

Detecting asset price bubbles using deep learning

Francesca Biagini* Lukas Gonon† Andrea Mazzon‡ Thilo Meyer-Brandis *

June 21, 2024

Abstract

In this paper we employ deep learning techniques to detect financial asset bubbles by using observed call option prices. The proposed algorithm is widely applicable and model-independent. We test the accuracy of our methodology in numerical experiments within a wide range of models and apply it to market data of tech stocks in order to assess if asset price bubbles are present. Under a given condition on the pricing of call options under asset price bubbles, we are able to provide a theoretical foundation of our approach for positive and continuous stochastic asset price processes. When such a condition is not satisfied, we focus on local volatility models. To this purpose, we give a new necessary and sufficient condition for a process with time-dependent local volatility function to be a strict local martingale.

1 Introduction

The study of methodologies for detecting asset price bubbles has attracted an increasing interest over the last years in order to prevent, or mitigate, the financial distress often following their burst. A financial bubble is defined as the temporary deviation of the market price of an asset from its fundamental value. In the context of the martingale theory of bubbles (see among others Cox and Hobson (2005), Loewenstein and Willard (2000), Jarrow et al. (2007), Jarrow et al. (2010), Jarrow et al. (2011b), Biagini et al. (2014), Protter (2013)), the market price of a (discounted) asset represents a bubble if it is given by a strict local martingale. In this framework, detecting an asset price bubble amounts to determine if the stochastic process modelling the discounted asset price is a true martingale or a strict local martingale.

Our contribution is part of a wide range of works on asset price bubbles detection. The problem is tackled within the framework of local volatility models in the papers Jarrow (2016),

*Workgroup Financial and Insurance Mathematics, Department of Mathematics, Ludwig-Maximilians Universität, Theresienstrasse 39, 80333 Munich, Germany. Emails: biagini@math.lmu.de, meyer-brandis@math.lmu.de.

†Department of Mathematics, Imperial College London, London, SW7 1NE UK. Email: l.gonon@imperial.ac.uk

‡Department of Economics, University of Verona, 37129 Verona, Italy. Email: andrea.mazzon@univr.it

Jarrow et al. (2011b) and Jarrow et al. (2011c), via volatility estimation techniques. In Fusari et al. (2020), under the joint assumption of *no free lunch with vanishing risk* (NFLVR) and *no-dominance*, the authors pursue this goal by looking at the differential pricing between put and call options, with the motivation that traded call and put options reflect price bubbles differently. On the other hand, a nonparametric estimator of asset price bubbles with the only assumption of NFLVR is proposed in Jarrow and Kwok (2021), using a cross section of European option prices. The approach is based on a nonparametric identification of the state-price distribution and the fundamental value of the asset from option data. In Piironen et al. (2018), the authors introduce a statistical indicator of asset price bubbles based on the bid and ask market quotes for the prices of put and call options. This methodology is in particular applied to the detection of asset price bubbles for SABR dynamics. Finally, an asymptotic expansion of the right wing of the implied volatility smile is used in Jacquier and Keller-Ressel (2018) to determine the strict local martingale property of the underlying.

In this paper, we propose a new approach that formulates bubble detection as a supervised learning problem. The learning problem is solved using neural networks that take as inputs observed call option prices. In this way we provide a theoretical foundation and a model-free deep learning approach for bubble detection, which is new in the literature. More precisely, we feed the network with a set of training data whose i -th data point consists of analytical call option prices for different strikes K_1, \dots, K_n and maturities T_1, \dots, T_m written on an underlying process X^i with known dynamics, together with an indicator specifying if X^i is a strict local martingale or a true martingale. To assess if a new stock (possibly corresponding to market data) has a bubble or not, we evaluate the trained neural network at (market) prices of call options on this underlying for strikes K_1, \dots, K_n and maturities T_1, \dots, T_m . The output of the network is in $(0, 1)$ and indicates the probability that the underlying exhibits a bubble. Slightly changing the algorithm, we are also able to estimate the size of the bubble.

In this manner the methodology does not attempt to detect a bubble based on the trajectory of the underlying, but instead uses the information contained in all call option prices available at a given point in time for different maturities and strikes to recover information on the asset's price probability distribution. In fact, the (theoretical) observation of call option prices for all strikes and maturities determines full information about the distribution of the underlying by the Breeden-Litzenberger formula. The method aims to exploit the information contained in call option prices in the best possible way to provide an estimate for the probability that an asset price bubble is present.

Formulating bubble detection as a supervised learning problem leads to a purely sample-based procedure. The only information used when training the neural network is a collection of call option prices associated to the process X^i and an indicator stating if these call option prices correspond to a price process with a bubble or not. Thus, we can arbitrarily combine different models for bubbles and could also enhance synthetic model training data by historical data.

One of the main advantages of our approach is that it does not require the direct estimation of any parameter or quantity related to the asset price process, and that it is model independent. In particular, we show via numerical experiments that our method works not only within a class of models, but also if the network is trained using a certain kind of stochastic processes (like local volatility models) and tested within another class (for example, stochastic volatility models). We also test the method with market data associated to assets that have been suspected to be involved in the new tech bubble (see among others Fusari et al. (2020), Piironen et al. (2018), Kolakowski (2019), Libich et al. (2021), Bercovici (2017), Ozimek (2017), Serla (2017), Sharma (2017)). According to the neural network, output bubbles were present in these assets with high probability at certain points in time, and these dates match retrospectively expected results.

Our motivation for this method is manifold: on the one hand, the price of call options on a bubbly underlying has an additional term which is added to the usual risk neutral valuation due to a collateral requirement represented by a constant $\alpha \in [0, 1]$, see Cox and Hobson (2005) and Theorem 2.4 in our paper. Looking at call option prices, it is then theoretically possible to assess if the underlying has a bubble by identifying the presence of this term. On the other hand, in the case when the underlying price follows a local volatility model, a modified version of Dupire’s formula stated in Theorem 2.3 of Ekström and Tysk (2012) permits to theoretically recover the local volatility function, which is crucial to determine if the process is a strict local martingale or a true martingale, from the observation of call option prices.

In this way we are able to provide a theoretical foundation for our method. In particular, in Theorems 2.10 and 2.24 we prove the existence of a sequence of neural networks approximating a theoretical “bubble detection function” F , under the assumption $\alpha < 1$ for underlyings given by continuous and positive stochastic processes and in the general case $\alpha \in [0, 1]$ for local volatility models. The function F maps from the space \mathcal{X} of call option prices as functions of strike and maturity to $\{0, 1\}$, being 1 if and only if the underlying process is a strict local martingale. Several steps are required in order to prove Theorems 2.10 and 2.24. First, we show the existence of the function F introduced above and show that it is measurable with respect to a natural topology on \mathcal{X} , see Propositions 2.6 and 2.20. In the case when $\alpha = 1$, we consider the class of local volatility models under a fairly general assumption on the local volatility function. Specifically, we use some new findings providing a sufficient and necessary condition for a stochastic process with time dependent local volatility function to be a strict local martingale. The main result of this analysis is stated in Theorem 2.18 and is of independent interest, since it is a generalisation of Theorem 8 of Jarrow et al. (2011b). In Propositions 2.7 and 2.23 we then construct a sequence of measurable functions $(F^n)_{n \geq 1}$ approximating F pointwise, where $F^n : \mathbb{R}^n \rightarrow [0, 1]$ takes the prices of a call option with fixed underlying for n different strikes and maturities and outputs the likelihood that the

underlying has a bubble according to the observation of these prices. These propositions are then used together with the universal approximation theorem, see Hornik (1991), to prove Theorems 2.10 and 2.24, respectively.

The paper is structured as follows. Section 2 is devoted to the theoretical foundation of our approach. In particular, in Section 2.1 we briefly recall the martingale theory of asset price bubbles, in Section 2.2 we introduce call option pricing in presence of financial asset bubbles, and in 2.3 we prove existence of the test function F and of its neural network approximation when $\alpha < 1$, when the asset price is given by a continuous and positive stochastic process. On the other hand, in Section 2.4 we focus on local volatility models to cover the general case $\alpha \in [0, 1]$: in Section 2.4.1 we provide sufficient and necessary conditions for the strict local martingale property of local volatility models, whereas in Section 2.4.2 we give results analogous to the ones in Section 2.3 under this new setting.

In Section 3 we test the performance of our methodology via numerical experiments. Specifically, we introduce our methodology in Section 3.1 and in Section 3.2 we both feed and test the network within local volatility models, with different, randomly chosen parameters, whereas in Section 3.3 we feed the network with data coming from stochastic volatility models and test it with local volatility models, and vice-versa. In Section 3.4 we comment on the choice of the network architecture. In Section 4 we apply our approach to the analysis of four stocks that have been claimed to be affected by a bubble in recent years, namely Nvidia, Meta, GameStop and Tesla. In Section 5 we present some conclusions.

2 Theoretical foundation

2.1 Asset price bubbles

In the literature there have been many approaches proposed for the mathematical modelization of asset price bubbles.

From an economic point of view, the main challenge consists in explaining how bubbles are generated at the microeconomic level by the interaction of market participants; see Tirole (1982), Harrison and Kreps (1978), DeLong et al. (1990), Scheinkman and Xiong (2003), Abreu and Brunnermeier (2003), Föllmer et al. (2005) and the references therein. From a mathematical point of view, asset price bubbles are mainly studied by using the *local martingale framework*: Loewenstein and Willard (2000), Cox and Hobson (2005), Jarrow and Madan (2000), Jarrow et al. (2007), Jarrow et al. (2010), Jarrow et al. (2011a), Jarrow and Protter (2009), Jarrow and Protter (2011), Jarrow and Protter (2013), Pal and Protter (2010), Kardaras et al. (2015), Biagini et al. (2014). Another definition has been proposed in Herdegen and Schweizer (2016), where the fundamental value is assumed to be the superreplication price of the asset. In Jarrow et al. (2012), Biagini and Nedelcu (2015) and Biagini et al. (2018) constructive models are proposed to take in account triggering factors for bubble formation

and bubble bursting. For a survey, we refer to Protter (2013).

Here we follow the approach of Jarrow and Protter (2009), Jarrow et al. (2011a). The notion of an asset price bubble has two ingredients: the observed *market price* S of a given financial asset and the asset's *fundamental value*. Assume discounting equal 1. Given an equivalent martingale measure Q for S , the fundamental value S^Q is usually defined as the expected sum of future dividends under Q . The bubble perceived under Q is defined as the difference between the two:

$$\beta^Q = S - S^Q.$$

In this setting, we have a bubble if and only if S is a strict local martingale under Q . In the sequel, we present a new method to detect financial asset prices' bubbles by testing for strict local martingale properties based on observations of call option prices and by applying machine learning techniques.

2.2 European call options under asset price bubbles

Let (Ω, P, \mathcal{F}) be a probability space endowed with a filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ satisfying the usual hypothesis. It is well known that the risk-neutral valuation of European options when the underlying asset price has a bubble (i.e., when it is a strict local martingale) does not satisfy the put-call parity, see for example Jarrow et al. (2010). An alternative way to price call options for a bubbly underlying, also supported by market data, has been proposed in Cox and Hobson (2005).

We start with the following definition.

Definition 2.1. *Introduce an underlying given by a continuous, positive stochastic process $X = (X_t)_{t \geq 0}$, and fix a pricing measure $Q \sim P$ under which X is a local martingale. The price $C^{\alpha, X}(T, K)$ at time $t = 0$ of a call option of maturity T and strike K written on X with collateral constant $\alpha \in [0, 1]$ is given by*

$$C^{\alpha, X}(T, K) = \mathbb{E}^Q[(X_T - K)^+] + \alpha m(T), \quad (1)$$

where

$$m(T) := X_0 - \mathbb{E}^Q[X_T]. \quad (2)$$

The term $m(T)$ in (2) is called the *martingale defect* of X . Note that $m(T) > 0$ if X is a strict local martingale bounded from below and $m(T) = 0$ if X is a martingale.

Remark 2.2. *Note that if for $t \in [0, T]$ we set*

$$C_t^{\alpha, X}(T, K) := \mathbb{E}^Q[(X_T - K)^+ | \mathcal{F}_t] + \alpha m_t(T)$$

for $m_t(T) = X_t - \mathbb{E}^Q[X_T | \mathcal{F}_t]$, then $C_t^{\alpha, X}(T, K)$, $t \in [0, T]$, is a non-negative local martingale, which equals $(X_T - K)^+$ at final time. Hence $C_t^{\alpha, X}(T, K)$ represents an arbitrage-free price at time t for the call option $(X_T - K)^+$.

If the market is complete, a financial interpretation of (1)-(2) where $\alpha \in [0, 1]$ assumes the meaning of *collateral requirement* comes from Cox and Hobson (2005), where the following definition is given.

Definition 2.3. *The price at time $t = 0$ of a call option $C = (X_T - K)^+$ of maturity T and strike K written on a continuous, positive underlying $X = (X_t)_{t \geq 0}$ is the smallest initial value V_0 of a superreplicating portfolio $V = (V_t)_{t \geq 0}$ satisfying the admissibility condition*

$$V_t \geq \alpha(X_t - K)^+ \quad (3)$$

for all $t \in [0, T]$, where $\alpha \in [0, 1]$ is a constant representing the collateral requirement.

Condition (3) is always satisfied when the underlying asset has no bubble. For complete market models we have the following result, given in Cox and Hobson (2005).

Theorem 2.4. *Introduce an underlying given by a continuous, positive stochastic process $X = (X_t)_{t \geq 0}$, and assume the existence of a unique probability measure $Q \sim P$ under which X is a local martingale. The price according to Definition 2.3 at time $t = 0$ of a call option of maturity T and strike K written on X and with collateral requirement $\alpha \in [0, 1]$ is given by equation (1).*

Such results can be extended to càdlàg processes, see Cox and Hobson (2005).

Remark 2.5. *Under a condition known as no dominance, the price of a call option on a bubbly asset is given by formula (1) with $\alpha = 1$ and satisfies the put-call parity, see Theorem 6.2 of Jarrow et al. (2010). Our setting, where $\alpha \in [0, 1]$, includes the framework of Jarrow et al. (2010) as a particular case. Moreover, our method does not need to know which value of α is taken to price the option.*

2.3 Neural network approximation for collateral $\alpha \in [0, 1]$

In this section we fix a collateral constant $\alpha \in [0, 1]$ and suppose that the price of a call option is given by (1), with continuous and positive underlying X . For the sake of simplicity, from now on we write C^X for the call option price. We treat the case $\alpha = 1$ in Section 2.4 in the specific setting of local volatility models. Our aim is to find a way of detecting asset price bubbles by using neural networks, as we explain in the sequel. From now on, we fix the “pricing measure” $Q \sim P$.

We start our analysis with the following result.

Proposition 2.6. *Fix a constant $x_0 > 0$. Consider the spaces*

$$\mathcal{M}_{loc} = \{X = (X_t)_{t \geq 0} \text{ continuous and positive local martingale on } (\Omega, \mathcal{F}, Q, \mathbb{F}) \text{ with } X_0 = x_0\}$$

and

$$\mathcal{M}_L = \{X \in \mathcal{M}_{loc}: X \text{ is a strict local martingale}\}. \quad (4)$$

Introduce the set

$$\mathcal{X} := \{C^X : X \in \mathcal{M}_{loc}\} \cap C([0, \infty) \times (0, \infty), [0, \infty)) \quad (5)$$

endowed with the Borel sigma-algebra of $C([0, \infty) \times (0, \infty), [0, \infty))$. Then there exists a measurable function $F: \mathcal{X} \rightarrow \{0, 1\}$ such that $F(C^X) = \mathbb{1}_{\{X \in \mathcal{M}_L\}}$.

In particular,

$$F(f) = \mathbb{1}_{\{\text{there exists } t \geq 0 \text{ such that } \lim_{x \rightarrow 0} f(T, x) < X_0 \text{ for any } T \geq t\}}, \quad f \in \mathcal{X}. \quad (6)$$

Proof. Let $X \in \mathcal{M}_{loc}$. From (1) and (2) we get

$$\begin{aligned} \lim_{K \rightarrow 0} C^X(T, K) &= \lim_{K \rightarrow 0} \mathbb{E}^Q[(X_T - K)^+] + \alpha m(T) = \mathbb{E}^Q[X_T] + \alpha(X_0 - \mathbb{E}^Q[X_T]) \\ &= \alpha X_0 + \mathbb{E}^Q[X_T](1 - \alpha), \end{aligned} \quad (7)$$

by the Dominated Convergence Theorem since X is a supermartingale and therefore integrable for any $t \geq 0$. In particular, since X is a supermartingale, equation (7) implies that it is a strict local martingale if and only if there exists $t > 0$ such that

$$\lim_{K \rightarrow 0} C^X(T, K) < X_0 \text{ for any } T \geq t.$$

Therefore, we have

$$F(C^X) = \mathbb{1}_{\{X \in \mathcal{M}_L\}} \text{ for any } C^X \in \mathcal{X}$$

for F given in (6). In order to prove that F is measurable, note that

$$F(f) = 1 - \prod_{t \in \mathbb{N}} \bar{F}_t(f), \quad f \in \mathcal{X},$$

with $\bar{F}_t: \mathcal{X} \rightarrow \{0, 1\}$ defined by

$$\bar{F}_t(f) := \mathbb{1}_{\{\lim_{x \rightarrow 0} f(t, x) \geq X_0\}} = \mathbb{1}_{\{\lim_{x \rightarrow 0} \bar{G}_{t, x}(f) \geq 0\}}, \quad t \in \mathbb{N}, \quad f \in \mathcal{X},$$

where $\bar{G}_{t, x}: \mathcal{X} \rightarrow \mathbb{R}$ is given by $\bar{G}_{t, x}(f) := f(t, x) - X_0$. As \bar{G} is measurable, then also $\lim_{x \rightarrow 0} \bar{G}_{t, x}$ is measurable, and hence \bar{F}_t and consequently F are also measurable. \square

We now prove the existence of a sequence of functions $(F^n)_{n \in \mathbb{N}}$, $F^n: \mathbb{R}^n \rightarrow [0, 1]$, such that F^n approximates, in some way we specify, the function F from (6).

Proposition 2.7. *Define the set*

$$A := [0, \infty) \times (0, \epsilon] \quad (8)$$

for some fixed $\epsilon > 0$.

Let $(T_n, K_n)_{n \in \mathbb{N}}$ be a sequence of maturities and strikes such that the set $\{(T_n, K_n), n \in \mathbb{N}\}$ is dense in A . Let μ be a probability measure on \mathcal{X} such that there exists a probability measure ν on \mathcal{X} with

$$\nu(\{f \in \mathcal{X} \text{ such that } f(T_i, K_i) = g(T_i, K_i) \text{ for all } i \in \mathbb{N}\}) > 0 \quad \text{for any } g \in \text{supp}(\mu). \quad (9)$$

Then there exists a sequence of functions $(F^n)_{n \in \mathbb{N}}$, $F^n: \mathbb{R}^n \rightarrow [0, 1]$ such that

$$\int_{\mathcal{X}} |F^n(g(T_1, K_1), \dots, g(T_n, K_n)) - F(g)|^p d\mu(g) \xrightarrow{n \rightarrow \infty} 0, \quad (10)$$

for any $p \in [1, \infty)$, where $F: \mathcal{X} \rightarrow \{0, 1\}$ is the function introduced in (6). For fixed $n \in \mathbb{N}$, the function F^n can be chosen as follows:

$$F^n(c^n) := \begin{cases} \frac{1}{\nu(S_n(c^n))} \int_{S_n(c^n)} F(f) d\nu(f) & \text{if } \nu(S_n(c^n)) > 0, \\ 0 & \text{otherwise,} \end{cases} \quad (11)$$

$c^n \in \mathbb{R}^n$, where

$$S_n(c^n) := \{f \in \mathcal{X} \text{ such that } f(T_i, K_i) = c_i^n \text{ for all } i = 1, \dots, n\}, \quad c^n \in \mathbb{R}^n. \quad (12)$$

Proof. Let $(T_n, K_n)_{n \in \mathbb{N}}$ be a sequence of maturities and strikes such that the set $\{(T_n, K_n), n \in \mathbb{N}\}$ is dense in A given in (8). Let μ be the probability measure on \mathcal{X} introduced above and fix $g \in \text{supp}(\mu)$. Also let ν be a probability measure on \mathcal{X} which satisfies (9). For any $n \in \mathbb{N}$, consider the function $F^n: \mathbb{R}^n \rightarrow [0, 1]$ defined in (11) and set

$$c^{g,n} := (g(T_1, K_1), \dots, g(T_n, K_n)).$$

We first want to show that

$$F^n(c^{g,n}) \xrightarrow{n \rightarrow \infty} F^\infty(c^{g,\infty}), \quad (13)$$

with

$$c^{g,\infty} := \{g(T_i, K_i), \quad i \in \mathbb{N}\}$$

and

$$F^\infty(c^{g,\infty}) := \frac{1}{\nu(S_\infty(c^{g,\infty}))} \int_{S_\infty(c^{g,\infty})} F(f) d\nu(f),$$

where

$$S_\infty(c^{g,\infty}) := \{f \in \mathcal{X} \text{ such that } f(T_i, K_i) = g(T_i, K_i) \text{ for all } i \in \mathbb{N}\}$$

satisfies $\nu(S_\infty(c^{g,\infty})) \neq 0$ by (9). Since $S_{n+1}(c^{g,n+1}) \subset S_n(c^{g,n})$ for any $n \in \mathbb{N}$ and $S_\infty(c^{g,\infty}) = \bigcap_{n \in \mathbb{N}} S_n(c^{g,n})$, we have

$$\mathbb{1}_{S_n(c^{g,n})} \xrightarrow{n \rightarrow \infty} \mathbb{1}_{S_\infty(c^{g,\infty})} \quad \text{pointwise} \quad (14)$$

and

$$\nu(S_n(c^{g,n})) \xrightarrow{n \rightarrow \infty} \nu(S_\infty(c^{g,\infty})). \quad (15)$$

By (14) and since $|F| \leq 1$, we can apply the Dominated Convergence Theorem and get

$$\int_{S_n(c^{g,n})} F(f) d\nu(f) \xrightarrow{n \rightarrow \infty} \int_{S_\infty(c^{g,\infty})} F(f) d\nu(f). \quad (16)$$

Putting together (15) and (16), we obtain (13).

Moreover, as the set $\{(T_n, K_n), n \in \mathbb{N}\}$ is dense in A by assumption, g is continuous and every function in \mathcal{X} is continuous as well, we have that

$$S_\infty(c^{g,\infty}) = \{f \in \mathcal{X} \text{ such that } f(T, K) = g(T, K) \text{ for all } (T, K) \in A\}.$$

By the definition of F in (6), this implies that $F(f) = F(g)$ for any $f \in S_\infty(c^{g,\infty})$, and so $F^\infty(c^{g,\infty}) = F(g)$. By (13), it thus follows that

$$F^n(g(T_1, K_1), \dots, g(T_n, K_n)) \xrightarrow{n \rightarrow \infty} F(g).$$

We then obtain (10) by the Dominated Convergence Theorem, since $(F^n)_{n \in \mathbb{N}}$ is bounded by 1. \square

Remark 2.8. Any discrete probability measure μ on \mathcal{X} , i.e. any probability measure μ such that $\text{supp}(\mu) \subseteq \mathcal{X}_0$ for a countable set $\mathcal{X}_0 \subset \mathcal{X}$, satisfies the requirement of Proposition 2.7 that there exists a probability measure ν on \mathcal{X} for which (9) holds. An example is given by

$$\nu := \frac{1}{\sum_{n \in \mathbb{N}} n^{-2}} \sum_{n \in \mathbb{N}} \delta_{g_n} n^{-2},$$

where $\mathcal{X}_0 := \{g_n, n \in \mathbb{N}\}$. Indeed, for any $g \in \text{supp}(\mu)$ there exists $n \in \mathbb{N}$ such that $g = g_n$, thus

$$\nu(\{f \in \mathcal{X} \text{ such that } f(T_i, K_i) = g(T_i, K_i) \text{ for all } i \in \mathbb{N}\}) \geq \nu(\{g_n\}) > 0.$$

We now want to approximate the sequence of functions $F^n : \mathbb{R}^n \rightarrow [0, 1]$ defined in (11), and thus the function F , by a sequence of neural networks. We start with the formal definition of a neural network, see for example Hornik (1991).

Definition 2.9. Let $I \subseteq \mathbb{R}$. A neural network $\hat{F} : \mathbb{R}^n \rightarrow I$ is a function

$$\hat{F}(x) = (\psi \circ W_L \circ (\sigma \circ W_{L-1}) \circ \dots \circ (\sigma \circ W_1))(x), \quad x \in \mathbb{R}^n,$$

with given $L \in \mathbb{N}$, $L \geq 2$, where the inner activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is bounded, continuous and non-constant, where the outer activation function $\psi : \mathbb{R} \rightarrow I$ is invertible and Lipschitz continuous and where $W_\ell : \mathbb{R}^{N_{\ell-1}} \rightarrow \mathbb{R}^{N_\ell}$ are affine functions with $N_1, \dots, N_{L-1} \in \mathbb{N}$, $N_0 = n$, $N_L = 1$.

We then have the following result.

Theorem 2.10. Fix $p \in [1, \infty)$ and let $(T_n, K_n)_{n \in \mathbb{N}}$ be the sequence of maturities and strikes introduced in Proposition 2.7. Also let μ be a probability measure on \mathcal{X} such that there exists a probability measure ν on \mathcal{X} which satisfies (9). Then for any $\epsilon > 0$ there exists an $n \in \mathbb{N}$ and a neural network $\hat{F}^n: \mathbb{R}^n \rightarrow [0, 1]$ such that

$$\left(\int_{\mathcal{X}} |\hat{F}^n(g(T_1, K_1), \dots, g(T_n, K_n)) - F(g)|^p d\mu(g) \right)^{1/p} < \epsilon.$$

Proof. Fix $p \in [1, \infty)$ and $\epsilon > 0$. By Proposition 2.7 there exists $n \in \mathbb{N}$ such that

$$\int_{\mathcal{X}} |F^n(g(T_1, K_1), \dots, g(T_n, K_n)) - F(g)|^p d\mu(g) < \left(\frac{\epsilon}{2}\right)^p, \quad (17)$$

with F^n defined in (11). Let $\psi: \mathbb{R} \rightarrow [0, 1]$, invertible and Lipschitz continuous with Lipschitz constant $K > 0$. By the universal approximation theorem, see Hornik (1991), there exists a neural network $\phi^n: \mathbb{R}^n \rightarrow \mathbb{R}$ with identity as outer activation function and such that

$$\int_{\mathcal{X}} |\phi^n(g(T_1, K_1), \dots, g(T_n, K_n)) - (\psi^{-1} \circ F^n)(g(T_1, K_1), \dots, g(T_n, K_n))|^p d\mu(g) < \left(\frac{\epsilon}{2K}\right)^p.$$

Set now $\hat{F}^n := \psi \circ \phi^n$. Then we have

$$\begin{aligned} & \int_{\mathcal{X}} |\hat{F}^n(g(T_1, K_1), \dots, g(T_n, K_n)) - F^n(g(T_1, K_1), \dots, g(T_n, K_n))|^p d\mu(g) \\ & \leq K^p \int_{\mathcal{X}} |\phi^n(g(T_1, K_1), \dots, g(T_n, K_n)) - (\psi^{-1} \circ F^n)(g(T_1, K_1), \dots, g(T_n, K_n))|^p d\mu(g) < \left(\frac{\epsilon}{2}\right)^p. \end{aligned} \quad (18)$$

Therefore it holds

$$\begin{aligned} & \left(\int_{\mathcal{X}} |\hat{F}^n(g(T_1, K_1), \dots, g(T_n, K_n)) - F(g)|^p d\mu(g) \right)^{1/p} \\ & \leq \left(\int_{\mathcal{X}} |\hat{F}^n(g(T_1, K_1), \dots, g(T_n, K_n)) - F^n(g(T_1, K_1), \dots, g(T_n, K_n))|^p d\mu(g) \right)^{1/p} \\ & \quad + \left(\int_{\mathcal{X}} |F^n(g(T_1, K_1), \dots, g(T_n, K_n)) - F(g)|^p d\mu(g) \right)^{1/p} \\ & < \epsilon \end{aligned}$$

by (17) and (18). □

Remark 2.11. When $\alpha = 1$, we have

$$\lim_{K \rightarrow 0} C^X(T, K) = \lim_{K \rightarrow 0} \mathbb{E}^Q[(X_T - K)^+] + m(T) = X_0,$$

so the computations illustrated in the proof of Proposition 2.6 do not help detecting the strict local martingale property of X . In this case, we prove a neural network approximation focusing on local volatility models.

Remark 2.12. *The function F^n is defined on \mathbb{R}^n with n possibly very large. In addition to their universal approximation properties (see Hornik (1991)) neural networks have been shown to be particularly suitable for learning high-dimensional functions due to their ability to overcome the curse of dimensionality in many situations, see for example the seminal work Barron (1993) and Elbrächter et al. (2022), Grohs et al. (2018), Reisinger and Zhang (2020), Gonon and Schwab (2021), Gonon (2021) in the context of option pricing. These properties are a key motivation for our choice of neural networks to approximate F^n .*

2.4 The general case $\alpha \in [0, 1]$ in the setting of local volatility models

We now include the case $\alpha = 1$ in equation (1). In this setting, we focus on the class of local volatility models, i.e., stochastic processes $X = (X_t)_{t \geq 0}$ such that there exists a risk neutral measure $Q \sim P$ under which X has dynamics

$$dX_t = \sigma(t, X_t)dW_t, \quad t \geq 0, \quad X_0 = x > 0, \quad (19)$$

where $W = (W_t)_{t \geq 0}$ is a one-dimensional (\mathbb{F}, Q) -Brownian motion and the function σ satisfies the following conditions. In the sequel, *martingale* and *(strict) local martingale* are always meant with respect to the probability measure Q .

Assumption 2.13. *The function $\sigma : [0, \infty) \times (0, \infty) \rightarrow (0, \infty)$ is continuous. Moreover, it is locally Hölder continuous with exponent $1/2$ with respect to the second variable, i.e., for any $L > 0$ and any compact set $[0, L] \times [L^{-1}, L]$ there exists a constant $C \in [0, \infty)$ such that*

$$|\sigma(t, x) - \sigma(t, y)| \leq C|x - y|^{1/2}$$

for all $(t, x, y) \in [0, L] \times [L^{-1}, L]^2$. Furthermore, there exists a constant $A \in [0, \infty)$ such that $\sigma(t, x) \leq A$ for any $(t, x) \in [0, \infty) \times (0, 1]$.

Assumption 2.13 provides existence and uniqueness of a strong solution of (19), see Chapter IX.3 in Revuz and Yor (2013). In particular, we do not require the volatility function to be at most of linear growth in the spatial variable. As stated in the next result, which is Theorem 8 of Jarrow et al. (2011b), this allows X to be a strict local martingale.

Theorem 2.14. *Let $\sigma : [0, \infty) \times (0, \infty) \rightarrow (0, \infty)$ be a function satisfying Assumption 2.13. Assume that there exist two functions $\underline{\sigma}, \bar{\sigma} : (0, \infty) \rightarrow (0, \infty)$, which are continuous and locally Hölder continuous with exponent $1/2$, such that $\underline{\sigma}(x) \leq \sigma(t, x) \leq \bar{\sigma}(x)$ for any $(t, x) \in [0, \infty) \times (0, \infty)$. Let $X = (X_t)_{t \geq 0}$ be the process in (19). The following holds:*

1. *If $\int_c^\infty \frac{y}{\bar{\sigma}^2(y)} dy = \infty$ for every $c > 0$, then X is a martingale.*
2. *If there exists $c > 0$ such that $\int_c^\infty \frac{y}{\bar{\sigma}^2(y)} dy < \infty$, then X is a strict local martingale.*

We now give a generalization of Theorem 2.14 that is useful to provide the theoretical foundation of our methodology in this setting.

2.4.1 Strict local martingale property of local volatility models

Before stating the main result, we give two lemmas.

Lemma 2.15. *Fix $c > 0$. Let $\sigma : [0, \infty) \times (0, \infty) \rightarrow (0, \infty)$ be a function satisfying Assumption 2.13 and $X = (X_t)_{t \geq 0}$ be the process in (19). Then X is a strict local martingale if there exist s, T with $0 \leq s < T$ and a function $\underline{\sigma} : (0, \infty) \rightarrow (0, \infty)$, continuous and locally Hölder continuous with exponent $1/2$, such that:*

$$\int_c^\infty \frac{y}{\underline{\sigma}^2(y)} dy < \infty, \quad (20)$$

$$\underline{\sigma}(x) \leq \sigma(t, x) \text{ for any } (t, x) \in [s, T] \times (0, \infty). \quad (21)$$

Proof. Let $0 \leq s < T$ and $\underline{\sigma}$ as above. For $x \in (0, \infty)$, define the process $X^{\underline{\sigma}, T, s, x}$ to be the unique strong solution of

$$\begin{aligned} dX_t^{\underline{\sigma}, T, s, x} &= \underline{\sigma}(X_t^{\underline{\sigma}, T, s, x}) dW_t, \quad s \leq t \leq T, \\ X_s^{\underline{\sigma}, T, s, x} &= x. \end{aligned}$$

Lemma 9 of Jarrow et al. (2011b) together with the assumption in (21) implies that

$$\mathbb{E}^Q[X_T | X_s = x] \leq \mathbb{E}^Q[X_T^{\underline{\sigma}, T, s, x}] < x,$$

where the last inequality comes from the fact that the process $X^{\underline{\sigma}, T, s, x}$ is a strict local martingale for any $(s, x) \in [0, \infty) \times (0, \infty)$ by Theorem 7 of Jarrow et al. (2011b) and (20). \square

Lemma 2.16. *Fix $c > 0$ and $(r, x) \in [0, \infty) \times (0, \infty)$. Let $\sigma : [0, \infty) \times (0, \infty) \rightarrow (0, \infty)$ be a function satisfying Assumption 2.13. Denote by $X^{r, x} = (X_t^{r, x})_{t \geq r}$ the unique strong solution of the SDE*

$$\begin{aligned} dX_t^{r, x} &= \sigma(t, X_t^{r, x}) dW_t, \quad t \geq r, \\ X_r^{r, x} &= x. \end{aligned}$$

Then $X^{r, x}$ is a true martingale if for all $T > r$ there exists a function $\bar{\sigma}^{(T)} : (0, \infty) \rightarrow (0, \infty)$, continuous and locally Hölder continuous with exponent $1/2$, such that:

$$\int_c^\infty \frac{x}{(\bar{\sigma}^{(T)}(x))^2} dx = \infty, \quad (22)$$

$$\bar{\sigma}^{(T)}(x) \geq \sigma(t, x) \text{ for any } (t, x) \in [r, T] \times (0, \infty). \quad (23)$$

Proof. Consider s, T such that $T > s \geq r$, and $y \in (0, \infty)$. Define the process $X^{\bar{\sigma}^{(T)}, T, s, y}$ to be the unique strong solution of

$$\begin{aligned} dX_t^{\bar{\sigma}^{(T)}, T, s, y} &= \bar{\sigma}^{(T)}(X_t^{\bar{\sigma}^{(T)}, T, s, y}) dW_t, \quad s \leq t \leq T, \\ X_s^{\bar{\sigma}^{(T)}, T, s, y} &= y. \end{aligned}$$

Lemma 9 of Jarrow et al. (2011b) together with the assumption in (23) implies that

$$\mathbb{E}^Q[X_T^{r,x} | X_s^{r,x} = y] \geq \mathbb{E}^Q[X_T^{\bar{\sigma}, T, s, y}] = y,$$

where the last equality holds because the process $X^{\bar{\sigma}, T, s, x}$ is a martingale by Theorem 7 of Jarrow et al. (2011b) and (22). Since $X^{r,x}$ is a local martingale and then a supermartingale, we obtain $\mathbb{E}^Q[X_T^{r,x} | X_s^{r,x} = y] = y$. Therefore, $X^{r,x}$ is a martingale because s, T and y have been chosen arbitrarily. \square

In Theorem 2.18 we now provide a sufficient and necessary condition for X to be a true martingale in terms of the local volatility function σ . To this purpose we need requirements on σ which are more restrictive than in Assumption 2.13. However they are satisfied in the most commonly employed models in applications.

Assumption 2.17. *The function $\sigma : [0, \infty) \times (0, \infty) \rightarrow (0, \infty)$ satisfies Assumption 2.13 and there exist time points $(T_n)_{n \geq 0}$, with $T_0 = 0$ and $T_{n+1} > T_n$ for any $n \geq 0$, such that $\sigma(\cdot, x)$ restricted to any interval $[T_n, T_{n+1}]$, $n \geq 0$, is monotone increasing for all $x \in (0, \infty)$ or monotone decreasing for all $x \in (0, \infty)$.*

Note that Assumption 2.17 includes the case when σ is independent of time or is monotone increasing for all $x \in (0, \infty)$ or monotone decreasing for all $x \in (0, \infty)$ with respect to the time variable.

Theorem 2.18. *Fix $c > 0$, and let $\sigma : [0, \infty) \times (0, \infty) \rightarrow (0, \infty)$ satisfy Assumption 2.17. Let $X = (X_t)_{t \geq 0}$ be the process in (19). Then X is a true martingale if and only if*

$$\int_c^\infty \frac{x}{\sigma^2(T, x)} dx = \infty \text{ for any } T \in (0, \infty) \setminus B^\sigma, \quad (24)$$

with

$$B^\sigma := \{T > 0 : \sigma(\cdot, x) \text{ has a local maximum at } T \text{ for all } x \in (0, \infty)\}. \quad (25)$$

Proof. We start proving that if X is a martingale, then (24) holds. Suppose there exists $T \in (0, \infty) \setminus B^\sigma$ such that

$$\int_c^\infty \frac{x}{\sigma^2(T, x)} dx < \infty.$$

Assume first that $T \in (T_n, T_{n+1})$ for some $n \in \mathbb{N} \cup \{0\}$. Then $\underline{\sigma}(x) := \sigma(T, x)$ is a lower bound for σ on the interval $[T, T + \delta]$ (on the interval $[T - \delta, T]$) for some $\delta > 0$ if $\sigma(\cdot, x)|_{[T_n, T_{n+1}]}$ is monotone increasing (decreasing) for all $x \in (0, \infty)$. If $T = T_n$ for some $n \in \mathbb{N}$, then $\sigma(\cdot, x)|_{[T_{n-1}, T_n]}$ must be decreasing $\forall x \in (0, \infty)$ and $\sigma(\cdot, x)|_{[T_n, T_{n+1}]}$ increasing $\forall x \in (0, \infty)$, as $T \notin B^\sigma$. Hence, $\underline{\sigma}(x) := \sigma(T, x)$ is a lower bound on $[T_{n-1}, T_{n+1}]$. In any case, σ satisfies conditions (20) and (21), thus Lemma 2.15 implies that X is a strict local martingale.

We now prove that if (24) holds, then X is a martingale. For the sake of simplicity, we only consider the case when $\sigma(\cdot, x)|_{[0, T_1]}$ is increasing $\forall x \in (0, \infty)$ and $\sigma(\cdot, x)|_{[T_1, \infty)}$ is decreasing $\forall x \in (0, \infty)$. The proof can be analogously generalized.

Under this assumption, we have that (24) is equivalent to

$$\int_c^\infty \frac{x}{\sigma^2(T, x)} dx = \infty \text{ for any } T > 0, T \neq T_1. \quad (26)$$

We get

$$\mathbb{E}^Q[X_t | \mathcal{F}_s] = \mathbb{E}^Q[X_t^{s, X_s} | \mathcal{F}_s] = X_s, \quad T_1 < s < t, \quad (27)$$

where the last equality comes from Lemma 2.16, choosing $r = s$ and $\bar{\sigma}^{(T)}(x) := \sigma(s, x)$ for any $T > s$. We now show that the same holds for $s = T_1$. Theorem 2.2 of Ekström and Tysk (2012) states that the functions $C, P : [0, \infty) \times (0, \infty) \rightarrow \mathbb{R}^+$ defined by

$$C(T, K) := \mathbb{E}^Q[(X_T - K)^+], \quad P(T, K) := \mathbb{E}^Q[(K - X_T)^+], \quad (T, K) \in [0, \infty) \times (0, \infty),$$

are continuous with respect to T for any fixed $K > 0$. From this and from the put-call parity formula

$$C(T, K) = P(T, K) - K + \mathbb{E}^Q[X_T], \quad (T, K) \in [0, \infty) \times (0, \infty),$$

it follows that

$$\lim_{t \downarrow T_1} \mathbb{E}^Q[X_t] = \lim_{t \downarrow T_1} C(t, K) - \lim_{t \downarrow T_1} P(t, K) + K = C(T_1, K) - P(T_1, K) + K = \mathbb{E}^Q[X_{T_1}].$$

Since in our case $\mathbb{E}^Q[X_t]$ is constant for any $t > T_1$ by (27), we get $\mathbb{E}^Q[X_t] = \mathbb{E}^Q[X_{T_1}]$ for any $t > T_1$. This implies that

$$\mathbb{E}^Q[X_t | \mathcal{F}_{T_1}] = X_{T_1}, \quad T_1 < t, \quad (28)$$

because X is a supermartingale. Moreover, again by Lemma 2.16 and using that $\sigma(\cdot, x)|_{[0, T_1]}$ is increasing $\forall x \in (0, \infty)$, we get

$$\mathbb{E}^Q[X_t | \mathcal{F}_s] = X_s, \quad s < t < T_1. \quad (29)$$

Finally, we have that

$$\mathbb{E}^Q[X_t | \mathcal{F}_s] = \mathbb{E}^Q[\mathbb{E}^Q[X_t | \mathcal{F}_{T_1}] | \mathcal{F}_s] = \mathbb{E}^Q[X_{T_1} | \mathcal{F}_s] = X_s, \quad 0 < s < T_1 \leq t, \quad (30)$$

where the second equality follows by (28). The last equality holds because X is a supermartingale and $\mathbb{E}^Q[X_{T_1}] = \mathbb{E}^Q[X_s]$ for all $s \leq T_1$, since the expectation is continuous and constant by (29) for any $s < T_1$.

Putting together equations (27), (28), (29) and (30) we get that X is a true martingale. \square

The results above show that, under some conditions, the behaviour of the volatility for large values of the spatial variable determines if the process X is a strict local martingale, i.e., if it has bubbly dynamics.

In the sequel we use these findings as well as a modified version of Dupire's equation in Ekström and Tysk (2012) in order to prove the existence of a “bubble detection function” in local volatilities models.

2.4.2 Neural network approximation for local volatility models

We now provide a neural network approximation also for the case $\alpha = 1$. To this purpose we first state the general Dupire equation for call options for local volatility models, see Theorem 2.3 of Ekström and Tysk (2012).

Theorem 2.19. *Let X be the process in (19). The price $C^{\alpha,X}(T, K)$ in (1) of a call option on X with strike $K > 0$ and maturity $T > 0$ with collateral requirement $\alpha \in [0, 1]$ is the unique bounded classical solution $C^{\alpha,X} \in C^{1,2}((0, \infty) \times (0, \infty)) \cap C([0, \infty) \times [0, \infty))$ of the equation*

$$\begin{cases} \partial_T C^{\alpha,X} = \mathcal{L}C^{\alpha,X} - (1 - \alpha)\partial_T m & \text{for } (T, K) \in (0, \infty) \times (0, \infty) \\ C^{\alpha,X}(0, K) = (x - K)^+ \\ C^{\alpha,X}(T, 0) = \mathbb{E}^Q[X_T] + \alpha m(T), \end{cases} \quad (31)$$

where \mathcal{L} is the second order differential operator

$$\mathcal{L} = \frac{\sigma^2(T, K)}{2} \partial_{KK}$$

and m is the martingale defect of X introduced in (2).

Note that, since the martingale defect m is increasing with respect to time if X is a strict local martingale, then $(1 - \alpha)\partial_T m > 0$ for $\alpha < 1$.

In particular, since

$$C^{\alpha,X}(T, 0) = \mathbb{E}^Q[X_T] + \alpha (X_0 - \mathbb{E}^Q[X_T]) = \alpha X_0 + (1 - \alpha)\mathbb{E}^Q[X_T],$$

if $\alpha < 1$ we can write

$$\partial_T m(T) = \partial_T (X_0 - \mathbb{E}^Q[X_T]) = -\frac{\partial_T C^{\alpha,X}(T, 0)}{1 - \alpha}.$$

Therefore, for any $\alpha \in [0, 1]$ it holds

$$\partial_T C^{\alpha,X}(T, K) = \mathcal{L}C^{\alpha,X}(T, K) - (1 - \alpha)\partial_T m(T, K) = \frac{\sigma^2(T, K)}{2} \partial_{KK} C^{\alpha,X}(T, K) + \partial_T C^{\alpha,X}(T, 0),$$

from which we get

$$\sigma(T, K) = \sqrt{\frac{2\partial_T (C^{\alpha,X}(T, K) - C^{\alpha,X}(T, 0))}{\partial_{KK} C^{\alpha,X}(T, K)}} \quad (32)$$

if $\partial_{KK}C^{\alpha,X}(T, K) > 0$ for any $(T, K) \in [0, \infty) \times (0, \infty)$. Namely, the PDE in (31) allows to recover the function σ from prices of (possibly collateralized) call options. Together with Theorem 2.14, we use the PDE (31) and (32) for detecting asset price bubbles by looking at prices of call options, by means of neural network approximations. We formalize this in the following.

We now assume that call prices are given by (1) with $\alpha = 1$, which is the case that was not covered by the results in Section 2.3. We denote by C^X the price given in such a way, with underlying X defined by (19), where σ satisfies Assumption 2.17 and

$$\partial_{KK}C^X(T, K) > 0 \text{ for any } (T, K) \in [0, \infty) \times (0, \infty). \quad (33)$$

Since $C^X(T, 0) = X_0$, equation (32) becomes the classical Dupire formula

$$\sigma(T, K) = \sqrt{\frac{2\partial_T C^X(T, K)}{\partial_{KK}C^X(T, K)}}. \quad (34)$$

We prove an equivalent of Proposition 2.6 in the current setting.

Introduce the spaces

$$\mathcal{M}_{loc}^{locvol} = \{X = (X_t)_{t \geq 0} \text{ with dynamics (19), where } \sigma \text{ satisfies Assumption 2.17 and (33)}\},$$

and

$$\mathcal{M}_L^{locvol} = \{X \in \mathcal{M}_{loc}^{locvol}, X \text{ strict local martingale}\}.$$

Consider the set

$$\mathcal{X}^{locvol} := \{C^X : X \in \mathcal{M}_{loc}^{locvol}\} \subset C^{1,2}([0, \infty) \times (0, \infty), (0, \infty)) \quad (35)$$

endowed with the Borel sigma-algebra of the Fréchet space $C^{1,2}([0, \infty) \times (0, \infty), (0, \infty))$ with topology induced by the countable family of seminorms

$$\|f\|_{i,j,n} := \sup \left\{ \left| \partial_{t,x}^{(i,j)} f(t, x) \right|, (t, x) \in [0, n] \times [1/n, n] \right\}, \quad (36)$$

for $n \in \mathbb{N}$ and $(i, j) = (0, 0), (0, 1), (0, 2), (1, 0)$.

Proposition 2.20. *Under Assumption 2.17 there exists a measurable function $F: \mathcal{X}^{locvol} \rightarrow \{0, 1\}$ such that $F(C^X) = \mathbb{1}_{\{X \in \mathcal{M}_L^{locvol}\}}$. In particular,*

$$F(f) = \mathbb{1} \left\{ \text{there exists } t \in (0, \infty) \setminus B^{(\partial_t f, \partial_{xx} f)} \text{ such that } \int_c^\infty x \frac{\partial_{xx} f(t, x)}{\partial_t f(t, x)} dx < \infty \right\}, \quad f \in \mathcal{X}^{locvol}, \quad (37)$$

for arbitrary fixed $c > 0$ and with

$$B^{(\partial_t f, \partial_{xx} f)} := \left\{ T > 0 : \sigma := \sqrt{\frac{2\partial_t f}{\partial_{xx} f}} \text{ has a local maximum at } T \text{ for all } x \in (0, \infty) \right\}. \quad (38)$$

Proof. We have that $\mathcal{X}^{locvol} \subset C^{1,2}([0, \infty) \times (0, \infty), (0, \infty))$ by Theorem 2.19. Fix $c > 0$, and consider $X \in \mathcal{M}_{loc}^{locvol}$ with $C^X \in \mathcal{X}^{locvol}$. Theorem 2.18 implies that for any process X of the form (19) where the local volatility function σ satisfies Assumption 2.17 it holds

$$X \in \mathcal{M}_L^{locvol} \text{ if and only if } \exists t \in (0, \infty) \setminus B^\sigma : \int_c^\infty \frac{x}{\sigma^2(t, x)} dx < \infty, \quad (39)$$

with B^σ in (25). Consider now $F: \mathcal{X}^{locvol} \rightarrow \{0, 1\}$ as in (37). From (34) and (39), it is clear that $F(C^X) = \mathbb{1}_{\{X \in \mathcal{M}_L^{locvol}\}}$ for any $C^X \in \mathcal{X}^{locvol}$. We now prove that F is measurable. Fix a volatility function σ satisfying Assumption 2.17. First note that, if

$$D^{\sigma, \mathbb{Q}} := \left(\left\{ t > 0 : \int_c^\infty \frac{x}{\sigma^2(t, x)} dx < \infty \right\} \cap \mathbb{Q} \right) \setminus B^\sigma = \emptyset, \quad (40)$$

then

$$D^\sigma := \left\{ t > 0 : \int_c^\infty \frac{x}{\sigma^2(t, x)} dx < \infty \right\} \setminus B^\sigma = \emptyset.$$

Indeed, suppose $D^\sigma \neq \emptyset$ and consider $t \in D^\sigma$. Then either t is a local minimum for $\sigma(\cdot, x)$ for all $x \in (0, \infty)$ or there exists an interval I_t with $t \in I_t$ on which $\sigma(\cdot, x)$ is monotone increasing for all $x \in (0, \infty)$ or monotone decreasing for all $x \in (0, \infty)$: in all cases, it is possible to find a positive $\bar{t} \in \mathbb{Q} \setminus B^\sigma$ such that $\int_c^\infty \frac{x}{\sigma^2(\bar{t}, x)} dx < \infty$, hence $D^{\sigma, \mathbb{Q}} \neq \emptyset$.

Then we can write

$$F(f) = 1 - \prod_{t \in \mathbb{Q} \cap ((0, \infty) \setminus B(\partial_t f, \partial_{xx} f))} (1 - F_t(f)) = 1 - \prod_{t \in \mathbb{Q} \cap (0, \infty)} (1 - F_t(f) \mathbb{1}_{\{t \notin B(\partial_t f, \partial_{xx} f)\}}), \quad (41)$$

$f \in \mathcal{X}^{locvol}$, with $B(\partial_t f, \partial_{xx} f)$ as in (38) and where $F_t: \mathcal{X}^{locvol} \rightarrow \{0, 1\}$ is defined by

$$F_t(f) := \mathbb{1}_{\left\{ \int_c^\infty x \frac{\partial_{xx} f(t, x)}{\partial_t f(t, x)} dx < \infty \right\}}, \quad f \in \mathcal{X}^{locvol}, \quad (42)$$

for any $t > 0$. Note that for any $t \in \mathbb{Q} \cap (0, \infty)$ and any $\delta \in (0, t)$ we have

$$\mathbb{1}_{\{t \notin B(\partial_t f, \partial_{xx} f)\}} = 1 - \prod_{x \in \mathbb{Q} \cap (0, \infty)} \prod_{s \in \mathbb{Q} \cap (0, \delta)} \mathbb{1}_{\left\{ \frac{\partial_t f(t-s, x)}{\partial_{xx} f(t-s, x)} < \frac{\partial_t f(t, x)}{\partial_{xx} f(t, x)} \right\}} \mathbb{1}_{\left\{ \frac{\partial_t f(t+s, x)}{\partial_{xx} f(t+s, x)} < \frac{\partial_t f(t, x)}{\partial_{xx} f(t, x)} \right\}},$$

which is measurable. Moreover, we can write $F_t = \hat{F} \circ \tilde{F}_t$, where the operator $\tilde{F}_t: \mathcal{X}^{locvol} \rightarrow C((0, \infty), (0, \infty))$ is given by

$$\tilde{F}_t(f) := I_d \cdot \frac{\partial_{xx} f(t, \cdot)}{\partial_t f(t, \cdot)}$$

for any $t > 0$, with $I_d(x) = x$, and $\hat{F}: C((0, \infty), (0, \infty)) \rightarrow \{0, 1\}$ equal to

$$\hat{F}(g) = \mathbb{1}_{\left\{ \int_c^\infty g(x) dx < \infty \right\}}.$$

The space $C((0, \infty), (0, \infty))$ is a Fréchet space equipped with the Borel sigma-algebra with respect to the topology induced by the family of seminorms

$$\|g\|_m := \sup \{|g(x)|, x \in [1/m, m]\} \quad (43)$$

for $m \in \mathbb{N}$. Note that \tilde{F}_t is well defined for any $t > 0$ by (33) and by (34), since σ is strictly positive. Moreover, \tilde{F}_t is measurable for all $t > 0$ because it is continuous with respect to the uniform convergence on compact sets away from zero induced by the topologies of the Fréchet spaces $C^{1,2}([0, \infty) \times (0, \infty), (0, \infty))$ on \mathcal{X}^{locvol} and $C((0, \infty), (0, \infty))$ with the seminorms in (36) and (43), respectively.

Coming now to \hat{F} , it can be written as

$$\hat{F}(g) = \mathbb{1}_{\{I(g) < \infty\}}$$

where $I : C((0, \infty), (0, \infty)) \rightarrow [0, \infty]$ is the integral operator

$$I(g) = \int_c^\infty g(x)dx, \quad g \in C((0, \infty), (0, \infty)).$$

Clearly, $I = \lim_{k \rightarrow \infty} I_k$, where the functions $I_k : C((0, \infty), (0, \infty)) \rightarrow [0, \infty)$ with

$$I_k(g) := \int_c^k g(x)dx, \quad g \in C((0, \infty), (0, \infty)), \quad k \in \mathbb{N},$$

are measurable because the operator $g \rightarrow \int_c^k g(x)dx$ is continuous for any finite k . Then I is measurable and \hat{F} as well.

The function F_t in (42) is thus measurable for every $t > 0$ since it is the composition of measurable functions. Hence, the function F is measurable by (41). □

Remark 2.21. *In Proposition 2.20, we focus on $\alpha = 1$ because the case $\alpha < 1$ has already been covered in Section 2.3. Alternatively, if $\alpha < 1$, using equation (32) we can define a new detection function $F : \mathcal{X}^{locvol} \rightarrow \{0, 1\}$ similar to (37) whose measurability can be proved using the same arguments as in the proof of Proposition 2.20.*

Remark 2.22. *The assumption that $\partial_{KK}C^X(T, K) > 0$ for any $(T, K) \in [0, \infty) \times (0, \infty)$ in (33) is equivalent to*

$$\partial_T C^X(T, K) > 0, \quad (T, K) \in [0, \infty) \times (0, \infty),$$

which is usually satisfied for the observed market prices.

The above conditions are satisfied, for instance, if X_T admits a strictly positive density $g_T^X : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ for any $T \geq 0$. Indeed, under this assumption it holds that

$$\partial_{KK}C^X(T, K) = \partial_{KK} \int_K^\infty (x - K)g_T^X(x)dx = g_T^X(K) > 0$$

for any $K > 0$.

We now give a result analogous to Proposition 2.7. The proof follows the same steps.

Proposition 2.23. *Define the set*

$$A = [0, \infty) \times [C, \infty)$$

for some fixed $C > 0$. Let $(T_n, K_n)_{n \in \mathbb{N}}$ be a sequence of maturities and strikes such that the set $\{(T_n, K_n), n \in \mathbb{N}\}$ is dense in A . Let μ be a probability measure on the space \mathcal{X}^{locvol} introduced in (35) such that there exists a probability measure ν on \mathcal{X}^{locvol} with

$$\nu \left(\left\{ f \in \mathcal{X}^{locvol} \text{ such that } f(T_i, K_i) = g(T_i, K_i) \text{ for all } i \in \mathbb{N} \right\} \right) > 0 \quad \text{for any } g \in \text{supp}(\mu). \quad (44)$$

Then there exists a sequence of functions $(F^n)_{n \in \mathbb{N}}$, $F^n: \mathbb{R}^n \rightarrow [0, 1]$ such that

$$\int_{\mathcal{X}^{locvol}} |F^n(g(T_1, K_1), \dots, g(T_n, K_n)) - F(g)|^p d\mu(g) \xrightarrow{n \rightarrow \infty} 0,$$

for any $p \in [1, \infty)$, where $F: \mathcal{X}^{locvol} \rightarrow \{0, 1\}$ is the function introduced in (37). For fixed $n \in \mathbb{N}$, the function F^n can be chosen as follows:

$$F^n(c^n) := \begin{cases} \frac{1}{\nu(S_n(c^n))} \int_{S_n(c^n)} F(f) d\nu(f) & \text{if } \nu(S_n(c^n)) > 0, \\ 0 & \text{otherwise,} \end{cases}$$

$c^n \in \mathbb{R}^n$, where

$$S_n(c^n) := \left\{ f \in \mathcal{X}^{locvol} \text{ such that } f(T_i, K_i) = c_i^n \text{ for all } i = 1, \dots, n \right\}, \quad c^n \in \mathbb{R}^n. \quad (45)$$

We then have the following result. The proof follows the same steps as for Theorem 2.10.

Theorem 2.24. *Fix $p \in [1, \infty)$ and let $(T_n, K_n)_{n \in \mathbb{N}}$ be the sequence of maturities and strikes introduced in Proposition 2.23. Also let μ be a probability measure on \mathcal{X}^{locvol} such that there exists a probability measure ν on \mathcal{X} which satisfies (44). Then for any $\epsilon > 0$ there exists an $n \in \mathbb{N}$ and a neural network $\hat{F}^n: \mathbb{R}^n \rightarrow [0, 1]$ such that*

$$\int_{\mathcal{X}^{locvol}} |\hat{F}^n(g(T_1, K_1), \dots, g(T_n, K_n)) - F(g)|^p d\mu(g) < \epsilon.$$

3 Numerical experiments

3.1 The methodology

In Section 2 we have provided a theoretical foundation for the bubble detection methodology explained in the introduction. The method does not require any parameter estimation related to the asset price, and it is model independent when $\alpha < 1$. Fixing $\alpha = 1$, we now show via numerical experiments that our approach works not only within a class of models, but also if the network is trained using a certain class of stochastic processes (like local volatility models) and tested within another class (for example, stochastic volatility models). In Section

4, we also test the method with market data associated to assets involved in the new tech bubble burst at the beginning of 2022 (see among others Bercovici (2017), Ozimek (2017), Serla (2017), Sharma (2017)) finding a close match between the output of the network and the expected results.

Remark 3.1. *We show the results for a collateral constant $\alpha = 1$ because this is the most common choice in the literature, see for example Dias et al. (2020), Heston et al. (2007) and Jarrow et al. (2010). For completeness, we have repeated the same tests for $\alpha = \frac{1}{2}$ finding not inferior accuracies.*

Remark 3.2. *The theoretical results proved in Section 2 provide an important justification for our approach: they show that in our setting bubbles can indeed be detected using a neural network with call option price inputs up to an error that tends to 0 as the number of call option prices and the size of the neural network become arbitrarily large. The functional form of the bubble indicator F in (6) and (37) is not explicitly used in the numerical algorithm, but instead an approximation of F is learnt from data. An alternative approach for bubble detection could aim to directly approximate the functional form of F in (6) and (37) via an extrapolation of the calibrated local volatility function as in Jarrow et al. (2011b). One of the key features of our approach is that no choice of extrapolation needs to be made explicitly, but instead learning is only based on the available call option prices.*

For the sake of clarity, before showing our results we specify how our methodology works:

- We first fix a constant $x_0 \in \mathbb{R}$ and a set of strikes and maturities $(T_j, K_j)_{1 \leq j \leq m}$, $m \in \mathbb{N}$.
- We take a set of local martingales $(X^i)_{1 \leq i \leq n_{train}}$, $n_{train} \in \mathbb{N}$, with possibly different dynamics and such that $X_0^i = x_0$ for any $1 \leq i \leq n_{train}$. Half of them are true martingales, half strict local martingales.
- We consider a binary classification problem with m -dimensional input and output in $\{0, 1\}$. Each of the n_{train} input data points consists of the call option prices with underlying X^i for strikes and maturities $(T_j, K_j)_{1 \leq j \leq m}$. The label associated to data point i is the indicator $\mathbb{1}_{\{X^i \text{ is a strict local martingale}\}}$. The call option prices are computed either via Monte-Carlo techniques or, when possible, analytically.
- Based on these data points we train a neural network in order to make it learn the indicator function from the prices. Since this is a supervised learning, binary classification problem, we choose binary cross entropy loss function and Sigmoid activation function for the output layer.

We test the network on a set of new local martingales $(Y_i)_{1 \leq i \leq n_{test}}$, $n_{test} \in \mathbb{N}$, with $Y_0^i = x_0$ for any $1 \leq i \leq n_{test}$, again with the property that half of them are strict local martingales

and half true martingales. In particular, our choice for the output layer’s activation function implies that, for any stock, the network returns a number $p \in (0, 1)$ which can be interpreted as the probability that the underlying under consideration is a strict local martingale. If $p > 0.5$, we label the underlying to be a strict local martingale according to the network. As performance metrics, both on the training and test set, we consider *accuracy* a and F_1 *score*, defined as

$$a := \frac{\text{true positives} + \text{true negatives}}{\text{true positives} + \text{true negatives} + \text{false positives} + \text{false negatives}}$$

and

$$F_1 := 2 \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}},$$

respectively, where

$$\text{precision} := \frac{\text{true positives}}{\text{true positives} + \text{false positives}}, \quad \text{recall} := \frac{\text{true positives}}{\text{true positives} + \text{false negatives}}.$$

In our framework, *true positives* (*true negatives*) are strict local martingales (true martingales) correctly labelled as strict local martingales (true martingales), and *false positives* (*false negatives*) are true martingales (strict local martingales) incorrectly labelled as strict local martingales (true martingales).

For both accuracy and F_1 score, and for all the cases under investigation, we compute the average over 10 runs of the network’s training, for fixed training and test data ¹.

Unless differently specified, we consider a neural network having a first layer with 60 nodes, two hidden layers with 30 and 7 nodes, respectively, and ReLu activation function for the input and hidden layers. We use an Adam optimization algorithm with learning rate equal to 0.001 and 20 epochs. These choices strike a balance between accuracy and efficiency, enabling the algorithm to yield high performance indicators in a significant number of cases without excessive computational overhead. In Section 3.3, where the stochastic models in the training and test set are significantly different, we find that other architectures and choices of the learning rate deliver better performances in some cases.

3.2 Local volatility models

First, we consider a displaced CEV process $(X_t)_{t \geq 0}$ with dynamics

$$dX_t = \sigma \cdot (X_t + d)^\beta dW_t, \quad t \geq 0, \quad X_0 = x_0 > 0, \quad (46)$$

where $(W_t)_{t \geq 0}$ is a one-dimensional standard Brownian motion and $\sigma, \beta, d > 0$.

Proposition 3.3. *The SDE (46) has a unique strong solution, which is a true martingale if and only if $\beta \leq 1$.*

¹Since the training of the network is not deterministic, different outputs might be obtained for different runs.

Proof. When $d = 0$, the SDE (46) admits a unique strong solution since the function $\sigma(x) = x^\beta$, $x > 0$, satisfies Assumption 2.13. A process X with dynamics given by (46) for general $d > 0$ is of the form $X_t = Y_t - d$, $t \geq 0$, where Y is the unique strong solution of (46) with $d = 0$ and $Y_0 = X_0 + d$. By Theorem 2.18, Y is a true martingale if and only if $\beta \leq 1$, and thus the same holds for X . \square

The proof of the above proposition shows that all the results of Section 2 also apply to (46), even if X is in general not positive but only bounded from below. Based on this result, we proceed to show our numerical experiments in Sections 3.2.1 and 3.2.2.

3.2.1 Random displacement

As a first experiment, we work with $n_{train} \in \mathbb{N}$ different underlyings, all following (46) with fixed $\sigma > 0$ and different, randomly chosen parameters d and β . In particular, for any $i = 1, \dots, n_{train}$, the underlying X^i has displacement d_i uniformly distributed in an interval $[0, D]$ for given $D > 0$ and strictly positive exponent β_i . We let β_i be uniformly distributed in an interval $(a, 1]$ with $a \in (0, 1)$ in $n_{train}/2$ cases and in an interval $(1, A)$ with $A > 1$ for the other $n_{train}/2$ cases, so that we consider $n_{train}/2$ true martingales and $n_{train}/2$ strict local martingales. We also fix strikes $0 < K_1 < \dots < K_m$ and maturities $0 < T_1 < \dots < T_l$ for $l, m \in \mathbb{N}$. Note that the values of the calls can be computed analytically, as indicated for example in Lindsay and Brecher (2012).

We choose $n_{train} = 100000$, $n_{test} = 5000$, $a = 0.2$, $A = 2$, $D = 0.2$, $\sigma = 1$, $X_0 = 2$, 100 equally spaced maturities from $T_1 = 1$ to $T_l = 2$. We consider two sets of 51 equally spaced strikes: one from $K_1 = 1.8$ to $K_m = 2.3$ (“at-the-money strikes”) and one from $K_1 = 3.5$ to $K_m = 4$ (“out-of-the-money strikes”). For strikes in $[1.8, 2.3]$, we obtain:

- Accuracies of 99.1% for both the training and test data;
- F_1 scores of 0.992 and 0.991 for the training and test data, respectively;

For strikes in $[3.5, 4]$, we obtain:

- Accuracies of 99.5% and 99.4% for the training and test data, respectively;
- F_1 scores of 0.995 and 0.994 for the training and test data, respectively.

As a further experiment, we give the network the same matrix as before but where the last element of the i -th row is now given by the martingale defect of X^i , $1 \leq i \leq n_{train}$. In this case, the network has to learn the martingale defect from the option prices. We choose linear activation function for the output layer and the R^2 coefficient of determination score of the computation of the martingale defect as a performance indicator. This is given by

$$R^2 := 1 - \frac{\sum_{i=1}^{n_{test}} (y_i - f_i)^2}{\sum_{i=1}^{n_{test}} (y_i - \bar{y})^2},$$

where y_i and f_i , $i = 1, 2, \dots, n$, are the true and predicted values, respectively, and \bar{y} is the average of the true values. The best possible score is therefore 1. In our case, on the test set we get an R^2 coefficient of determination score of 0.994 for strikes in $[1.8, 2.3]$ and of 0.999 for strikes in $[3.5, 4]$.

All the metrics show better results for strikes out-of-the-money. A possible explanation is the following: looking at equation (1) we note that the price of a call option in presence of a bubble diverges from the one without bubble by a term proportional to the martingale defect $m(T) = X_0 - \mathbb{E}^Q[X_T]$, and this term has a big impact on the price of the option when $C(T, K)$ in (1) is small, i.e., for large values of the strike K .

3.2.2 Train with no displacement, test with a strictly positive displacement

We now repeat the experiment described in Section 3.2.1 with the difference that this time the displacements are not random, and they differ from training and testing. In particular, we train and test the network with data generated by an underlying in (46) with $d = 0$ and $d = 0.5$, respectively. In this way, we can check if the network is able to detect bubbles also when it is trained and tested with data coming from processes having shifted dynamics. All the other parameters are the same as in Section 3.2.1. For strikes in $[1.8, 2.3]$, we obtain:

- Accuracies of 99.9% and 86.9% for the training and test data, respectively;
- F_1 scores of 0.999 and 0.883 for the training and test data, respectively;

For strikes in $[3.5, 4]$ we get:

- Accuracies of 99.9% and 91.1% for the training and test data, respectively;
- F_1 scores of 0.999 and 0.918 for the training and test data, respectively;

As before, strikes out-of-the-money yield better results. However, we observe smaller accuracies and F_1 -score in the test set than in the first case: this can be expected because of the shifted dynamics of the underlying processes that produce the prices to train and test the network, respectively. Note that the displacement we use to test the network is bigger than the biggest possible displacement in the first experiment.

Very similar results are observed if we instead train the network with displacement $d = 0.5$ and test with $d = 0$.

In all the experiments, the accuracy of the results gets smaller if we decrease the number of strikes or the number of maturities, keeping the intervals the same.

3.3 Stochastic volatility models

As an extension of the setting of Section 3.2, we now include stochastic volatility models into our numerical experiments. In particular, we check how our methodology works when the

network is fed with data coming from stochastic volatility processes and tested looking at local volatility models, and vice-versa. Since the nature of the two classes of models is deeply different, this analysis constitutes a robust test for our methodology.

In particular, we consider the SABR model and the class of stochastic volatility models introduced in Sin (1998). The SABR volatility model $F = (F_t)_{t \geq 0}$ is defined as the unique strong solution of the SDE

$$dF_t = \sigma_t(F_t)^\gamma dW_t, \quad t \geq 0, \quad (47)$$

$$d\sigma_t = \alpha\sigma_t dZ_t, \quad t \geq 0, \quad (48)$$

with initial value $F_0 > 0$, initial volatility $\sigma_0 > 0$, exponent $\gamma > 0$ and volatility of the volatility $\alpha > 0$. Here W and Z are Brownian motions with correlation $\rho \in [-1, 1]$.

A class of stochastic volatility models of the form

$$dY_t = \sigma_1 v_t^\alpha Y_t dB_t^1 + \sigma_2 v_t^\alpha Y_t dB_t^2, \quad t \geq 0, \quad (49)$$

$$dv_t = a_1 v_t dB_t^1 + a_2 v_t dB_t^2 + \kappa(L - v_t)dt, \quad t \geq 0, \quad v_0 = 1, \quad (50)$$

is considered in Sin (1998), where $\alpha, \kappa, L \in \mathbb{R}^+$, $\sigma_1, \sigma_2, a_1, a_2 \in \mathbb{R}$ and (B_1, B_2) is a two-dimensional Brownian motion. It is well known that Y is a strict local martingale if and only if $a_1\sigma_1 + a_2\sigma_2 > 0$, see Sin (1998). From now on, we call (49)-(50) the Sin model. It can be noted that, when $\gamma = 1$, the SABR model is a particular case of the Sin model, and is a strict local martingale if and only if $\rho > 0$.

The call prices for the SABR model admit an analytical approximation², see for example Piiroinen et al. (2018). For the Sin model, on the other hand, we are not aware of any way to analytically compute or approximate the prices, so we rely on Monte-Carlo approximations with 200000 simulations and time step 0.01.

We take 61 maturities equally spaced from $T_1 = 2$ to $T_l = 5$ and, as in Section 3.2, two groups of 100 equally spaced strikes: one from $K_1 = 1$ to $K_m = 3.5$ and one from $K_1 = 3$ to $K_m = 5.5$. We consider three different classes of processes for the underlying, in all cases with initial value $X_0 = 2$:

- A set of displaced CEV models with same specifications as in Section 3.2, that is, in particular, with randomly varying exponent and displacement;
- A set of SABR models (47)-(48) with parameters $\alpha = 0.5$, $\gamma = 1$ and $\sigma_0 = 0.5$. We let the correlation ρ between the process and the volatility be randomly chosen and a priori different for any underlying: in particular, we set it to be uniformly distributed in $[-0.8, 0]$ in half of the cases (the ones where the underlying is a true martingale) and in $(0, 0.8]$ in the other half of the cases (the ones where the underlying is a strict local martingale).

²We compute these approximations using the library available at https://github.com/google/tf-quant-finance/tree/master/tf_quant_finance/models/sabr..

- A set of Sin models (49)-(50) with parameters $v_0 = 0.5$, $\alpha = 1$, $\kappa = 0$, $\sigma_2 = -0.5$, $a_1 = 1.8$ and $a_2 = 1.2$. We let σ_1 be randomly chosen and a priori different for any underlying: specifically, it is uniformly distributed in $[0, 1/3]$ in half of the cases (the ones where the underlying is a true martingale) and in $(1/3, 1]$ in the other half of the cases (the ones where the underlying is a strict local martingale).

For all the experiments, we consider 50000 underlyings for the training set and 5000 for the test set, and investigate accuracy and F_1 score of our methodology.

As a benchmark for our following investigation, for which training and test sets will consist of prices generated by underlying coming from different classes of models, Table 1 shows the accuracy and the F_1 score in the training and test phases when we both train and test the neural network with any of the three models above. Half of the them are true martingales and half strict local martingales.

	Displaced CEV	Sin	SABR
Strikes in $[1, 3.5]$	$a(\text{train}) = 99.8\%$, $a(\text{test}) = 98.1\%$ $F_1(\text{train}) = 0.988$ $F_1(\text{test}) = 0.927$	$a(\text{train}) = 99.8\%$, $a(\text{test}) = 99.6\%$ $F_1(\text{train}) = 0.995$ $F_1(\text{test}) = 0.934$	$a(\text{train}) = 99.4\%$, $a(\text{test}) = 94.8\%$ $F_1(\text{train}) = 0.997$ $F_1(\text{test}) = 0.876$
Strikes in $[3, 5.5]$	$a(\text{train}) = 99.9\%$ $a(\text{test}) = 99.2\%$ $F_1(\text{train}) = 0.991$ $F_1(\text{test}) = 0.941$	$a(\text{train}) = 99.8\%$, $a(\text{test}) = 99.7\%$ $F_1(\text{train}) = 0.998$ $F_1(\text{test}) = 0.947$	$a(\text{train}) = 99.3\%$ $a(\text{test}) = 98.4\%$ $F_1(\text{train}) = 0.995$ $F_1(\text{test}) = 0.914$

Table 1: Accuracy ($a(\text{train})$ and $a(\text{test})$) and F_1 score ($F_1(\text{train})$ and $F_1(\text{test})$) on the training set and on the test set when the algorithm to identify bubbles is both trained and tested with one single family of models: the three investigated cases are displaced CEV, Sin model and SABR model.

We now mix the models for training and testing. Since training and test data are now drawn from very different distributions, in some cases using the same network architecture and learning rate specified in Section 3.1 is not optimal, and better results can be obtained with different specifications. In particular, we let the learning rate λ vary in $\{0.01, 0.001, 0.0001\}$, and also consider a bigger network constituted by an input layer with 100 nodes, two hidden layers with 70 and 30 nodes, and possibly $L1$ or $L2$ -regularization with parameter in the set $\{0.1, 0.01, 0.001\}$. For some choices of the training and test data, all reported in Tables 2 and 3, we find that changing the learning rate or taking the bigger, regularized network lead to bigger accuracy and F_1 score in the test set. All the results shown in Tables 2 and 3 are relative to the best choice of network and learning rate. When the choice of the network and of the learning rate is different from the specifications in Section 3.1, we indicate it.

In particular, when we write “Big, Reg” this means that we use the bigger network with L_2 -regularization with parameter 0.01: this turns out to be the best regularization choice in terms of performance in the test set in all cases when regularization is required. Looking at Tables 2 and 3, it can be noted that only two variants of the setting described in Section 3.1 are needed to always get the highest test accuracy and F_1 score: the bigger, regularized network and the “default” network with smaller learning rate.

First, to investigate the performance of our method when trained on stochastic volatility processes and tested on local volatility models, we test the algorithm looking at the prices generated by the set of displaced CEV models, for three different training sets: one constituted exclusively from data coming from the SABR model, one only from the Sin model and one from both of them. When both the Sin and the SABR models are used for training, the training set is equally split between the two cases. Table 2 shows the accuracy and the F_1 score we get on the training set and on the test set in the different cases.

	Only SABR	Only Sin (Big, Reg)	SABR and Sin (Big, Reg)
Strikes in $[1, 3.5]$	$a(\text{train}) = 99.8\%$ $a(\text{test}) = 68.7\%$ $F_1(\text{train}) = 0.998$ $F_1(\text{test}) = 0.761$	$a(\text{train}) = 99.7\%$ $a(\text{test}) = 82.8\%$ $F_1(\text{train}) = 0.997$ $F_1(\text{test}) = 0.790$	$a(\text{train}) = 98.7\%$ $a(\text{test}) = 95.2\%$ $F_1(\text{train}) = 0.986$ $F_1(\text{test}) = 0.955$
Strikes in $[3, 5.5]$	$a(\text{train}) = 99.9\%$ $a(\text{test}) = 86.2\%$ $F_1(\text{train}) = 0.999$ $F_1(\text{test}) = 0.875$	$a(\text{train}) = 99.9\%$ $a(\text{test}) = 86.3\%$ $F_1(\text{train}) = 0.999$ $F_1(\text{test}) = 0.841$	$a(\text{train}) = 97.3\%$ $a(\text{test}) = 86.7\%$ $F_1(\text{train}) = 0.977$ $F_1(\text{test}) = 0.879$

Table 2: Accuracy ($a(\text{train})$ and $a(\text{test})$) and F_1 score ($F_1(\text{train})$ and $F_1(\text{test})$) on the training set and on the test set when the algorithm to identify bubbles is tested looking at call prices generated by displaced CEV models, for two different sets of strikes and three different training sets.

We see that looking at prices of out-of-the-money options generally improves the performance, as in Section 3.2, apart from the case when the network is trained with prices generated by both SABR and Sin processes. Moreover, splitting the training set into the two sets of SABR and Sin processes results in an higher accuracy and F_1 score on the test set with respect to both cases, when we train with only one set of processes, at the expenses of the performance on the training set.

We then move to test our methodology looking at the prices generated by the class of SABR models, again for three different training sets: the first two sets consist of call prices generated by the displaced CEV models and by the Sin models, respectively, and the third one of prices generated by both of them. In this way, we complete our investigation: the first case regards training with local volatility and testing with stochastic volatility, the second one

concerns both training and testing with stochastic volatility, whereas the third one studies if a combination of training from local and stochastic volatility can be beneficial when looking at a stochastic volatility model. The results are summarized in Table 3.

	Only disp. CEV (Big, reg)	Only Sin	Disp. CEV and Sin ($\lambda = 0.0001$)
Strikes in $[1, 3.5]$	$a(\text{train}) = 99.8\%$ $a(\text{test}) = 50.0\%$ $F_1(\text{train}) = 0.998$ $F_1(\text{test}) = 0$	$a(\text{train}) = 99.9\%$ $a(\text{test}) = 73.0\%$ $F_1(\text{train}) = 0.999$ $F_1(\text{test}) = 0.787$	$a(\text{train}) = 99.8\%$ $a(\text{test}) = 88.5\%$ $F_1(\text{train}) = 0.997$ $F_1(\text{test}) = 0.898$
Strikes in $[3, 5.5]$	$a(\text{train}) = 98.1\%$ $a(\text{test}) = 79.2\%$ $F_1(\text{train}) = 0.982$ $F_1(\text{test}) = 0.736$	$a(\text{train}) = 99.4\%$ $a(\text{test}) = 93.4\%$ $F_1(\text{train}) = 0.997$ $F_1(\text{test}) = 0.930$	$a(\text{train}) = 99.7\%$ $a(\text{test}) = 96.0\%$ $F_1(\text{train}) = 0.997$ $F_1(\text{test}) = 0.959$

Table 3: Accuracy ($a(\text{train})$ and $a(\text{test})$) and F_1 score ($F_1(\text{train})$ and $F_1(\text{test})$) on the training set and on the test set when the algorithm to identify bubbles is tested looking at data coming from the SABR model, for two different sets of strikes and three different training sets.

Looking at bigger strikes we get bigger accuracy and F_1 -score in the test set in all cases. Moreover, for both the groups of strikes, training with a combination of data coming from a stochastic volatility model and from a local volatility model increases both accuracy and F_1 -score in the test set for the SABR model at the expenses of the accuracy in the training set: this suggests again that training with higher variety of data makes our methodology more robust. On the other hand, for the at-the-money-strikes, training only with data coming from the displaced CEV processes does not help guessing the strict local martingale property of the SABR processes: in particular, the algorithm always identifies all the processes as true martingales. Comparing this case with Table 2 for this choice of strikes we see that our methodology works fairly good if we train with the SABR models and try to label the CEV models, but not vice-versa.

In the Appendix we show the Receiver Operating Characteristic (ROC) curves for all the cases taken into consideration in Tables 2 and 3. ROC curves plot the true positive rate (TPR) against the false positive rate (FPR) at various values of the acceptance threshold, see for example Carter et al. (2016), Fan et al. (2006), Gonçalves et al. (2014) for a deeper analysis. In particular, Figures 6 and 7 display the ROC curves for the at-the-money and the out-of-the-money strikes, respectively.

We note that, when we train with CEV and test with SABR for at-the-money strikes, the ROC curve is just the diagonal line: this is due to the fact that the network always predicts true martingales in the test set.

Moreover, in seven cases, the ROC curve is (very close to) the combination of the vertical line

at zero and the upper line in the plot. For these examples, the network exhibits two distinct patterns:

- a) It consistently identifies strict local martingales correctly: for any strict local martingale in the test set, it outputs a probability close to 1 that it is a strict local martingale. This implies that $\text{TPR} = 1$ for all thresholds smaller than a value proximal to 1. In particular, the thresholds for which $\text{TPR} < 1$ are big enough so that $\text{FPR} = 0$.
- b) It consistently identifies true martingales correctly: for any true martingale in the test set, it outputs a probability close to 0 that it is a strict local martingale. This implies that $\text{FPR} = 0$ for all thresholds bigger than a value close to 0. In particular, the thresholds for which $\text{FPR} > 0$ are small enough so that $\text{TPR} = 1$.

In both the scenarios above, for all thresholds we have either $\text{TPR} = 1$ or $\text{FPR} = 0$, and this explains the observed ROC curves. In particular, scenario a) occurs when:

- we train with CEV and Sin processes and test with SABR (for both groups of strikes);
- we train with Sin and test with SABR (for out-of-the-money strikes);
- we train with SABR and Sin and test with CEV (for at-the-money strikes).

Scenario b) instead is obtained when:

- we train with CEV and test with SABR (for out-of-the-money strikes);
- we train with Sin and test with CEV (for both groups of strikes).

Finally, in order to show the role that epochs play during training, we investigate the evolution of the accuracy in the set constituted by the prices generated by SABR and Sin processes, with same parameters specified before, as a function of the epoch number. In this case, a fraction of the training data is used as validation data. The model sets apart this fraction of the training data, does not train on it, and evaluates the accuracy on this data at the end of each epoch. In Figure 1 we show the training set accuracy (that is, the accuracy on the data that are actually used to train the model) and validation set accuracy letting the validation set size be 10% of the total training set size.

Summarizing, the above results show that, in particular when combining different classes of processes in the training phase, our methodology still works when including stochastic volatility models, even if the results are (not surprisingly) less accurate than when we consider local volatility models for both training and testing.




Figure 1: Training set accuracy and validation set accuracy with respect to epochs, when training with SABR and Sin models. The validation set size is 10% of the total training set size.

3.4 The choice of the network architecture

We conclude Section 3 by commenting on the chosen neural network architecture. The network described in Section 3.1 (which we call network A in the remainder of the current subsection) has been selected since it guarantees good performances for most of cases, for a fairly small amount of layers and nodes.

In order to test the sensitivity of our approach to the choice of the network architecture, we also consider other possible choices, initially keeping other hyperparameters such as learning rate and number of epochs unchanged. In particular, we introduce:

- Network B : same number of layers as A but more nodes: input layer with 100 nodes, two hidden layers with 70 and 30 nodes;
- Network C : same number of nodes and layers as B , but with $L2$ regularization with parameter 0.01;
- Network D : higher number of layers and more nodes than A : input layer with 200 nodes, four hidden layers with 150, 100, 70 and 30 nodes.

We then repeat the experiments of Sections 3.2 and 3.3 for any of these networks. We note that:

- The average training accuracy is above 99.5% for all networks and all cases;
- Networks B and D never give (strictly) the best results in terms of either accuracy or F_1 score in the test set. In particular, choosing D or B decreases the performance in the test set when we train with Sin and test with SABR, train with Sin and CEV and test with SABR, train with Sin and test with CEV, train with Sin and SABR and test

with CEV. In all other cases (among which the local volatility models of Section 3.2) accuracies and F_1 score in the testing set remain unchanged.

- The regularized network C is (strictly) the best choice in terms of performance in the test set when we train with CEV and test with SABR, when we train with Sin and SABR and test with CEV and when we train with Sin and test with CEV (see also Tables 2 and 3); in particular, in the first case and for out-of-the-money strikes, it represents a big improvement with respect to the non regularized network B , which gives a test accuracy of only 54%. For the examples of Section 3.2, instead, regularization yields worse results in terms of both test accuracy and F_1 score.
- Network A strictly gives the best results in the test set when we train with Sin and test with SABR, train with Sin and CEV and test with SABR, and when we train with SABR and test with CEV.

For network D , which has the highest number of parameters to optimize, we also evaluate whether increasing the number of epochs from 20 to 40 or 60 improves the average testing accuracy. We observe that training with more epochs significantly enhances performance only when training with Sin and testing with displaced CEV for at-the-money strikes. In this scenario, the average test accuracy rises from 80.2% at 20 epochs to 84.1% at 40 epochs. This latter value exceeds the accuracy reported for network C in Table 2. However, since the results are similar and network D with 40 epochs is computationally more costly, we retain network C for consistency.

4 Analysis on market data: detecting tech stocks bubbles

4.1 The methodology

We now use our approach to detect if some tech stocks have been affected by an asset price bubble in last years. In particular, we focus on Meta Platforms Inc. (META), GameStop Corp. (GME), Nvidia Corporation (NVDA) and Tesla Inc. (TSLA). These companies have experienced a boom in the last years, followed by a contraction at the beginning of 2022, as it can also be seen by the time series of prices shown in Figures 2, 3, 4 and 5. This has brought many financial analysts to claim the presence of a new tech bubble, after the *dot com mania* of the late 1990s (see for example Bercovici (2017), Fusari et al. (2020), Ozimek (2017), Serla (2017), Sharma (2017)).

Let $S_{t_i}^i$ be the value at time t_i of a stock $i \in \{\text{META}, \text{GME}, \text{NVDA}, \text{TSLA}\}$. In order to assess if a bubble is present at t_i , we consider the call option market prices $C_1^i, \dots, C_{n_i}^i$ at t_i for maturities $T_1^i, \dots, T_{n_i}^i$ and strikes $K_1^i, \dots, K_{n_i}^i$. We use call option prices for a total of 50000 different underlyings for maturities $T_1^i, \dots, T_{n_i}^i$, strikes $K_1^i, \dots, K_{n_i}^i$ and initial value of the underlying S^i to train a neural network which has a first layer with 60 nodes, two hidden

layers with 30 and 7 nodes, respectively, ReLu activation function for the input and hidden layers, Sigmoid activation function for the output layer and $L2$ -regularization with parameter 0.01. For training we use the Adam optimization algorithm with learning rate equal to 0.001 and 20 epochs.

As potential candidates to generate the prices used to train the network, we only consider displaced CEV models in (46) and SABR models in (47)-(48) with $\gamma = 1$, as they admit analytic formulas for the call prices. In particular:

- For the SABR models, we let σ_0 and α be both uniformly distributed in $[0.1, 0.9]$. Moreover, we let ρ be uniformly distributed in $[-0.9, 0]$ in half cases (those where the underlying is a true martingale) and in $(0, 0.9]$ in the other half of cases (those where the underlying is a strict local martingale).
- For the displaced CEV models, we let the volatility σ and the displacement d be uniformly distributed in $[0, 5]$ and $[0.8, 1.5]$, respectively. Moreover, we let the exponent β be uniformly distributed in $[0.5, 1]$ in half cases (those where the underlying is a true martingale) and in $(1, 1.2]$ in the other half of cases (those where the underlying is a strict local martingale).

Note that in the models specification above we take randomly chosen parameters in order to increase the variety of the training data and then the robustness of our methodology.

After the training phase, we feed the network with the realized call market prices $C_1^i, \dots, C_{n_i}^i$, so that it outputs a number in $(0, 1)$ which can be interpreted as the probability that the underlying has a bubble based on the training and on the market prices. We repeat this procedure 10 times, and take the average of the probabilities returned by the networks.

For any date and any stock we decide whether to use the network trained with data from displaced CEV models only, SABR models only or both, based on the following test. We train the network for each of these alternatives, and for the maturities, strikes and initial values observed in the market. After running the prediction 10 times with same training data, we select the training choice (displaced CEV only, SABR only or both) which ranks first in most of the metrics below:

- Lowest variance of the output probabilities on the observed market data;
- Highest average accuracy on the training set;
- Highest average accuracy on a test set constituted by option prices generated by both displaced CEV and SABR models;
- Lowest variance of the output probabilities on the test set of displaced CEV and SABR models.

If two alternatives rank first in two of the above metrics, we take the average prediction between them.

According to this criterion we select 6 times both displaced CEV and SABR models, 5 times only SABR models and 15 times only displaced CEV models. Moreover, in 7 cases we take the average predictions between the network trained with both displaced CEV and SABR models and the one trained with displaced CEV models only. In more detail:

- The network trained with both displaced CEV and SABR models gives the lowest variance of returned probabilities for the market data in 14 cases, the highest accuracy on the training set in 1 case, the highest accuracy on the test set in 31 cases and the lowest variance of returned probabilities for the underlyings in the test set in 2 cases;
- The network trained with only SABR models gives the lowest variance of returned probabilities for the market underlying in 6 cases and the highest accuracy on the training set in 9 cases;
- The network trained with only displaced CEV models gives the lowest variance of returned probabilities for the market underlying in 12 cases, the highest accuracy on the training set in 21 cases, the highest accuracy on the CEV/SABR test set in 2 cases and the lowest variance of returned probabilities for the underlyings in the test set in 30 cases.

Note that, even if the training set with both displaced CEV and SABR models provides in most cases the highest accuracy on the test set, still training with the displaced CEV alone generally provides good test accuracies for the strikes and maturities available in the market. Moreover, for some strikes and maturities, it scores a lot better (with test accuracies around 97%) when the test set is generated by only displaced CEV models.

We now state our findings for the four underlyings. Note that only American call option prices are available on the market. However, all the stocks that we consider here do not pay dividends or have a very small dividend yield (Nvidia Corporation) that we neglect. In this case, American call option prices are equal to European option prices and can be used for our tests.

4.2 Nvidia Corporation

The advanced semiconductor manufacturer Nvidia Corp. has seen its shares rocketing to huge heights in last years. From January 2016 to its peak in late November 2021, Nvidia's closing stock prices boosted their value by about 4600%. This seems to point to a bubble in Nvidia stock, see for example Kolakowski (2019).

We apply our methodology in order to check for the presence of a bubble from January 2016 to July 2022, on a six-monthly basis. In Figure 2 we plot the price evolution of the asset with the dates when we use our methodology to detect if a bubble is present. In particular, for each date we consider the average probability P_b for the presence of a bubble we get running our methodology 10 times, as described above. Different colors of the lines correspond to different values of P_b . We list these findings in more details in the following, in terms of the evolution of P_b .

- No bubble is detected from January 2016 to July 2017: in particular, P_b is equal to 4.8% in January 2016, 0.6% in July 2016, 8.7% in January 2017 and 1.5% in July 2017.
- The network outputs a probability of a bubble $P_b = 71.2\%$ in January 2018 and $P_b = 41.4\%$ in July 2018.
- In the second half of 2018 we observe a decrease of the asset price, and the network outputs bubble probabilities $P_b = 19.8\%$ in January 2019, $P_b = 20.3\%$ in July 2019 and $P_b = 41.2\%$ in January 2020.
- A second increase of the price leads to a second bubble according to our methodology. Indeed, the algorithm outputs high probabilities of having a bubble from July 2020: we have $P_b = 98.9\%$ in July 2020, $P_b = 99.5\%$ in January and July 2021, $P_b = 95.2\%$ slightly before the peak of November 2021.
- A second decline of the price can be observed from the end of 2021: after the beginning of this decline, our methodology outputs $P_b = 44.3\%$.

4.3 GameStop Corporation

The stock value of the video games and gaming merchandise retailer GameStop has experienced a meteoric rise over the few last years. From the beginning of January 2021, GameStop's stock price increased from around 17\$ to an all-time high of nearly 483\$ on January 28, 2021, marking an astonishing gain of over 2700% in a matter of weeks. After this surge, the price experienced a rapid decline, to rise again and stabilize, albeit with high volatility, at levels much higher than those of early 2021. Observing this behaviour, most financial analysts have claimed the stock clearly had a bubble, see for example Libich et al. (2021). We here assess this statement using our methodology. Figure 3 shows the results. The network detects a bubble with high probability already in January and September 2020. In the middle of the huge rise of January 2021 a bubble is detected with average probability $P_b = 86.7\%$, and later on always with probabilities higher than 99% apart from June 2021 when the network outputs $P_b = 94.7\%$.

4.4 Meta Platforms Inc.

The stock price of Meta Platforms, Inc., formerly known as Facebook, has experienced a substantial surge over the last few years, more than doubling from January 2018 to July 2021. We use our methodology in order to assess if this rise indicates the presence of a bubble, see Figure 4. Before 2020, the network outputs small probabilities for Meta to have a bubble (average probabilities are 0.2% in November 2018, 0.3% in March 2019, 5% in September 2019). Things change after the substantial rise from early 2020: P_b equals 82.1% in late July 2020, 92.8% in early April 2021 and 93.4% in September 2021. Finally, the decrease after the peak is interpreted as the burst of the bubble, as $P_b = 0.7%$ in November 2022. Note that the presence of temporary bubbles on Facebook over last years has been detected also with the method of Fusari et al. (2020).

4.5 Tesla Inc.

As a third example, we focus on the case of Tesla. The Tesla stock price has seen an extremely steep increase in last years, of about 30000% from early 2016 to the peak in late 2021. Because of such a huge boost, many experts claimed that the stock is affected by a bubble, see for example Libich et al. (2021). By using our methodology we assess if a bubble is present or not four dates, which are shown in Figure 5. The color of the lines which identify the dates corresponds to the probability P_b assigned by our methodology for the presence of a bubble. The algorithm outputs $P_b = 1.3%$ the 3rd of January 2016, $P_b = 70.1%$ the 1st of November 2021 (right before the peak of the stock price), $P_b = 95%$ the 1st of February 2022 (right after the peak) and $P_b = 98.9%$ the 17th of July 2022.

4.6 Discussion

Overall, for each of these assets our method finds that bubbles were present with high probability at certain points in time. These dates coincide with points shortly before a rapid decrease of the asset price. Similarly, points in time at which the method outputs low bubble probabilities coincide with times which are followed neither by a rapid increase nor a rapid decrease of the price.

5 Conclusion

In this article we propose a deep learning-based methodology to detect financial asset price bubbles, where a deep neural network is trained using synthetic option price data generated from a collection of models. We apply our method to local and stochastic volatility models, in which the presence of a bubble is determined by the value of certain parameters. This trained neural network is then used for bubble detection in more general situations. The deep




Figure 2: Nvidia Corporation stock prices (USD). The lines illustrate the dates considered to detect a bubble with our methodology. The colors change based on the average probability P_b that the asset has a bubble according to our methodology. In particular: for the red lines we have $P_b > 95\%$, for the yellow line $P_b = 71.2\%$, for the black lines $40\% < P_b < 45\%$, for the green lines $19\% < P_b < 21\%$ and for the blue lines $P_b < 10\%$.




Figure 3: GameStop corporation. stock prices (USD). The lines illustrate the dates considered to detect a bubble with our methodology. The colors change based on the average probability P_b that the asset has a bubble according to our methodology. In particular, for the red lines we have $P_b > 98\%$, for the orange line $P_b = 94.7\%$ and for the yellow one $P_b = 86.7\%$.




Figure 4: Meta Platforms Inc. stock prices (USD). The lines illustrate the dates considered to detect a bubble with our methodology. The colors change based on the average probability P_b that the asset has a bubble according to our methodology. In particular, for the orange lines we have $92 < P_b < 94\%$, for the yellow line $P_b = 82.1\%$ and for the blue ones $P_b \leq 5\%$.




Figure 5: Tesla Inc. stock prices (USD). The lines illustrate the dates considered to detect a bubble with our methodology. The colors change based on the average probability P_b that the asset has a bubble according to our methodology. In particular: for the red lines we have $P_b \geq 95\%$, for the yellow line $P_b = 70.1\%$ and for the blue one $P_b = 1.3\%$.

neural network takes as input call option prices for a given underlying and returns as output a number in $(0, 1)$, which indicates the probability that the underlying exhibits a bubble. We have assessed the performance of the methodology on both synthetic and real data. First we have tested various settings in which the network is trained using a certain class of local or stochastic volatility models and tested within another class. In particular, a high out-of-sample prediction accuracy has been observed when two model classes are used for training and the trained network is used for bubble detection in another class of models. We have also tested the methodology with market data of four tech stocks at various dates and have observed a close match between output of the network and expected results. Furthermore, these numerical experiments have been complemented by a mathematical proof for the method in a local volatility setting. Extensions of these theoretical results to more general model classes are also possible, provided an analogue of the “bubble detection function” F can be constructed. On the other hand, the proposed algorithm could be directly modified to take into account additional market factors as inputs for the neural network, provided that these factors can be incorporated into the models used to generate the training data.

References

- D. Abreu and M.K. Brunnermeier. Bubbles and crashes. *Econometrica*, 71(1):173–204, 2003.
- Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Trans. Inform. Theory*, 39(3):930–945, 1993.
- Jeff Bercovici. Yes, it’s a Tech Bubble. *Inc.*, 2017. URL <https://www.inc.com/magazine/201509/jeff-bercovici/are-we-in-a-tech-bubble.html>.
- F. Biagini and S. Nedelcu. The formation of financial bubbles in defaultable markets. *SIAM Journal on Financial Mathematics*, 6(1):530–558, 2015.
- F. Biagini, A. Mazzon, and T. Meyer-Brandis. Liquidity induced asset bubbles via flows of ELMs. *SIAM Journal on Financial Mathematics*, 9(2):800–834, 2018.
- Francesca Biagini, Hans Föllmer, and Sorin Nedelcu. Shifting martingale measures and the slow birth of a bubble. *Finance and Stochastics*, 18(2):297–326, 2014.
- Jane V Carter, Jianmin Pan, Shesh N Rai, and Susan Galandiuk. Roc-ing along: Evaluation and interpretation of receiver operating characteristic curves. *Surgery*, 159(6):1638–1645, 2016.
- Alexander MG Cox and David G Hobson. Local martingales, bubbles and option prices. *Finance and Stochastics*, 9(4):477–492, 2005.

- J.B. DeLong, A. Shleifer, L. Summers, and R. Waldmann. Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703–738, 1990.
- José Carlos Dias, João Pedro Vidal Nunes, and Aricson Cruz. A note on options and bubbles under the CEV model: implications for pricing and hedging. *Review of Derivatives Research*, 23(3):249–272, 2020.
- Erik Ekström and Johan Tysk. Dupire’s equation for bubbles. *International Journal of Theoretical and Applied Finance*, 15(06):1250041, 2012.
- Dennis Elbrächter, Philipp Grohs, Arnulf Jentzen, and Christoph Schwab. DNN expression rate analysis of high-dimensional PDEs: application to option pricing. *Constr. Approx.*, 55(1):3–71, 2022.
- Jerome Fan, Suneel Upadhye, and Andrew Worster. Understanding receiver operating characteristic (roc) curves. *Canadian Journal of Emergency Medicine*, 8(1):19–20, 2006.
- H. Föllmer, U. Horst, and A. Kirman. Equilibria in financial markets with heterogeneous agents: A probabilistic perspective. *Journal of Mathematical Economics*, 41(1-2):123–155, 2005.
- Nicola Fusari, Robert A. Jarrow, and Sujan Lamichhane. Testing for asset price bubbles using options data. *Johns Hopkins Carey Business School Research Paper*, (20-12), 2020.
- Luzia Gonçalves, Ana Subtil, M Rosário Oliveira, and Patricia de Zea Bermudez. Roc curve estimation: An overview. *REVSTAT-Statistical journal*, 12(1):1–20, 2014.
- Lukas Gonon. Random feature neural networks learn Black-Scholes type PDEs without curse of dimensionality. *Preprint, arXiv 2106.08900*, 2021.
- Lukas Gonon and Christoph Schwab. Deep ReLU network expression rates for option prices in high-dimensional, exponential Lévy models. *Finance Stoch.*, 25(4):615–657, 2021.
- Philipp Grohs, Fabian Hornung, Arnulf Jentzen, and Philippe von Wurstemberger. A proof that artificial neural networks overcome the curse of dimensionality in the numerical approximation of Black-Scholes partial differential equations. *To appear in Mem. Am. Math. Soc.; arXiv:1809.02362*, 2018.
- J.M. Harrison and D.M. Kreps. Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics*, 92(2):323–336, 1978.
- Martin Herdegen and Martin Schweizer. Strong bubbles and strict local martingales. *International Journal of Theoretical and Applied Finance*, 19(04):1650022, 2016.

- Steven L Heston, Mark Loewenstein, and Gregory A Willard. Options and bubbles. *The Review of Financial Studies*, 20(2):359–390, 2007.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(1989):251–257, 1991.
- Antoine Jacquier and Martin Keller-Ressel. Implied volatility in strict local martingale models. *SIAM Journal on Financial Mathematics*, 9(1):171–189, 2018.
- R. Jarrow and D. Madan. Arbitrage, martingales and private monetary value. *Journal of Risk*, 3(1), 2000.
- R. Jarrow and P. Protter. Foreign currency bubbles. *Review of Derivatives Research*, 14(1): 67–83, 2011.
- R. Jarrow and P. Protter. Relative asset price bubbles. Preprint, Available at SSRN: <http://ssrn.com/abstract=2265465> or <http://dx.doi.org/10.2139/ssrn.2265465>, 2013.
- R. Jarrow, P. Protter, and A. Roch. A Liquidity Based Model for Asset Price. *Quantitative Finance*, 12(1):1339–1349, 2012.
- Robert A. Jarrow. Testing for asset price bubbles: three new approaches. *Quantitative Finance Letters*, 4(1):4–9, 2016.
- Robert A. Jarrow and Simon S. Kwok. Inferring financial bubbles from option data. *Journal of Applied Econometrics*, 36(7):1013–1046, 2021.
- Robert A. Jarrow and Philip Protter. Forward and futures prices with bubbles. *International Journal of Theoretical and Applied Finance*, 12(7):901–924, 2009.
- Robert A. Jarrow, Philip Protter, and Kazuhiro Shimbo. Asset price bubbles in complete markets. *Advances in Mathematical Finance*, In Honor of Dilip B. Madan:105–130, 2007.
- Robert A. Jarrow, Philip Protter, and Kazuhiro Shimbo. Asset price bubbles in incomplete markets. *Mathematical Finance*, 20(2):145–185, 2010.
- Robert A. Jarrow, Younes Kchia, and Philip Protter. How to detect an asset bubble. *SIAM Journal on Financial Mathematics*, 2:839–865, 2011a.
- Robert A. Jarrow, Younes Kchia, and Philip Protter. How to detect an asset bubble. *SIAM Journal on Financial Mathematics*, 2(1):839–865, 2011b.
- Robert A. Jarrow, Younes Kchia, and Philip Protter. Is there a bubble in LinkedIn’s stock price? *The Journal of Portfolio Management*, 38(1):125–130, 2011c.

- C. Kardaras, D. Kreher, and A. Nikeghbali. Strict local martingales and bubbles. *The Annals of Applied Probability*, 25(4):1827–1867, 2015.
- Mark Kolakowski. Nvidia’s stock signals techs near bubble like 2000. 2019. URL <https://www.investopedia.com/news/nvidia/>.
- Jan Libich, Liam Lenten, et al. Bitcoin, Tesla and GameStop bubbles as a flight to focal points. *World Economics*, 22(1):83–108, 2021.
- Alan E Lindsay and DR Brecher. Simulation of the CEV process and the local martingale property. *Mathematics and Computers in Simulation*, 82(5):868–878, 2012.
- Mark Loewenstein and Gregory A. Willard. Rational equilibrium asset-pricing bubbles in continuous trading models. *Journal of Economic Theory*, 91(1):17–58, 2000.
- Adam Ozimek. There obviously is a tech bubble, but hopefully that doesn’t matter. *Forbes*, 2017. URL <https://www.forbes.com/sites/modeledbehavior/2017/07/16/there-obviously-is-a-tech-bubble-but-hopefully-that-doesnt-matter/#737270d9d426>.
- S. Pal and P. Protter. Analysis of continuous strict local martingales via h-transforms. *Stochastic processes and their applications*, 120(8):1424–1443, 2010.
- Petteri Piiroinen, Lassi Roininen, Tobias Schoden, and Martin Simon. Asset price bubbles: An option-based indicator. *arXiv preprint arXiv:1805.07403*, 2018.
- P. Protter. A mathematical theory of financial bubbles. In V. Henderson and R. Sincar, editors, *Paris-Princeton Lectures on Mathematical Finance*, volume 2081 of *Lecture Notes in Mathematics*. Springer, 2013.
- Christoph Reisinger and Yufei Zhang. Rectified deep neural networks overcome the curse of dimensionality for nonsmooth value functions in zero-sum games of nonlinear stiff systems. *Anal. Appl. (Singap.)*, 18(6):951–999, 2020.
- Daniel Revuz and Marc Yor. *Continuous martingales and Brownian motion*, volume 293. Springer Science & Business Media, 2013.
- J. Scheinkman and W. Xiong. Overconfidence and speculative bubbles. *Journal of political economy*, 111(6):1183–1219, 2003.
- Rusli Serla. The tech bubble: how close is it to bursting? *Telegraph*, 2017. URL <http://www.telegraph.co.uk/technology/2017/06/22/tech-bubble-close-bursting/>.
- Ruchir Sharma. When will the tech bubble burst? *New York Times*, 2017. URL <https://www.nytimes.com/2017/08/05/opinion/sunday/when-will-the-tech-bubble-burst.html>.

Carlos A. Sin. Complications with stochastic volatility models. *Advances in Applied Probability*, 30(1):256–268, 1998.

J. Tirole. On the possibility of speculation under rational expectations. *Econometrica*, 53(6): 1163–1182, 1982.

6 Appendix

We here provide the ROC curves relative to the experiments of Section 3.3.

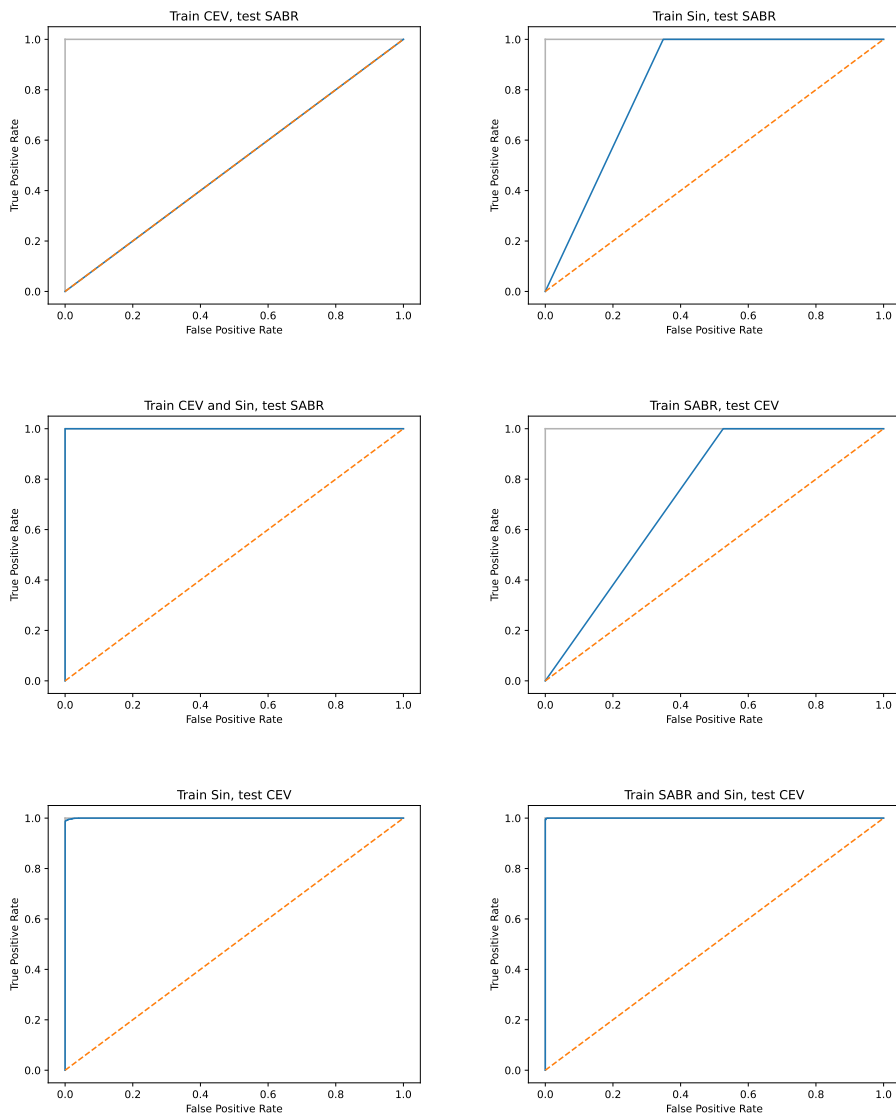


Figure 6: ROC curves obtained taking underlyings with initial value $x_0 = 2$ and dynamics specified in Section 3.3, equally spaced maturities in (2, 5) and equally spaced strikes in (1, 3.5).




Figure 7: ROC curves obtained taking underlyings with initial value $x_0 = 2.0$ and dynamics specified in Section 3.3, equally spaced maturities in $(2, 5)$ and equally spaced strikes in $(3, 5.5)$.