underlying prices share a linear relationship. This allows the statistical test to signal for strict local martingale behaviour in the stock price process. Therefore, only call options that abide by the ND condition are selected for estimating and testing of bubbles. Finally, the fundamental value from call options, is computed by averaging modelled prices, across the different seeds.

## 6.3 Call Option Bubbles: Most Liquid Cross Section

A bubble exists when the market price of the call option exceeds its fundamental value, and is considered significant, if  $\widehat{\mathbb{B}}_t > \alpha \widehat{\sigma}_t$  (see 3.1.3). Under the ND condition, call option and underlying price bubbles share a linear relationship, with identification in the former implying strict local martingale tendencies in the latter. Daily bubbles within each of the stock prices are detected over the sample period ranging from August 1, 2022, to April 30, 2024. The time-varying threshold,  $\alpha \widehat{\sigma}_t$ , given the relatively small sample size is computed over 30- and 60-day window sizes, and three levels of significance. Recall, the difference between the market prices of American call options and the fundamental value of their European-styled counterpart, capture the bubble similar to revealing the early exercise value. It is under this capacity, in correspondence to the local martingale theory of bubbles, the neural network enhanced three-step approach is utilised for identifying exuberance in the aforementioned stocks. The  $\theta_{HCV}$  and  $\theta_{Surface}$  parameters are applied for detecting bubbles in *HCV* call options. This section begins by examining the key statistics, prior to conducting a market analysis and evaluation of model performance. Each of these steps are crucial towards overcoming any potential joint-hypothesis related issues.

### 6.3.1 Key Statistics

The significant call options bubbles, indicating towards strict local martingale tendencies in each of the selected stocks, are examined using two window sizes, and three levels of significance. The *HCV* call options are priced using  $\theta_{HCV}$  and  $\theta_{Surface}$  parameters obtained from the application of the deep calibration framework. The key statistics corresponding to exuberance, for each parameter set, window size, and levels of significance, are presented

in Table 6.3 for the respective stocks. The performances are measured by  $RMSE_{\sigma,t}$  metrics, across the various window sizes and levels of significance. The  $\theta_{HCV}$  reveal better performances, in comparison to those from  $\theta_{Surface}$ , on days displaying exuberance. Additionally, across all stocks, average  $RMSE_{\sigma,t}$  from  $\theta_{HCV}$ , is lower than the entire sample (Table 6.2).

**Table 6.3: GSVJD Performances relevant to the Case Study.**

Note:  $RMSE_{\sigma,t}$  (%) error metric, along with number of significant bubbles ( $N_B$ ), and percentage of the sample period (excluding burn-in) revealing exuberance  $N_B(\%)$ , are reported for days displaying  $\widehat{\mathbb{B}}_t > \alpha \hat{\sigma}_t$  according to HCV and Entire Surface put options, in respective panels for various stocks.

$\theta_{HCV}$					$\theta_{Surface}$		
Window	$\alpha$	$\mathbb{N}_B$	$\mathbb{N}_B$ (%)	$RMSE_{\sigma,t}$	$\mathbb{N}_B$	$\mathbb{N}_B$ (%)	$RMSE_{\sigma,t}$
Panel A: MSFT							
30.00	10.00	27.00	6.62%	1.63%	31.00	7.60%	2.13%
	5.00	19.00	4.66%	1.77%	26.00	6.37%	2.15%
	1.00	12.00	2.94%	1.64%	16.00	3.92%	1.90%
60.00	10.00	7.00	1.85%	1.96%	6.00	1.59%	2.44%
	5.00	6.00	1.59%	1.91%	5.00	1.32%	2.12%
	1.00	1.00	0.26%	0.73%	3.00	0.79%	1.81%
Panel B: NVDA							
30.00	10.00	13.00	3.19%	1.52%	12.00	2.94%	4.04%
	5.00	10.00	2.45%	1.71%	5.00	1.23%	3.77%
	1.00	5.00	1.23%	1.09%	-	-	-
60.00	10.00	8.00	2.12%	0.93%	3.00	0.79%	5.70%
	5.00	4.00	1.06%	0.72%	2.00	0.53%	6.67%
	1.00	1.00	0.26%	1.14%	-	-	-
Panel C: AMZN							
30.00	10.00	23.00	5.64%	1.34%	8.00	1.96%	4.98%
	5.00	19.00	4.66%	1.10%	5.00	1.23%	4.18%
	1.00	13.00	3.19%	1.19%	1.00	0.25%	4.47%
60.00	10.00	20.00	5.29%	1.36%	6.00	1.59%	4.65%
	5.00	15.00	3.97%	1.02%	2.00	0.53%	2.95%
	1.00	12.00	3.17%	1.01%	2.00	0.53%	2.95%

Table 6.3 continued...

Table 6.3 continued...

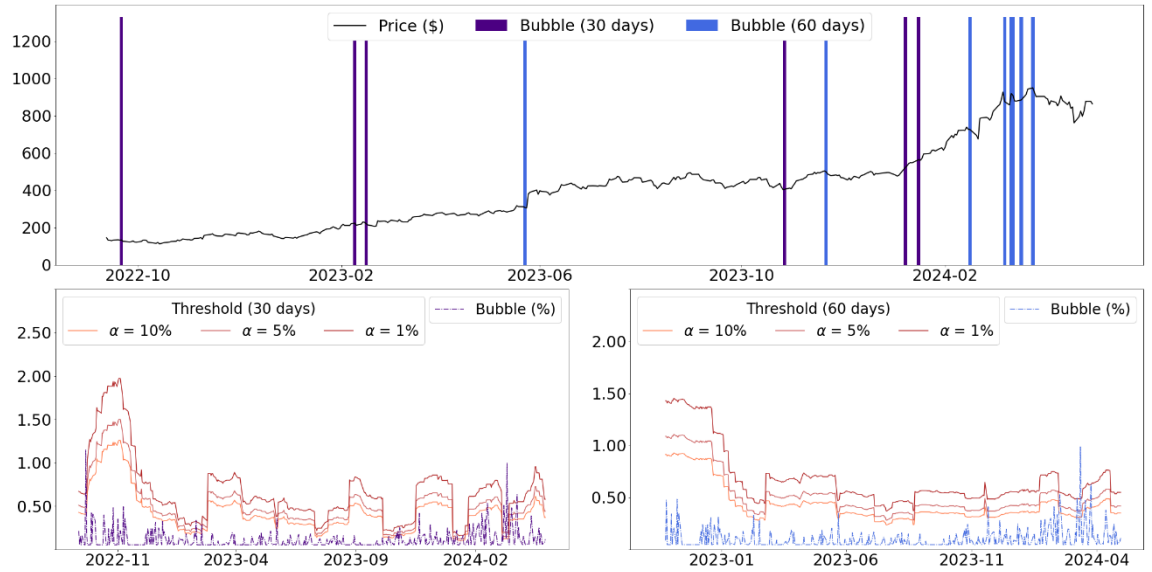
Panel D: META							
	10.00	17.00	4.17%	0.89%	14.00	3.43%	3.30%
30.00	5.00	11.00	2.70%	0.88%	10.00	2.45%	2.86%
	1.00	4.00	0.98%	0.91%	8.00	1.96%	2.23%
	10.00	11.00	2.91%	0.96%	7.00	1.85%	4.76%
60.00	5.00	8.00	2.12%	0.97%	5.00	1.32%	4.98%
	1.00	5.00	1.32%	1.01%	-	-	-
Panel E: AMD							
	10.00	26.00	6.37%	1.34%	10.00	2.45%	3.23%
30.00	5.00	18.00	4.41%	1.36%	7.00	1.72%	2.86%
	1.00	12.00	2.94%	1.72%	4.00	0.98%	2.93%
	10.00	12.00	3.17%	1.38%	5.00	1.32%	3.05%
60.00	5.00	11.00	2.91%	1.48%	4.00	1.06%	2.67%
	1.00	6.00	1.59%	1.03%	2.00	0.53%	2.31%

Calibration to the entire surface does not detect bubble episodes at the 1% significance level in *NVDA* price. Similarly, absence of exuberance, at this level, when employing a 60-day window, is documented for *META*. The  $RMSE_{\sigma,t}$  performances being inferior for  $\theta_{Surface}$  parameters, is consistent with the trend observed during calibrations to the S&P 500 put options. Recall, the *Entire Surface* dataset comprises of daily option contracts over a wide range of maturities, a large amount of which are significantly less liquid in contrast to those in the *HCV* dataset. In comparison to the average metric from the entire sample, the  $RMSE_{\sigma,t}$  for *MSFT* and *META* is lower across all window sizes and significant level combinations, with application of a 60-day window, at the respective 10%; and 10% and 5% levels of significance acting as exceptions. Whereas  $RMSE_{\sigma,t}$  corresponding to days revealing exuberance in *NVDA*, *AMZN*, and *AMD*, prices are significantly higher in comparison to average calibration performance across the entire sample. In regard to *AMD*, this holds throughout, barring the application of a 60-day window at the 1% level. Concerns over joint-hypothesis related issues are raised; however, it is also precisely the reason behind requiring a market analysis to examine validity of the detected bubbles.

## 6.3.2 Market Analysis

A brief chronological analysis of the exuberances detected, in correspondence to the respective applications of  $\theta_{HCV}$  and  $\theta_{Surface}$  parameters, follows. In contrast to the market analysis of S&P 500 bubbles, it is difficult to reason the formation of exuberance in individual stocks. Hence, in the next section of the chapter, an analysis regarding the price implication of bubbles, similar Fusari et al. (2024) is conducted. Bubbles observed in the *NVDA* stock, are displayed in *Figure 6.1*, with each panel corresponding to the set of parameters utilised for computing the fundamental value of call options. The days of exuberance are timestamped (*top*), along with the underlying price. Moreover, bubble magnitude and time-varying thresholds over different window sizes and levels of significances are revealed (*bottom*). Exuberance within each of the selected stocks are presented individually, with all illustrations having likewise interpretation.

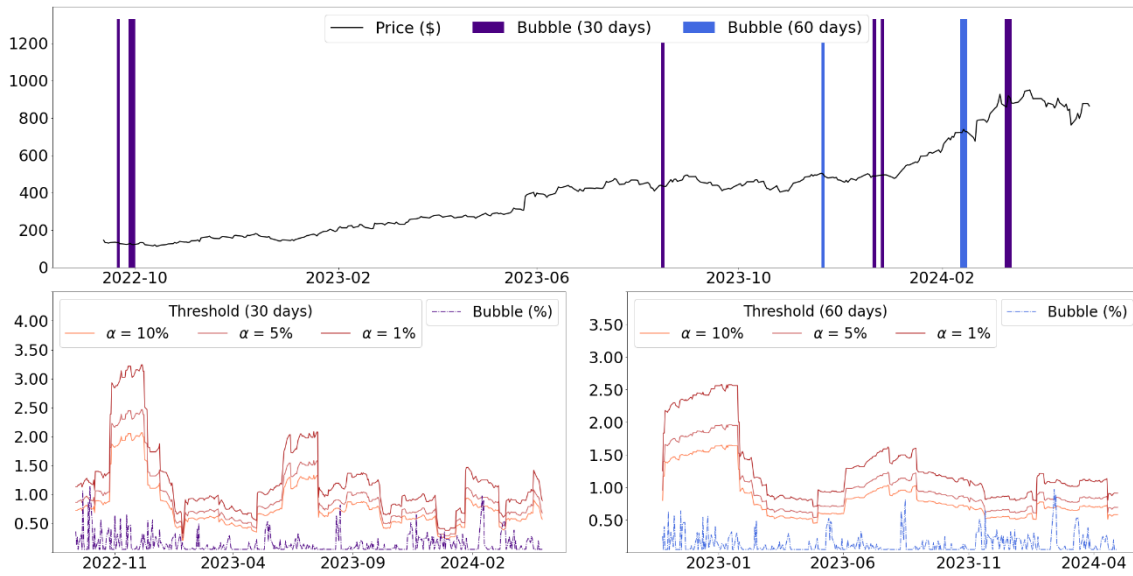
**Figure 6.1: NVDA Bubbles in HCV Call Options.**



Panel A:  $\theta_{HCV}$

*Figure 6.1 continued...*

Figure 6.1 continued...

Panel B:  $\theta_{Surface}$ 


Note: Daily bubbles detected in NVDA call options using  $\theta_{HCV}$  and  $\theta_{Surface}$ , are displayed in Panel A and B, respectively. Bubbles are timestamped (top) and assessed at each combination of window sizes (30, and 60 days) and levels of significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ).

After eliminating the initial burn-in period, between September 12, 2022, and April 30, 2024, the NVDA stock gained 495.67%. Despite the monumental gains, within approximately 20 months, very few instances of exuberance were identified. Both sets of parameters revealed an instance on September 21, 2022, with  $\theta_{Surface}$  also highlighting occurrences on September 28 and 30. Interestingly, the stock lost 5.28% and 4.05% on respective days, following the bubbles on September 21 and 29. Next,  $\theta_{HCV}$  parameters spotted exuberance on February 9, 16 and May 23, 2023. The NVDA stock lost 6.53% between February 8 and 21, before rebounding by 14.02% on February 23, 2023. A loss of 1.53% was observed on May 23, after which, a rally of 31.34% was experienced during May 25 and 30, 2023.

The NVDA stock peaked at \$950.02 on March 25, 2024, gaining 91.83%, since the start of the year. During this period, exuberance was first revealed on January 8 and 16, by  $\theta_{HCV}$  parameters. These dates, respectively experienced price rises of 6.43% and 3.06%.

In addition, calibration to the most liquid maturity exhibited a bubble on February 16, whereas  $\theta_{Surface}$  acknowledged instances on February 13 and 15. By this time, since the start of the new year the stock has risen by 46.63%. Next, a cluster of 5 bubbles, between March 8 and 25, 2024, were captured by  $\theta_{HCV}$ . Furthermore, three-consecutive days of exuberance, between March 11 and 13, were identified by using  $\theta_{Surface}$ . All but that on March 11, were exhibited by the former calibration approach, as well. Ironically, during the cluster, the stock only experienced a hike of 2.52%. After a 7.16% rise on March 12, 2024, the next day coincided with a 1.12% drop. The stock lost 7.44% of its value two days prior to the hike, including 5.55% on March 8, 2024.

**Figure 6.2: META Bubbles in HCV Call Options.**




Panel A:  $\theta_{HCV}$

Figure 6.2 continued...



Figure 6.2 continued...



Panel B:  $\theta_{Surface}$

Note: Daily bubbles detected in META call options using  $\theta_{HCV}$  and  $\theta_{Surface}$ , are displayed in Panel A and B, respectively. Bubbles are timestamped (top) and assessed at each combination of window sizes (30, and 60 days) and levels of significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ).

The META stock experienced a 154.60% price increase over the sample period, with detected exuberance displayed in Figure 6.2. First, a bubble was detected on April 21, 2023, using  $\theta_{HCV}$  parameters, 3 trading days prior to a 13.93% hike. The same set of parameters went on to observe occurrences on October 13, 20 and 23, 2023. Losses of 2.92% and 1.33%, were experienced on the respective first two occasions, which were followed by gains of 2.05% and 1.74%. The exuberance on October 20 was even captured by  $\theta_{Surface}$ .

Moreover,  $\theta_{Surface}$  parameters report a cluster of 8 bubbles between November 28 and December 11, 2023, with the final episode being acknowledged by  $\theta_{HCV}$ . During this cluster, the stock price reduced by 2.81%, with each day closing in red, barring those on December 7 and 8. In respective order, the two exceptions experienced spikes of 2.88% and 1.89%, with only the former exhibiting exuberance. At the turn of the year, both sets of parameters signalled towards exuberance in January 2024, however neither acknowledged the 20.32% hike experienced on February 2, 2024. Next,  $\theta_{HCV}$  detected bubbles on March 11, 13, and 18, 2024, similar to those experienced in NVDA prices. The former coincided with losses of

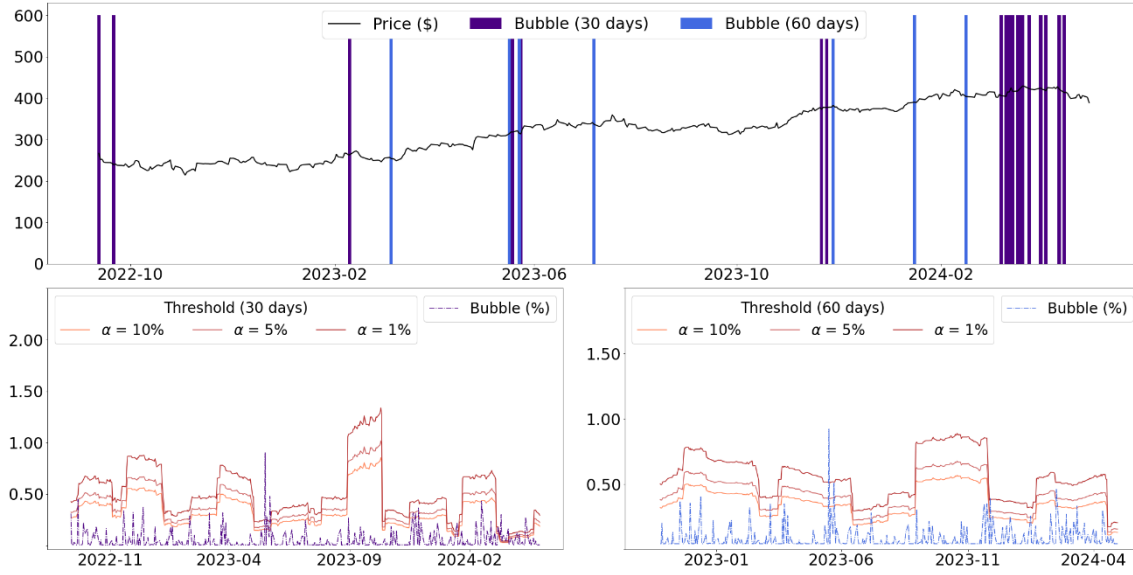
4.42% and 0.84% respectively, whereas the latter witnessed a 2.66% gain. In the following month,  $\theta_{HCV}$  and  $\theta_{Surface}$  captured instances on April 12 and 18, as the stock respectively changed by -2.15% and 1.54%. Additionally,  $\theta_{Surface}$  revealed bubbles associated with losses on April 16 and 17, whereas  $\theta_{HCV}$  exhibited occurrences coinciding with gains, on April 10 and 23. The *META* stock witnessed a 2.98% hike on April 23, marking a 49.03% rise from the beginning of 2024. However, it crashed by 10.56% on April 25, as a direct consequence of comments from CEO Mark Zuckerberg regarding the need to increase spending on AI products prior to making revenue from them (Milmo, 2024).

Next, exuberance in *MSFT* prices, as seen in *Figure 6.3*, is examined. Over the sample period, the stock gained 46.01%. Though highly promising, the returns are significantly lower than those experienced by *NVDA* and *META*. Initially,  $\theta_{HCV}$  revealed bubbles on September 12, and 21, 2022, with the latter, and that on September 22, being spotted by calibrating to the entire surface. Moreover,  $\theta_{HCV}$  captured occurrences on February 10, and March 7, 2023. In all instances, barring that on September 22, the stock closed in red. Each parameter set addressed instances on May 17, 23, and 24, with  $\theta_{HCV}$  also acknowledging that on May 19, 2023. The two trading days following May 24, experienced respective spikes of 3.85% and 2.14%. Additionally, according to  $\theta_{HCV}$ , exuberance was exhibited on November 21, 24, and 28, with each day revealing losses, barring the latter. Alternatively,  $\theta_{Surface}$  identified a cluster of 13 bubbles, between November 21, and December 26, 2023, during which *MSFT* prices fell by 0.74%.

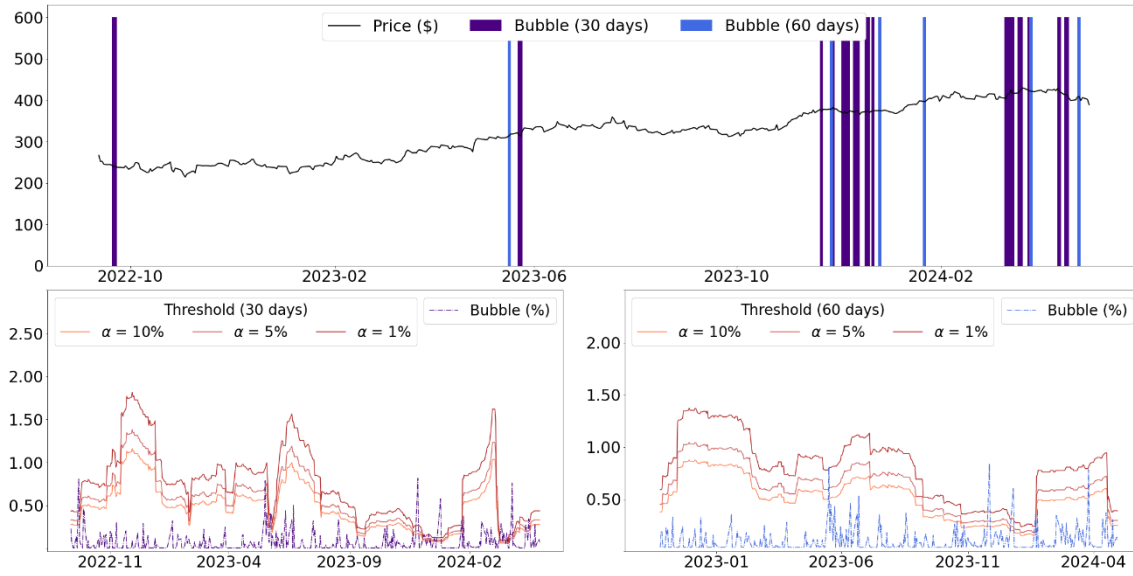
In 2024, *MSFT* prices exhibited exuberance on January 16, and February 16; and January 22, according to  $\theta_{HCV}$  and  $\theta_{Surface}$  parameters, respectively. The former parametric set further acknowledges bubbles on each day between March 8 and 21, 2024, in addition to that on March 25. Similarly, all episodes, but for those on March 18, and 21, are addressed by the latter. Moreover,  $\theta_{Surface}$  captured a bubble on March 26, 2024. During this period, the *MSFT* stock increased by 3.06%. On March 12, it rose by 2.66%, after closing in red during 5 out of the last 6 days, losing approximately 2.61%. The stock went on to gain 2.44% on March 14, prior to dropping by 2.07% on March 15, 2024. Recall, that bubbles in the early stages of the cluster, were similarly timestamped in *NVDA* and *META* prices. Finally,  $\theta_{HCV}$  revealed bubbles on April 1, 4, 12, and 15, whereas  $\theta_{Surface}$  captured instances on

April 12, 16 and 17. Overall, from the start of the year, till the final bubble, on April 17, 2024, *MSFT* prices increased by 9.52%.

**Figure 6.3: *MSFT* Bubbles in *HCV* Call Options.**



Panel A:  $\theta_{HCV}$




Panel B:  $\theta_{Surface}$

*Note: Daily bubbles detected in MSFT call options using  $\theta_{HCV}$  and  $\theta_{Surface}$ , are displayed in Panel A and B, respectively. Bubbles are timestamped (top) and assessed at each combination of window sizes (30, and 60 days) and levels of significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ).*

The prices of *AMZN* experienced growth of 28.25% during the sample, the lowest amongst the selected stocks. Initially, as witnessed in *Figure 6.4*, both forms of calibrations revealed bubbles on January 30 and May 23, 2023, coinciding with losses. In regard to the former, after the bubble, the stock gained 10.44%, prior to crashing by 8.43% and kick-starting 5-day losing streak on February 3. The  $\theta_{HCV}$  parameters further revealed four consecutive bubbles between October 23 and 27, including the occurrence on October 26 being captured by  $\theta_{Surface}$ . The *AMZN* stock finished 5.58% and 1.50% below on October 25 and 26, respectively, prior to gaining 6.83% the next day.

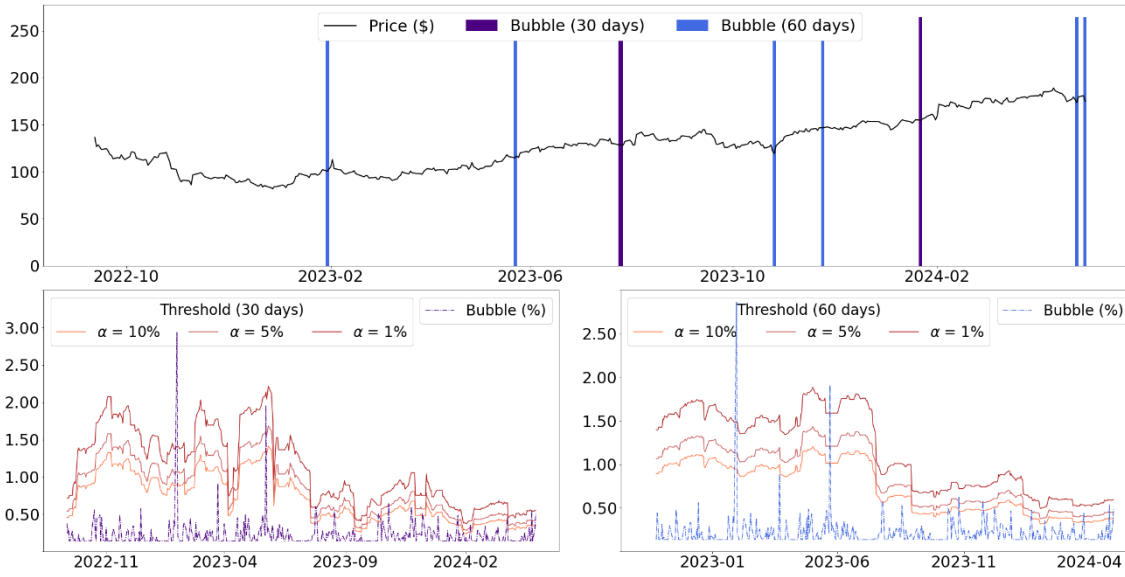
**Figure 6.4: *AMZN* Bubbles in HCV Call Options.**



Panel A:  $\theta_{HCV}$

*Figure 6.4 continued...*

Figure 6.4 continued...




Panel B:  $\theta_{Srface}$


Note: Daily bubbles detected in AMZN call options using  $\theta_{HCV}$  and  $\theta_{Srface}$ , are displayed in Panel A and B, respectively. Bubbles are timestamped (top) and assessed at each combination of window sizes and levels of significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ).

Nevertheless, all days closed in red, between December 22, 2023, and January 4, 2024, expect December 28, 2023. Calibration to the most liquid maturity allowed for capturing exuberance on December 27 and 28, 2023; and January 3, and 8, 2024. The latter corresponded to a 2.66% hike, and was followed by bubbles on January 16, 19, 23 and 24, 2024, with the occurrence on January 22, which was revealed by  $\theta_{Srface}$ . Moreover,  $\theta_{HCV}$  addressed exuberance on April 12, corresponding to a 1.54% loss. However, as of April 11, 2024, share prices grew year-on-year by 89.20%, and marked a 26.09% increase, since the start of 2024. From this day, till April 30, 7.49% of the value was lost, despite a spike of 3.43% on April 26, 2024. The last two bubbles were identified on April 25 and 30, 2024, by  $\theta_{Srface}$ , coinciding with losses of 1.65% and 3.29%, respectively. Although heavy losses were experienced towards the close of April 2024, AMZN prices provided a 71.48% year-on-year return.

**Figure 6.5: AMD Bubbles in HCV Call Options.**



Panel A:  $\theta_{HCV}$



Panel B:  $\theta_{Surface}$

*Note: Daily bubbles detected in AMD call options using  $\theta_{HCV}$  and  $\theta_{Surface}$ , are displayed in Panel A and B, respectively. Bubbles are timestamped (top) and assessed at each combination of window sizes (30, and 60 days) and levels of significance ( $\alpha = 10\%$ ,  $5\%$ , and  $1\%$ ).*

The exuberance within prices of *AMD*, which gained 87.12% over the sample period are illustrated in *Figure 6.5*. The stock lost 21.67% from the start of the sample, till September 26, 2022; with  $\theta_{Surface}$  detecting exuberance on September 21. Alternatively, after a 6/16% loss on November 9, share prices rose by 14.27% and 7.39% over the next consecutive days, with a bubble being timestamped on November 14, 2022. A 5-day winning streak ended on May 3, 2023, with a 9.22% crash, as exuberance was identified on May 1. Furthermore, sandwiched between crashes of 5.64% and 1.35%, a bubble was observed on June 1, coinciding with a 1.07% rise. Occurrences on August 11 and 22, witnessed respective losses of 2.41% and 2.37%, but were followed by gains in excess of 3.55%. The latter of the two was spotted by  $\theta_{HCV}$ .

A collective of episodes, similar to those in *AMZN* prices was spotted by  $\theta_{HCV}$ , between October 24 and 27, 2023. Additionally, bubbles were revealed on October 13 and 31, with the former, along with that on October 25, addressed by  $\theta_{Surface}$ . During this period, the stock lost 9.46%, prior to rebounding by 9.69% on November 1, 2023. At the turn of the year, exuberance on January 8, 19, and 23, 2024, coinciding with respective returns of 5.48%, 7.11% and 0.14%, were detected by  $\theta_{HCV}$ . Whereas  $\theta_{Surface}$  parameters acknowledged occurrences on January 24, and 25, corresponding to the following gains: 5.86% and 1.14%. Interestingly, a similar cluster to that seen in *NVDA*, *META*, and *MSFT* prices was revealed by  $\theta_{HCV}$ , as all days between March 8 and 19, 2024 displayed bubbles. Despite the cluster period witnessing a 14.17% drop in value, hikes of 2.20% and 2.14% were observed on March 12 and 15, respectively. Furthermore,  $\theta_{HCV}$  addressed a bubble on April 12, as the share price plummeted by 4.23%. Finally,  $\theta_{Surface}$  identified bubbles on April 24, and 25, 2024. Even though the stock experienced a large growth over the sample period, *AMD* prices rose by only 7.44% since the start of 2024.

### 6.3.3 Model Performance

The deep calibration framework was applied to the respective *HCV* and *Entire Surface* datasets of *NVDA*, *META*, *MSFT*, *AMZN*, and *AMD*. The identification of call option bubbles revealed presence of strict local martingale tendencies within the equity prices, across several days between August 1, 2022, and April 30, 2024. Furthermore, bubbles in stock

prices occurred, more frequently over shorter periods of time, in comparison to those witnessed in the S&P 500 index. Once again, statistical validity of the observed bubbles, is examined by  $RMSE_{\sigma,t}$ . The performances of the GSVJD model across the entire sample (initial window period is burned), days with positive and negative bubble magnitudes, are presented in *Table 6.4*.

**Table 6.4: Model Performance relevant to the Case Study.**

*Note:  $RMSE_{\sigma,t}$  (%) metrics with respect to  $\theta_{HCV}$  and  $\theta_{Surface}$  from each stock, corresponding to the entire sample (excluding burn-in), along days with negative ( $\widehat{\mathbb{B}}_t < 0$ ) and positive ( $\widehat{\mathbb{B}}_t > 0$ ) bubble magnitudes, across all window sizes.*

	Window	Entire	$\widehat{\mathbb{B}}_t < 0$	$\widehat{\mathbb{B}}_t > 0$
Panel A: MSFT				
$\theta_{HCV}$	30.00	2.15%	2.16%	2.14%
	60.00	2.16%	2.19%	2.13%
$\theta_{Surface}$	30.00	2.34%	2.36%	2.32%
	60.00	2.35%	2.37%	2.32%
Panel B: NVDA				
$\theta_{HCV}$	30.00	2.03%	2.16%	1.91%
	60.00	1.97%	2.10%	1.85%
$\theta_{Surface}$	30.00	3.01%	3.20%	2.82%
	60.00	3.02%	3.22%	2.82%
Panel C: AMZN				
$\theta_{HCV}$	30.00	1.40%	1.33%	1.47%
	60.00	1.36%	1.28%	1.43%
$\theta_{Surface}$	30.00	2.61%	2.80%	2.42%
	60.00	2.61%	2.82%	2.40%
Panel D: META				
$\theta_{HCV}$	30.00	1.37%	1.38%	1.37%
	60.00	1.35%	1.35%	1.35%
$\theta_{Surface}$	30.00	3.58%	4.00%	3.11%
	60.00	3.56%	4.03%	3.06%
Panel E: AMD				
$\theta_{HCV}$	30.00	1.60%	1.56%	1.64%
	60.00	1.55%	1.51%	1.58%
$\theta_{Surface}$	30.00	2.43%	2.51%	2.36%
	60.00	2.40%	2.49%	2.33%



According to the local martingale theorem, occurrence of negative bubbles, can be interpreted as potential model misspecification. Since, the fundamental value of the stock, acts as the lower bound to its market price, negative bubbles cannot exist. In the sample, days associated with  $\widehat{\mathbb{B}}_t < 0$ , on average should coincide with larger error metrics. This ensures most of the error within the model is concentrated away from  $\widehat{\mathbb{B}}_t > 0$  or  $\widehat{\mathbb{B}}_t > \alpha\widehat{\sigma}_t$ , hence eliminating concerns over the occurrence of a potential joint-hypothesis related issue. Although,  $RMSE_{\sigma,t}$  metrics, across the entire sample are relatively larger, than those associated with the S&P 500 index, when using  $\theta_{Surface}$ , the average errors are lower for  $\widehat{\mathbb{B}}_t > 0$ . Alternatively, application of  $\theta_{HCV}$  provide larger  $RMSE_{\sigma,t}$ , when  $\widehat{\mathbb{B}}_t > 0$ , for *AMZN*, and *AMD*, but by very small margins across both windows.

## 6.4 Factors Influencing the Formation of Bubbles

Throughout this study, presence of exuberance in the S&P 500 index (*SPX*), and 5 stocks; *MSFT*, *NVDA*, *AMZN*, *META*, and *AMD*, have been thoroughly examined. Next, in abidance with Fusari et al. (2024), implications of bubbles detected in *HCV* call options are explored, providing economic intuition to the exhibition of exuberance. Additionally, this acts a robustness check for applying the deep calibration framework constructed to calibrate GSVJD parameters within the three-step approach. The impacts of underlying price volatility, trading volume, earnings announcements, and GSVJD parameters ( $\rho$  and  $p$ ) on the formation of bubbles are examined. Given smaller sample sizes for each individual stock, the analysis is conducted by pooling together data of all the assets. Furthermore, detection of stock price bubbles, though not displayed in 6.3.2, are considered over the 90- and 180-day, to maintain consistency with S&P 500 observations.

The presence of bubbles is governed by explosive behaviour in underlying volatility, with respect to the drift, which can be measured by taking the expectation of the risk-neutral

quadratic variations<sup>51</sup>. The correlation between the quadratic variation and call option bubble magnitudes are documented in *Table 6.5*, over all window sizes, in correspondence to parameters from each dataset. In all instances, the correlation coefficient is positive, and much larger when associated with  $\theta_{Surface}$ . The strength of positive correlation between the quadratic variations and bubble magnitudes is not a strong one, which could possibly hint at the presence of explosiveness in underlying price (Jarrow & Kwok, 2024).

**Table 6.5: Correlation between Quadratic Variation and Call Option Bubbles.**

*Note: Correlation coefficients (%) between risk-neutral quadratic variation and call option bubbles, across various window sizes (days), when using  $\theta_{HCV}$  and  $\theta_{Surface}$ .*

Window	$\theta_{HCV}$	$\theta_{Surface}$
30.00	12.80%	16.14%
60.00	13.16%	15.53%
90.00	13.50%	14.92%
180.00	17.18%	19.37%

The importance of explosive behaviour in volatility, stems from Delbaen & Shirakawa (2002), who indicate a bubble to exist if and only if,  $\int_{\epsilon}^{\infty} \frac{s}{\sigma(s)} < \infty$ , for  $\epsilon > 0$ . This intuition is similar to the strict local martingale condition on parameters  $\rho$  and  $p$  (Andersen & Piterbarg, 2007; Fusari et al., 2024). Rising volatility is typically accompanied by increased trading activity, and subsequently a greater likelihood of exuberance being exhibited. Several factors govern the trading activity associated with the underlying and its options, however, earning announcements, which are prescheduled, present a greater opportunity to comprehend the motive of solely purchasing an asset to resell it at a higher price. Trading activity analogous to the options and underlying are driven by expectations regarding earnings, prior to the date of announcement, which contribute towards explosive behaviour in volatility. Therefore, inspired by Fusari et al. (2024) the impact of trading volume and earning announcements, on the formation of bubbles, is explored through (51) and (52).

$$\mathbb{B}_{i,t}^+ = \gamma_0 + \gamma_1 \ln(Vol_{i,t}) + \gamma_2 \ln(Call Vol_{i,t}) + \epsilon_t \quad (51)$$

<sup>51</sup> The estimation of quadratic variation,  $QV_{t,t+\tau} = \frac{1}{\tau} \left( \int_t^{t+\tau} V_s ds + \sum_{t < s \leq t+\tau} \mathcal{J}_s^2 \right) \approx V_{0,t} + 2\lambda_t \left( e^{\mu_{y,t} + 0.5\sigma_{y,t}^2} - \mu_{y,t} - 1 \right)$ , depends spot variances as the volatility tends to be strongly persistent (Fusari et al. 2024).

$$\begin{aligned}\widehat{\mathbb{B}}_{i,t}^+ = & \omega_0 + \omega_1 D_{1-14,i,t} + \omega_2 D_{15-28,i,t} + \omega_3 D_{29-42,i,t} \\ & + \omega_4 \ln(Vol_{i,t}) + \omega_5 \ln(Call Vol_{i,t}) + \epsilon_t\end{aligned}\quad (52)$$

The impact of daily trading activity on the formation of positive call option bubbles,  $\widehat{\mathbb{B}}_{i,t}^+$ , is examined in (51). The coefficient,  $\gamma_1$  and  $\gamma_2$ , respectively, document change in  $\widehat{\mathbb{B}}_{i,t}^+$ , for a 1% increase in daily traded volume of the underlying ( $Vol_{i,t}$ ), and within the accumulated volume of the most liquid call options ( $Call Vol_{i,t}$ ). The regression is estimated over two versions of the *HCV* and *Entire Surface* pooled dataset, with the former comprising of  $i = \{MSFT, NVDA, AMZN, META, AMD\}$ , whereas the latter is inclusive of the *SPX*. The S&P 500 index, unlike its constituents, is not applicable for examining the impacts of earning announcements on bubble formations. Hence, its exclusion ensures the possibility to reveal the robust impact of volume by controlling for earnings announcements, and vice versa, on the presence of stock bubbles. In equation (52),  $D_{t_1-t_2} = 1$ , if the announcement is within the next  $t + t_1$  and  $t + t_2$  days, and, 0, otherwise. Since,  $\widehat{\mathbb{B}}_{i,t}^+$  represents positive bubble magnitudes, normalised by the underlying price, the coefficients of (51) and (52) are interpreted as percentages (%).

**Table 6.6: Results for Regression (51).**

*Note: Results from regression (51), corresponding to  $\theta_{HCV}$  and  $\theta_{Surface}$  parameters, along with estimated coefficients ( $\gamma_0, \gamma_1, \gamma_2$ ), and their levels of significance (in parentheses) are reported. The following, ‘\*’, ‘\*\*’, and ‘\*\*\*’, represent significance at the respective, 1%, 5%, and 10%, levels. Additionally, the model  $p$ -value ( $F$ -stat), and  $R^2$  values are displayed. Results reported in Panel A are associated with  $i = \{MSFT, NVDA, AMZN, META, AMD\}$ , whereas those Panel B are inclusive of *SPX*.*

Window Size	30.00		60.00		90.00		180.00	
	$\theta_{HCV}$	$\theta_{Surface}$	$\theta_{HCV}$	$\theta_{Surface}$	$\theta_{HCV}$	$\theta_{Surface}$	$\theta_{HCV}$	$\theta_{Surface}$
Panel A								
$\gamma_0$	0.105	0.185	0.097	0.160	0.101	0.193	0.115	0.198
	(**)	(*)	(**)	(*)	(**)	(*)	(**)	(*)
$\gamma_1$	0.009	0.018	0.007	0.014	0.007	0.012	0.011	0.016
	(*)	(*)	(*)	(*)	(*)	(*)	(*)	(*)
$\gamma_2$	-0.020	-0.040	-0.017	-0.031	-0.016	-0.031	-0.023	-0.038
	(*)	(*)	(*)	(*)	(*)	(*)	(*)	(*)
$R^2$	0.019	0.019	0.019	0.019	0.019	0.019	0.019	0.019
$p$ -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 6.6 continued...

Table 6.6 continued...

Panel B								
$\gamma_0$	0.110 (*)	0.176 (*)	0.092 (**)	0.132 (**)	0.093 (**)	0.151 (*)	0.119 (*)	0.173 (*)
$\gamma_1$	0.008 (*)	0.017 (*)	0.006 (*)	0.014 (*)	0.006 (*)	0.013 (*)	0.007 (*)	0.015 (*)
$\gamma_2$	-0.018 (*)	-0.037 (*)	-0.014 (*)	-0.029 (*)	-0.014 (*)	-0.029 (*)	-0.018 (*)	-0.033 (*)
$R^2$	0.016	0.031	0.010	0.020	0.008	0.018	0.012	0.024
$p$ -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

The estimated coefficients of regression (51) are presented in Table 6.6, with Panel A displaying results corresponding to the consideration of only stocks, whereas Panel B illustrates those associated with the inclusion of SPX. In both cases, across all window sizes and datasets utilised for calibrating GSVJD parameters,  $\gamma_2$  is negative at 1% significance. This translates to a drop in the bubble magnitude, when daily accumulated call option volume across the most liquid maturity increases. Falling volume, could indicate bearish sentiments, signalling towards the market price of the stock exceeding its fundamental value. The examination of Panel A reveals  $\gamma_1$  to be positive at 1% significance, in all situations. Hence, the size of exuberance is found to rise with trading activity of the underlying asset. This is consistent with the intuition that rising trading volumes are associated with growing volatility, and subsequently the formation of bubbles. The robustness of the findings is maintained in Panel B, with respect to the impact and significance of  $\gamma_1$  and  $\gamma_2$  coefficients. Moreover, observe that the magnitude impact associated with both underlying and call options volumes, is greater when detecting bubbles using  $\theta_{Surface}$ .

Next, results from (52) are depicted in Table 6.7. The coefficients reported in Panel A correspond solely towards examining the impact of earnings announcements, whereas those in Panel B are associated with controlling for underlying and options trading activity. In both panels,  $\omega_1$  has a positive impact on the formation of bubbles, whereas  $\omega_2$  and  $\omega_3$  have a negative effect. The latter coefficients are statistically insignificant for exuberance observed using  $\theta_{HCV}$ , with the exception of the 90-day window length in Panel A. Alternatively, the negative impact of  $D_{15-28,t}$  is significant at 1% in association with  $\theta_{Surface}$ .



It can be comprehended that as the earnings announcement nears, especially when being between 1 to 14 days away, the motive to purchase the stock for the sole purpose of reselling at a higher price, and subsequent bubble magnitude, increases. This reflects in the trading volume of the underlying, which signals towards potential explosive behaviour in volatility. The impact of earning announcements remains the same, even when controlling for volume. Notice, the signs of coefficients and significance of  $\omega_4$  and  $\omega_5$  are similar to their counterparts observed in *Panel A* of *Table 6.6*. It implies that increased underlying and reduced accumulated call option volumes, have a positive impact on bubble magnitudes, even when controlling for earnings announcements. To further, assess for this claim, the impact of GSVJD parameters  $\rho$  and  $p$  are examined.

$$\mathbb{B}_{i,t}^+ = \beta_0 + \beta_1 \rho_{i,t} + \beta_2 D_{p,i,t} + \beta_3 (D_{p,i,t} \times \rho_{i,t}) + \epsilon_t \quad (53)$$

$$\mathbb{P} \left( 1_{\{\mathbb{B}_{i,t}^+ > 0\}} \right) = \beta_0 + \beta_1 \rho_{i,t} + \beta_2 D_{p,i,t} + \beta_3 (D_{p,i,t} \times \rho_{i,t}) + \epsilon_t \quad (54)$$

The intuition behind the underlying price process revealing strict local martingale tendencies, when  $0.5 < p < 1.5$ , and  $\rho > 0$ , stems from the exuberance being observed due to explosive behaviour in volatility (Delbaen & Shirakawa, 2002; Fusari et al., 2024). The motive behind estimating regressions (53) and (54), inclusive of the S&P 500 findings, is to provide robustness for the test utilized to detect significant call option bubbles. In both equations, following Fusari et al. (2024) the 21-day moving average of  $\mathbb{B}_{i,t}^+$ ,  $\rho$  and  $p$  are considered, with  $D_{p,i,t} = 1$ , if  $0.5 < p < 1.5$ , or 0, otherwise. The variable  $D_{p,i,t} \times \rho_{i,t}$  captures the interaction between the parameters, and subsequent effect on call option bubbles. Furthermore, in (54), the probability of a bubble occurring is examined, hence a logit regression is implemented, whereas for (53) an OLS model is used.

The examination of an increase in  $\rho$  is crucial, as by itself,  $0.5 < p < 1.5$ , is not sufficient enough for volatility to exhibit explosive behavior. The explosiveness is only possible when, in addition to the condition on  $p$ ,  $\rho > 0$ . Therefore, the practitioner is requested to thread carefully, in the aforementioned scenario. Overall, during instances when the coefficients have a significant impact on the exhibition of bubbles, the average effect of a 0.1 increase in  $\rho$ , for when  $0.5 < p < 1.5$ , can be examined. In *Panel A*,  $\Delta \mathbb{B}$  indicates the average change in bubble magnitude, whereas  $\Delta \mathbb{P}$  (*Panel B*) reveals the variation in the probability of

identifying a bubble. Across, all instances, the overall impact of an increase in  $\rho$  is positive, resulting in a greater bubble magnitude and probability of detection. To illustrate this, consider the bubbles observed using  $\theta_{Surface}$  parameters, within the 30-day window. For a 0.1 increase, in  $\rho$ , the magnitude of the bubble rises by close to 0.80%, whereas the average probability<sup>52</sup> of detection spikes by approximately 1.50%.

**Table 6.8: Results for Regression (53) and (54).**

*Note: Results from regression (53) and (54), for  $i = \{MSFT, NVDA, AMZN, META, AMD, SPX\}$ , are reported in Panel A and B, respectively. The estimated coefficients  $(\beta_0, \beta_1, \beta_2, \beta_3)$ , and corresponding levels of significance (in parentheses), are reported in accordance with  $\theta_{HCV}$  and  $\theta_{Surface}$  parameters. The following, ‘\*’, ‘\*\*’, and ‘\*\*\*’, represent significance at, 1%, 5%, and 10% levels respectively. Furthermore, Panel A displays model  $p$ -value ( $F$ -stat), and  $R^2$  metrics, whereas Panel B reveals  $p$ -value (LLR), pseudo  $R^2$  figures. In accordance, the average change in bubble magnitude ( $\Delta\mathbb{B}$ ) and probability of occurrence ( $\Delta\mathbb{P}$ ), for a 0.1 units increase in  $\rho$  are documented.*

Window Size	30		60		90		180	
	$\theta_{HCV}$	$\theta_{Surface}$	$\theta_{HCV}$	$\theta_{Surface}$	$\theta_{HCV}$	$\theta_{Surface}$	$\theta_{HCV}$	$\theta_{Surface}$
Panel A: OLS								
$\beta_0$	-0.085 (*)	0.052 (*)	-0.122 (*)	0.049 (*)	-0.137 (*)	0.044 (*)	-0.188 (*)	0.030 (*)
$\beta_1$	-0.093 (*)	0.160 (*)	-0.154 (*)	0.143 (*)	-0.175 (*)	0.114 (*)	-0.283 (*)	-0.004 -
$\beta_2$	0.123 (*)	-0.069 (*)	0.153 (*)	-0.077 (*)	0.164 (*)	-0.077 (*)	0.207 (*)	-0.080 (*)
$\beta_3$	0.200 (*)	-0.076 (*)	0.241 (*)	-0.090 (*)	0.246 (*)	-0.082 (*)	0.323 (*)	-0.024 -
$R^2$	0.095	0.078	0.077	0.074	0.064	0.064	0.070	0.091
$p$ -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\Delta\mathbb{B}$	0.011	0.008	0.009	0.005	0.007	0.003	0.004	-0.003
Panel B: Logit								
$\beta_0$	-0.912 (***)	1.694 (*)	-1.685 (*)	1.698 (*)	-2.157 (*)	1.632 (*)	-3.886 (*)	1.146 (*)
$\beta_1$	-0.660 -	3.340 (*)	-1.927 (**)	3.273 (*)	-2.537 (*)	2.988 (*)	-6.319 (*)	0.336 -

Table 6.8 continued...

<sup>52</sup> It is obtained by computing the average difference between probabilities obtained from considering  $\rho = [-1, 1)$ , at increments of 0.1, and adding 0.1 to each  $\rho$  value.

Table 6.8 continued...

$\beta_2$	1.717 (*)	-2.357 (*)	2.407 (*)	-2.354 (*)	2.831 (*)	-2.342 (*)	4.451 (*)	-1.957 (*)
$\beta_3$	2.860 (*)	-2.682 (*)	3.885 (*)	-2.678 (*)	4.325 (*)	-2.606 (*)	7.592 (*)	-0.447 -
<i>Pseudo R</i> <sup>2</sup>	0.041	0.080	0.034	0.080	0.030	0.075	0.031	0.082
<i>p-value</i>	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\Delta\mathbb{P}$	0.039	0.015	0.037	0.013	0.035	0.008	0.027	-0.002

Recall, under the ND condition, the put-call parity holds, therefore bubbles in call options have a linear relationship, with those in the underlying price. Even though the entire size of the bubble is not captured, the exuberance within call options is capable of signalling towards strict local martingale tendencies in the underlying price process. This is the foundation of the three-step approach, which relies upon a statistical test for identifying call option bubbles. The calibration of parameters are prone to non-unique solutions (Wshah et al., 2020) and high standard errors (Fusari et al., 2024). Thus, the statistical test was relied upon for acknowledging exuberance in the underlying prices. However, bubbles identified by the three-step approach, in combination with the deep neural network framework, align with the conditions associated with parameters  $\rho$  and  $p$ . Both the probability of observing, and magnitude of exuberance increase with  $\rho$ , for  $0.5 < p < 1.5$ , indicating to the existence of bubbles due to explosive behaviour in volatility. Hence, providing robustness towards the implementation of the three-step approach, especially over the *Entire Surface* dataset, for bubble detection in call option and underlying prices.

## 6.5 Summary and Discussion

The current chapter builds on the previous exploration of bubbles in the S&P 500 index, by applying the deep calibration framework to five technology stocks: *AMZN*, *AMD*, *META*, *MSFT*, and *NVDA*. The options written on these stocks are American-styled and cannot reveal exuberance within their prices. However, in abidance with the local martingale theory of bubbles, exuberance can be detected by examining the difference between market price of American-styled call options, and the fundamental value of their European-styled counterparts. Since, American put options cannot reveal bubbles, the deep calibration framework is applicable for computing GSVJD model parameters from market observations.



Parameters were calibrated from the *HCV* and *Entire Surface* datasets of each stock, during the sample period between August 1, 2022, and April 30, 2024.

The exuberance in stock prices occurred more frequently over shorter time, reflecting the existence of such instances in clustered periods. Despite substantial increases witnessed in *NVDA* and *META* prices, relatively, very few days of exuberance were noticed. It highlighted that the appearance of bubbles was not necessarily associated with rapid price rises and also acknowledged the benefits of using smaller time frequencies during detection. The latter prevents the false classification of stocks exhibiting bubbles over larger periods, due to overestimation of explosiveness in the price process. Majority of bubble episodes coincided with days preceding/succeeding those that experienced relatively large losses/gains. Additionally, days witnessing monumental returns were not always dated for exuberance.

A similar cluster of bubbles was observed for *NVDA*, *META*, *AMD* and *MSFT*, between March 8 and 19, 2024. Despite strong performances over the course of 2024, technology stocks remained volatile. Overall, exuberance in 2024, coincided more frequently with profitable days. However, in regard to the cluster, traders sought to book profits, as rising inflation and treasury yields sparked concerns regarding overvaluation, with semiconductor stocks being hit the most. This is consistent with the local martingale theory of bubbles, which states bubbles to capture the willingness of trades to purchase for the sole purpose of reselling the asset at a higher price. Especially, in the case of *NVDA*, massive price hikes do not necessarily reveal exuberance, as they could simply resemble strong fundamentals, and traders wanting to purchase the stock for reciprocating gains from holding it.

Next, in abidance with Fusari et al. (2024), factors influencing the formation of bubbles were explored. Exuberance was discovered to be positively influenced by increased trading activity in the underlying, and proximity to earning announcements. The latter indicated towards motives of investors purchasing the stock, for the sole purpose of reselling, in anticipation of a price hike. Whereas, trading activity in the underlying is associated with high levels of volatility, following the same intuition as Delbaen & Shirakawa (2002). When observations in the S&P 500 index, were considered, the impact of trading volume remained robust. Moreover, it was revealed that reduced call option volume, which signal towards

bearish sentiment, and subsequently, the market price exceeding the perceived fundamental value, increased the bubble magnitude.

Finally, a robustness check for bubbles identified by the enhanced three-step approach was conducted. Recall, given the non-unique nature of the differential evolution optimiser, it was difficult to rely on the strict local martingale conditions associated with the  $p$  and  $\rho$  parameters. Instead, signals of exuberance in the underlying price were identified from call option prices, with the assistance of a statistical test. Hence, upon estimating the OLS and Logit regressions from Fusari et al. (2024), the magnitude and detection probability of bubbles, was revealed to respectively increase, with  $\rho$  increases, given that  $0.5 < p < 1.5$ . Therefore, bubbles identified by the three-step approach, with application of the deep calibration framework, were consistent with the strict local martingale conditions, associated with the parameters. This provided strong support to the implementation of the enhanced three-step approach, especially with respect to using the *Entire Surface* dataset for detecting bubbles in call options and underlying prices.



# Chapter 7:

# Conclusion

This thesis thoroughly explores the phenomenon of asset price bubble detection, and the importance of timely interventions to mitigate the devastating financial and socioeconomic consequences associated with their bursts. Given the increasing integration of global financial markets, and the frequent occurrence of exuberance in recent times, this research highlights the detrimental need for an accurate, robust and efficient detection method. The current chapter summarises the research, whilst highlighting contributions and discussing future implications.

## 7.1 Summary, Contributions and Future Implications

Traditional bubble detection approaches are largely focused on estimating the fundamental value of the asset and identifying deviations in market prices. Whilst recursive regression techniques (Phillips et al., 2011, 2015) and the LPPLS model (Johansen et al., 1999, 2000) have gained traction, they suffer from the joint-hypothesis issue, making results inconclusive. Furthermore, these methods require large time series data, making them

vulnerable to structural breaks, with the latter, failing to consider exogenous factors that could influence the formation of bubbles. To overcome these challenges, preference for the local martingale theory of bubbles was revealed. It focuses on the detection of type III bubbles, which are short-term, and observed in the prices of assets with bounded lives. By examining for strict local martingale tendencies in the asset price, the local martingale theory does not rely on the estimation of fundamental values and hence avoids the joint-hypothesis problem. Moreover, it provides a mathematical and economic framework for the detection of asset price bubbles from option prices, which allows for capturing the forward looking expectations of market participants.

One of the key features of this research is the adaptation of the three-step approach developed by Fusari et al. (2024), under the tutelage of the local martingale theory. The method significantly improves detection by using option prices for capturing forward looking expectations regarding the formation of bubbles. It relies on the calibration of the GSVJD model to market put options, prior to detecting bubbles in the prices of call options for overcoming the joint-hypothesis issue. Put options have bounded payoffs, and cannot reveal bubbles, implying that their market price and fundamental values must align. Alternatively, call options have unbounded payoffs, making them desirable to purchase for selling at a higher price. Under the ND condition, exuberance in call option prices share a linear relationship with those in the underlying and hence signal towards strict local martingale tendencies in its process.

Despite its robustness, the three-step approach relies on the computationally expensive Monte Carlo simulations during calibrations, that limit its application during real-time bubble detection. Fusari et al. (2024) were forced to comprise by considering only the most liquid cross section of daily put options, for striking a balance between robustness and efficiency. Such calibrations are unable to capture the full extent of forward looking information prevalent across the variety of option maturities within the entire surface. The abundant information within the surface, especially with respect to changes in the underlying prices, significantly improves the quality of bubble detection. Furthermore, even a highly sophisticated model such as the GSVJD with capabilities of capturing mean-reverting volatility, jump and strict local martingale tendencies in the underlying price process, can be easily overlooked by practitioners due to its poor tractability. Compromising accuracy for

tractability during bubble detection can make the method suffer from a joint-hypothesis issue, and provide inadequate and untimely risk management, which could be devastating. To address this limitation, a deep calibration framework was proposed to answer the following research question: *How to improve the efficiency and tractability of the three-step approach and subsequently extend its application for bubble detection?*

Neural networks possess universal approximation capabilities and can estimate any function to a certain level of accuracy, making them ideal for pricing options, and handling the computational demands of the three-step approach. They mimic the functions of an animal's nervous system, to comprehend the complex non-linear relationship between input and output variables. There are two types of neural network based approaches for calibrating parameters of a given stochastic process. First, the one-step approach, which directly estimates parameters from market data. Second, the two-step approach, which trains the network to learn dynamics of the process and uses an optimiser to calibrate parameters from market observations. The latter is favoured, as it provides real-time validation and ensures robustness across different market regimes. Furthermore, given the provision of real-time validation, it is favourable amongst regulators for dismissing the 'black-box' nature of neural networks. By leveraging the computational power of neural networks, the proposed framework makes the need for Monte Carlo simulations redundant. Hence, the three-step approach is made more attractive for real-time applications, allowing for efficient bubble detection.

The deep calibration framework is the main contribution of this research. It significantly improves the computational efficiency and accuracy of bubble detection, making the three-step approach suitable for real-time applications. Given the lack of a closed form solution, the network was trained using Monte Carlo simulations. Even though 10,000 simulations, with a timestep of  $\frac{1}{365 \times 5}$  were considered, the accuracy and efficiency of calibration could potentially benefit from the availability of a closed form solution. Furthermore, it could improve the reliance on parameters  $p$  and  $\rho$  for detecting strict local martingale tendencies. Nevertheless, deriving a closed-form solution is beyond the scope of the current study and left for future research. However, most importantly, the boost in efficiency and accuracy provided by neural networks, enables the three-step approach to utilise the entire surface of

daily option prices, rather than being limited to a single volatility smile. This improves the ability to capture forward-looking market expectations across various maturities and better the overall accuracy of bubble detection.

Prior to applying the deep calibration framework, the GSVJD model was tested against various stochastic processes, over daily S&P 500 index data. The comparison provides another contribution, especially with respect to market turmoil periods, such as the COVID-19 induced crash. It reveals the GSVJD process to be superior, as it possesses capabilities of capturing price jumps and strict local martingale tendencies, despite being the most inefficient amongst the considered models. This cemented the importance of considering the GSVJD model in the three-step approach, and further motivated the requirement for improved efficiency, without sacrificing accuracy. For achieving the latter, the deep calibration framework is required to develop a comprehensive and robust understanding of the GSVJD dynamics. This is accomplished by conducting an extensive random search over a broad range of hyperparameters, with each architecture undergoing 3-fold cross validation. Next, the top performing networks are selected for optimal training and tested on market data and compared to calibration results obtained from Monte Carlo simulations. The most optimal architecture was found to nearly double the accuracy and improve efficiency by a  $2^{54}$  order of magnitude.

When comparing to the benchmark, the GSVJD was calibrated to the most liquid daily S&P 500 index put options. However, the boost in efficiency obtained from the deep calibration framework, enabled parametric estimation from the entire surface. Initially, the respective parameters,  $\theta_{HCV}$  and  $\theta_{Surface}$ , were utilised to detect exuberance from the most liquid cross section of call options. Notably, calibrating to the daily surface excelled at identifying exuberance in the S&P 500 index during the COVID-19 induced crash, and subsequent recovery phase. It captured speculative-driven bubbles, coinciding with circuit breakers and major price fluctuations. These findings were consistent with the local martingale theory, especially with respect to the NFLVR condition. The  $\theta_{HCV}$  documented similar occurrences, however failed to acknowledge exuberance on March 16, as the S&P 500 index collapsed by 11.98%, and experienced a circuit breaker. Moreover, it was not able to acknowledge instances during the immediate recovery period. Both sets of parameters addressed

exuberance in mid-2021, amidst rising inflation and spread of coronavirus cases. Overall,  $\theta_{Surface}$  demonstrated superior ability in detecting bubbles linked to speculative behaviour and economic triggers, by leveraging information from several volatility smiles.

Calibrating to the entire surface allowed for capturing a broader range of market expectations. It enhanced the identification of bubbles, by extending the application of the three-step approach to examining their existence and formation across various call option maturity groups, and over the lifetime of certain contracts. Similar instances of exuberance were identified as in the most liquid cross section, however, in larger clusters. Most notably in regard to the COVID-19 induced crash, and subsequent recovery period in 2020. The examination of bubbles over the lifetime of options provided important revelations about their characteristics. Options with greater maturities, in consistency with the local martingale theory, displayed larger bubbles, providing further insight into the potential size of the phenomenon in the underlying prices. Bubbles, though smaller, were more pronounced near the end of contract lifetimes, highlighting the desires of traders to purchase the options for a short-term profit. This reflected bubble formations to be more frequent than previously perceived. Moreover, it reflects differing investor perceptions across maturities, with long-term holders viewing bubbles to be larger, and a consequence of more prominent events.

The detected bubbles in the S&P 500 bare resemblance to the findings from Song et al. (2022) and Shu & Song (2024), especially for those observed during the post-crash recovery period, when employing  $\theta_{Surface}$  parameters. However, the neural network enhanced three-step approach does not reveal a strong presence of exuberance in 2021, as opposed to the application of the LPPLS model in the aforementioned works. Recall, type III bubbles are short-lived, in contrast to the type II episodes observed by the LPPLS method. Hence, as suggested by Fusari et al. (2024), methods such as the LPPLS and recursive regression techniques can be used to complement the three-step approach in future research. This would allow for capturing bubbles by modelling both the drift and diffusion components of the asset price process. Nevertheless, it is crucial to highlight the application of the deep calibration framework, discovered exuberance during the COVID-19 crash period, when the S&P 500 experienced multiple circuit breakers. This novel revelation is a major contribution



of the thesis, as it is consistent with the local martingale theory of bubbles and strongly justifies the requirement for a deep calibration framework.

The proposed framework captures forward-looking market expectations more effectively and provides a deeper understanding of speculative behaviour in financial markets, by expanding the calibration process to include the entire option price surface. The ability to detect bubbles across different option maturities and over the lifetime of options offers valuable insights into the formation and evolution of exuberance. As the frequency of occurrences increase, with growing interconnectedness of global financial markets, the need for robust detection systems becomes ever more critical. The deep calibration framework for enhancing the three step approach, provides a powerful tool for detecting asset price bubbles in real-time, enabling financial institutions and regulators to implement timely interventions for mitigating the potential damage of future market crashes, associated with exuberance.

The computational efficiency provided by the deep calibration framework makes the three-step approach more feasible for practitioners who seek to implement timely risk management practices. Furthermore, the flexibility of the neural networks allows it to handle a wide range of option maturities, providing a deeper understanding of bubbles formations across different investment horizons. The robustness of the enhanced three-step approach was tested by applying it to five major technology stocks: *AMZN*, *AMD*, *META*, *MSFT*, and *NVDA*. The case study revealed that stock bubbles, occur more frequently over short durations, however, their appearance is not necessarily associated with massive prices increases. This finding underscores the importance of distinguishing between rapid price rises driven by strong fundamentals and speculative hikes. Furthermore, it was established that stock bubbles occurred in clusters, often around periods of heightened trading activity and with close proximity to earnings announcements. Bubble magnitudes and call option volume were found to share a negative relationship, suggesting the presence of exuberance to signal bearish investor sentiments.

The case study revealed the flexibility of the deep calibration framework for bubble detection, using different option contract styles. Furthermore, it can be applied over several asset classes, even if each requires a different stochastic process. For instance, the GSVJD

model, which already excels at capturing mean-reverting volatility, jumps and strict local martingale tendencies in the underlying price; can be extended to incorporate attention factors, as in Cretarola et al. (2018, 2020), and applied to various cryptocurrencies. This enables novel identification of exuberance in cryptocurrencies, by capturing forward looking expectations to the fullest extent, from respective derivative markets. As witnessed in Liu et al. (2019) and Horvath et al. (2021), the two-step network calibration approach can be trained to learn several stochastic processes, even within a single architecture. Hence, a robust deep calibration framework could be designed, for implementing the three-step approach over various financial instruments by using the most favourable stochastic processes. The work of Biagini et al. (2024) explores this with several processes, capable of capturing strict local martingale tendencies. The deep calibration framework can be further extended in a similar manner, for parametric estimation of various stochastic process under the three-step approach.

A key component of the stochastic processes must be their ability to capture strict local martingale tendencies in the underlying process. This research revealed that bubbles occur more frequently than previously perceived, and do not necessarily coincide with large price hikes. Instead, they capture the eagerness of traders seeking to make a quick profit. It is important to take this observation into consideration. Recall, stochastic processes addressing strict local martingale features, outperformed their direct counterparts, and better captured market characteristics. Therefore, it is suggested to explore the pricing of options or generation of implied volatility surfaces, with a model that possess the ability to acknowledge strict local martingale characteristics. The recommendation is not necessarily implied for detecting bubbles, but rather to incorporate their possible existence when pricing instruments or computing key market metrics. From a regulatory perspective, it can boost the ability to implement timely, and most importantly accurately devised risk management practices.

The timely implementation of risk management techniques is crucial, relying heavily on the ability to accurately detect bubbles. Kwong & Wong (2022) revealed the benefits of pairing of Value-at-Risk (VaR) models, with recursive-regression bubble detection techniques. However, such methods suffer from a joint hypothesis issue and are unable to capture

forward looking expectations of market participants. The three-step approach overcomes such issues associated with traditional methods, and when enhanced with the deep calibration framework, enables capturing forward looking expectations to the full extent. Therefore, it is recommended to explore the pairing of VaR models, with the neural network enhanced three-step approach, possibly applicable across various asset classes and desired stochastic processes, to improve the timely implementation of risk management measures. It is on this note that the reader's patience and interest is deeply appreciated, and the thesis is concluded.

# References

- Abreu, D., & Brunnermeier, M. K. (2003). Bubbles and Crashes. *Econometrica*, 71(1), 173–204. <https://doi.org/10.1111/1468-0262.00393>
- Acharya, D. (2024). Comparative Analysis of Stock Bubble in S&P 500 Individual Stocks: A Study Using SADF and GSADF Models. *Journal of Risk and Financial Management*, 17(2), 59. <https://doi.org/10.3390/jrfm17020059>
- Ackerer, D., Tagasovska, N., & Vatter, T. (2020). Deep Smoothing of the Implied Volatility Surface. *Advances in Neural Information Processing Systems*, 33, 11552–11563. <https://proceedings.neurips.cc/paper/2020/hash/858e47701162578e5e627cd93ab0938a-Abstract.html>
- Almeida, C., Fan, J., Freire, G., & Tang, F. (2023). Can a Machine Correct Option Pricing Models? *Journal of Business & Economic Statistics*, 41(3), 995–1009. <https://doi.org/10.1080/07350015.2022.2099871>
- Andersen, L. B. G., & Piterbarg, V. V. (2007). Moment explosions in stochastic volatility models. *Finance and Stochastics*, 11(1), 29–50. <https://doi.org/10.1007/s00780-006-0011-7>
- Bashchenko, O., & Marchal, A. (2020). Deep Learning for Asset Bubbles Detection. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3531154>
- Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *The Review of Financial Studies*, 9(1), 69–107.

- Bayer, C., Horvath, B., Muguruza, A., Stemper, B., & Tomas, M. (2019). *On deep calibration of (rough) stochastic volatility models* (arXiv:1908.08806). arXiv. <http://arxiv.org/abs/1908.08806>
- Bayer, C., & Stemper, B. (2018). *Deep calibration of rough stochastic volatility models* (arXiv:1810.03399). arXiv. <http://arxiv.org/abs/1810.03399>
- Ben Osman, M., Galariotis, E., Guesmi, K., Hamdi, H., & Naoui, K. (2024). Are markets sentiment driving the price bubbles in the virtual? *International Review of Economics & Finance*, 89, 272–285. <https://doi.org/10.1016/j.iref.2023.10.041>
- Bergstra, J., & Bengio, Y. (2012). *Random Search for Hyper-Parameter Optimization*. 10, 281–305.
- Biagini, F., Föllmer, H., & Nedelcu, S. (2014). Shifting martingale measures and the birth of a bubble as a submartingale. *Finance and Stochastics*, 18(2), 297–326. <https://doi.org/10.1007/s00780-013-0221-8>
- Biagini, F., Gonon, L., Mazzon, A., & Meyer-Brandis, T. (2024). Detecting asset price bubbles using deep learning. *Mathematical Finance*, n/a(n/a). <https://doi.org/10.1111/mafi.12443>
- Black, F., & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81(3), 637–654. <https://doi.org/10.1086/260062>
- Bonetto, R., & Latzko, V. (2020). Chapter 8—Machine learning. In F. H. P. Fitzek, F. Granelli, & P. Seeling (Eds.), *Computing in Communication Networks* (pp. 135–167). Academic Press. <https://doi.org/10.1016/B978-0-12-820488-7.00021-9>
- Brée, D. S., & Joseph, N. L. (2013). Testing for financial crashes using the Log Periodic Power Law model. *International Review of Financial Analysis*, 30, 287–297. <https://doi.org/10.1016/j.irfa.2013.05.005>
- Broadie, M., & Kaya, Ö. (2006). Exact Simulation of Stochastic Volatility and Other Affine Jump Diffusion Processes. *Operations Research*, 54(2), 217–231. <https://doi.org/10.1287/opre.1050.0247>
- Büchel, P., Kratochwil, M., Nagl, M., & Rösch, D. (2022). Deep calibration of financial models: Turning theory into practice. *Review of Derivatives Research*, 25(2), 109–136. <https://doi.org/10.1007/s11147-021-09183-7>

- Cao, J., Chen, J., & Hull, J. (2020). A neural network approach to understanding implied volatility movements. *Quantitative Finance*, 20(9), 1405–1413. <https://doi.org/10.1080/14697688.2020.1750679>
- Carr, P., & Madan, D. (1999). Option valuation using the fast Fourier transform. *The Journal of Computational Finance*, 2(4), 61–73. <https://doi.org/10.21314/JCF.1999.043>
- Chaim, P., & Laurini, M. P. (2019). Is Bitcoin a bubble? *Physica A: Statistical Mechanics and Its Applications*, 517, 222–232. <https://doi.org/10.1016/j.physa.2018.11.031>
- Choi, S. H., & Jarrow, R. (2024). Is the Current Bull Market A Bubble? An Empirical Investigation. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4789419>
- Choi, S. H., & Jarrow, R. A. (2022). APPLYING THE LOCAL MARTINGALE THEORY OF BUBBLES TO CRYPTOCURRENCIES. *International Journal of Theoretical and Applied Finance*, 25(03), 2250013. <https://doi.org/10.1142/S0219024922500133>
- Christensen, J., & Bastien, C. (2016). Chapter | three—Introduction to General Optimization Principles and Methods. In J. Christensen & C. Bastien (Eds.), *Nonlinear Optimization of Vehicle Safety Structures* (pp. 107–168). Butterworth-Heinemann. <https://doi.org/10.1016/B978-0-12-417297-5.00003-1>
- Clevert, D.-A., Unterthiner, T., & Hochreiter, S. (2016). *Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)* (arXiv:1511.07289). arXiv. <http://arxiv.org/abs/1511.07289>
- Cont, R. (2010). Model Calibration. In R. Cont (Ed.), *Encyclopedia of Quantitative Finance* (1st ed.). Wiley. <https://doi.org/10.1002/9780470061602.eqf08002>
- Cont, R., & da Fonseca, J. (2002). Dynamics of implied volatility surfaces. *Quantitative Finance*, 2(1), 45–60. <https://doi.org/10.1088/1469-7688/2/1/304>
- Cont, R., Fonseca, J. da, & Durrleman, V. (2002). Stochastic Models of Implied Volatility Surfaces. *Economic Notes*, 31(2), 361–377. <https://doi.org/10.1111/1468-0300.00090>
- Corbet, S., Lucey, B., & Yarovaya, L. (2018). Datestamping the Bitcoin and Ethereum bubbles. *Finance Research Letters*, 26, 81–88. <https://doi.org/10.1016/j.frl.2017.12.006>
- Cox, A. M., & Hobson, D. G. (2005). Local martingales, bubbles and option prices. *Finance and Stochastics*, 9(4), 477–492.

- Cretarola, A., Figà-Talamanca, G., & Patacca, M. (2018). A Continuous Time Model for Bitcoin Price Dynamics. In M. Corazza, M. Durbán, A. Grané, C. Perna, & M. Sibillo (Eds.), *Mathematical and Statistical Methods for Actuarial Sciences and Finance* (pp. 273–277). Springer International Publishing. [https://doi.org/10.1007/978-3-319-89824-7\\_49](https://doi.org/10.1007/978-3-319-89824-7_49)
- Cretarola, A., Figà-Talamanca, G., & Patacca, M. (2020). Market attention and Bitcoin price modeling: Theory, estimation and option pricing. *Decisions in Economics and Finance*, 43(1), 187–228. <https://doi.org/10.1007/s10203-019-00262-x>
- Culkin, R., & Das, S. (2017). Machine Learning in Finance: The Case of Deep Learning for Option Pricing. *Journal of Investment Management*, 15(4), 92–100.
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Math. Control Signals and Systems*, 2, 303–314. <https://doi.org/10.1007/BF02551274>
- Darken, C., Chang, J., & Moody, J. (1992). Learning rate schedules for faster stochastic gradient search. *Neural Networks for Signal Processing II Proceedings of the 1992 IEEE Workshop*, 3–12. <https://doi.org/10.1109/NNSP.1992.253713>
- Dauphin, Y., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., & Bengio, Y. (2014). *Identifying and attacking the saddle point problem in high-dimensional non-convex optimization* (arXiv:1406.2572). arXiv. <http://arxiv.org/abs/1406.2572>
- De Spiegeleer, J., Madan, D. B., Reyners, S., & Schoutens, W. (2018). Machine learning for quantitative finance: Fast derivative pricing, hedging and fitting. *Quantitative Finance*, 18(10), 1635–1643.
- Delbaen, F., & Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300(1), 463–520. <https://doi.org/10.1007/BF01450498>
- Delbaen, F., & Shirakawa, H. (2002). No Arbitrage Condition for Positive Diffusion Price Processes. *Asia-Pacific Financial Markets*, 9(3), 159–168. <https://doi.org/10.1023/A:1024173029378>
- Demmler, M., & Fernández, A. O. (2024). Explosive behavior in historic NASDAQ market prices. *The North American Journal of Economics and Finance*, 71, 102095. <https://doi.org/10.1016/j.najef.2024.102095>
- Diba, B. T., & Grossman, H. I. (1988a). Explosive rational bubbles in stock prices? *The American Economic Review*, 78(3), 520–530.

- Diba, B. T., & Grossman, H. I. (1988b). The theory of rational bubbles in stock prices. *The Economic Journal*, 98(392), 746–754.
- Dimitroff, G., Röder, D., & Fries, C. P. (2018). Volatility Model Calibration With Convolutional Neural Networks. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3252432>
- Driffill, J., & Sola, M. (1998). Intrinsic bubbles and regime-switching. *Journal of Monetary Economics*, 42(2), 357–373.
- Duchi, J., Hazan, E., & Singer, Y. (2011). Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *The Journal of Machine Learning Research*, 12, 2121–2159.
- Evans, G. W. (1991). Pitfalls in Testing for Explosive Bubbles in Asset Prices. *The American Economic Review*, 81(4), 922–930. JSTOR.
- Fang, F., & Oosterlee, C. W. (2009). A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions. *SIAM Journal on Scientific Computing*, 31(2), 826–848. <https://doi.org/10.1137/080718061>
- Filimonov, V., & Sornette, D. (2013). A stable and robust calibration scheme of the log-periodic power law model. *Physica A: Statistical Mechanics and Its Applications*, 392(17), 3698–3707. <https://doi.org/10.1016/j.physa.2013.04.012>
- Froot, K. A., & Obstfeld, M. (1991). Intrinsic bubbles: The case of stock prices (National Bureau of Economic Research). *Forthcoming in American Economic Review*.
- Funahashi, H. (2023). SABR equipped with AI wings. *Quantitative Finance*, 23(2), 229–249. <https://doi.org/10.1080/14697688.2022.2150561>
- Fusari, N., Jarrow, R., & Lamichhane, S. (2024). Testing for Asset Price Bubbles using Options Data. *Journal of Business & Economic Statistics*, 1–46. <https://doi.org/10.1080/07350015.2024.2429470>
- Gatheral, J. (2012). Volatility Surface Asymptotics. In *The Volatility Surface* (pp. 87–100). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781119202073.ch7>
- Glorot, X., Bordes, A., & Bengio, Y. (2011). Deep Sparse Rectifier Neural Networks. In G. Gordon, D. Dunson, & M. Dudík (Eds.), *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics* (Vol. 15, pp. 315–323). PMLR. <https://proceedings.mlr.press/v15/glorot11a.html>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT press.



- Guidolin, M., & Wang, K. (2023). The empirical performance of option implied volatility surface-driven optimal portfolios. *Physica A: Statistical Mechanics and Its Applications*, 618, 128496. <https://doi.org/10.1016/j.physa.2023.128496>
- Gürkaynak, R. S. (2008). ECONOMETRIC TESTS OF ASSET PRICE BUBBLES: TAKING STOCK. *Journal of Economic Surveys*, 22(1), 166–186. <https://doi.org/10.1111/j.1467-6419.2007.00530.x>
- Hagan, M. T., Demuth, H. B., Beale, M. H., & De Jesús, O. (2014). *Neural Network Design*.
- Hagan, P. S., Kumar, D., Lesniewski, A., & Woodward, D. (2002). Managing Smile Risk. *Wilmott Magazine*, 1, 84–108.
- Hagan, P. S., & West, G. (2006). Interpolation Methods for Curve Construction. *Applied Mathematical Finance*, 13(2), 89–129. <https://doi.org/10.1080/13504860500396032>
- Han, C.-H., & Kuo, C.-L. (2017). Monte Carlo calibration to implied volatility surface under volatility models. *Japan Journal of Industrial and Applied Mathematics*, 34(3), 763–778. <https://doi.org/10.1007/s13160-017-0270-z>
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). Linear Methods for Regression. In T. Hastie, R. Tibshirani, & J. Friedman (Eds.), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (pp. 43–99). Springer. [https://doi.org/10.1007/978-0-387-84858-7\\_3](https://doi.org/10.1007/978-0-387-84858-7_3)
- Hernandez, A. (2016). Model Calibration with Neural Networks. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.2812140>
- Heston, S. L. (1993). 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.
- Heston, S. L., Loewenstein, M., & Willard, G. A. (2007). Options and Bubbles. *Review of Financial Studies*, 20(2), 359–390. <https://doi.org/10.1093/rfs/hhl005>
- Homm, U., & Breitung, J. (2012). Testing for speculative bubbles in stock markets: A comparison of alternative methods. *Journal of Financial Econometrics*, 10(1), 198–231.
- Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2), 251–257. [https://doi.org/10.1016/0893-6080\(91\)90009-T](https://doi.org/10.1016/0893-6080(91)90009-T)

- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5), 359–366. [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8)
- Horvath, B., Muguruza, A., & Tomas, M. (2021). Deep learning volatility: A deep neural network perspective on pricing and calibration in (rough) volatility models. *Quantitative Finance*, 21(1), 11–27. <https://doi.org/10.1080/14697688.2020.1817974>
- Hoshisashi, K., Phelan, C. E., & Barucca, P. (2023). *No-Arbitrage Deep Calibration for Volatility Smile and Skewness* (arXiv:2310.16703). arXiv. <http://arxiv.org/abs/2310.16703>
- Hull, J. C. (2014). FUNDAMENTALS OF FUTURES AND OPTIONS MARKETS (8th Edition). *New York, NY: Pearson*.
- Hutchinson, J. M., Lo, A. W., & Poggio, T. (1994). A Nonparametric Approach to Pricing and Hedging Derivative Securities Via Learning Networks. *The Journal of Finance*, 49(3), 851–889. <https://doi.org/10.1111/j.1540-6261.1994.tb00081.x>
- Itkin, A. (2019). *Deep learning calibration of option pricing models: Some pitfalls and solutions* (arXiv:1906.03507). arXiv. <http://arxiv.org/abs/1906.03507>
- Ivaşcu, C.-F. (2021). Option pricing using Machine Learning. *Expert Systems With Applications*, 7.
- Izadkhah, H. (2022). Chapter 4—Basic structure of neural networks. In H. Izadkhah (Ed.), *Deep Learning in Bioinformatics* (pp. 67–93). Academic Press. <https://doi.org/10.1016/B978-0-12-823822-6.00011-1>
- Jang, H., & Lee, J. (2019). Machine learning versus econometric jump models in predictability and domain adaptability of index options. *Physica A: Statistical Mechanics and Its Applications*, 513, 74–86. <https://doi.org/10.1016/j.physa.2018.08.091>
- Jarrow, R. (2016). Testing for asset price bubbles: Three new approaches. *Quantitative Finance Letters*, 4(1), 4–9. <https://doi.org/10.1080/21649502.2015.1165838>
- Jarrow, R. A. (2015). Asset Price Bubbles. *Annual Review of Financial Economics*, 7(1), 201–218. <https://doi.org/10.1146/annurev-financial-030215-035912>
- Jarrow, R. A. (2021). *Continuous-time asset pricing theory*. Springer.

- Jarrow, R. A., Jin, X., & Madan, D. B. (1999). The Second Fundamental Theorem of Asset Pricing. *Mathematical Finance*, 9(3), 255–273. <https://doi.org/10.1111/1467-9965.00070>
- Jarrow, R. A., & Kwok, S. S. (2021). Inferring financial bubbles from option data. *Journal of Applied Econometrics*, 36(7), 1013–1046. <https://doi.org/10.1002/jae.2862>
- Jarrow, R. A., & Kwok, S. S. (2024). A study on asset price bubble dynamics: Explosive trend or quadratic variation? *Quantitative Finance*, 24(5), 613–626. <https://doi.org/10.1080/14697688.2024.2342897>
- Jarrow, R. A., Protter, P., & Shimbo, K. (2007). Asset Price Bubbles in Complete Markets. In M. C. Fu, R. A. Jarrow, J.-Y. J. Yen, & R. J. Elliott (Eds.), *Advances in Mathematical Finance* (pp. 97–121). Birkhäuser Boston. [https://doi.org/10.1007/978-0-8176-4545-8\\_7](https://doi.org/10.1007/978-0-8176-4545-8_7)
- Jarrow, R. A., Protter, P., & Shimbo, K. (2010). ASSET PRICE BUBBLES IN INCOMPLETE MARKETS\*. *Mathematical Finance*, 20(2), 145–185. <https://doi.org/10.1111/j.1467-9965.2010.00394.x>
- Jarrow, R., Kchia, Y., & Protter, P. (2011). How to Detect an Asset Bubble. *SIAM J. Financial Math.*, 2, 839–865. <https://doi.org/10.2139/ssrn.1621728>
- Jindal, S. K., Banerjee, S., Patra, R., & Paul, A. (2022). Deep learning-based brain malignant neoplasm classification using MRI image segmentation assisted by bias field correction and histogram equalization. In *Brain Tumor MRI Image Segmentation Using Deep Learning Techniques* (pp. 135–161). Elsevier. <https://doi.org/10.1016/B978-0-323-91171-9.00008-9>
- Johansen, A., Ledoit, O., & Sornette, D. (2000). Crashes as critical points. *International Journal of Theoretical and Applied Finance*, 3(02), 219–255.
- Johansen, A., Sornette, D., & Olivier, L. (1999). Predicting Financial Crashes Using Discrete Scale Invariance. *Journal of Risk*, 1.
- Karatas, T., Oskoui, A., & Hirsa, A. (2019). *Supervised Deep Neural Networks (DNNs) for Pricing/Calibration of Vanilla/Exotic Options Under Various Different Processes* (arXiv:1902.05810). arXiv. <http://arxiv.org/abs/1902.05810>
- Karimov, A. (2017). *Identifying Stock Market Bubbles*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-65009-8>

- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., & Tang, P. T. P. (2017). *On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima* (arXiv:1609.04836). arXiv. <http://arxiv.org/abs/1609.04836>
- Khan, K., Su, C.-W., Umar, M., & Yue, X.-G. (2021). Do crude oil price bubbles occur? *Resources Policy*, 71, 101936. <https://doi.org/10.1016/j.resourpol.2020.101936>
- Kingma, D. P., & Ba, J. (2017). *Adam: A Method for Stochastic Optimization*.
- Kotani, S. (2006). On a Condition that One-Dimensional Diffusion Processes are Martingales. In M. Émery & M. Yor (Eds.), *In Memoriam Paul-André Meyer* (Vol. 1874, pp. 149–156). Springer Berlin Heidelberg. [https://doi.org/10.1007/978-3-540-35513-7\\_12](https://doi.org/10.1007/978-3-540-35513-7_12)
- Kwong, R., & Wong, H. (2022). Value-at-risk in the presence of asset price bubbles. *Journal of Applied Economics*, 25(1), 361–384. <https://doi.org/10.1080/15140326.2021.1927441>
- Laurini, M. P., & Chaim, P. (2021). Brazilian stock market bubble in the 2010s. *Sn Business & Economics*, 1(1), 8. <https://doi.org/10.1007/s43546-020-00005-w>
- LaValle, S. M., Branicky, M. S., & Lindemann, S. R. (2004). On the Relationship between Classical Grid Search and Probabilistic Roadmaps. *The International Journal of Robotics Research*, 23(7–8), 673–692. <https://doi.org/10.1177/0278364904045481>
- Lavine, B. K., & Blank, T. R. (2009). Feed-Forward Neural Networks. *Chemical and Biochemical Data Analysis*, 3, 571–586.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Lek, S., & Park, Y. S. (2008). Multilayer Perceptron. In S. E. Jørgensen & B. D. Fath (Eds.), *Encyclopedia of Ecology* (pp. 2455–2462). Academic Press. <https://doi.org/10.1016/B978-008045405-4.00162-2>
- LeRoy, S. F., & Porter, R. D. (1981). The present-value relation: Tests based on implied variance bounds. *Econometrica: Journal of the Econometric Society*, 555–574.
- Li, C.-Q., & Yang, W. (2023). 2—Essential reliability methods. In C.-Q. Li & W. Yang (Eds.), *Time-Dependent Reliability Theory and Its Applications* (pp. 51–119). Woodhead Publishing. <https://doi.org/10.1016/B978-0-323-85882-3.00006-4>

- Liashchynskiy, P., & Liashchynskiy, P. (2019). *Grid Search, Random Search, Genetic Algorithm: A Big Comparison for NAS* (arXiv:1912.06059). arXiv. <http://arxiv.org/abs/1912.06059>
- Lions, P.-L., & Musiela, M. (2007). Correlations and bounds for stochastic volatility models. *Annales de l'Institut Henri Poincaré C, Analyse Non Linéaire*, 24(1), 1–16. <https://doi.org/10.1016/j.anihpc.2005.05.007>
- Liu, S., Borovykh, A., Grzelak, L. A., & Oosterlee, C. W. (2019). A neural network-based framework for financial model calibration. *Journal of Mathematics in Industry*, 9(1), 9. <https://doi.org/10.1186/s13362-019-0066-7>
- Liu, S., Leitao, Á., Borovykh, A., & Oosterlee, C. W. (2021). On a Neural Network to Extract Implied Information from American Options. *Applied Mathematical Finance*, 28(5), 449–475. <https://doi.org/10.1080/1350486X.2022.2097099>
- Liu, S., Oosterlee, C. W., & Bohte, S. M. (2019). Pricing options and computing implied volatilities using neural networks. *Risks*, 7(1), 16.
- Loewenstein, M., & Willard, G. A. (2000a). Local martingales, arbitrage, and viability. *Economic Theory*, 16(1), 135–161. <https://doi.org/10.1007/s001990050330>
- Loewenstein, M., & Willard, G. A. (2000b). Rational Equilibrium Asset-Pricing Bubbles in Continuous Trading Models. *Journal of Economic Theory*, 91(1), 17–58. <https://doi.org/10.1006/jeth.1999.2589>
- Lorenzo, P. R., Nalepa, J., Kawulok, M., Ramos, L. S., & Pastor, J. R. (2017). Particle swarm optimization for hyper-parameter selection in deep neural networks. *Proceedings of the Genetic and Evolutionary Computation Conference*, 481–488. <https://doi.org/10.1145/3071178.3071208>
- Mackay, C. (1841). *Extraordinary popular delusions and the madness of crowds*. Simon and Schuster (2012).
- Malliaris, M., & Salchenberger, L. (1993). Beating the best: A neural network challenges the Black-Scholes formula. *Proceedings of 9th IEEE Conference on Artificial Intelligence for Applications*, 445–449. <https://doi.org/10.1109/CAIA.1993.366633>
- Malliaris, M., & Salchenberger, L. (1996). Using neural networks to forecast the S&P 100 implied volatility. *Neurocomputing*, 10(2), 183–195. [https://doi.org/10.1016/0925-2312\(95\)00019-4](https://doi.org/10.1016/0925-2312(95)00019-4)

- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4), 115–133. <https://doi.org/10.1007/BF02478259>
- McGhee, W. A. (2018). *An Artificial Neural Network Representation of the SABR Stochastic Volatility Model*.
- Mckay, M., Beckman, R., & Conover, W. (1979). A Comparison of Three Methods for Selecting Vales of Input Variables in the Analysis of Output From a Computer Code. *Technometrics*, 21, 239–245. <https://doi.org/10.1080/00401706.1979.10489755>
- Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 141–183.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1–2), 125–144.
- Mijatović, A., & Urusov, M. (2012). On the martingale property of certain local martingales. *Probability Theory and Related Fields*, 152(1–2), 1–30. <https://doi.org/10.1007/s00440-010-0314-7>
- Milmo, D. (2024, April). *Meta value falls \$190bn as investors react to plan to increase spending on AI*. Guardian News and Media. <https://www.theguardian.com/technology/2024/apr/25/meta-value-falls-190bn-as-investors-react-to-plan-to-increase-spending-on-ai>
- Nguyen, Q. N., & Waters, G. A. (2022). Detecting periodically collapsing bubbles in the S&P 500. *The Quarterly Review of Economics and Finance*, 83, 83–91. <https://doi.org/10.1016/j.qref.2021.11.005>
- Obayashi, Y., Protter, P., & Yang, S. (2017). The lifetime of a financial bubble. *Mathematics and Financial Economics*, 11(1), 45–62. <https://doi.org/10.1007/s11579-016-0170-z>
- Park, S., & Yang, J.-S. (2024). Machine learning models based on bubble analysis for Bitcoin market crash prediction. *Engineering Applications of Artificial Intelligence*, 135, 108857. <https://doi.org/10.1016/j.engappai.2024.108857>
- Perifanis, T. (2019). Detecting west Texas intermediate (WTI) prices’ bubble periods. *Energies*, 12(14), 2649.
- Phillips, P. C., Shi, S., & Yu, J. (2015). Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. *International Economic Review*, 56(4), 1043–1078.

- Phillips, Wu, Y., & Yu, J. (2011). EXPLOSIVE BEHAVIOR IN THE 1990s NASDAQ: WHEN DID EXUBERANCE ESCALATE ASSET VALUES?\*. *International Economic Review*, 52(1), 201–226. <https://doi.org/10.1111/j.1468-2354.2010.00625.x>
- Piironen, P., Roininen, L., Schoden, T., & Simon, M. (2018). Asset Price Bubbles: An Option-based Indicator. *arXiv:1805.07403 [Math, q-Fin]*. <http://arxiv.org/abs/1805.07403>
- Poggio, T., & Liao, Q. (2018). Theory I: Deep networks and the curse of dimensionality. *Bulletin of the Polish Academy of Sciences Technical Sciences*, 761–773. <https://doi.org/10.24425/bpas.2018.125924>
- Pothina, H., & Nagaraja, K. V. (2023). Artificial Neural Network and Math Behind It. In Y.-D. Zhang, T. Senjyu, C. So-In, & A. Joshi (Eds.), *Smart Trends in Computing and Communications* (Vol. 396, pp. 205–221). Springer Nature Singapore. [https://doi.org/10.1007/978-981-16-9967-2\\_21](https://doi.org/10.1007/978-981-16-9967-2_21)
- Poulsen, R., Schenk-Hoppé, K. R., & Ewald, C.-O. (2009). Risk minimization in stochastic volatility models: Model risk and empirical performance. *Quantitative Finance*, 9(6), 693–704. <https://doi.org/10.1080/14697680902852738>
- Protter, P. (2013). A Mathematical Theory of Financial Bubbles. In *Paris-Princeton Lectures on Mathematical Finance 2013: Editors: Vicky Henderson, Ronnie Sircar* (pp. 1–108). Springer International Publishing. [https://doi.org/10.1007/978-3-319-00413-6\\_1](https://doi.org/10.1007/978-3-319-00413-6_1)
- Protter, P. (2016). Mathematical models of bubbles. *Quantitative Finance Letters*, 4(1), 10–13. <https://doi.org/10.1080/21649502.2015.1165863>
- Protter, P. E. (2005). *Stochastic Integration and Differential Equations* (Vol. 21). Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-662-10061-5>
- Quinn, W., & Turner, J. D. (2020). *Boom and Bust: A Global History of Financial Bubbles* (1st ed.). Cambridge University Press. <https://doi.org/10.1017/9781108367677>
- Quinn, W., & Turner, J. D. (2023). Bubbles in history. *Business History*, 65(4), 636–655. <https://doi.org/10.1080/00076791.2020.1844668>
- Ravikumar, N., Zakeri, A., Xia, Y., & Frangi, A. F. (2024). Chapter 16—Deep learning fundamentals. In A. F. Frangi, J. L. Prince, & M. Sonka (Eds.), *Medical Image*

- Analysis* (pp. 415–450). Academic Press. <https://doi.org/10.1016/B978-0-12-813657-7.00041-8>
- Refaeilzadeh, P., Tang, L., & Liu, H. (2009). Cross-Validation. In L. LIU & M. T. ÖZSU (Eds.), *Encyclopedia of Database Systems* (pp. 532–538). Springer US. [https://doi.org/10.1007/978-0-387-39940-9\\_565](https://doi.org/10.1007/978-0-387-39940-9_565)
- Roeder, D., & Dimitroff, G. (2020). *Volatility model calibration with neural networks a comparison between direct and indirect methods* (arXiv:2007.03494). arXiv. <http://arxiv.org/abs/2007.03494>
- Romo, J. M. (2014). Dynamics of the implied volatility surface. Theory and empirical evidence. *Quantitative Finance*, 14(10), 1829–1837. <https://doi.org/10.1080/14697688.2012.686668>
- Ruder, S. (2017). *An overview of gradient descent optimization algorithms* (arXiv:1609.04747). arXiv. <http://arxiv.org/abs/1609.04747>
- Ruf, J., & Wang, W. (2020). *Neural networks for option pricing and hedging: A literature review* (arXiv:1911.05620). arXiv. <http://arxiv.org/abs/1911.05620>
- Scherbina, A., & Schlusche, B. (2014). Asset price bubbles: A survey. *Quantitative Finance*, 14(4), 589–604. <https://doi.org/10.1080/14697688.2012.755266>
- Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71(June), 421–436.
- Shobha, G., & Rangaswamy, S. (2018). Machine Learning. In *Handbook of Statistics* (Vol. 38, pp. 197–228). Elsevier. <https://doi.org/10.1016/bs.host.2018.07.004>
- Shu, M., & Song, R. (2024). Detection of financial bubbles using a log-periodic power law singularity (LPPLS) model. *WIREs Computational Statistics*, 16(2), e1649. <https://doi.org/10.1002/wics.1649>
- Shu, M., Song, R., & Zhu, W. (2021). The ‘COVID’Crash of the 2020 US Stock Market. *The North American Journal of Economics and Finance*, 101497.
- Sin, C. A. (1998). Complications with stochastic volatility models. *Advances in Applied Probability*, 30(1), 256–268. <https://doi.org/10.1239/aap/1035228003>
- Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical Bayesian Optimization of Machine Learning Algorithms. In F. Pereira, C. J. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems* (Vol. 25). Curran Associates, Inc.



[https://proceedings.neurips.cc/paper\\_files/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf)

- Song, R., Shu, M., & Zhu, W. (2022). The 2020 global stock market crash: Endogenous or exogenous? *Physica A: Statistical Mechanics and Its Applications*, 585, 126425. <https://doi.org/10.1016/j.physa.2021.126425>
- Sornette, D., Demos, G., Zhang, Q., Cauwels, P., Filimonov, V., & Zhang, Q. (2015). Real-time prediction and post-mortem analysis of the Shanghai 2015 stock market bubble and crash. *Swiss Finance Institute Research Paper*, 15–31.
- Stahl, P., & Blauth, J. (2024). Martingale defects in the volatility surface and bubble conditions in the underlying. *Review of Derivatives Research*. <https://doi.org/10.1007/s11147-023-09200-x>
- Stone, H. (2020). Calibrating rough volatility models: A convolutional neural network approach. *Quantitative Finance*, 20(3), 379–392. <https://doi.org/10.1080/14697688.2019.1654126>
- Storn, R., & Price, K. (1997). Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *Journal of Global Optimisation*, 11, 341–359.
- Su, C.-W., Li, Z.-Z., Chang, H.-L., & Lobont, O.-R. (2017). When Will Occur the Crude Oil Bubbles? *Energy Policy*, 102, 1–6. <https://doi.org/10.1016/j.enpol.2016.12.006>
- Subah, F. Z., & Deb, K. (2023). Chapter 13—A comprehensive study on atlas-based classification of autism spectrum disorder using functional connectivity features from resting-state functional magnetic resonance imaging. In A. S. El-Baz & J. S. Suri (Eds.), *Neural Engineering Techniques for Autism Spectrum Disorder* (pp. 269–296). Academic Press. <https://doi.org/10.1016/B978-0-12-824421-0.00021-7>
- Telgarsky, M. (2016). *Benefits of depth in neural networks*. 1517–1539.
- Tharsanee, R. M., Soundariya, R. S., Kumar, A. S., Karthiga, M., & Sountharajan, S. (2021). 7—Deep convolutional neural network-based image classification for COVID-19 diagnosis. In U. Kose, D. Gupta, V. H. C. de Albuquerque, & A. Khanna (Eds.), *Data Science for COVID-19* (pp. 117–145). Academic Press. <https://doi.org/10.1016/B978-0-12-824536-1.00012-5>

- Tieleman, T., & Hinton, G. (2012). Lecture 6.5-rmsprop: Divide the Gradient by a Running Average of Its Recent Magnitude. *COURSERA: Neural Networks for Machine Learning*, 4, 26–31.
- Tiwari, K., & Young Chong, N. (2020). 5—Preliminaries: A primer. In K. Tiwari & N. Young Chong (Eds.), *Multi-robot Exploration for Environmental Monitoring* (pp. 41–52). Academic Press. <https://doi.org/10.1016/B978-0-12-817607-8.00018-6>
- Tomasini, U., & Wyart, M. (2024). *How Deep Networks Learn Sparse and Hierarchical Data: The Sparse Random Hierarchy Model* (arXiv:2404.10727). arXiv. <http://arxiv.org/abs/2404.10727>
- Ulrich, M., & Walther, S. (2020). Option-implied information: What’s the vol surface got to do with it? *Review of Derivatives Research*, 23(3), 323–355. <https://doi.org/10.1007/s11147-020-09166-0>
- West, K. D. (1987). A specification test for speculative bubbles. *The Quarterly Journal of Economics*, 102(3), 553–580.
- Wojtowysch, S., & Weinan, E. (2020). Can Shallow Neural Networks Beat the Curse of Dimensionality? A Mean Field Training Perspective. *IEEE Transactions on Artificial Intelligence*, 1(2), 121–129. <https://doi.org/10.1109/TAI.2021.3051357>
- Wshah, S., Shadid, R., Wu, Y., Matar, M., Xu, B., Wu, W., Lin, L., & Elmoudi, R. (2020). Deep Learning for Model Parameter Calibration in Power Systems. *2020 IEEE International Conference on Power Systems Technology (POWERCON)*, 1–6. <https://doi.org/10.1109/POWERCON48463.2020.9230531>
- Yang, L., & Shami, A. (2020). On hyperparameter optimization of machine learning algorithms: Theory and practice. *Neurocomputing*, 415, 295–316. <https://doi.org/10.1016/j.neucom.2020.07.061>
- Zagoruyko, S., & Komodakis, N. (2017). *Wide Residual Networks* (arXiv:1605.07146). arXiv. <http://arxiv.org/abs/1605.07146>
- Zhang, Y.-J., & Yao, T. (2016). Interpreting the movement of oil prices: Driven by fundamentals or bubbles? *Economic Modelling*, 55, 226–240.
- Zhou, Q., Shen, Z., Yong, B., & Zhi, P. (2022). Chapter 5—Introduction of machine learning and neural networks. In Q. Zhou, Z. Shen, B. Yong, R. Zhao, & P. Zhi (Eds.), *Theories and Practices of Self-Driving Vehicles* (pp. 147–175). Elsevier. <https://doi.org/10.1016/B978-0-323-99448-4.00005-9>



# Appendix

The Appendix is dedicated to providing supplementary material, supporting the findings of this research. Section *A* provides detailed information of the training and testing procedures applied to the deep calibration framework, such that the optimal architecture is determined for bubble detection.

## A. Deep Calibration Framework: Training and Testing

Previously in 4.3, superior calibration performances from the GSVJD model, in comparison to other stochastic processes were revealed. However, the task was associated with a large computation burden, which further exemplifies over the sample period. During this operation, only options corresponding to the most liquid maturities were considered. Despite such attempts towards improving efficiency, on average, daily calibration required 3.56 hours, creating a major bottleneck for employing the three-step approach. Additionally, quality of the method is hampered, as daily calibrations are conducted over a single volatility smile, rather than the entire available surface. Therefore, it is crucial to improve the speed of GSVJD parameter estimation, without sacrificing accuracy.

The two-step calibration framework comprises of an MLP neural network. During the first stage, *forward pass*, the network is trained to learn dynamics of the GSVJD model. The training phase consists of a random search over 720 hyperparameter combinations, with each undergoing a 3-fold validation. Next, top performing architectures are further trained over a much larger dataset and carried to the *backward pass*. The latter step utilises the trained hidden layers, and calibrates GSVJD parameters from market observations, with the assistance of the Differential Evolution optimizer. Parameters from the best architecture are selected for pricing call options for bubble detection. This section proceeds with examining the training and testing phase of the *forward pass*, followed by analyzing calibrations to market data during the *backward pass*.

## A.1 Forward Pass

In the *forward pass* phase, the neural network is trained to learn the dynamics of the GSVJD model. The goal of the extensive random search, across the vast range of hyperparametric combinations, is to seek the best performing architectures, for further training. First, each network architecture is trained on a smaller *Random Search Dataset*. Next, the top performing hyperparametric combinations are selected to be trained on a significantly larger *Optimal Dataset*. Once trained, each of the optimal architectures are carried to that *backward pass* phase, to test calibration performances on S&P 500 data. Both training datasets are synthetically generated and therefore is important to understand the reasoning and method behind their creation.

### A.1.1 Synthetic Training Data

The implementation of neural networks for option pricing are plagued with the following interlinked concerns; unavailability of sufficient historical data, and a lack of economic interpretation (Hutchinson et al., 1994; Bayer et al., 2019; Ivaşcu, 2021). The latter refers to the lack of generalisation during output estimation, catering to the notorious ‘black-box’ nature of machine learning models. A major source of such problems is identified as the training dataset quality. These concerns are widely observed with the one-step calibration

method<sup>53</sup>. The approach trains networks on stochastic model parameters corresponding to the historical market observations. Such information is not readily available; hence, the practitioner is required to conduct calibrations off-line for creating the training dataset. An additional burden is created during this computationally intensive stage. Furthermore, a network trained on such a dataset is vulnerable to overfitting, and weak performances during unseen market scenarios (Hernandez, 2016; Dimitroff et al., 2018). Therefore, one must frequently indulge in the computationally cumbersome task of retraining the network to accommodate the latest market trend. This reveals a loss of generalization during the estimation of output, given the absence of economic interpretation<sup>54</sup>.

Such issues, even though on a lesser extent, could arise when training the two-step calibration approach on historical data. Therefore, creating a synthetic training dataset is favourable. This ensures that the inclusion of various market scenarios is controlled by the practitioner deciding the distribution range of input and output variables. In addition, there are no size restrictions, making it extremely beneficial when catering to instruments lacking sufficient historical data. Given flexibility in construction, inclusion of market scenarios are only restricted by the practitioner's imagination. Training on such datasets is robust, reducing the chances of overfitting, and most crucially making the need for re-training, redundant (Bayer et al., 2019). The term 'synthetic' arises from the random generation of inputs, which resemble their behaviour under various market conditions. The selected asset pricing model is applied to determine implied volatilities corresponding to each combination of inputs. This ensures the availability of sufficient amounts of data, across various possible market scenarios, such that training is robust.

The generation of the synthetic training dataset, as illustrated in *Figure A.1*, is a fusion of utilising randomly generated parameters and option contracts related input data, to estimate option prices and subsequently implied volatilities. The contract characteristics data,  $\Omega = \{mon, \tau\}$ , comprises of information on stock prices ( $S$ ), risk-free rates ( $r$ ), dividend yields


---

<sup>53</sup> These explanations are relevant and can be extended for to the general utilization of neural networks in financial derivative markets.

<sup>54</sup> This issue can also be attributed to the absence of real-time validation, which is dealt with by the two-step calibration approach.

( $q$ ) and option maturities ( $\tau$ ). Forward prices,  $F = Se^{(r-q)\tau}$ , are obtained from all constituents  $\Omega$ , to compute strike prices,  $\mathbb{K}$ . First, a respective lower and upper bound of 0.4, and 1.6, was set for moneyness,  $mon = (\mathbb{K}/F)$ . Next, from the range and given forward prices,  $\mathbb{K}$  values were generated at increments of \$25.00.

*Figure A.1: Creation of Synthetic Dataset.*



In the pointwise two-step approach, since moneyness and maturity are employed as inputs, the generation of  $\Omega$ , does not impact the robustness of the network. The strike price and expiry date of an option are fixed and specified as constituents of the contract and can be similar for options written on other underlying instruments. There might be concern with the generation of  $S$ , and subsequently, computing  $F$ , however this is eliminated by considering moneyness as an input. Moneyness describes the degree of extent to which an option is profitable and therefore is not particular to individual contracts. The GSVJD parameter values, are not calibrated but sampled at random. The components of  $\Omega$  merely represent contract details of options on a given day, or even over a time period. The motivation behind creating such a dataset is derived from wanting to train the neural network over various market regimes, representing unique ELMs, defined by each synthetic parametric combination.

For to each  $\Omega$ , 50 different parameter combinations,  $\theta = \{V_0, \bar{v}, \kappa, \sigma_v, \rho, p, \mu_y, \sigma_y, \lambda\}$ , abiding by Feller's condition ( $2\kappa v \geq \sigma_v^2$ ) were generated. The spawning of parameters requires a sampling technique with good space-filling properties (Liu et al., 2019). Latin Hypercube Sampling (LHS), introduced by McKay et al. (1979), is considered due to its ability to well represent the parameter space by simulating random multidimensional distributions. Additionally, the approach is independent to number of dimensions, and has a memory, providing the closest representation of realistic variability (Li & Yang, 2023). Once

the number of samples required are decided, LHS remembers the location of the taken point, to prevent repetition<sup>55</sup>, analogous to having a chessboard with  $N$  rooks that do not threaten each other.

Traditionally,  $\theta$  is calibrated to recognise the existing ELMM, and accordingly used for pricing options. By assigning each  $\Omega$  set, 50 unique parametric combinations, the neural network will be trained on a large number of market regimes, to prevent overfitting. It will be equipped with the understanding of accurately identifying the prevalent ELMM for a given day, during the *backward pass*. Since the neural network will be trained to learn the dynamics of the GSVJD model, the stochastic process is used to price put ( $P(\theta, \Omega)$ ) and call ( $C(\theta, \Omega)$ ) options, which will be used to compute implied volatilities. Given, the unavailability of a closed form solution, Monte Carlo simulations are relied upon, for pricing.

The lack of a closed form solution can be problematic, giving rise to two error sources, when training the network (Liu et al., 2019). First, errors arise from Monte Carlo simulations, implemented for pricing. As a solution, Horvath et al. (2021) resorted to the spot martingale control variate method for reducing noise. However, to remain consistent with the practice implemented during the benchmark model calibration, variance reduction or control variance techniques were not adopted. Instead, Monte Carlo simulations, consisting of 10,000 sample paths, with time steps of  $\frac{1}{365 \times 5}$  are used. A time-step equivalent to one-fifth of a day helps in significantly reducing the noise. Furthermore, options revealing prices less than \$0.01, in addition to those with implied volatilities greater (less) than 100% (1%), were dropped. The second error occurs during optimisation of the network in the training and calibrating phases. Regarding the former, presence of some error is appreciated, as it implies that overtraining has not taken place (Lek & Park, 2008; Lavine & Blank, 2009). For the latter, concerns over the error are relevant, which is precisely why the two-step calibration approach includes an optimizer, to ensure real-time validation. In context of bubble detection, it is equivalent to the real-time testing of a potential joint-hypothesis issue.

---

<sup>55</sup> In contrast, random sampling generates outputs without taking previous points under consideration (Li & Yang, 2023).



**Table A.1: Summary Statistics of the Random Search and Optimal Training Datasets.**

Note: Summary statistics of the Random Search (left) and Optimal Training (right) datasets.  $V(\theta, \Omega)/\mathbb{K}$  represents option (call and put) prices, scaled by  $\mathbb{K}$ , across a large range of  $mon = (\mathbb{K}/F)$  and  $\tau$  (days), for L.H.S. generated random  $\theta = \{V_0, \bar{v}, \kappa, \sigma_v, \rho, p, \mu_J, \sigma_J, \lambda_J\}$  vectors.

	Random Search Dataset					Optimal Training Dataset				
	Mean	SD	P25	P50	P75	Mean	SD	P25	P50	P75
$\kappa$	7.07	3.12	4.65	7.31	9.75	7.07	3.12	4.65	7.31	9.75
$\sigma_v$	1.00	0.66	0.43	0.90	1.49	1.00	0.66	0.43	0.90	1.49
$\bar{v}$	0.38	0.18	0.24	0.39	0.53	0.38	0.18	0.24	0.39	0.53
$V_0$	0.46	0.28	0.22	0.43	0.68	0.46	0.28	0.21	0.43	0.68
$\rho$	-0.01	0.58	-0.51	-0.02	0.50	-0.01	0.58	-0.51	-0.02	0.50
$p$	1.00	0.57	0.50	0.99	1.50	1.00	0.57	0.50	0.99	1.50
$\lambda_J$	1.47	1.29	0.43	1.05	2.22	1.47	1.29	0.43	1.05	2.21
$\mu_J$	-0.13	0.44	-0.45	-0.13	0.17	-0.13	0.44	-0.45	-0.13	0.17
$\sigma_J$	0.36	0.26	0.14	0.30	0.53	0.36	0.26	0.14	0.30	0.53
$\tau$	163.63	99.01	80.00	150.00	242.00	163.71	98.99	80.00	150.00	242.00
$mon$	1.04	0.33	0.77	1.04	1.32	1.04	0.33	0.77	1.04	1.32
$V(\theta, \Omega)/\mathbb{K}$	0.25	0.25	0.08	0.19	0.34	0.25	0.25	0.08	0.19	0.34
$\sigma(\theta, \Omega)$	76.81%	13.61%	68.27%	78.82%	87.58%	76.82%	13.62%	68.27%	78.84%	87.58%

It is well documented that the performances of neural networks improve, with the size of training datasets, given the greater range of distribution among input and output variables. During their respective calibration attempts with neural networks, Liu et al. (2019) and Büchel et al. (2022), resorted to creating a large training datasets, comprising millions of data points. The generation of such sizeable datasets can be computationally cumbersome, especially when using the GSVJD model with Monte Carlo simulations. It would be preferred to endure the one-time computational expense of simulating millions of synthetic option prices, and corresponding implied volatilities.

Synthetic datasets provide the ability to train the network on various market regimes, allowing for robust performances. Adhering to this stream of reasoning, a large number of option prices and implied volatilities were simulated, out of which 10 million, were selected at random for constructing a dataset. It is tasked with the purpose of training optimal neural networks, and hence will be referred to as the ‘*Optimal Training Dataset.*’ The optimal architectures are determined by initiating a random search, in union with the 3-fold cross

validation approach, over a wide range of hyperparameter combinations. Each architecture undergoes training and validation, and it would be efficient to conduct these practices the ‘*Random Search Dataset*.’ The secondary dataset is one-tenth of the size of its counterpart, comprising of 1 million option contracts. Hence, even when split into 3 equivalent folds, during cross validation, sufficient data exist to ensure robust training.

As observed from *Table A.1*, all components, across both datasets have very similar distributions. This ensures that the quality of the *Random Search* dataset has not been compromised, despite being reduced for efficiency. Since these datasets are synthetically created, there are no concerns regarding the overlapping/leaking of timeseries data between training and testing phases. After conducting a comprehensive analysis of literature on the application of neural networks for option pricing, Ruf & Wang (2020) highlight the ease of working with such datasets. An error exists from the pricing method used to create the dataset, however, on the plus side, training can be targeted towards learning a specified relationship between input and output variables. Recall the goal is to calibrate the GSVJD model, therefore such targeted training is preferable. It adds economic interpretation to the neural network and reduces regulatory concerns over ‘black-box’ operations.

### A.1.2 Random Search

The optimal architecture of the *forward pass* is determined by a random search over 720 hyperparameter combinations. Each one undergoes a 3-fold cross validation, to ensure robustness during the selection process. Performances, over the *Random Search* dataset, are evaluated using the average of error metrics, across the 3 folds. Given its superior ability to deal with outliers, the RMSE metric will be the key determinant, along with MSE and MAE. Instead of proceeding with the singular best architecture, the top 10 from the random search qualify for further training on the *Optimal Training* dataset.

The architectures in *Table A.2* are arranged in accordance with the  $Val_{RMSE}$  metric. All reported combinations reveal a strong preference for  $\lambda = 0$ , and *ELU* activation in the hidden layers. The performance of the neural network is adversely impacted by the  $L_2$  regularisation penalty term. This can be highlighted by the 50<sup>th</sup> ranked architecture, which is the best for  $\lambda = 1.0 \times 10^{-3}$ . Keeping the remaining hyperparameters unchanged, and

only adjusting the  $L_2$  regularisation term to  $\lambda = 0$ , the performance of this particular architecture would significantly improve, climbing up to the 2<sup>nd</sup> rank. Therefore,  $L_2$  regularisation is not employed during the optimal training of neural networks.

**Table A.2: Top Performing Architectures from the random search.**


*Note: Top 10 performing hyperparametric combinations from the random search over 720 neural network architectures, each uniquely identified by Arch ID. For the given values of hyperparameters, Layers, Nodes, Batch Size, Act. Fun. (hidden layer activation function),  $\lambda$  ( $L_2$  regularisation term), and  $\eta$  (learning rate), MSE and MAE metrics from the training ( $Train'_{MSE}, Train'_{MAE}$ ) and validation ( $Val'_{MSE}, Val'_{MAE}$ ) of the normalised output, along with RMSE from the validation of the original output ( $Val_{RMSE}$ ), are reported.*

Arch ID	Layers	Nodes	Batch Size	Act. Fun.	$\lambda$	$\eta$	$Train'_{MSE}$	$Train'_{MAE}$	$Val'_{MSE}$	$Val'_{MAE}$	$Val_{RMSE}$
610	4	30	2048	ELU	0	0.01	0.011	0.075	0.010	0.073	0.014
430	3	40	2048	ELU	0	0.01	0.011	0.075	0.011	0.076	0.014
670	4	40	2048	ELU	0	0.01	0.010	0.072	0.011	0.077	0.014
710	4	40	8192	ELU	0	0.01	0.012	0.079	0.011	0.076	0.014
450	3	40	4096	ELU	0	0.01	0.011	0.075	0.011	0.077	0.015
370	3	30	2048	ELU	0	0.01	0.012	0.078	0.012	0.078	0.015
470	3	40	8192	ELU	0	0.01	0.013	0.085	0.012	0.079	0.015
671	4	40	2048	ELU	0	0.001	0.012	0.079	0.012	0.078	0.015
650	4	30	8192	ELU	0	0.01	0.013	0.084	0.013	0.081	0.015
690	4	40	4096	ELU	0	0.01	0.011	0.076	0.013	0.083	0.015

A collective preference for deeper and wider networks, with smaller batch sizes is revealed in Table A.2. Amongst the top 10 architectures, a close to even split is witnessed between choices for 3 or 4 hidden layers, with the larger value being slightly favoured. A similar observation can be made over preference for 30 or 40 neurons. For batch sizes, half the architectures, including the top 3 prefer 2,048, whereas the rest favour 8,192. Smaller batch sizes improve the generalisation of the network by bettering convergence to flat sections of the error surface, but also increase computational inefficiency and the variance during gradient estimation (Goodfellow et al., 2016; Keskar et al., 2017). Furthermore, recall that a balance between the depth and width of the neural network is crucial, given the trade-off between accuracy and efficiency. Selecting shallow and narrow networks would worsen the performance but boost efficiency by reducing the computations required during the optimizations phase.

There exists a trade-off between estimation accuracy and computational efficiency, during the selection of hidden layers, neurons, and batch sizes. Therefore, it is important to look beyond the top 10 architectures and examine the extent to which neural network performances are impacted due to adjustments within these structural features. The extended examination analyses the distribution range of such hyperparameters, amongst the top 10<sup>th</sup> percentile of architectures, in *Figure A.2*. Furthermore, the average Val<sub>RMSE</sub> metric, across the hyperparameter ranges, is investigated in *Figure A.3*. Examining number of neurons and hidden layers, display a similar picture to that in *Table A.2*. Amongst the 72 architectures, *Figure A.2* reveals only 11 networks that have 2 hidden layers, whereas the rest are split 27 to 33, in favour of possessing 4. A preference for deeper networks is observed in *Figure A.3*, as those with 4, 3, and 2 layers, provide average Val<sub>RMSE</sub> metrics of 1.74%, 1.77%, and 1.90%, respectively. Similarly, there is an inclination towards wider networks, as those comprising of 20, 30, and 40 neurons, appearing 7, 24, and 40 times, give respective average performances of 1.87%, 1.77%, and 1.76%.

**Figure A.2: Hyperparameter Range of Architectures in the top 10<sup>th</sup> percentile.**

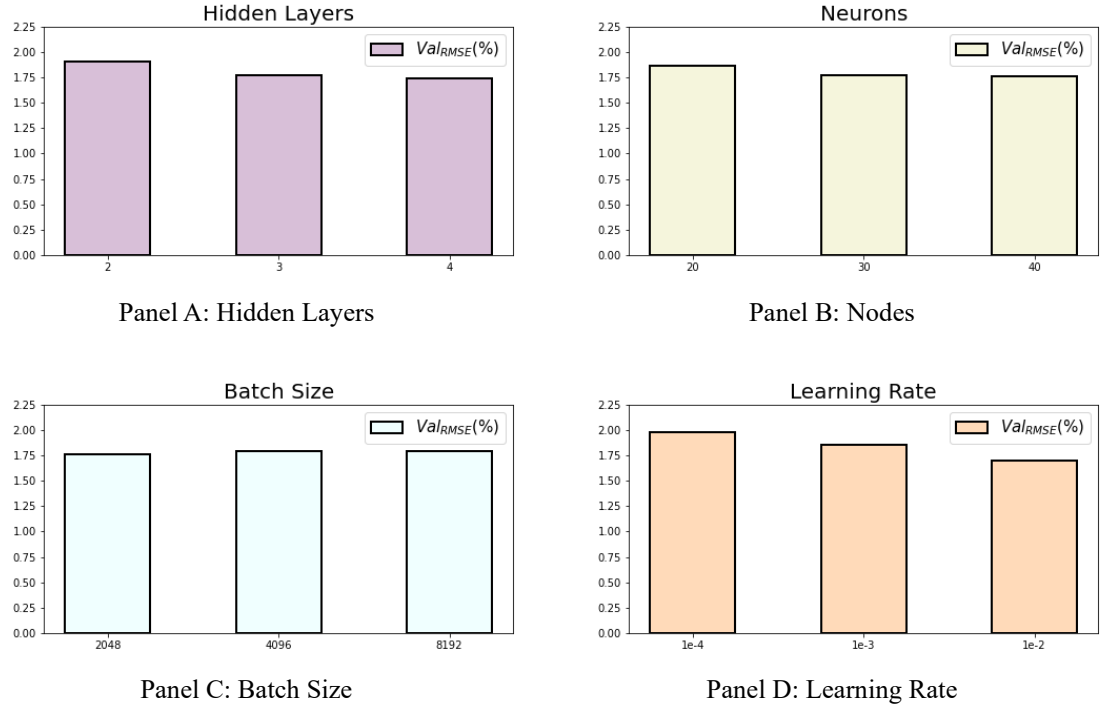


*Note: The count of Hidden Layers, Nodes, Batch Sizes and Learning Rate values, amongst the top 10<sup>th</sup> percentile (72) performing architectures from the random search, are displayed in Panel A, B, C, and D respectively.*

Neural networks, using 2048, 4096 and 8192 batch sizes during the optimisation, appear 18, 24, 29 times, providing mean  $Val_{RMSE}$  metrics of 1.75%, 1.79%, and 1.80%, respectively. Smaller batch sizes provide superior convergence to flat section of the error surface, hence revealing better average performances. Their presence is dominant within the highly ranked networks. Alternatively, they are vulnerable to high variances when estimating gradients, hence reducing favourability in comparison to the largest batch size. However, this problem is overcome, by pairing small batch sizes with deeper and wider networks, as witnessed with architectures in *Table A.2*. In contrast to the other hyperparameters, the learning rate,  $\eta$ , can be altered during the training phase. It controls the pace of optimisation by deciding the number of units moved, along the error surface. Initially, during the random search,  $\eta$  ranged from  $10^{-6}$  to  $10^{-2}$ , at an interval of  $10^1$ . However, in the top 10<sup>th</sup> percentile of architectures, the lower bound of the range is cut-off at  $10^{-4}$ . In ascending order, from *Figure A.3*, corresponding average  $Val_{RMSE}$  metrics, to  $\eta$  values of  $1e-4$ ,  $1e-3$ , and  $1e-2$ , are disclosed as 1.97%, 1.85%, and 1.70%, respectively. Regarding frequency of appearance, *Figure A.2* exhibits 2 architectures, amongst the 72 to possess a learning rate of  $1e-4$ , whereas 32 and 37 encompass  $1e-3$ , and  $1e-2$ , respectively.

In practice, it is difficult to select an optimal learning rate, given the trade-off between small and large values. A smaller  $\eta$ , increases the computational burden during training, however, is extremely beneficial when narrowing down to a minimum on the error surface. Alternatively, larger rates speed up the optimisation of the loss function but could cause the network to fluctuate around a minimum, or even worse, diverge from it. An argument to improve performance is put forward for scheduled learning rates, those decaying after certain number of epochs. Starting with a larger  $\eta$ , can improve efficiency in narrowing down towards a minimum. Furthermore, the large movements along the error surface would also prevent the network from being stuck as a local minima. To prevent fluctuations or divergence,  $\eta$  can be reduced, therefore improving the ability of the network to navigate towards the minimum. The computational burden would not be a source of concern, as majority of the navigation towards the minimum are done with a larger  $\eta$ .

**Figure A.3: Performances of top 10<sup>th</sup> percentile architectures across various Hyperparameters.**



Note: Average  $Val_{RMSE}$  metric across Hidden Layers, Nodes, Batch Sizes and Learning Rate values, amongst the top 10<sup>th</sup> percentile (72) performing architectures from the random search, are displayed in Panel A, B, C, and D respectively.

Though appealing, two flaws prevented the use of scheduled learning rates during the random search. First, the  $\eta$  values must be defined prior to training, and hence could fail to adapt to characteristics of the dataset. Second, regardless of the changes during training, the same learning rate, within each epoch, would be applied to all network parameter updates, even for those that occur rarely. The latter concerns are dealt with by employing the Adam optimizer. It is a galvanized version of the SGD, that employs adaptive learning rates, with respect to the weights and biases, whilst updating them at each epoch. The doubts over pre-determination of learning rates for the scheduled decline, are overcome by the information obtained from the random search. The extensive search was conducted on the *Random Search* dataset, which shares strong resemblances with the *Optimal Training* dataset. Hence, the scheduled learning rate is obtained from examining performances of the top 10<sup>th</sup> percentile of architectures. For optimal training decaying learning rate is proposed, starting at  $10^{-2}$  and halved after every 50 epochs, such that throughout the course of training,  $\eta$  will

be within the  $10^{-2}$  to  $10^{-4}$  range. This range is consistent with that observed amongst the top 10<sup>th</sup> percentile of architectures from the random search. It must be acknowledged that,  $\eta$  will approach, but never equal  $10^{-4}$ , given the lower frequency of appearances and relatively worse error metrics, designated with this value.

### A.1.3 Optimal Neural Network Training

The random search concluded a preference for deeper and wider network architectures, with smaller batch sizes and learning rates. Upon a 3-fold cross validation investigation, employed on each of the 720 hyperparametric combinations trained over the *Random Search* dataset, the 10 best performing networks, from *Table A.2*, are selected for optimal training.

The optimal training is conducted on a much larger, *Optimal Training* dataset. It shares the same characteristics as the *Random Search* dataset but is 10 times larger, comprising of 10 million synthetically priced options, approximately half of which correspond to put contracts. For boosting performances, input and output variables were scaled using (45) and (46), respectively. Such operations make the loss function more symmetrical, causing simpler and quicker convergence, and protecting the network from the vanishing gradient issue. Hence, requiring a large number of epochs is made redundant, with the optimal training being conducted over only 200. Efficiency and convergence are further boosted by the implementation of a scheduled decaying learning rate, with an initial value of  $10^{-2}$  halved, after 50 epochs. The network is further shielded from being stuck at a local minima, by the larger learning rate. When narrowing down to a minimum, to avoid fluctuations or divergence, the learning rate is reduced to better convergence.

The aim is to seek out a neural network architecture, that can best learn the GSVJD model for estimating  $\sigma_P(\theta, \Omega)$ , from inputs  $(\theta, mon, \tau)$ . The architectures considered for optimal training vary with respect to three hyperparameters: hidden layers, neurons, and batch sizes. The rest are set constant, in accordance with findings from the random search. The activation function in the hidden and output layers are fixed at ‘ELU’ and ‘linear’, respectively. The decaying learning rate that halves after 50 epochs, with an initial value of  $10^{-2}$ , applied all architectures, in addition with the decision to opt against  $L_2$  regularisation, by setting  $\lambda =$

0. Under these circumstances, neural networks with Arch IDs 670 and 671, are identical in *Table A.2*. They were differentiated between, solely based on their respective learning rates of  $10^{-2}$  and  $10^{-3}$ . To avoid repetition, the network with Arch ID<sup>56</sup> 671 is selected to undergo optimal training.

In the random search a 3-fold cross validation was conducted, which split the *Random Search* dataset into 3 components. While being trained, the networks were validated on each component, during the respective fold. Recall, the optimal training procedure is undertaken to determine the *forward pass*, such that its hidden layers are utilised in the *backward pass*. If the 3-fold cross validation technique were to be implemented, then there would be 3 sets of weights and biases, each corresponding to a respective fold of the *Optimal Training* dataset, and the hidden layers from best performing fold will be selected. This practice would be inadequate and computationally inefficient. The latter is a consequence of 3 training cycles required to complete the cross validation. The former stems from the best performing fold being selected without maximizing the utility of the entire dataset. Given hidden layers from only one-fold can be considered in the *backward pass*, splitting the dataset into 3-folds reduces the distributive range of the training dataset. Alternatively, the training dataset can be split into a greater number of folds, but this would reduce training efficiency. Therefore, it is preferred not to continue with the implementation of k-fold cross validation for optimal training.

The *Optimal Training* dataset undergoes a 90:10 random split to create training of validation subsets. As opposed to conventional practices, a third subset for additional testing is not considered. This is justified by the selection process for the optimal architecture, continuing into the *backward pass*. The purpose of this research does not end at calibrating the GSVJD model but rather at detecting bubbles from option prices. Hence, the architecture selected for bubble detection must display the best performances to fit market put options. This provides a strong economic background when selecting an optimal neural network for calibration, which assists in putting regulatory concerns regarding the ‘black-box’ nature of the model to bed.

---

<sup>56</sup> Arch ID 671 was allocated a learning rate of  $10^{-3}$  during the random search and allows to better capture the new implementation of the scheduled decaying learning rates, starting at  $10^{-2}$ .



The performance from each architecture considered for optimal training are reported in *Table A.3*. Similar to the random search,  $Val_{RMSE}$  plays a crucial role in determining the best architecture. The ranking of network architectures varies in comparison to those from the random search. Observe, the finest of margins separate the trained networks, during estimation of  $\sigma_p(\theta, \Omega)$ . For instance,  $Val_{RMSE}$  of the best and worst performing models, are separated by 0.104%. In regard to  $Val_{MAE}$ , the proximity of reported values is even closer, with near identical values obtained across all architectures. Therefore, to analyse the performances of these models, given the minute differences amongst them,  $Val_{MSE}$  becomes useful. Furthermore, a closer look should be taken at examining the difference between  $Train'_{MSE}$  and  $Val'_{MSE}$  metrics to ensure the networks do not suffer from overfitting.

According to  $Val_{MSE}$  from *Table A.3*, the first and second ranked candidates, comprise of the deepest and widest networks, across the whole hyperparameter search. These networks also provide an exception to the preference for smaller batch sizes. In all other instances, when comparing neural networks with the same width and depth, those with larger batch sizes perform worse. The next three ranking positions are occupied by networks comprising of 40 nodes within their hidden layers. A batch size of 2048 makes Arch ID 430 triumph over 710, and 450, even though the former is deeper. Overall, wider architectures are noticeably ranked higher, barring Arch ID 470, possibly due to a larger batch size. There is a preference for smaller batch sizes, but the depth of the neural network also plays a vital role in its performance. Consider networks with Arch ID 710 and 450, which have the same width, but the deeper one performs better despite having a larger batch size. A similar conclusion can be arrived at following the comparison of Arch IDs 650 and 690, with 370 and 430, respectively.

The optimal training reveals preference for wider networks, amongst which, those that are deeper with smaller batch sizes, perform even better, consistent with the findings of the random search. A strong case for exploring even deeper and wider neural network architectures can be put forward. However, proceeding down this route would make the network suffer from the curse of dimensionality, as the number of weights and bias parameters would increase exponentially, worsening the efficiency of the model. The

increase in computational burden of training the *forward pass*, is not a major source of concern, compared to potential additional seconds required during calibrations in the *backward pass*. The calibration task can be simplified to matrix operations, which would become even higher-dimensional in nature, if the network were to become wider or deeper. This might discourage the practitioner from implementing the framework, given the frequent requirements of calibrating parameters due to volatile essence of markets.

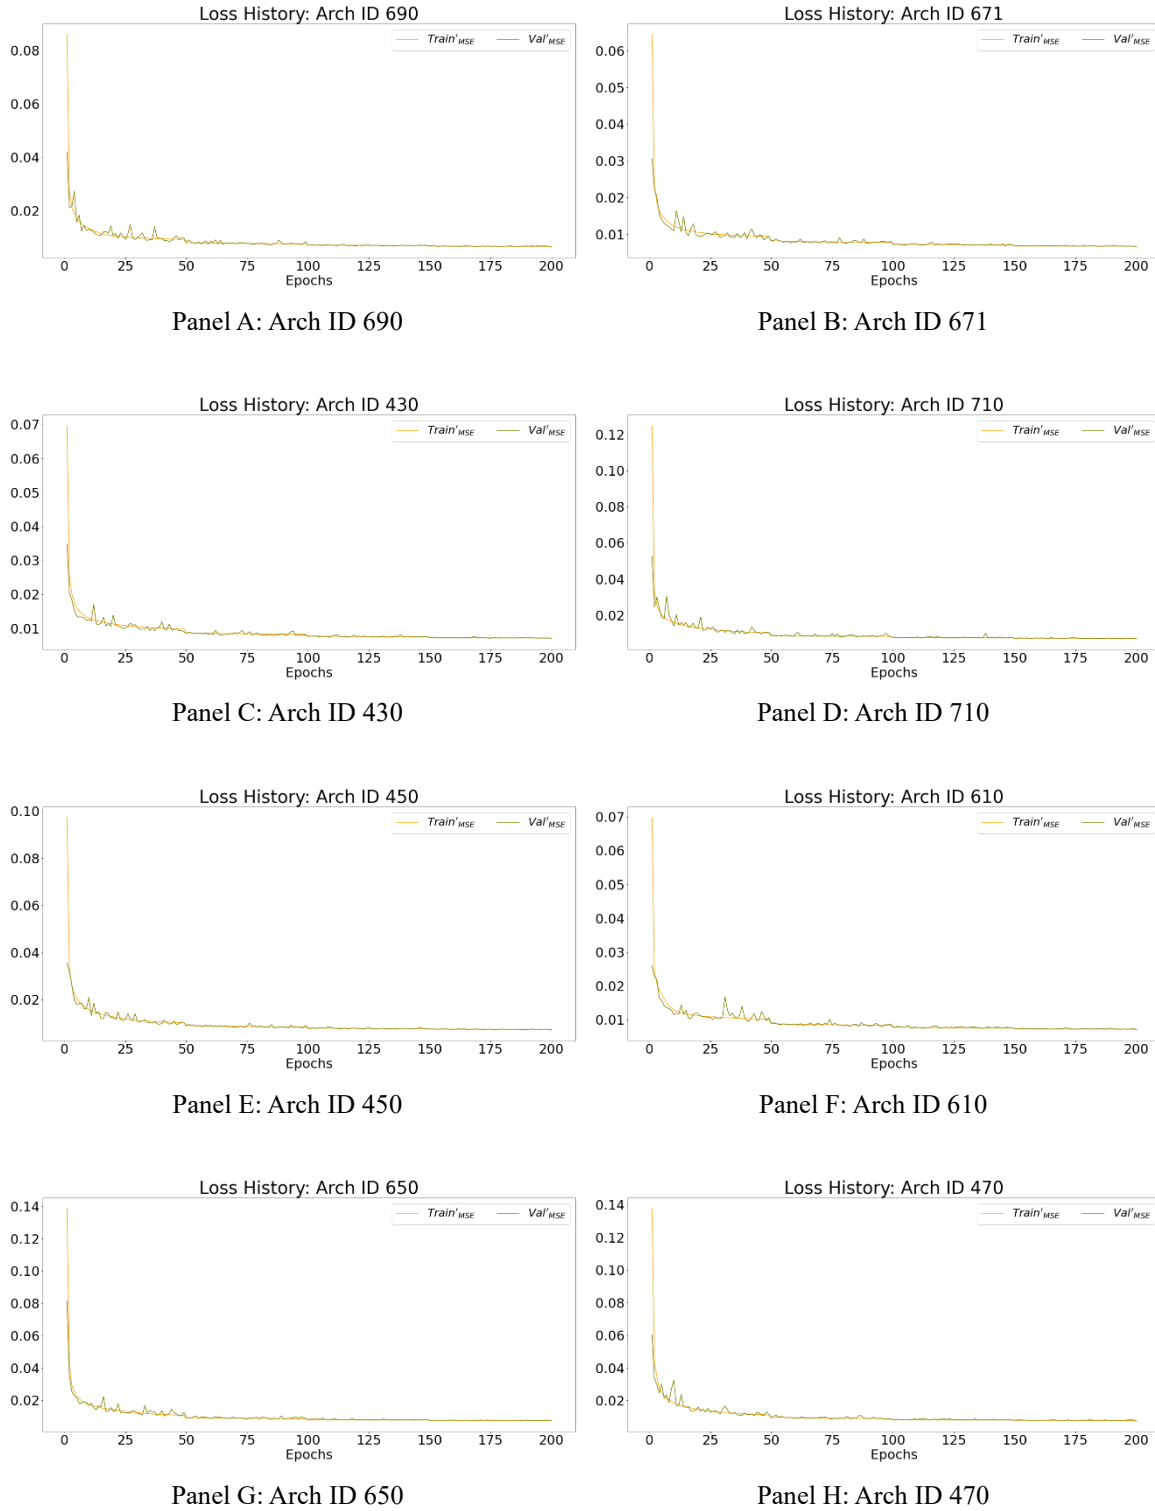
**Table A.3: Optimal Training Performance.**

*Note: Optimal training summary of the best performing hyperparametric combinations from the random search, each uniquely identified by Arch ID. For each combination of Layers, Nodes, and Batch Size, MSE and MAE metrics from the training ( $Train'_{MSE}$ ,  $Train'_{MAE}$ ) and validation ( $Val'_{MSE}$ ,  $Val'_{MAE}$ ) of the normalised output, along with MSE, RMSE, and MAE metrics from the validation of the original output ( $Val_{MSE}$ ,  $Val_{RMSE}$ ,  $Val_{MAE}$ ), are reported.*

Arch ID	Layers	Nodes	Batch Size	$Train'_{MSE}$	$Train'_{MAE}$	$Val'_{MSE}$	$Val'_{MAE}$	$Val_{MSE}$	$Val_{RMSE}$	$Val_{MAE}$
690	4	40	4096	0.007	0.058	0.007	0.058	1.22e-04	0.011	0.008
671	4	40	2048	0.007	0.058	0.007	0.059	1.25e-04	0.011	0.008
430	3	40	2048	0.007	0.060	0.007	0.059	1.29e-04	0.011	0.008
710	4	40	8192	0.007	0.059	0.007	0.059	1.29e-04	0.011	0.008
450	3	40	4096	0.007	0.060	0.007	0.060	1.33e-04	0.012	0.008
610	4	30	2048	0.007	0.060	0.007	0.060	1.33e-04	0.012	0.008
650	4	30	8192	0.007	0.061	0.007	0.061	1.38e-04	0.012	0.008
470	3	40	8192	0.008	0.062	0.008	0.062	1.40e-04	0.012	0.008
370	3	30	2048	0.008	0.063	0.008	0.064	1.46e-04	0.012	0.009


The  $Val_{MSE}$  metric allows for assessing the ranking of network architectures, given negligible differences observed within their  $Val_{RMSE}$  and  $Val_{MAE}$  values. Each of these 3 error metrics reveal performances on the validation dataset, ensuring robustness, and diminishing concerns over overfitting. This can be further established by examining the plots in *Figure A.4*, that reveal the history of training ( $Train'_{MSE}$ ) and validation ( $Val'_{MSE}$ ) loss functions over 200 epochs for each network. It is important to witness convergence of both loss functions at the final epoch, as this would resemble similar performances of the neural network on two different datasets. In other words, convergence confirms the network having learned the complex non-linear dynamics of the GSVJD model.

**Figure A.4: History of Loss during Optimal Training.**



*Figure A.4 continued...*

Figure A.4 continued...



Panel I: Arch ID 370

*Note: The history of loss over 200 epochs, measured by the  $Train'_{MSE}$  and  $Val'_{MSE}$  metrics, for all architectures from the training and validation of the forward pass, are respectively displayed in each panel.*

All architectures experience convergence between  $Train'_{MSE}$  and  $Val'_{MSE}$  metrics. The values of these metrics at the last epoch, along with those of  $Train'_{MAE}$  and  $Val'_{MAE}$ , are documented in *Table A.3*. This convergence reveals an economic foundation for an approach that is heavily criticized by regulators for its ‘black box’ nature and over reliance on data. The improved robustness of the neural network should prevent it from facing difficulty in mimicking such performances on unseen market data and boost its ability to calibrate GSVJD parameters. Recall, the three-step approach required the GSVJD model to accurately price put options, such that their fundamental and market values align, indicating the absence of a bubble. This is crucial for overcoming the joint-hypothesis related issue, when attempting to detect bubbles in call option/underlying asset prices. Therefore, each optimally trained architecture is carried to the *backward pass*, such that their hidden layers are frozen and utilised for calibrating from daily market put option implied volatilities.

## A.2 Backward Pass

The optimal training of the *forward pass* illustrated a preference for deeper and wider networks, with smaller batch sizes. The selection of a single best neural network was hindered by the discovery that all provided nearly identically strong performances. Furthermore, the approval of the network architecture for bubble detection also stems from its ability to accurately calibrate the GSVJD parameters to market put option prices/implied volatilities. Subsequently, the superior network is one that is least likely to suffer from the

joint hypothesis issue, by possessing best capabilities in identifying the true market regime. Therefore, the *backward pass* must not only be viewed as the calibration phase, but also one for testing the quality of the network architecture for bubble detection.

### A.2.1 Backward Pass: Optimal Architecture

The *backward pass* is designed to incorporate the trained hidden layers from the optimal *forward pass* and improve the computational efficiency of calibrating the GSVJD model. The benchmark for accuracy is set by using Monte Carlo simulation for calibrating the GSVJD model. Since both approaches, utilise the Differential Evolution optimiser to minimise the  $RMSE_{\sigma,t}$  between market and estimated implied volatilities, comparisons between daily accuracy and computational time are straightforward. However, prior to comparisons being made, the optimal architecture of the *forward pass* must be determined.


The secondary purpose of the *backward pass* is to assess the performance of each architecture that has undergone optimal training. This is conducted by calibrating the GSVJD model, on the same dataset as employed during the benchmark model, *HCV* put options (see *Table 4.2*), to maintain consistency with 4.2 and 4.3. Daily S&P 500 put options, corresponding to the most liquid maturity between January 2019 and December 2022, are considered. The Differential Evolution optimiser is utilised to minimise the  $RMSE_{\sigma,t}$ , and for additional robustness, the  $SSE_{\sigma,t}$  loss functions, between  $\sigma_P(\theta, \Omega)'$  and  $\sigma_P(mon, \tau)'$ . Therefore, the deep calibration framework is further validated over market data, prior to being utilised for bubble detection. The additional procedure reduces regulatory concerns, by eliminating the presence of joint-hypothesis related issues, that could arise from selecting a suboptimal architecture. Finally, given that the choice of asset pricing models in industrial practice is heavily governed by tractability, it is vital to investigate the time required by each architecture for daily calibrations.

The  $RMSE_{\sigma,t}$  values, between the estimated and market put option implied volatilities, across the entire sample period, are plotted for each of the architectures in *Figure A.5*. These plots illustrate the performance of the *backward pass*, when using  $RMSE_{\sigma,t}$  for calibrating the GSVJD model. The performances of the architectures are assessed by the average MSE,

RMSE, and MAE, error metric values. Along with the mean time for daily calibrations, depicted in *Table A.4*, with respect to  $RMSE_{\sigma,t}$  and  $SSE_{\sigma,t}$ . Overall, the *backward pass* delivers strong, stable, and most remarkably rapid calibrations of the GSVJD model to market data, across the entire study period, for all architectures.


The performance trends over the sample, of all neural network architectures, are similar to those displayed by using Monte Carlo. There are two noticeable instances, consistent with the benchmark, over the considered period, during which concerns regarding the calibration to daily S&P 500 options, are raised. First, a clustered rise in the  $RMSE_{\sigma,t}$ , can be observed at the start of March 2020, across all architecture plots in *Figure A.5*. Several instances of relatively larger errors, peaking close to 2%, are witnessed in the respective plots of each network architecture. This period corresponds to the start of the COVID-19 induced lockdown and subsequent crash when great amounts of turmoil and uncertainty clouded market expectations. Second, all neural networks deliver their worst performance on January 27, 2021, a day corresponding to the massive price hike experienced by the GameStop stock due to a short squeeze. A singular, yet towering  $RMSE_{\sigma,t}$  values, close to 4%, is noticed in all panels of *Figure A.5*. Barring these two scenarios corresponding to exceptional circumstances, all neural networks, deliver strong performances.

**Figure A.5: Backward Pass Calibrations using Optimal Architectures.**




*Figure A.5 continued...*


Figure A.5 continued...




Panel C: Arch ID 430




Panel D: Arch ID 710




Panel E: Arch ID 450




Panel F: Arch ID 610



Panel G: Arch ID 650



Panel H: Arch ID 470



Panel I: Arch ID 370

Note: Backward pass calibration of the GSVJD model by minimising  $RMSE_{\sigma,t}$  between  $\sigma_p(\theta, \Omega)$  and  $\sigma_p(mon, \tau)$ , using all architectures that underwent optimal forward pass training.

The performance of architectures during the optimal training phase of the *forward pass* were very similar. However, *Table A.4* reveals noticeable differences, with respect to accuracy and computational time. The network architectures are reported in order of their average performances over the study period, measured by the  $\sigma_{RMSE}$  metric. Upon closer inspection, a clear favorability for deeper and wider architectures is revealed, which automatically highlights the dreaded trade-off between accuracy and computational efficiency. The *backward pass* experimented by calibrating all architectures, using two loss functions,  $RMSE_{\sigma,t}$  and  $SSE_{\sigma,t}$ . The difference between the impact on accuracy is very minor, as both loss functions present the same ranking of architectures, with former being marginally superior in the estimation of put option implied volatilities. However, regarding computational efficiency, the latter is on average 9.5 seconds faster during the daily calibration of the GSVJD parameters from market data. The inclusion of the penalty term,  $\lambda = 1 \times 10^{-6}$  in the  $SSE_{\sigma,t}$  function, improves convergence, and hence computational speed, without having a major impact on calibration performances. It is on these grounds that the practitioner is strongly encouraged to utilise the  $SSE_{\sigma,t}$  loss function with a penalty term, when attempting to calibrate high-dimensional asset pricing models, using neural networks. However, despite the additional computational cost, the  $RMSE_{\sigma,t}$  loss function is selected for bubble detection, in order to maintain consistency with the benchmark.

In general, regardless of the loss function, considering all architectures, on average, calibration is sped up by a magnitude ranging between 253-542, when employing the neural network-based two-step calibration approach, compared to Monte Carlo simulations. Furthermore, excluding the Arch ID 610, the remaining architectures display superior performances, when examined against the benchmark. Thus, the application of neural networks not only boosts computational efficiency, but also calibration accuracy. As a result, a large potential for enhancing the precision of three-step approach, along with its tractability for real time bubble detection exists, such that it can be exploited to make the method favourable amongst practitioners. The selection of the optimal architecture for bubble detection is conducted by analysing the  $\sigma_{RMSE}$  metric. The ranking of architectures is similar



across the  $RMSE_{\sigma,t}$  and  $SSE_{\sigma,t}$  loss function, and in order to maintain consistency with the benchmark, will be examined only in accordance with the former.

**Table A.4: Calibration Performances during Backward Pass.**

*Note: Average performance and time (seconds) of the backward pass, for each optimal architecture from the forward pass phase, are displayed with respect to  $RMSE_{\sigma,t}$  (Panel A) and  $SSE_{\sigma,t}$  (Panel B). The performances are reported by MSE, RMSE and MAE metrics between estimated and observed implied volatilities ( $\sigma$ ).*

	$\sigma_{MSE}$	$\sigma_{RMSE}$	$\sigma_{MAE}$	Time
Panel A: $RMSE_{\sigma,t}$				
671	7.94E-06	0.17%	0.0011	50.371
690	9.39E-06	0.19%	0.0014	48.169
710	9.95E-06	0.22%	0.0016	50.556
650	1.12E-05	0.23%	0.0015	44.626
470	1.53E-05	0.28%	0.0020	35.086
370	1.69E-05	0.28%	0.0021	35.679
450	1.69E-05	0.29%	0.0021	34.154
430	1.65E-05	0.30%	0.0022	34.803
610	2.64E-05	0.38%	0.0028	34.090
Panel B: $SSE_{\sigma,t}$				
671	8.01E-06	0.17%	0.0011	43.177
690	9.64E-06	0.20%	0.0014	41.879
710	9.93E-06	0.22%	0.0016	36.039
650	1.12E-05	0.23%	0.0015	32.224
470	1.52E-05	0.28%	0.0020	29.466
370	1.67E-05	0.28%	0.0021	24.355
450	1.71E-05	0.29%	0.0021	27.062
430	1.69E-05	0.31%	0.0022	24.531
610	2.79E-05	0.39%	0.0029	23.606

It can be observed from *Table A.4* that the Arch ID 671, with 4 hidden layers, each comprising of 40 neurons, and a batch size of 2048, is the best performing architecture. It provides  $\sigma_{RMSE}$  value of 0.17%. Furthermore, it is the deepest and widest network, with the smallest batch size, giving a vindication to the random search findings. In fact, the top three networks measured by the  $\sigma_{RMSE}$ , are equally as deep and wide, with the following two, in order of performance, Arch ID 690 and 710, having respective batch sizes of 4096 and 8192. Arch ID 610, is the worst performing network, even being inferior to the benchmark.

Ironically, it is also the most computationally efficient architecture for the *backward pass*. Even though practitioners favour tractability, accuracy must not be overlooked when determining the optimal architecture for bubble detection. The application of neural networks for calibrating the GSVJD has significantly boosted computational accuracy, and even improved efficiency, in comparison to the Monte Carlo simulations. From *Table A.4*, preference for speed, comes at a cost of accuracy, with the best performing network being the slowest, whereas the fastest, Arch ID 610, being inferior to the benchmark.

Naturally, one can arrive at the conclusion that deeper and wider networks require more calibration time. This is a fair perception, given the curse of dimensionality results in an exponential increase in the number of weights and bias parameters within the hidden layers. However, Arch ID 610 comprised of 4 hidden layers and 30 neurons and was even faster than shallower and narrower networks. The possibility of the optimiser being stuck at a local minimum, and stopping given the lack of improvement, due to the drawbacks of the architecture, must not be ignored. Therefore, it is preferable to curb the greed for speed and select a network with better accuracy. Hence, the optimal architecture selected for bubble detection is Arch ID 671, comprising of 4 hidden layers, 40 neurons and a batch size of 2048. It provides speed up with a magnitude of 254, in addition to being approximately twice as accurate as the benchmark. Finally, all calibrations during the *backward pass* stage were computed on a personal computer, with an i7-processor, and 16 GB RAM, over a single CPU unit.