

A deep learning approach to data-driven model-free pricing and to martingale optimal transport

Ariel Neufeld and Julian Sester

Abstract

We introduce a novel and highly tractable supervised learning approach based on neural networks that can be applied for the computation of model-free price bounds of, potentially high-dimensional, financial derivatives and for the determination of optimal hedging strategies attaining these bounds. In particular, our methodology allows to train a single neural network offline and then to use it online for the fast determination of model-free price bounds of a whole class of financial derivatives with current market data. We show the applicability of this approach and highlight its accuracy in several examples involving real market data. Further, we show how a neural network can be trained to solve martingale optimal transport problems involving fixed marginal distributions instead of financial market data.

I. INTRODUCTION

FINANCIAL derivatives are financial contracts between the corresponding seller, typically a bank, and a buyer, typically another financial institution or a private person, with a future uncertain payoff depending on another (typically simpler) financial instrument, often a stock, to which we refer as the underlying security. Options are a large class of financial derivatives which allow, but do not oblige the owner of the option to buy or sell the underlying securities involved in the contract. The most common types of traded financial derivatives are call and put options which allow to buy and sell, respectively, the underlying single security at a future maturity at a predetermined price, the so called *strike* of the option. Due to the uncertainty involved in the future cashflow, today's price of the financial derivative is a priori unclear and subject to a high degree of ambiguity. The classical paradigm in mathematical finance, which is commonly applied to determine the *fair* value of some financial derivative, consists in capturing the developments of the real underlying market by a sophisticated financial market model¹. This model is then calibrated to observable market parameters such as current spot prices, prices of liquid options, interest rates, and dividends, and is thus believed to capture the reality appropriately, see e.g. [53] for details of this procedure. However, such an approach evidently involves the uncontrollable risk of having a priori chosen the wrong type of model - this refers to the so called *Knightian uncertainty* ([42]).

To reduce this apparent model risk the research in the area of mathematical finance recently developed a strong interest in the computation of model-independent² and robust price bounds for financial derivatives (compare among many others [7], [13], [16], [17], [20], [27], [35], [36], [38], [46], and [47]). We speak of model-independent price bounds if realized prices within these bounds exclude any arbitrage opportunities³ under usage of liquid market instruments independent of any model assumptions related to potential underlying stochastic models, whereas robust price bounds refer to the exclusion of model-dependent arbitrage within a range of models that are deemed to be admissible.

We present an approach enabling the fast and reliable computation of model-independent price bounds of financial derivatives. This approach is mainly based on supervised deep learning ([43], [52]) and proposes how a deep⁴ feed-forward neural network⁵ can be trained to learn the relationship between observed financial data and associated model-independent price bounds of any potentially high-dimensional financial derivative from an entire parametric class of exotic⁶ options. The great advantage of the presented methodology is that, in contrast to computational intensive and therefore potentially time-consuming pricing methods which have to be reapplied for each new set of observed financial data and each derivative one wants to value, it allows to use a *sole pre-trained* neural network for *real time* pricing of every financial derivative from a pre-specified class of payoff functions. Let us consider some family of financial derivatives defined through payoff functions

$$\Phi_{\theta} : \mathbb{R}_+^{nd} \rightarrow \mathbb{R}, \quad \theta \in \Theta,$$

A. Neufeld is with the Division of Mathematical Sciences, NTU Singapore. e-mail: ariel.neufeld@ntu.edu.sg

J. Sester is with the Department of Mathematics, NUS Singapore. e-mail: jul_ses@nus.edu.sg

¹Examples for sophisticated financial market models include among many others the Heston model (compare [33]) and Dupire's local volatility model (compare [22]).

²The terms model-free and model-independent are used synonymously in the literature.

³*Arbitrage* refers to a profit that can be realized without taking any risk. Prices of a derivative that allow for arbitrage are considered as not reasonable as the arbitrage profit would be immediately exploited by *arbitrageurs*.

⁴We speak of *deep* neural networks if there are at least 2 hidden layers involved.

⁵As a convention we refer to feed-forward simply as neural networks throughout the paper.

⁶Every option that is neither a call nor a put option is called *exotic*.

which determines the payoff an investor receives at time t_n in case he bought the derivative Φ_θ at initial time t_0 . The payoff depends on the values of $d \in \mathbb{N}$ underlying securities at $n \in \mathbb{N}$ future times $t_1 < t_2 < \dots < t_n$, i.e., the derivative depends on each security $\mathbf{S}^k = (S_{t_1}^k, \dots, S_{t_n}^k)$ for $k = 1, \dots, d$. The goal is then to determine all possible today's prices for each Φ_θ such that a potential investor cannot profit from one of these prices to exploit *model-independent arbitrage*. This notion refers to strategies that involve trading in underlying securities and/or in liquid options which are cost-free and lead to a profit independent of any model assumptions, i.e., for any possible future evolution of the underlying security, see also [2]. As a canonical example we consider the class of payoffs associated to basket options, which are financial derivatives that allow (but not oblige) at a future time to buy a weighted sum of financial assets (with weights denoted by $(w_i^k)_{i,k}$) at a predetermined strike L . Such an option is only executed if it is favorable for the option-holder to do so, which is the case if the difference between the weighted sum and the strike is positive and therefore the set of payoffs is given by

$$\{\Phi_\theta(\mathbf{S}^1, \dots, \mathbf{S}^d); \theta \in \Theta\} := \left\{ \max \left\{ \sum_{i=1}^n \sum_{k=1}^d w_i^k S_{t_i}^k - L, 0 \right\} \text{ where } \theta := \left((w_i^k)_{i,k}, L \right) \in \mathbb{R}^{nd} \times \mathbb{R} \right\}. \quad (\text{I.1})$$

To find the arbitrage-free upper price bounds of a financial derivative Φ_θ , we consider model-independent super-replication strategies (also called super-hedging strategies) of Φ_θ , i.e., trading strategies that lead for every possible evolution of the underlying securities to a greater or equal outcome than the payoff of Φ_θ , which is referred to as the trading strategy super-replicating Φ_θ . Prices of such strategies need to be at least as high as the price of Φ_θ , since otherwise the market would admit model-independent arbitrage, which can indeed be seen by buying the strategy and by selling the derivative Φ_θ at initial time. Thus, the smallest price among all model-independent super-replication strategies leads to the arbitrage-free upper price bound of Φ_θ . Analogue, the greatest price among sub-replication strategies yields the arbitrage-free lower price bound. Moreover, it is a consequence of (adaptions of) the fundamental theorem of asset pricing (see [2] and [19]) that there exist a dual method to approach the valuation problem: One may also consider all martingale models⁷ which are consistent with bid and ask prices of liquidly traded option prices written on the underlying securities $\mathbf{S} = (\mathbf{S}^1, \dots, \mathbf{S}^d)$ and expire at the future maturities $(t_i)_{i=1, \dots, n}$ as candidate models. Then, minimizing and maximizing the expectations $\mathbb{E}_{\mathbb{Q}}[\Phi_\theta(\mathbf{S})]$ among all associated martingale / risk-neutral measures \mathbb{Q} of potential models leads to the desired price bounds, compare e.g. [2] and [15] for such results in the discrete time model-independent setting.

Given a payoff function Φ_θ from a (parametric) set of payoff functions $\{\Phi_\theta, \theta \in \Theta\}$, for example from the set of basket options as in (I.1), we use the sub/super-replication method to compute the lower and upper price bound of Φ_θ for various different sets of financial data and for several choices of $\theta \in \Theta$, i.e., we compute the bounds in dependence of different observed financial data. The observable market parameters comprised in the financial data include prices of the underlying securities as well as bid and ask prices of liquidly traded call and put options and its associated strikes. After having computed the price bounds for various different sets of financial data, we let, in accordance with the universal approximation theorem from [37], a specially designed neural network learn the relationship between observed financial data and the corresponding model-independent price bounds for a parametric family of payoff functions, compare also Figure 1a.

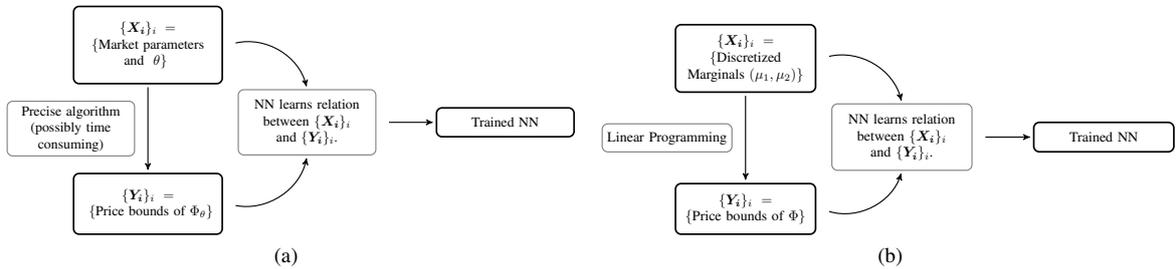


Fig. 1. (a): Illustration of the presented approach, that is described in detail in Algorithm 1, in order to train a neural network (NN) to learn the model-independent price bounds of a derivative Φ_θ from a family $\{\Phi_\theta, \theta \in \Theta\}$ in dependence of given market prices.

(b): Illustration of Algorithm 2 which is applied to learn price bounds of MOT problems from marginals. The price bounds contained in \mathbf{Y}_i correspond to the solutions of MOT problems, i.e., to $\inf_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi(S_{t_1}, S_{t_2})]$ and $\sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi(S_{t_1}, S_{t_2})]$, where $\mathcal{M}(\mu_1, \mu_2)$ denotes the set of martingale measures with fixed marginal distributions μ_1 and μ_2 , compare also equation (III.2).

While there exist several numerical routines to compute model-free price bounds of financial derivatives, the only numerical routine that allows to compute model-free price bounds in a purely data-driven approach without imposing any probabilistic assumptions on the market is the approach from [47] which we therefore use to construct a training set of price bounds. Indeed, while [23], [24], [29], [31], [32] all provide methods to compute price bounds of financial derivatives, they all rely fundamentally on the assumption that the marginal distributions of each single asset are known exactly. Moreover, each established methodology so far requires for every new set of financial data to employ a potentially time-consuming valuation method to find price bounds for every financial derivative of interest. Our approach circumvents this problem as it enables to train offline a *single* neural network for a *whole family* of related payoff functions, such as e.g. basket options with different weights and strikes, and then to determine model-free price bounds in *real time* by using the already trained neural network. Thus, in practice, it suffices to train a couple of neural networks (one for each relevant family of payoffs) and then to use the pre-trained neural networks for valuation-purposes. We refer to Remark II.9 (e) for further possible examples of parametric families $\{\Phi_\theta, \theta \in \Theta\}$, where we highlight that only a few neural networks are necessary to cover the most relevant payoff functions of financial derivatives.

⁷One often refers to martingale models as *risk-neutral models*, in which an investor is indifferent of either investing in the underlying security or keeping her money in a bank account with constant interest rate (where typically one assumes the interest rate to be zero for simplicity).

In Section II, we first present our approach in a very general setting including multiple assets, multiple time steps, as well as market frictions. To justify our methodology, we show, by proving a continuous relationship between market data and resultant model-free price bounds, that the universal approximation theorem from [37] is indeed applicable, see Theorem II.2. More precisely, under some continuity assumptions on the parametric family of payoff functions $\theta \mapsto \Phi_\theta$ which are typically satisfied for the relevant families of payoff functions in finance (see Section II-C), we prove that both upper and lower arbitrage-free price bound depend continuously on the relevant inputs: strike prices and bid-ask prices of the liquid call and put options, the current prices of the underlying stocks, as well as the parameter determining the payoff function. This, together with the universal approximation property of neural networks proves that a single neural network can indeed learn the arbitrage-free price bounds of a parametric family of payoff functions. To the best of our knowledge, Theorem II.2 (as well as Theorem III.1) is the first result which proves a continuous relationship between the respective inputs and outputs. This result justifies to learn model-free price bounds by neural networks, and hence provides an important novel contribution to the field. In particular, this means that it is possible to train a single neural network offline on past market data and then to use it online with current market data to compute price bounds of each financial derivative $\Phi_\theta, \theta \in \Theta$. Additionally, we show accuracy and tractability of our presented approach in various high-dimensional relevant examples involving real market data.

In Section III, we show that the methodology can also be applied to compute two-marginal martingale optimal transport (MOT) problems, see also Figure 1b for an illustration of the approach where instead of market data entire marginal distributions are the input to the neural network, we refer to Theorem III.1 for the novel theoretical justification of that approach. The knowledge about marginal distributions can be motivated by the findings from [11] which lead to the insight that complete information about the marginal distributions is equivalent to the knowledge of prices of call options written on the underlying securities for a whole continuum of strikes.

We further show within several examples the applicability of the presented approach. Mathematical proofs of the theoretical results are provided in Section IV.

II. APPROXIMATING MODEL-FREE PRICE BOUNDS WITH NEURAL NETWORKS

In this section we present an arbitrage-free approach to determine model-free price bounds of a possibly high-dimensional financial derivative when real market data is given. In addition to prices of underlying securities we observe bid and ask prices of call and put options written on these securities, where bid and ask prices refer to the quotations for which the options can be sold and bought. Moreover, we explain how model-independent price bounds can be approximated through neural networks.

A. Model-independent valuation of derivatives

We consider at the present time $t_0 \in [0, \infty)$ $d \in \mathbb{N}$ underlying securities and $n \in \mathbb{N}$ future times $t_0 < t_1 < \dots < t_n < \infty$, i.e., the underlying process is given by

$$\mathbf{S} := (\mathbf{S}^1, \dots, \mathbf{S}^d) = (\mathbf{S}_{t_1}, \dots, \mathbf{S}_{t_n}) = (S_{t_i}^k)_{i=1, \dots, n}^{k=1, \dots, d}$$

with $\mathbf{S}^k := (S_{t_i}^k)_{i=1, \dots, n}$ denoting the k -th underlying security and $\mathbf{S}_{t_i} := (S_{t_i}^k)_{k=1, \dots, d}$ denoting the values of the underlying securities at time t_i . The process \mathbf{S} is modelled as the canonical process on \mathbb{R}_+^{nd} equipped with the Borel σ -algebra denoted by $\mathcal{B}(\mathbb{R}_+^{nd})$, i.e., for all $i = 1, \dots, n, k = 1, \dots, d$ we have

$$S_{t_i}^k(\mathbf{s}) = s_i^k \text{ for all } \mathbf{s} = (s_1^1, \dots, s_1^d, \dots, s_n^1, \dots, s_n^d) \in \mathbb{R}_+^{nd}.$$

As we want to consider real market data, we cannot - as usual in a vast majority of the mathematical literature on model-independent finance - neglect bid-ask spreads as well as transaction costs. Thus, we assume that option prices do not necessarily coincide for buyer and seller, instead we take into account a bid price and an ask price. Let $k \in \{1, \dots, d\}$, $i \in \{1, \dots, n\}$, $j \in \{1, \dots, n_{ik}^{\text{opt}}\}$, where n_{ik}^{opt} denotes the amount of tradable put and call options⁸ with maturity t_i written on $S_{t_i}^k$. Then a call option on \mathbf{S}^k with maturity t_i for strike $K_{ijk}^{\text{call}} \in \mathbb{R}_+$ can be bought at price $\pi_{\text{call}, i, j, k}^+$ and be sold at price $\pi_{\text{call}, i, j, k}^-$. As a call option entitles the owner of the option to *buy* the underlying security at price K_{ijk}^{call} at time t_i it is only exercised if the difference between underlying security and strike K_{ijk}^{call} is positive, and therefore possesses the payoff $\max\{S_{t_i}^k - K_{ijk}^{\text{call}}, 0\}$. Similarly, bid and ask prices for traded put options are denoted by $\pi_{\text{put}, i, j, k}^-$ and $\pi_{\text{put}, i, j, k}^+$, respectively. Put options give the right to *sell* the underlying security at price K_{ijk}^{put} at time t_i . Hence, put options are only exercised if the difference between strike K_{ijk}^{put} and underlying security is positive, leading to the payoff $\max\{K_{ijk}^{\text{put}} - S_{t_i}^k, 0\}$.

Moreover, we assume proportional transaction costs, similar to the approaches in [15, Section 3.1.1.] and [21]. This means, at each time t_i , after having observed the values $\mathbf{S}_{t_1}, \dots, \mathbf{S}_{t_i}$, rearranging a dynamic self-financing⁹ trading position in the underlying security from¹⁰ $\Delta_{i-1}^k \in B(\mathbb{R}_+^{(i-1)d}, \mathbb{R})$, which was the trading position after having observed only the values $\mathbf{S}_{t_1}, \dots, \mathbf{S}_{t_{i-1}}$, to a new trading position $\Delta_i^k \in B(\mathbb{R}_+^{id}, \mathbb{R})$, causes transaction costs¹¹ of

$$\kappa |S_{t_i}^k| \left| \Delta_i^k(\mathbf{S}_{t_i}, \dots, \mathbf{S}_{t_1}) - \Delta_{i-1}^k(\mathbf{S}_{t_{i-1}}, \dots, \mathbf{S}_{t_1}) \right|$$

⁸We assume the same amount of traded put and call options. This simplifies the presentation, but can without difficulties be extended to a more general setting.

⁹Self-financing means that at any time there is neither consumption nor any money injection. The profit of the trading strategy is purely a consequence of the trading in the underlying security.

¹⁰For $m, n \in \mathbb{N}$ and some set $K \subseteq \mathbb{R}^m$, we denote by $B(K, \mathbb{R}^n)$ the set of all functions $f : K \rightarrow \mathbb{R}^n$ which are $B(K)/\mathcal{B}(\mathbb{R}^n)$ -measurable, whereas $C(K, \mathbb{R}^n)$ denotes the set of all continuous functions $f : K \rightarrow \mathbb{R}^n$.

¹¹Here also different approaches to measure transaction costs would have been possible. Compare for example the presentations in [12] and [15].

for some fixed $\kappa \geq 0$. We denote for each $k = 1, \dots, d$ by $S_{t_0}^k \in \mathbb{R}_+$ the observable and therefore deterministic current value of the k -th security, also called the *spot price* of S^k . Then, given spot prices $\mathbf{S}_{t_0} = (S_{t_0}^1, \dots, S_{t_0}^d)$ and strikes $\mathbf{K} := ((K_{ijk}^{\text{call}})_{i,j,k}, (K_{ijk}^{\text{put}})_{i,j,k})$, we consider trading strategies with profits of the form¹²

$$\begin{aligned} \Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\mathbf{S}) := & a + \sum_{i=1}^n \sum_{k=1}^d \sum_{j=1}^{n_{ik}^{\text{opt}}} (c_{ijk}^+ - c_{ijk}^-) \max\{S_{t_i}^k - K_{ijk}^{\text{call}}, 0\} + \sum_{i=1}^n \sum_{k=1}^d \sum_{j=1}^{n_{ik}^{\text{opt}}} (p_{ijk}^+ - p_{ijk}^-) \max\{K_{ijk}^{\text{put}} - S_{t_i}^k, 0\} \\ & + \sum_{k=1}^d \sum_{i=0}^{n-1} \left(\Delta_i^k(\mathbf{S}_{t_i}, \dots, \mathbf{S}_{t_1}) (S_{t_{i+1}}^k - S_{t_i}^k) - \kappa |S_{t_i}^k| \left| \Delta_i^k(\mathbf{S}_{t_i}, \dots, \mathbf{S}_{t_1}) - \Delta_{i-1}^k(\mathbf{S}_{t_{i-1}}, \dots, \mathbf{S}_{t_1}) \right| \right) \end{aligned} \quad (\text{II.1})$$

for an amount of cash $a \in \mathbb{R}$, non-negative long positions $c_{ijk}^+, p_{ijk}^+ \in \mathbb{R}_+$ and non-negative short positions $c_{ijk}^-, p_{ijk}^- \in \mathbb{R}_+$ in call and put options, respectively, for $j = 1, \dots, n_{ik}^{\text{opt}}$, $i = 1, \dots, n$, $k = 1, \dots, d$. In equation (II.1) and for the rest of the paper we use the abbreviations $\mathbf{c}_{ijk} = (c_{ijk}^+, c_{ijk}^-) \in \mathbb{R}_+^2$ and $\mathbf{p}_{ijk} = (p_{ijk}^+, p_{ijk}^-) \in \mathbb{R}_+^2$. Further, the strategies involve self-financing trading positions $\Delta_i^k \in B(\mathbb{R}_+^{id}, \mathbb{R})$ with the convention $\Delta_0^k \in \mathbb{R}$, i.e., to be deterministic, as well as $\Delta_{-1}^k := 0$. The costs for setting up the position $\Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}$ with respect to the bid-ask prices

$$\boldsymbol{\pi} := \left((\pi_{\text{call}, i, j, k}^-)_{i, j, k}, (\pi_{\text{call}, i, j, k}^+)_{i, j, k}, (\pi_{\text{put}, i, j, k}^-)_{i, j, k}, (\pi_{\text{put}, i, j, k}^+)_{i, j, k} \right)$$

are given by

$$\mathcal{C} \left(\Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}, \boldsymbol{\pi} \right) := a + \sum_{i=1}^n \sum_{k=1}^d \sum_{j=1}^{n_{ik}^{\text{opt}}} (c_{ijk}^+ \pi_{\text{call}, i, j, k}^+ - c_{ijk}^- \pi_{\text{call}, i, j, k}^-) + \sum_{i=1}^n \sum_{k=1}^d \sum_{j=1}^{n_{ik}^{\text{opt}}} (p_{ijk}^+ \pi_{\text{put}, i, j, k}^+ - p_{ijk}^- \pi_{\text{put}, i, j, k}^-). \quad (\text{II.2})$$

For a strategy $\Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}$ with parameters $(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)_{i, j, k}$ we introduce the function

$$\Sigma(\mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k) := \sum_{i=1}^n \sum_{k=1}^d \sum_{j=1}^{n_{ik}^{\text{opt}}} (c_{ijk}^+ + c_{ijk}^- + p_{ijk}^+ + p_{ijk}^-) + \sum_{k=1}^d |\Delta_0^k| + \sum_{i=1}^{n-1} \sum_{k=1}^d \|\Delta_i^k\|_\infty,$$

where $\|\cdot\|_\infty$ denotes the supremum norm. Imposing a universal upper bound on the function Σ , i.e., $\Sigma(\cdot) \leq \mathfrak{B} < \infty$ for some $\mathfrak{B} \in \mathbb{R}_+$, relates to a restriction on the maximal position an investor is willing/allowed to invest. We want to value a derivative with payoff $\Phi \in B(\mathbb{R}_+^{nd}, \mathbb{R})$. Hence, given strikes \mathbf{K} , spot prices \mathbf{S}_{t_0} , and bid-ask prices $\boldsymbol{\pi}$, our goal is to solve the following super-hedging problem

$$\begin{aligned} \overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi) := & \inf_{\substack{a \in \mathbb{R}, \\ \mathbf{c}_{ijk}, \mathbf{p}_{ijk} \in \mathbb{R}_+^2, \\ (\Delta_i^k) \in B(\mathbb{R}_+^{id}, \mathbb{R})}} \left\{ \mathcal{C} \left(\Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}, \boldsymbol{\pi} \right) \text{ s.t. } \Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\mathbf{s}) \geq \Phi(\mathbf{s}) \text{ for all } \mathbf{s} \in [0, B]^{nd}, \right. \\ & \left. \text{and } \Sigma(\mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k) \leq \mathfrak{B} \right\}, \end{aligned} \quad (\text{II.3})$$

for some bounds $\mathfrak{B}, B \in (0, \infty]$, where the bound B corresponds to a restriction of the form $S_{t_i}^k \leq B$ for all i, k . It is economically reasonable to assume a large but finite B for the securities under consideration, since it imposes no severe restriction and reduces artificial high prices which were not realistic in practice.¹³ A solution of (II.3) defines the largest model-independent arbitrage-free price of the derivative Φ and simultaneously comes with a strategy that enables to exploit arbitrage if prices for Φ lie above this price bound.

In analogy to (II.3), the smallest model-independent arbitrage-free price of Φ is given by the corresponding sub-hedging problem

$$\begin{aligned} \underline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi) := & \sup_{\substack{a \in \mathbb{R}, \\ \mathbf{c}_{ijk}, \mathbf{p}_{ijk} \in \mathbb{R}_+^2, \\ (\Delta_i^k) \in B(\mathbb{R}_+^{id}, \mathbb{R})}} \left\{ \mathcal{C} \left(\Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}, \boldsymbol{\pi} \right) \text{ s.t. } \Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\mathbf{s}) \leq \Phi(\mathbf{s}) \text{ for all } \mathbf{s} \in [0, B]^{nd}, \right. \\ & \left. \text{and } \Sigma(\mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k) \leq \mathfrak{B} \right\}. \end{aligned}$$

B. Training a neural network for option valuation

Next, we focus on the supervised learning approach we pursue in this paper. This approach is implemented using neural networks, thus we start this section with a short exposition on neural networks which can be found in similar form in [18], [5], [12], [23], [24], [26], or in every standard textbook on the topic (e.g. [9], [28], or [30]).

¹²To simplify the presentation we assume zero interest rates and zero dividend yields.

¹³We still allow a priori $\mathfrak{B}, B = \infty$ in case one does not want to make restrictions on the trading strategies or exclude unbounded price paths.

1) *Neural networks*: In the following we consider a fully-connected neural network which is for input dimension $d_{\text{in}} \in \mathbb{N}$, output dimension $d_{\text{out}} \in \mathbb{N}$, and number of layers $l \in \mathbb{N}$ defined as a function of the form

$$\begin{aligned} \mathbb{R}^{d_{\text{in}}} &\rightarrow \mathbb{R}^{d_{\text{out}}} \\ \mathbf{x} &\mapsto \mathbf{A}_l \circ \varphi_l \circ \mathbf{A}_{l-1} \circ \dots \circ \varphi_1 \circ \mathbf{A}_0(\mathbf{x}), \end{aligned} \quad (\text{II.4})$$

where $(\mathbf{A}_i)_{i=0,\dots,l}$ are functions of the form

$$\mathbf{A}_0 : \mathbb{R}^{d_{\text{in}}} \rightarrow \mathbb{R}^{h_1}, \quad \mathbf{A}_i : \mathbb{R}^{h_i} \rightarrow \mathbb{R}^{h_{i+1}} \text{ for } i = 1, \dots, l-1, \text{ (if } l > 1), \quad \mathbf{A}_l : \mathbb{R}^{h_l} \rightarrow \mathbb{R}^{d_{\text{out}}}, \quad (\text{II.5})$$

and where for $i = 1, \dots, l$ we have $\varphi_i(x_1, \dots, x_{h_i}) = (\varphi(x_1), \dots, \varphi(x_{h_i}))$, with $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ being a non-constant function called *activation function*. Here $\mathbf{h} = (h_1, \dots, h_l) \in \mathbb{N}^l$ denotes the dimensions (the number of neurons) of the hidden layers, also called *hidden dimension*. Moreover, for all $i = 0, \dots, l$, the function \mathbf{A}_i is assumed to have an affine structure of the form

$$\mathbf{A}_i(\mathbf{x}) = \mathbf{M}_i \mathbf{x} + \mathbf{b}_i$$

for some matrix $\mathbf{M}_i \in \mathbb{R}^{h_{i+1} \times h_i}$ and some vector $\mathbf{b}_i \in \mathbb{R}^{h_{i+1}}$, where $h_0 := d_{\text{in}}$ and $h_{l+1} := d_{\text{out}}$. We then denote by $\mathfrak{N}_{d_{\text{in}}, d_{\text{out}}}^{l, \mathbf{h}}$ the set of all neural networks with input dimension d_{in} , output dimension d_{out} , l hidden layers, and hidden dimension \mathbf{h} . Moreover, we consider the set of all neural networks with input dimension d_{in} , output dimension d_{out} , a fixed amount of l hidden layers, but unspecified hidden dimension

$$\mathfrak{N}_{d_{\text{in}}, d_{\text{out}}}^l := \bigcup_{\mathbf{h} \in \mathbb{N}^l} \mathfrak{N}_{d_{\text{in}}, d_{\text{out}}}^{l, \mathbf{h}},$$

as well as the set of all neural networks mapping from $\mathbb{R}^{d_{\text{in}}}$ to $\mathbb{R}^{d_{\text{out}}}$ with an unspecified amount of hidden layers

$$\mathfrak{N}_{d_{\text{in}}, d_{\text{out}}} := \bigcup_{l \in \mathbb{N}} \mathfrak{N}_{d_{\text{in}}, d_{\text{out}}}^l.$$

One fundamental result that is of major importance for the approximation of functions through neural networks is the universal approximation theorem from e.g. [37, Theorem 2], stating that, given some mild assumption on the activation function φ , every continuous function can be approximated arbitrarily well by neural networks on compact subsets.

Proposition II.1 (Universal approximation theorem for continuous functions [37]). *Assume that $\varphi \in C(\mathbb{R}, \mathbb{R})$ and that φ is not constant, then for any compact $\mathbb{K} \subset \mathbb{R}^{d_{\text{in}}}$ the set $\mathfrak{N}_{d_{\text{in}}, d_{\text{out}}} |_{\mathbb{K}}$ is dense in $C(\mathbb{K}, \mathbb{R}^{d_{\text{out}}})$ w.r.t. the topology of uniform convergence on $C(\mathbb{K}, \mathbb{R}^{d_{\text{out}}})$.*

Popular examples for activation functions are the *ReLU* function given by $\varphi(x) := \max\{x, 0\}$ or the logistic function $\varphi(x) := 1/(1 + e^{-x})$, which fulfil the assumptions of Proposition II.1. Further, we remark that the original statement from [37, Theorem 2] only covers output dimension $d_{\text{out}} = 1$, and $l = 1$ hidden layer, but can indeed be generalized to the above statement, compare e.g. [40, Theorem 3.2.].

2) *Approximation of the super-replication functional through neural networks*: We consider for $i = 1, \dots, S$, where $S \in \mathbb{N}$ denotes the number of samples, input data of the form

$$\mathbf{X}_i = (\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)$$

and we aim at predicting via an appropriately trained neural network the following target

$$\mathbf{Y}_i = \left(\underline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta), \overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta) \right),$$

for a parametrized family $\{\Phi_\theta, \theta \in \Theta\}$. If we are additionally interested in predicting the optimal super-replication strategy, then $\widetilde{\mathbf{Y}}_i$ contains instead the associated parameters of the strategies, i.e.,

$$\widetilde{\mathbf{Y}}_i = \left(a, (\mathbf{c}_{ijk})_{i,j,k}, (\mathbf{p}_{ijk})_{i,j,k}, (\Delta_i^k)_{i,k} \right),$$

for the minimal super-replication strategy $\Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}$ (and analogue also for the maximal sub-replication strategy), which implicitly also contains the minimal super-replication price by calculating the corresponding cost using (II.2). However, after having trained a neural network to predict $\widetilde{\mathbf{Y}}_i$ given market data \mathbf{X}_i , due to a different training error, the implied price bounds are expected to differ to a larger extent from \mathbf{Y}_i than those from a neural network which directly predicts the prices \mathbf{Y}_i , see also Example II.6 in which we compare both approaches.

According to Proposition II.1, a trained neural network can be used to predict price bounds and optimal strategies if price bounds and strategies, respectively, are continuous functions of the input, i.e., of \mathbf{X}_i . The following result stated in Theorem II.2 ensures that, under mild assumptions which we discuss subsequently in Remark II.3, this requirement is fulfilled. For this, we denote for all $k \in \mathbb{N}$ by $\|\cdot\|_k$ some norm on \mathbb{R}^k . Since all norms on Euclidean spaces are equivalent, the specific choice of the norm is irrelevant for the following assertions. The induced metric for $k \in \mathbb{N}$ is denoted by $d_k(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_k$. Moreover, we define for every $B \in (0, \infty)$ the norm $\|f\|_{\infty, B} := \sup_{x \in [0, B]^{n_d}} |f(x)|$ and $d_{\infty, B}(f, g) := \|f - g\|_{\infty, B}$ for $f, g \in C(\mathbb{R}_+^{n_d}, \mathbb{R})$. In the case $B = \infty$ we set

$$d_{\infty, \infty}(f, g) := \frac{\|f - g\|_{\infty}}{1 + \|f - g\|_{\infty}}.$$

To the best of our knowledge, the following Theorem II.2 proves for the first time a continuous relation between the market inputs and the corresponding price bounds. This novel result justifies to apply neural networks to determine model-independent price bounds and hence provides a significant contribution to the literature.

Theorem II.2. *Let $B \in (0, \infty]$, $\mathfrak{B} \in (0, \infty)$, and $M := \sum_{i=1}^n \sum_{k=1}^d n_{ik}^{\text{opt}}$. Let $\{\Phi_\theta, \theta \in \Theta\}$, for some $\Theta \subset \mathbb{R}^p$ and $p \in \mathbb{N}$, be a (parametric) family of functions in $C(\mathbb{R}_+^{nd}, \mathbb{R})$ such that*

$$\begin{aligned} (\Theta, d_p) &\rightarrow \left(C(\mathbb{R}_+^{nd}, \mathbb{R}), d_{\infty, B} \right) \\ \theta &\mapsto \Phi_\theta \end{aligned} \quad (\text{II.6})$$

is continuous and let $N_{\text{input}} := 2M + 4M + d + p$. Then, the following holds.

(a) Let $\mathbb{K}_1 \subset \mathbb{R}_+^{2M} \times \mathbb{R}^{4M} \times \mathbb{R}_+^d \times \Theta$ be a compact set such that both

$$\underline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta), \overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta) \in (-\infty, \infty) \quad (\text{II.7})$$

holds for all $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \in \mathbb{K}_1$. Then, the map

$$\begin{aligned} (\mathbb{K}_1, d_{N_{\text{input}}}) &\rightarrow (\mathbb{R}^2, d_2) \\ (\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) &\mapsto \left(\underline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta), \overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta) \right) \end{aligned}$$

is continuous.

(b) Let \mathbb{K}_1 be defined as in (a). Then, for all $\varepsilon > 0$ there exists a neural network $\mathcal{N}_1 \in \mathfrak{N}_{N_{\text{input}}, 2}$ such that for all $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \in \mathbb{K}_1$ it holds

$$\left\| \mathcal{N}_1(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) - \left(\underline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta), \overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta) \right) \right\|_2 < \varepsilon. \quad (\text{II.8})$$

(c) Let $n = 1$. Let $\mathbb{K}_2 \subset \mathbb{R}_+^{2M} \times \mathbb{R}^{4M} \times \mathbb{R}_+^d \times \Theta$ be a compact set such that for all $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \in \mathbb{K}_2$ we have that (II.7) holds and $\overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta)$ is attained by a unique strategy

$$\left(a^*, (c_{1jk}^*)_{j,k}, (p_{1jk}^*)_{j,k}, (\Delta_0^{k*})_k \right) (\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)$$

satisfying

$$\sup_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \in \mathbb{K}_2} |a^*(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)| < \infty. \quad (\text{II.9})$$

Then the map

$$\begin{aligned} (\mathbb{K}_2, d_{N_{\text{input}}}) &\rightarrow (\mathbb{R}^{1+4M+d}, d_{1+4M+d}) \\ (\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) &\mapsto \left(a^*, (c_{1jk}^*)_{j,k}, (p_{1jk}^*)_{j,k}, (\Delta_0^{k*})_k \right) (\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \end{aligned}$$

is continuous.

(d) Let $n = 1$ and let \mathbb{K}_2 be defined as in (c). Then, for all $\varepsilon > 0$ there exists a neural network $\mathcal{N}_2 \in \mathfrak{N}_{N_{\text{input}}, 1+4M+d}$ such that for all $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \in \mathbb{K}_2$ it holds

$$\left\| \mathcal{N}_2(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) - \left(a^*, (c_{1jk}^*)_{j,k}, (p_{1jk}^*)_{j,k}, (\Delta_0^{k*})_k \right) (\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \right\|_{1+4M+d} < \varepsilon. \quad (\text{II.10})$$

Proof. See Section IV. \square

Remark II.3. (a) Assumption (II.7) means that the market with its parameters $\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}$ is arbitrage-free, compare e.g. [47, Assumption 2.1. and Theorem 2.4.] for the case $n = 1$, [10, Definition 1.1. and Theorem 5.1.] for the multi-period case with traded options, and [15, Theorem 2.1.] for the general case with market frictions. Note that assuming an arbitrage-free market is a necessity to determine arbitrage-free price bounds of financial derivatives. Indeed, if the market offers arbitrage, then we can identify a trading strategy fulfilling $\Psi_{(a, c_{ijk}, p_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\mathbf{s}) \geq 0$ for some parameters $(a, c_{ijk}, p_{ijk}, \Delta_i^k)_{i,j,k}$ with price $\mathcal{C} \left(\Psi_{(a, c_{ijk}, p_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\boldsymbol{\pi}) \right) < 0$. Now, consider a super-replication strategy $(\tilde{a}, \tilde{c}_{ijk}, \tilde{p}_{ijk}, \tilde{\Delta}_i^k)_{i,j,k}$ of some derivative Φ satisfying $\Psi_{(\tilde{a}, \tilde{c}_{ijk}, \tilde{p}_{ijk}, \tilde{\Delta}_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\mathbf{s}) \geq \Phi(\mathbf{s})$. Then we have for all $\lambda > 0$ that

$$\Psi_{(\tilde{a}, \tilde{c}_{ijk}, \tilde{p}_{ijk}, \tilde{\Delta}_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\mathbf{s}) + \lambda \cdot \Psi_{(a, c_{ijk}, p_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\mathbf{s}) = \Psi_{(\tilde{a} + \lambda a, \tilde{c}_{ijk} + \lambda c_{ijk}, \tilde{p}_{ijk} + \lambda p_{ijk}, \tilde{\Delta}_i^k + \lambda \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\mathbf{s}) \geq \Phi(\mathbf{s})$$

meaning that $(\tilde{a} + \lambda a, \tilde{c}_{ijk} + \lambda c_{ijk}, \tilde{p}_{ijk} + \lambda p_{ijk}, \tilde{\Delta}_i^k + \lambda \Delta_i^k)$ is another super-replication strategy, whose price is given by

$$\mathcal{C} \left(\Psi_{(\tilde{a} + \lambda a, \tilde{c}_{ijk} + \lambda c_{ijk}, \tilde{p}_{ijk} + \lambda p_{ijk}, \tilde{\Delta}_i^k + \lambda \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\boldsymbol{\pi}) \right) = \mathcal{C} \left(\Psi_{(\tilde{a}, \tilde{c}_{ijk}, \tilde{p}_{ijk}, \tilde{\Delta}_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\boldsymbol{\pi}) \right) + \lambda \cdot \mathcal{C} \left(\Psi_{(a, c_{ijk}, p_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\boldsymbol{\pi}) \right). \quad (\text{II.11})$$

By scaling up $\lambda > 0$, we see from (II.11) that the corresponding price decreases, which in turn implies $\overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi) \ll 0$ (for \mathfrak{B} large enough).

With an analogue argument we conclude that if the market offers arbitrage, then we have $\underline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi) \gg 0$, preventing the computation of reasonable price bounds.

- (b) In an arbitrage-free market, a necessary requirement for the existence of a unique optimizer, as assumed in Theorem II.2 (c), is that the considered market instruments are non-redundant, i.e., that the payoffs of the market instruments are linear independent. In our case, this means that to avoid ambiguity of minimal super-replication strategies, one should only consider put options that are written on other strikes than the ones for the call options under consideration.
- (c) As a canonical example for a parametric family of payoff functions, we consider for example basket call options with payoffs

$$\left\{ \Phi_\theta = \max \left\{ \sum_{i=1}^n \sum_{k=1}^d w_i^k S_{t_i}^k - L, 0 \right\} \text{ where } \theta \in \Theta := \left\{ ((w_i^k)_{i,k}, L) \right\} = \mathbb{R}^{nd} \times \mathbb{R}_+ \right\},$$

i.e., the strike L and the weights $(w_i^k)_{i,k}$ are inputs to the trained neural network. For any $0 < B < \infty$, we have the continuity of the map

$$(\Theta, d_{nd+1}) \rightarrow \left(C(\mathbb{R}_+^{nd}, \mathbb{R}), d_{\infty, B} \right), \theta \mapsto \Phi_\theta.$$

Thus, we can find a neural network which fulfils (II.8) with respect to $\underline{D}^{\mathfrak{B}, B}$ and $\overline{D}^{\mathfrak{B}, B}$. We remark that assuming a uniform large bound B on the possible values of $S_{t_i}^k$ imposes no severe constraint for practical applications with real market data and allows to reduce the difference between the no-arbitrage price bounds by not considering unbounded prices which are unrealistic in practice. For further examples of parametric families of payoff functions, e.g., best-of-call options or call-on-max options, we refer to [47, Example 3.2. (i)–(vi)].

- (d) Theorem II.2 is also applicable to a single pre-specified continuous payoff function Φ when setting $\Phi_\theta = \Phi$ for all $\theta \in \Theta$.
- (e) Note that we restrict the assertion of Theorem II.2 (d) to $n = 1$ to make sure that Δ_i^k , which is the output of the neural network, is a number, not a function.
- (f) An analogue result as in Theorem II.2 (c) and Theorem II.2 (d) for optimal sub-hedging strategies can be obtained in the same way.
- (g) We highlight that our approach computes model-independent price bounds, i.e., no assumptions on underlying financial models are imposed and market prices of call and put options are considered as exogenous inputs. Nevertheless, in a frictionless market, it is possible for each arbitrage-free sample $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)$ to find a stochastic model (in which call option prices are given endogenously) that is consistent with this data, compare e.g. [34]. Therefore one could understand the model-independent price bounds obtained with respect to the given sample $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)$ also as the prices of hedging strategies in such a consistent model. However, note that such a model (expressed by a probability measure) would be different for each considered sample.

Finally, Algorithm 1 describes, relying on the results from Theorem II.2, how one can train a neural network which approximates these price bounds.

Algorithm 1: Training of a neural network via back-propagation for the computation of the price bounds $\underline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta)$, $\overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta)$ of a class of financial derivatives $\{\Phi_\theta\}_{\theta \in \Theta}$.

Data: Call and put option prices (bid and ask) on different securities and maturities; Associated strikes, maturities, and spot prices;

Input : Algorithm to compute price bounds of exotic derivatives; Family $\{\Phi_\theta, \theta \in \Theta\}$ of payoff functions $\Phi_\theta : \mathbb{R}_+^{nd} \rightarrow \mathbb{R}$ fulfilling the requirements of Theorem II.2; Hyper-parameters of the neural network; Number n_{subset} of considered functions from $\{\Phi_\theta, \theta \in \Theta\}$ for each sample; Transaction costs $\kappa \geq 0$; Bounds \mathfrak{B} and B ;

for each sample $(\mathbf{K}_i, \boldsymbol{\pi}_i, \mathbf{S}_{t_0_i})$ of data considering exactly n maturities and d securities do

$\widetilde{\mathbf{X}}_i \leftarrow (\mathbf{K}_i, \boldsymbol{\pi}_i, \mathbf{S}_{t_0_i})$;

end

$\mathcal{S} \leftarrow \#\{\widetilde{\mathbf{X}}_i\}$; // Assign number of samples and call it \mathcal{S} .

for i in $\{1, \dots, \mathcal{S}\}$ do

 Generate a (random) subset $\{\Phi_{\theta_j}, j = 1, \dots, n_{\text{subset}}\} \subset \{\Phi_\theta, \theta \in \Theta\}$;

for j in $\{1, \dots, n_{\text{subset}}\}$ do

 Compute for $\widetilde{\mathbf{X}}_i$ the corresponding price bounds $\left(\underline{D}_{(\mathbf{K}_i, \boldsymbol{\pi}_i, \mathbf{S}_{t_0_i})}^{\mathfrak{B}, B}(\Phi_{\theta_j}), \overline{D}_{(\mathbf{K}_i, \boldsymbol{\pi}_i, \mathbf{S}_{t_0_i})}^{\mathfrak{B}, B}(\Phi_{\theta_j}) \right)$;

$\mathbf{X}_{(i-1)n_{\text{subset}}+j} \leftarrow (\widetilde{\mathbf{X}}_i, \theta_j)$;

$\mathbf{Y}_{(i-1)n_{\text{subset}}+j} \leftarrow \left(\underline{D}_{(\mathbf{K}_i, \boldsymbol{\pi}_i, \mathbf{S}_{t_0_i})}^{\mathfrak{B}, B}(\Phi_{\theta_j}), \overline{D}_{(\mathbf{K}_i, \boldsymbol{\pi}_i, \mathbf{S}_{t_0_i})}^{\mathfrak{B}, B}(\Phi_{\theta_j}) \right)$;

 // In this step it would be possible to use any algorithm that can compute these bounds reliably, see Remark II.4

end

end

Train with back-propagation ([50]) a neural network $\mathcal{N} \in \mathfrak{N}_{d_{\text{in}}, 2}$ with a sufficient number of neurons and hidden layers such that $\mathcal{N}(\widetilde{\mathbf{X}}_i) \approx \mathbf{Y}_i$;

Output: Trained neural network $\mathcal{N} \in \mathfrak{N}_{N_{\text{input}}, 2}$ with N_{input} as in Theorem II.2;

Remark II.4. To compute model-independent price bounds given option prices, we can use e.g. a linear programming approach based on grid discretization as proposed in [23], [29], or [31]. If the payoff function only depends on one future maturity and is continuous piecewise affine (CPWA, see e.g. [47, Example 3.2.]) we can also use the numerically very efficient algorithm proposed in [47]. If the payoff function fulfils a so-called martingale Spence–Mirrlees condition¹⁴ one can apply the algorithm presented in [32]. For multiple time-steps, another possibility is to apply the penalization approach presented in [24]. The minimization is then performed using a stochastic gradient descent algorithm with some penalization parameter γ which enforces the optimizing strategy to be a super-hedge.

C. Examples

In this section we present, in selected examples, the results of our approach when applied to real market data.

1) *Training data:* We consider for the training of all neural networks financial market data received from *Thomson Reuters Eikon* that was observed on 10th June 2020. The data includes bid and ask prices on call options written on all 500 constituents of the American stock market index S&P 500. Note that, in Example II.6, we also predict the optimal super-hedging strategy and not only optimal price bounds. Thus, to avoid ambiguity of the optimal strategy, as explained in Remark II.3, we do not consider any put options there. We consider for each constituent and each available maturity of an option the 20 most liquid strikes, i.e., the bid and ask prices of options with the highest trading volume.

2) *Test data:* For testing the trained neural networks we consider - as for the training data - option prices on all constituents of the S&P 500. The data was observed on 23rd August 2020. We highlight that, in particular, the test data comes from a different dataset than the training data.

3) *Implementation:* The training of each of the neural networks is performed using the back-propagation algorithm ([50]) with an Adam optimizer ([41]) implemented in *Python* using *Tensorflow* ([1]). For the optimization with the Adam optimizer we use a batch size of 256. The architecture involves a L^2 -loss function, the neural networks comprise 3 hidden layers with 512 neurons each and *ReLU* activation functions. The samples are normalized before training with a min-max scaler. Moreover, we assign 10% of the training data to a validation set to be able to apply early stopping (compare [28, Chapter 7.8.]) to prevent overfitting to the training data. To reduce the internal covariate shift of the neural network and to additionally regularize it, we apply batch normalization ([39]) after each layer. All the codes related to the examples below¹⁵, as well as the trained neural networks are provided under https://github.com/juliansester/deep_model_free_pricing. For all examples we assume no transaction costs, i.e., we have $\kappa = 0$.

Example II.5 (Training of the valuation of call options given prices of other call options). We want to train the valuation of call options for arbitrary strikes, i.e., we consider payoff functions from the set

$$\{\Phi_\theta, \theta \in \Theta\} = \left\{ \Phi_L(S_{t_1}^1) := \max \{S_{t_1}^1 - L, 0\}, \text{ with } L \in \mathbb{R}_+ \right\}.$$

Note that the assumptions of Theorem II.2 are met for any $B \in (0, \infty]$, $\mathfrak{B} < \infty$, which we choose therefore large enough to not impose a restriction. Thus Theorem II.2 (b) ensures that we can train a single neural network for the above mentioned parameterized family of payoff functions.

In this example, a single sample \mathbf{X}_i consists of 62 total entries which comprise 20 bid prices, 20 ask prices, 20 associated strikes of call options, as well as the underlying spot price and the strike L of the call option Φ_L which we want to price.

We apply Algorithm 1 to train a neural network, in particular, for each of the prices from the training data, we create several different random strikes L , for which we compute a corresponding \mathbf{Y}_i which consists of lower and upper model-independent price bounds for all samples according to the algorithm from [47]. With this methodology we create a training set with 100000 samples. We then train, as described in Section II-C3, a neural network using back-propagation and test it on the test data, described in Section II-C2, which was observed at a later date (August 2020). The test set consists of 10000 samples.

The results of the training yield a mean absolute error of 2.2033 as well as a mean squared error of 25.8779 on the test set and are depicted in Figure 2a. To be able to compare the error independent of the size of the spot price of the underlying security, we report a mean absolute error of 0.0111 when dividing the predicted prices by the spot prices. We call this value the relative mean absolute error. The corresponding squared distance after division by the spot prices amounts to 0.0003 and is called relative mean squared error. Compare also Figure 2b, where we depict the relative error of each sample in the test set, i.e., the difference of each prediction from its target value, after division with the spot price.

Example II.6 (Optimal strategies of basket options). We consider payoff functions of basket options written on two assets, i.e., the class of payoff functions is defined through

$$\{\Phi_\theta, \theta \in \Theta\} = \left\{ \Phi_{w_1, w_2, L}(S_{t_1}^1, S_{t_1}^2) = \max \{w_1 S_{t_1}^1 + w_2 S_{t_1}^2 - L, 0\}, \text{ with } w_1, w_2, L \in \mathbb{R}_+ \right\}.$$

When $B, \mathfrak{B} < \infty$, the assumptions of Theorem II.2 (b), respectively those of Remark II.3 (c), are fulfilled. For each of the considered market prices from the training set and test set, respectively, we create in accordance with Algorithm 1 several different weights w_1, w_2 and some strike L . Thus, a sample \mathbf{X}_i consists of the spot prices $S_{t_0}^1, S_{t_0}^2$, the generated values w_1, w_2, L as well as of bid and ask prices with associated strikes of both assets, i.e., in total each \mathbf{X}_i consists of 125 numbers.

We aim at simultaneously predicting the minimal super-replication strategy and the maximal sub-replication strategy. Therefore, we compute the parameters of the strategies attaining the price bounds according to the algorithm from [47]. Thus, in our case

¹⁴This means that $\frac{\partial^3}{\partial x y^2} \Phi(x, y)$ exists and satisfies $\frac{\partial^3}{\partial x y^2} \Phi(x, y) > 0$ for all x, y .

¹⁵For copyright reasons we can only provide the used code, but we cannot provide the used data.

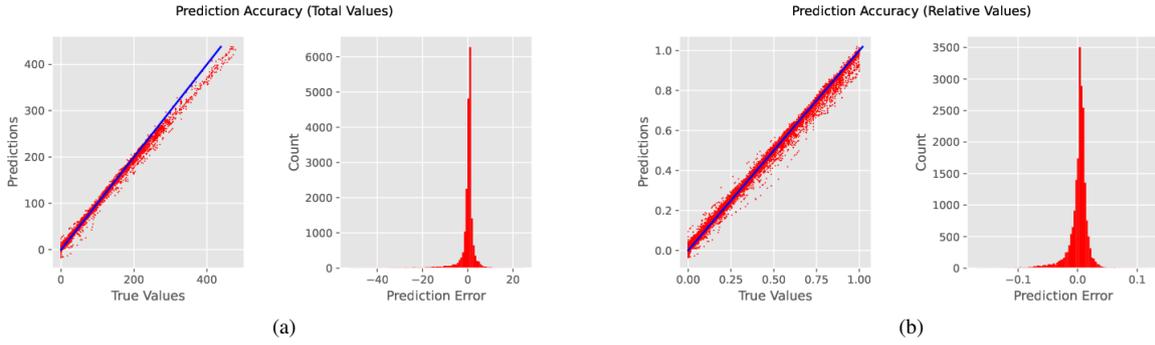


Fig. 2. (a): This figure illustrates the accuracy of the predictions on the test set. The left panel shows a plot of all target values (x-values) and its predictions (y-values), the right panel depicts a histogram of the prediction error, i.e., the error between target values and predicted values. (b): This figure shows the accuracy of the predictions of call option prices on the test set when considering the relative error, i.e., when dividing the predicted prices \mathbf{Y}_i by the corresponding spot prices.

each sample \mathbf{Y}_i comprises 86 values, which constitute, for both lower and upper bound, of the initial investment a (1 parameter), the buy and sell positions in call options¹⁶ $(c_{1jk})_{j,k}^{j=1,\dots,20, k=1,2}$ (40 parameters) and the investment positions in the underlying securities $(\Delta_0^k)_{k=1,2}$ (2 parameters).

Moreover, training a neural network to learn the relationship between market parameters and optimal super-replication strategy (without the price bound) is possible due to Theorem II.2 (d), which implies that the difference in absolute values between predicted parameters $(a, (c_{1jk})_{j,k}, (\Delta_0^k)_k)$ and true parameters $(a^*, (c_{1jk}^*)_{j,k}, (\Delta_0^{k*})_k)$ of the optimal strategy should not differ significantly after training. We train the neural network on 150000 samples, test it on 10000 samples, and obtain indeed a small relative mean absolute error of 0.0015. We highlight that the training set remains the same, in particular does not consist of trading strategies.

After having trained the neural network to predict the minimal super-hedging strategies, we are able to derive from these strategies the optimal price bounds using (II.2) and compare it with the predictions from a neural network which is trained on predicting lower and upper price bounds directly instead of predicting optimal strategies. In Figure 3 we show that however, as expected, the neural network that predicts prices directly performs by far better than the price bound predictions that are derived via (II.2) from the trained strategies, when evaluated on the test set. The relative¹⁷ mean absolute error of the direct prediction of the price bounds is 0.0319, whereas the relative mean absolute error of the prediction relying on the strategies is 0.1804. The corresponding relative mean squared errors are 0.0148 and 0.5045, respectively.

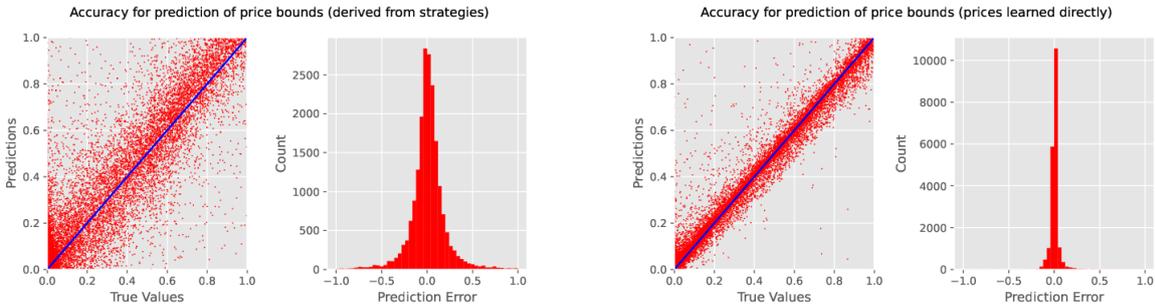


Fig. 3. This figure compares the accuracy of prediction of price bounds derived from the trained strategies using (II.2) (left) with those predictions from neural networks that are trained to predict the prices directly (right). We depict the relative error of the predictions by dividing the prediction error through the weighted sum of spot prices, where the weights are according to the weights in the payoff of the basket option.

The larger approximation error when approximating first the strategies by a neural network and then deriving price bounds from this approximation can be explained as follows.

When approximating the price bounds directly, then we have, after sufficient training of a neural network, according to Theorem II.2 (b), a maximal absolute approximation error of order ε_1 between the upper price bound and the output of the neural network, given a tolerance level of $\varepsilon_1 > 0$.

¹⁶ Here, by slight abuse of notation, we denote by $c_{ljk} \in \mathbb{R}$ the net position invested in the option, i.e., c_{ljk} is also allowed to attain negative values.

¹⁷Note that here the relative error refers to the error after division with the *weighted* sum of the spot prices, where the weights are determined by the weights in the payoff of the basket option.

In contrast, when approximating the optimal super-replication strategy $\left(a^*, (c_{1jk}^*)_{j,k}, (p_{1jk}^*)_{j,k}, (\Delta_0^{k*})_k\right) (\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)$ by the output of a neural network, denoted by $\left(a^{\mathcal{NN}}, (c_{1jk}^{\mathcal{NN}})_{j,k}, (p_{1jk}^{\mathcal{NN}})_{j,k}, (\Delta_0^{k\mathcal{NN}})_k\right) (\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)$ at some tolerance level $\varepsilon_2 > 0$, then the absolute error between the upper price bound $\overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathbb{B}, \mathbb{B}}(\Phi_\theta) = C\left(\Psi_{(a^*, c_{1jk}^*, p_{1jk}^*, \Delta_0^{k*})}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\boldsymbol{\pi})\right)$ and the price $C\left(\Psi_{(a^{\mathcal{NN}}, c_{1jk}^{\mathcal{NN}}, p_{1jk}^{\mathcal{NN}}, \Delta_0^{k\mathcal{NN}})}^{(\mathbf{K}, \mathbf{S}_{t_0})}(\boldsymbol{\pi})\right)$ derived from the approximated strategy computes by (II.2) as

$$\left| a^* - a^{\mathcal{NN}} + \sum_{k=1}^d \sum_{j=1}^{n_{1k}^{\text{opt}}} \left((c_{1jk}^+)^* - c_{1jk}^{+\mathcal{NN}} \right) \pi_{\text{call}, 1, j, k}^+ - (c_{1jk}^-)^* - c_{1jk}^{-\mathcal{NN}} \right) \pi_{\text{call}, 1, j, k}^- + \sum_{k=1}^d \sum_{j=1}^{n_{1k}^{\text{opt}}} \left((p_{1jk}^+)^* - p_{1jk}^{+\mathcal{NN}} \right) \pi_{\text{put}, 1, j, k}^+ - (p_{1jk}^-)^* - p_{1jk}^{-\mathcal{NN}} \right) \pi_{\text{put}, 1, j, k}^- \right|$$

which is, according to (II.10), as large as of order

$$\varepsilon_2 \cdot \max \left\{ \max_{j,k} \pi_{\text{call}, 1, j, k}^+, \max_{j,k} \pi_{\text{call}, 1, j, k}^-, \max_{j,k} \pi_{\text{put}, 1, j, k}^+, \max_{j,k} \pi_{\text{put}, 1, j, k}^- \right\}.$$

However, $C := \max \left\{ \max_{j,k} \pi_{\text{call}, 1, j, k}^+, \max_{j,k} \pi_{\text{call}, 1, j, k}^-, \max_{j,k} \pi_{\text{put}, 1, j, k}^+, \max_{j,k} \pi_{\text{put}, 1, j, k}^- \right\}$ is typically significantly larger than 1 (the largest call option price in the considered test set was 349\$). Therefore, to obtain similar approximation results for the prices derived from strategies as those when approximating prices directly, one needs to consider a tolerance level of $\varepsilon_2 = \frac{\varepsilon_1}{C}$ which usually is significantly smaller than ε_1 . In our case, the training set was not large enough to obtain such a high precision in the approximation of the strategies¹⁸.

The above outlined reasoning is supported by a comparatively large approximation error which is already observed on the training set. When predicting prices from strategies via (II.2), we have a relative mean absolute error of 0.1311 on the training set, which is only marginally smaller than the relative mean absolute error of 0.1804 which we reported for predictions on the test set. In comparison, the relative mean absolute error for directly predicting prices is 0.0160 on the training set, whereas we reported a relative mean absolute error of 0.0319 on the test set.

We conclude that, even though it is theoretically possible to derive price bounds with an arbitrarily high precision from approximated strategies, in practice it turns out to be more efficient (and requires a smaller training set) to train a neural network that approximates the price bounds directly if one is only interested in the prediction of these.

Example II.7 (Neural networks trained for basket options applied to call options). We reconsider the trained neural network from Example II.6 predicting the price bounds of basket options written on two assets.

- We now test this neural network on the same test set as in Example II.5 which takes into account call options instead of basket options. To be able to apply the neural network that takes inputs with 125 entries, we modify the original samples \mathbf{X}_i from Example II.5 (originally containing 62 entries) by duplicating the original entries (except for the strike of the call option) and by additionally adding weights of 0.5 and 0.5 as well as the original strike, leading to $2 \cdot 61 + 2 + 1 = 125$ entries. This means that we consider a basket option with payoff of the form $\max\{0.5S_{t_1}^1 + 0.5S_{t_1}^2 - K, 0\} = \max\{S_{t_1}^1 - K, 0\}$ which is a call option.
- We then observe a relative mean absolute error of 0.0230 and a relative mean squared error of 0.0023 (in comparison with 0.0111 and 0.0003, respectively for the neural network from Example II.5 that was only trained on call options).
- This shows that the more general neural network performs reasonably well also on the more specific payoffs, but is less specifically trained and therefore is outperformed by the neural network solely trained on call options.
- To improve the performance of this neural network, we retrain the neural network by adding in addition to the 150000 samples containing basket options from Example II.6 also the 100000 samples containing call options from Example II.5.
- Indeed, the result is a trained neural network that performs well on both basket options and call options. On the test set for call options we obtain a relative mean absolute error of 0.0112 and a relative mean squared error of 0.0004 (compared to 0.0111 and 0.0003 obtained in Example II.5). On the test set for basket options we compute a relative mean absolute error of 0.0292 and a relative mean squared error of 0.0374 (compared to 0.0319 and 0.0148 obtained in Example II.6). See also Figure 4, where we depict the accuracy of the predictions before and after adding the additional samples. This shows that training the neural network on an additional task (call option price bounds) did improve notably the approximation quality on this specific task while the approximation quality on the well-known task (basketed option price bounds) is barely affected.

This example showcases that it is possible to train a single neural network that approximates price bounds reasonably well for different types of payoff functions given that these payoff functions are represented appropriately in the training set. For an overview of different types of payoff functions that can be covered by a single neural network we refer to Table I. Moreover, the approach pursued in Example II.7 can be considered as an instance of transfer learning, since we used an existing neural network as a starting point to improve the approximation quality on a more specific class of payoff functions.

¹⁸One certainly could enlarge the training set to improve the approximation results. However, the goal of this example is to showcase that on a fixed training set predicting price bounds directly is more efficient than predicting price bounds via (II.2) after having approximated the optimal strategies.

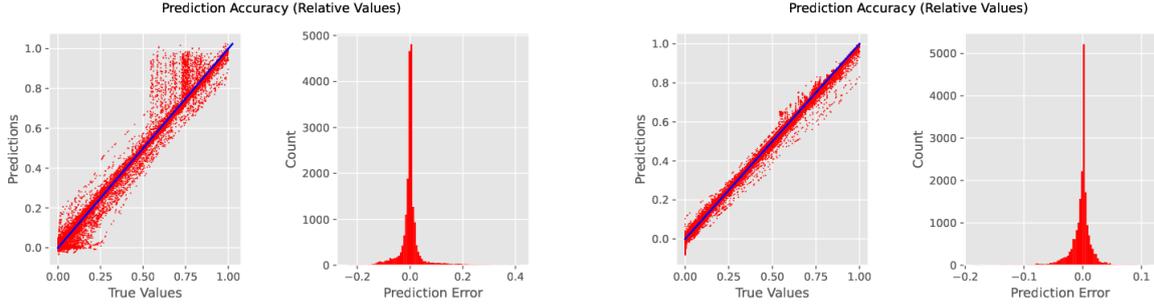


Fig. 4. This figure compares the accuracy of prediction of price bounds of call options using the neural network trained solely on basket options (left) with those predictions from a neural network that is trained additionally also on call options (right). We depict the relative error of the predictions by dividing the prediction error by the spot prices.

Example II.8 (High-dimensional payoff function). We train a neural network to predict prices of a basket option depending on 30 underlying securities, i.e., we consider a family of payoff functions of the form

$$\{\Phi_\theta, \theta \in \Theta\} = \left\{ \Phi_{w_1, \dots, w_{30}, L}(S_{t_1}^1, \dots, S_{t_1}^{30}) = \max \left\{ \sum_{k=1}^{30} w_k S_{t_1}^k - L, 0 \right\}, \text{ with } w_1, \dots, w_{30}, L \in \mathbb{R}_+ \right\}.$$

We train the neural network only on 8000 different samples, where each sample consists of bid and ask prices of call options for 20 different strikes of 30 different underlying securities from data observed in June 2020 with randomly generated weights w_k , $k = 1, \dots, 30$, and randomly generated strikes L , i.e., one single sample X_i consists of 1861 entries. These entries consist of $30 \cdot 20$ strikes of call options, $30 \cdot 20$ bid prices of call options, $30 \cdot 20$ ask prices of call options, as well as 30 spot prices and 1 strike L of the basket option and 30 associated weights w_k , $k = 1, \dots, 30$.

The corresponding prices are computed with the algorithm from [47] which enables to compute precise price bounds even in this high-dimensional setting¹⁹.

After having trained the neural network, we test on 2000 samples from data on options on the S&P 500 that were observed in August 2020, as described in Section II-C2. We test only for the lower bound of the basket option and achieve a relative²⁰ mean absolute error of 0.0101 and a relative mean squared error of 0.0002 on the test set. Compare Figure 5, where we depict the relative error, i.e., we divide predictions and prices by the weighted sum of the spot prices.

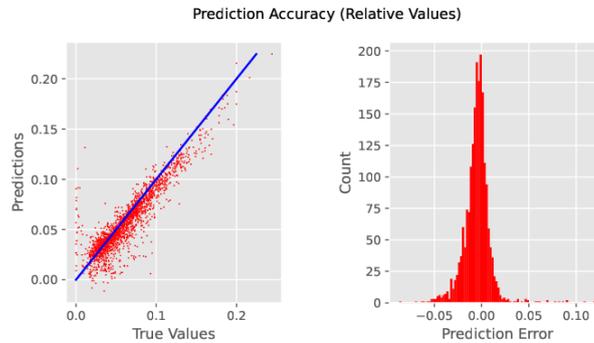


Fig. 5. This figure shows the relative error when predicting the lower price bound of a basket option that depends on 30 underlying securities. We divide the target prices (and predicted prices) by the weighted spot prices, where the weights are the ones in the payoff function under consideration.

Remark II.9. (a) It turns out that indeed our proposed approach can be executed significantly faster than comparable methods that can be applied to compute model-free price bounds. The computation of 100 price bounds in the setting of Example II.8 takes 225.31 seconds on a standard computer²¹ when using the LSIP approach²² from [47]. The execution of a trained neural network to predict 100 price bounds however only takes 0.00303 seconds. The execution of the neural network is therefore approximately 75000 times faster. This highlights the computational advantage of our proposed approach over comparable

¹⁹With this algorithm it is even possible to compute price bounds of basket options that depend on 60 securities.

²⁰Note that here, as in Example II.6, the relative error refers to the error after division with the weighted sum of the spot prices, where the weights are determined by the weights in the payoff of the basket option.

²¹We used for the computations a Gen Intel(R) Core(TM) i7-1165G7, 2.80 GHz processor with 40 GB RAM.

²²The Matlab-code for the execution of the LSIP approach is provided under <https://github.com/qikunxiang/ModelFreePriceBounds>.

numerical methods, and further indicates that our approach indeed allows almost in real time the model-free valuation of financial derivatives.

- (b) We report quite fast training times that are required to train the neural networks. The computationally most expensive setting considered in Example II.8 involves, after applying early stopping, the training of 284 Epochs on 7200 samples with a batch size of 256. This training procedure takes in total only 142.88 seconds²³.
- (c) Note that even though the underlying approach is presented in a great generality, it can be easily modified to meet the potentially more specific requirements of applicants.
If less call or put options are traded than the trained neural network contains, then one can simply set several strikes and prices to the same value and then apply the trained neural network on the smaller financial market. Moreover, if the payoff function considered in the trained neural network depends on more assets than the same-type payoff function which one wants to price, then this is possible within our approach by adjusting the parameters from the more general payoff function. For example it is possible to determine the price bounds of call options after having trained price bounds of basket options as it was shown in Example II.7, compare also the overview provided in Table I which clarifies which payoffs are of the same type. Moreover, if one wants to take into account asset-specific investment constraints, both universal bounds B and \mathfrak{B} can be replaced by asset-specific bounds $(B_i)_{i=1,\dots,n}$ and option specific bounds $(\mathfrak{B}_{i,j,k})_{i=1,\dots,n,k=1,\dots,d,j=1,\dots,n_{ik}^{\text{opt}}}$. The assertion of Theorem II.2 remains valid as the continuity mainly relies on a compactness argument which still can be applied.
- (d) One major implicit assumptions of the presented approach is that the considered call and put options can be traded liquidly. Even though this assumption is usually fulfilled in practice one should verify this assumption carefully. Moreover, given that other traded options are considered sufficiently liquid, the presented approach can be extended in a straightforward way by including these options in (II.1) and (II.2).
- (e) If one is not interested in imposing a proper trading restriction through the bound \mathfrak{B} , then setting \mathfrak{B} to a sufficiently large value does lead in practice to the same price bounds as in an unbounded setting. Therefore, also the optimal parameters of the neural networks that approximate these bounds are the same as in an unbounded setting. This holds true since both the trained parameters of the neural network and the corresponding trained trading strategy a posteriori turn out to remain bounded over the whole training period, compare e.g. [5, Fig. 4].
- (f) It is noteworthy that the presented approach allows to determine model-free price bounds of the most common types of traded financial derivatives by only training a couple of neural networks (one for each type of payoff function), compare the non-exhaustive Table I for an overview which relies partly on the presentation provided in [47]. Table I shows in particular which payoff functions are of the same type and can therefore be trained with a single neural network. Note that with Example II.7 we provide empirical evidence that this approach not only works in theory but indeed in practice. Recall that in Example II.7 we trained a neural network to predict price bounds of call options and basket options that are both of the same type.

Type	Name	Payoff	Parameters Θ
I.	Basket call option with weights $(w^k)_{k=1,\dots,d}$ and strike L	$\max\{\sum_{k=1}^d w^k S_{t_1}^k - L, 0\}$	$\Theta = \{w^1, \dots, w^d \in \mathbb{R}, L \in \mathbb{R}\}$
I.	Basket put option with weights $(w^k)_{k=1,\dots,d}$ and strike L	$\max\{L - \sum_{k=1}^d w^k S_{t_1}^k, 0\}$	$\Theta = \{w^1, \dots, w^d \in \mathbb{R}, L \in \mathbb{R}\}$
I.	Call option on the i -th asset with strike L	$\max\{S_{t_1}^i - L, 0\}$	$\Theta = \{L \in \mathbb{R}\}$
I.	Put option on the i -th asset with strike L	$\max\{L - S_{t_1}^i, 0\}$	$\Theta = \{L \in \mathbb{R}\}$
I.	Spread call options with weights w^i, w^j and strike L	$\max\{w^i S_{t_1}^i - w^j S_{t_1}^j - L, 0\}$	$\Theta = \{w^i, w^j \in \mathbb{R}, L \in \mathbb{R}\}$
I.	Spread put options with weights w_i, w_j and strike L	$\max\{L - (w^i S_{t_1}^i - w^j S_{t_1}^j), 0\}$	$\Theta = \{w^i, w^j \in \mathbb{R}, L \in \mathbb{R}\}$
II.	Call-on-max with strike L	$\max\{\max\{S_{t_1}^k, k = 1, \dots, d\} - L, 0\}$	$\Theta = \{L \in \mathbb{R}\}$
II.	Put-on-max with strike L	$\max\{L - \max\{S_{t_1}^k, k = 1, \dots, d\}, 0\}$	$\Theta = \{L \in \mathbb{R}\}$
III.	Call-on-min with strike L	$\max\{\min\{S_{t_1}^k, k = 1, \dots, d\} - L, 0\}$	$\Theta = \{L \in \mathbb{R}\}$
III.	Put-on-min with strike L	$\max\{L - \min\{S_{t_1}^k, k = 1, \dots, d\}, 0\}$	$\Theta = \{L \in \mathbb{R}\}$
IV.	Best-of-calls option with strikes L_1, \dots, L_d	$\max\{\max\{S_{t_1}^k - L_k, 0\}, k = 1, \dots, d\}$	$\Theta = \{L_1, \dots, L_d \in \mathbb{R}\}$
IV.	Best-of-puts option with strikes L_1, \dots, L_d	$\max\{\max\{L_k - S_{t_1}^k, 0\}, k = 1, \dots, d\}$	$\Theta = \{L_1, \dots, L_d \in \mathbb{R}\}$

TABLE I

THE TABLE DEPICTS THE MOST COMMON TYPES OF FINANCIAL DERIVATIVES THAT CAN BE TRAINED BY THE PRESENTED APPROACH. THE LEFTMOST COLUMN IDENTIFIES THE TYPE OF THE RESPECTIVE FINANCIAL DERIVATIVE, WHERE THE PRICE BOUNDS OF PAYOFFS OF THE SAME TYPE CAN BE LEARNED FROM A SINGLE NEURAL NETWORK.

III. MARTINGALE OPTIMAL TRANSPORT

The presented approach from Section II-B can easily be adjusted to modified market settings given that it is possible to establish a continuous relationship between the prevailing market scenario and resultant model-free price bounds of derivatives. In this section we show how the approach can be adapted to the setting used in martingale optimal transport (compare among many other relevant articles [7], [8], [14], and [20]), where instead of observing a finite amount of call and put option prices, one assumes that the entire one-dimensional marginal distributions of the underlying assets at future dates are known. This situation is according to the Breeden-Litzenberger result [11] equivalent to the case where one can observe call and put option prices for a continuum of strikes

²³See footnote 21,

on each of the associated maturities on which the financial derivative $\Phi \in B(\mathbb{R}_+^{nd}, \mathbb{R})$ depends. In the martingale optimal transport case, one wants to compute the arbitrage-free upper price bound²⁴ of Φ which leads to the maximization problem

$$\sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \dots, \mu_n)} \mathbb{E}_{\mathbb{Q}}[\Phi(S)], \quad (\text{III.1})$$

where²⁵

$$\mathcal{M}(\mu_1, \dots, \mu_n) := \left\{ \mathbb{Q} \in \mathcal{P}(\mathbb{R}_+^{nd}) \mid \mathbb{Q} \circ (S_{t_i}^k)^{-1} = \mu_i^k, \mathbb{E}_{\mathbb{Q}}[S_{t_{i+1}}^k \mid \mathbf{S}_{t_i}, \dots, \mathbf{S}_{t_1}] = S_{t_i}^k \text{ } \mathbb{Q}\text{-a.s. for all } i, k \right\} \quad (\text{III.2})$$

describes the set of all n -step martingale measures with fixed one-dimensional marginals $\mu_i = (\mu_i^1, \dots, \mu_i^d)$, $i = 1, \dots, n$, of all involved securities.

We show that, when considering a single asset and two future maturities, (III.1) can be approximated by a properly constructed neural network. To that end, let

$$\mathcal{P}_1(\mathbb{R}_+) := \left\{ \mathbb{Q} \in \mathcal{P}(\mathbb{R}_+) \mid \int_{\mathbb{R}_+} x \, d\mathbb{Q}(x) < \infty \right\}$$

denote the set of probability measures on \mathbb{R}_+ with existing first moment. Further, we introduce the 1-Wasserstein-distance $W(\cdot, \cdot)$ between two measures $\mu_1, \mu_2 \in \mathcal{P}_1(\mathbb{R}_+)$, which is defined through

$$W(\mu_1, \mu_2) := \inf_{\pi \in \Pi(\mu_1, \mu_2)} \int_{\mathbb{R}_+^2} |u - v| \, d\pi(u, v),$$

with $\Pi(\mu_1, \mu_2)$ denoting the set of all couplings²⁶ of μ_1 and μ_2 , compare e.g. [55].

We recall the construction of the \mathcal{U} -quantization from [6]. Given some probability measure $\mu \in \mathcal{P}_1(\mathbb{R}_+)$ and some $N \in \mathbb{N}$ we set for $i = 1, \dots, N$

$$x_i^{(N)}(\mu) := N \int_{(i-1)/N}^{i/N} F_{\mu}^{-1}(u) \, du,$$

where $F_{\mu}^{-1}(u) := \inf\{x \in \mathbb{R}_+ : F_{\mu}(x) \geq u\}$ denotes the u -quantile associated to the cumulative distribution function F_{μ} of μ , and we denote $\mathbf{x}^{(N)}(\mu) := (x_l^{(N)}(\mu))_{l=1, \dots, N}$. Then, by [6, Theorem 2.4.12.] it holds that

$$\mathcal{U}^{(N)}(\mu) := \frac{1}{N} \sum_{i=1}^N \delta_{x_i^{(N)}(\mu)} \quad (\text{III.3})$$

converges weakly to μ for $N \rightarrow \infty$. Since the mean of $\mathcal{U}^{(N)}(\mu)$ and μ coincide for all $N \in \mathbb{N}$ due to [6, Lemma 2.4.4.] and $\mu, \mathcal{U}^{(N)}(\mu) \in \mathcal{P}_1(\mathbb{R}_+)$, we further obtain convergence in the 1-Wasserstein-distance, compare [55, Definition 6.8]. This means particularly that for all $\mu \in \mathcal{P}_1(\mathbb{R}_+)$ and for all $\delta > 0$ there exists some $N \in \mathbb{N}$ such that $W(\mathcal{U}^{(N)}(\mu), \mu) < \delta$.

We derive the following novel result which asserts for the first time that two-marginal martingale optimal transport problems can be approximated arbitrarily well by neural networks.²⁷

Theorem III.1. *Let $\Phi : \mathbb{R}_+^2 \rightarrow \mathbb{R}$ be continuous such that $\sup_{x_1, x_2 \in \mathbb{R}_+} \frac{|\Phi(x_1, x_2)|}{1+x_1+x_2} < \infty$. Then, for all $\varepsilon > 0$, $N \in \mathbb{N}$, and compact sets $\mathbb{K} \subset \mathbb{R}_+^N$, there exists a neural network $\mathcal{N} \in \mathfrak{N}_{2N, 2}$ such that for all $(\mu_1, \mu_2) \in \mathcal{P}_1(\mathbb{R}_+) \times \mathcal{P}_1(\mathbb{R}_+)$ with $\mu_1 \preceq \mu_2$ ²⁸, there exists some $\delta > 0$ such that if $W(\mathcal{U}^{(N)}(\mu_1), \mu_1) < \delta$, $W(\mathcal{U}^{(N)}(\mu_2), \mu_2) < \delta$, and $\mathbf{x}^{(N)}(\mu_1), \mathbf{x}^{(N)}(\mu_2) \in \mathbb{K}$, then*

$$\left\| \mathcal{N}(\mathbf{x}^{(N)}(\mu_1), \mathbf{x}^{(N)}(\mu_2)) - \left(\inf_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) \right\|_2 < \varepsilon. \quad (\text{III.4})$$

Proof. See Section IV. □

Remark III.2. *The assumption in Theorem III.1 stating that the atoms of the approximating \mathcal{U} -quantizations are contained in some prespecified compact set \mathbb{K} can always be fulfilled if one only considers marginals with support in \mathbb{K} .*

If one does not want to restrict to compactly supported marginals, one can start with $\widehat{\mu}_1, \widehat{\mu}_2 \in \mathcal{P}_1(\mathbb{R}_+)$ and consider for every $r > 0$ the sets $B_r(\widehat{\mu}_i) := \{\mu_i \in \mathcal{P}_1(\mathbb{R}_+) : W(\mu_i, \widehat{\mu}_i) \leq r\}$ for $i = 1, 2$. Then there exists some compact set $\mathbb{K} \subset \mathbb{R}_+^N$ s.t. $x_l^{(N)}(\mu_i) \in \mathbb{K}$ for all $l = 1, \dots, N$, $\mu_i \in B_r(\widehat{\mu}_i)$, $i = 1, 2$.

Indeed, there exists some constant $C > 0$ such that

$$\left| x_l^{(N)}(\widehat{\mu}_i) \right| \leq C \text{ for all } l = 1, \dots, N, i = 1, 2. \quad (\text{III.5})$$

²⁴We implicitly assume absence of a bid-ask spread and of transaction costs.

²⁵ $\mathcal{P}(\mathbb{R}_+^{nd})$ denotes the set of all Borel probability measures on \mathbb{R}_+^{nd} .

²⁶More precisely, the set $\Pi(\mu_1, \mu_2)$ is defined as $\Pi(\mu_1, \mu_2) := \left\{ \pi \in \mathcal{P}_1(\mathbb{R}_+^2) : \pi \circ S_{t_i}^{-1} = \mu_i, i = 1, 2 \right\}$

²⁷We recall that $\|\cdot\|_2$ is an arbitrary norm on \mathbb{R}^2 .

²⁸Here \preceq denotes the convex order for measures $(\mu_1, \mu_2) \in \mathcal{P}_1(\mathbb{R}_+) \times \mathcal{P}_1(\mathbb{R}_+)$, i.e., $\mu_1 \preceq \mu_2$ means $\int_{\mathbb{R}_+} f \, d\mu_1 \leq \int_{\mathbb{R}_+} f \, d\mu_2$ for all convex functions $f : \mathbb{R}_+ \rightarrow \mathbb{R}$.

Moreover, for all $\mu_i \in B_r(\hat{\mu}_i)$, $i = 1, 2$, we have for all $l = 1, \dots, N$ that

$$\begin{aligned} \left| x_i^{(N)}(\hat{\mu}_i) - x_i^{(N)}(\mu_i) \right| &= \left| N \int_{(l-1)/N}^{l/N} F_{\hat{\mu}_i}^{-1}(u) - F_{\mu_i}^{-1}(u) \, du \right| \leq N \int_{(l-1)/N}^{l/N} |F_{\hat{\mu}_i}^{-1}(u) - F_{\mu_i}^{-1}(u)| \, du \\ &\leq N \int_0^1 |F_{\hat{\mu}_i}^{-1}(u) - F_{\mu_i}^{-1}(u)| \, du \\ &= N \cdot W(\hat{\mu}_i, \mu_i) \leq N \cdot r, \end{aligned} \quad (\text{III.6})$$

see for the last equality in (III.6) also [49, Equation 3.1.6] and [51, Equation 3.5]. This implies for all $i = 1, 2$ that $(x_1^{(N)}(\mu_i), \dots, x_N^{(N)}(\mu_i)) \in \mathbb{K} := \{(x_1, \dots, x_N) \in \mathbb{R}_+^N \mid |x_j| \leq C + N \cdot r \text{ for all } j\}$.

Remark III.3. If we are only interested in predicting the upper bound $\sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi]$, then one can see, by carefully reading the proof of Theorem III.1, that it suffices to assume that Φ is upper semi-continuous (and linearly bounded) to obtain the existence of a neural network $\mathcal{N} \in \mathfrak{N}_{2N,1}$ that approximates the bound as in (III.4) w.r.t. $\|\cdot\|_1$. Similarly, to derive the existence of a neural network approximating the lower bound $\inf_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi]$ it is only necessary to assume that Φ is lower semi-continuous (and linearly bounded).

In the case with two marginals, the training routine from Algorithm 1 modifies to the below presented Algorithm 2 which is also depicted in Figure 1b.

Algorithm 2: Training of a neural network for the computation of the price bounds of some financial derivative Φ in the MOT setting

Input : Payoff function $\Phi : \mathbb{R}_+^2 \rightarrow \mathbb{R}$ fulfilling the assumptions of Theorem III.1;

N describing the number of maximal supporting values of the approximating distribution; Number of samples S ;

Method to create samples such that the associated marginal distributions increase in convex order (cf.

Remark III.4 (a));

for i in $\{1, \dots, S\}$ **do**

 Create samples $\mathbf{X}_i = (x_1^i, \dots, x_N^i, y_1^i, \dots, y_N^i) \in \mathbb{R}_+^{2N}$ such that $\frac{1}{N} \sum_{j=1}^N \delta_{x_j^i} \preceq \frac{1}{N} \sum_{j=1}^N \delta_{y_j^i}$;

 Compute, e.g. via linear programming (compare [29]), the target value

$$\mathbf{Y}_i = \left(\inf_{\mathbb{Q} \in \mathcal{M}\left(\frac{1}{N} \sum_{j=1}^N \delta_{x_j^i}, \frac{1}{N} \sum_{j=1}^N \delta_{y_j^i}\right)} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}\left(\frac{1}{N} \sum_{j=1}^N \delta_{x_j^i}, \frac{1}{N} \sum_{j=1}^N \delta_{y_j^i}\right)} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) \in \mathbb{R}^2;$$

end

Train via back-propagation a neural network $\mathcal{N} \in \mathfrak{N}_{2N,2}$ to predict \mathbf{Y}_i given \mathbf{X}_i , i.e., such that $\mathcal{N}(\mathbf{X}_i) \approx \mathbf{Y}_i$ (which is possible due to Theorem III.1);

Output: Trained neural network $\mathcal{N} \in \mathfrak{N}_{2N,2}$;

Remark III.4.

- The critical point in Algorithm 2 is the method which is employed to create discrete samples of marginals that are increasing in convex order. It is a priori not obvious how to create a good sample set. One working methodology is to draw samples from marginals that are similar to marginals one wants to predict, then to apply \mathcal{U} -quantization and to discard measures that do not increase in convex order. Compare also Example III.5. Being similar can for example mean to be close in Wasserstein distance or being from the same parametric family of probability distributions.
- Since the proof of Theorem III.1 relies on the continuity of the MOT problem, as stated in [4], [48], and [56], for which - at the moment - no extension to the case with more than two marginals or in multidimensional settings is known, we stick to the two marginal case in the one-dimensional setting.
- The approach can be extended to the case with information on the variance as in [44], with Markovian assumptions as imposed in [25] and [54], or even more general constraints on the distribution (see e.g. [3]), whenever it is possible to establish a continuous relationship between marginals and prices. Compare also Example III.6, where we consider a constrained martingale optimal transport problem.

To compute the robust bounds $\inf_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi]$ and $\sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi]$ for arbitrary (possibly continuous) marginal distributions $\mu_1 \preceq \mu_2$, with the above explained approach, we approximate μ_1 and μ_2 in the Wasserstein distance by $\mathcal{U}^{(N)}(\mu_1) \preceq \mathcal{U}^{(N)}(\mu_2)$ and then compute the resultant price bound via $\mathcal{N}(\mathbf{x}^{(N)}(\mu_1), \mathbf{x}^{(N)}(\mu_2))$, where $\mathcal{N} \in \mathfrak{N}_{2N,2}$ denotes the neural network which was trained according to Algorithm 2.

A. Examples

We train a neural network according to Algorithm 2. Thus, the input features are, in contrast to the examples from Section II-C, not directly option prices but are given by marginal distributions which are discretized according to the \mathcal{U} -quantization. In the following we fix a payoff function $\Phi(S_{t_1}, S_{t_2}) = |S_{t_1} - S_{t_2}|$ and present the results when predicting $\inf_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi(S_{t_1}, S_{t_2})]$ and $\sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi(S_{t_1}, S_{t_2})]$ via neural networks in dependence of the marginals μ_1 and μ_2 .

1) *Implementation*: To train the neural network we create according to Algorithm 2 numerous artificial marginals with a fixed number $N \in \mathbb{N}$ of supporting values. In the examples below we choose $N = 20$. To discretize continuous marginal distributions we apply the \mathcal{U} -quantization as introduced in (III.3). Below, we describe which marginal distributions μ_1^i, μ_2^i we generate for $i = 1, \dots, S$ with S being the total number of samples.

1) **Log-Normally distributed marginals**

(if $i \bmod 4 = 0$)

$$\mu_1^i \sim \mathcal{LN}(\mu^i - (\sigma_1^i)^2/2, (\sigma_1^i)^2) \text{ with } \mu^i \sim \mathcal{U}([-2, 2]) \text{ and } \sigma_1^i \sim \mathcal{U}([0, 0.5]),$$

$$\mu_2^i \sim \mathcal{LN}(\mu^i - (\sigma_2^i)^2/2, (\sigma_2^i)^2) \text{ with } \sigma_2^i = \sigma_1^i \cdot \widetilde{\sigma}_2^i, \text{ where } \widetilde{\sigma}_2^i \sim \mathcal{U}([1, 2]);$$

2) **Uniform marginals**

(if $i \bmod 4 = 1$)

$$\mu_1^i \sim \mathcal{U}([\mu^i - a^i, \mu^i + a^i]),$$

$$\mu_2^i \sim \mathcal{U}([\mu^i - b^i, \mu^i + b^i]) \text{ with } \mu^i \sim \mathcal{U}([10, 20]), a^i \sim \mathcal{U}([0, 5]), b^i \sim \mathcal{U}([a^i, a^i + 5]);$$

3) **Continuous and discrete uniform marginals**

(if $i \bmod 4 = 2$)

$$\mu_1^i \sim \mathcal{U}([\mu^i - a^i, \mu^i + a^i]),$$

$$\mu_2^i \sim \mathcal{U}(\{\mu - a^i, \mu + a^i\}) \text{ with } \mu^i \sim \mathcal{U}([5, 10]), a^i \sim \mathcal{U}([0, 5]);$$

4) **Uniform and triangular marginals**

(if $i \bmod 4 = 3$)

$$\mu_1^i \sim \mathcal{U}([m^i - l^i/2, m^i + l^i/2]) \text{ for } l^i \sim \mathcal{U}([0, 5]), m^i \sim \mathcal{U}([l^i, l^i + 10]) \text{ and triangular marginals } \mu_2^i \text{ with lower limit } l^i, \text{ upper limit } u^i \sim \mathcal{U}([m^i, m^i + 10]) \text{ and mode } m^i.$$

If the generated marginals are in convex order²⁹, then we add the discretized values $\mathcal{U}^{(N)}(\mu_1^i), \mathcal{U}^{(N)}(\mu_2^i)$ to the sample set $(\mathbf{X}_i)_{i=1, \dots, S}$ and compute via linear programming $\inf_{\mathbb{Q} \in \mathcal{M}(\mu_1^i, \mu_2^i)} \mathbb{E}_{\mathbb{Q}}[\Phi]$ as well as $\sup_{\mathbb{Q} \in \mathcal{M}(\mu_1^i, \mu_2^i)} \mathbb{E}_{\mathbb{Q}}[\Phi]$ which we then add as corresponding target values to $(\mathbf{Y}_i)_{i=1, \dots, S}$.

2) *Architecture of the neural networks*: Given a set of samples $(\mathbf{X}_i)_{i=1, \dots, S}$ and a set of targets $(\mathbf{Y}_i)_{i=1, \dots, S}$ we train a neural network using the back-propagation algorithm with an Adam optimizer implemented in *Python* using *Tensorflow* similar to Section II-C. The loss function is a L^2 -loss function, the neural networks comprise 3 hidden layers with 512 neurons each, and with *ReLU* activation functions.

For further details of the code we refer to https://github.com/juliansester/deep_model_free_pricing.

Example III.5 (MOT without constraints). *We report the results for a neural network trained to $S = 100,000$ samples that are generated according to the procedure described above. We split the samples in training set, test set (10% of the samples) and validation set (20% of the training samples) and report a mean absolute error of 0.0082 as well as mean squared error of 0.0001 on the test set after 1000 epochs of training with early stopping. The accuracy of the trained neural network on the test set is displayed in Figure 6. Additionally, we test the neural network in the following specific situations.*

(a) $\mu_1 = \mathcal{LN}(0.5 - 0.5 \cdot 0.25^2, 0.25^2),$

$$\mu_2 = \mathcal{LN}(0.5 - 0.5 \cdot 0.5^2, 0.5^2)$$

(b) $\mu_1 = \mathcal{LN}(-0.05, 0.1), \mu_2 = \mathcal{LN}(-0.1, 0.2)$

(c) $\mu_1 = \mathcal{U}([8, 12]), \mu_2 = \mathcal{U}([5, 15])$

(d) $\mu_1 = \mathcal{U}([5, 10]), \mu_2 = \mathcal{U}(\{5, 10\})$

(e) $\mu_1 = \mathcal{U}([2, 4]),$

$$\frac{d\mu_2}{d\lambda}(x) = (x-1)/3\mathbf{1}_{[1,2]}(x) + (1/3)\mathbf{1}_{[2,4]}(x) + (5-x)/3\mathbf{1}_{[4,5]}(x), \text{ where } \lambda \text{ denotes the Lebesgue-measure.}$$

The results are displayed in Table II and indicate that the bounds can indeed be approximated very precisely.

TABLE II

THE TABLE DISPLAYS FOR DIFFERENT MARGINALS THE APPROXIMATED LOWER BOUNDS AND UPPER BOUNDS THAT ARE COMPUTED VIA TRAINED NEURAL NETWORKS (NN) AND VIA A LINEAR PROGRAMMING APPROACH (LP).

	Lower bound (LP)	Lower bound (NN)	Upper bound (LP)	Upper bound (NN)	Cumulative Error
(a)	0.2363	0.2573	0.4226	0.4210	0.0226
(b)	0.0814	0.0939	0.1870	0.1946	0.0202
(c)	1.7491	1.7503	2.6220	2.6082	0.0150
(d)	1.6688	1.6659	1.6687	1.6636	0.0080
(e)	0.3587	0.3626	0.7215	0.7151	0.0103

In the next example we show how the introduced methodology can be applied to constrained martingale optimal transport problems.

Example III.6 (MOT with variance constraints). *We train the neural network with the same artificial samples of marginals as in the previous Example III.5. In addition to the discretized marginals $\mathcal{U}^{(N)}(\mu_1^i), \mathcal{U}^{(N)}(\mu_2^i)$, we consider as an additional feature a pre-specified level of the variance of the returns as in [44]. This means we only consider martingale measures \mathbb{Q} which additionally*

²⁹This can be easily checked by verifying that $\int_{\mathbb{R}_+} (x-L)_+ d\mu_1(x) \leq \int_{\mathbb{R}_+} (x-L)_+ d\mu_2(x)$ for all atoms L as well as $\int_{\mathbb{R}_+} x d\mu_1(x) = \int_{\mathbb{R}_+} x d\mu_2(x)$, compare e.g. [45].

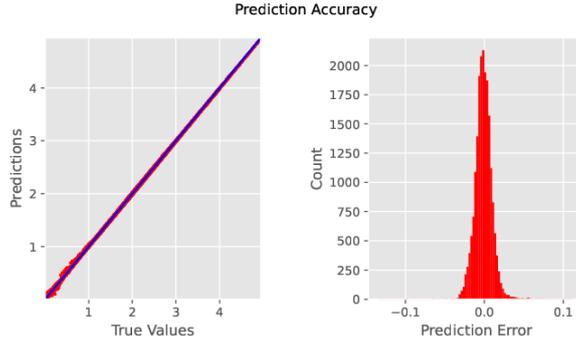


Fig. 6. This figure illustrates the accuracy of the predictions of the trained neural network in the MOT setting, evaluated on the test set.

fulfil $\text{Var}_{\mathbb{Q}}(S_{t_2}/S_{t_1}) = \sigma_{12}^2$. In particular, σ_{12} is thus, an additional feature of the samples \mathbf{X}_i . It was already indicated in Remark II.4 (c), that the approximation through neural networks, as stated in Theorem III.1, can also be obtained in this case due to [44, Theorem A.6.] ensuring continuity of the map

$$(\mu_1, \mu_2) \mapsto \sup \{ \mathbb{E}_{\mathbb{Q}}[\Phi] \mid \mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2) \text{ and } \text{Var}_{\mathbb{Q}}(S_{t_2}/S_{t_1}) = \sigma_{12}^2 \}$$

when Φ is Lipschitz-continuous and when the marginals have compact support. Then, an adaption of Theorem III.1 is straightforward. The results of a neural network approximation for the payoff function $\Phi(S_{t_1}, S_{t_2}) = |S_{t_2} - S_{t_1}|$ for the marginal distributions (c) and (e) from Example III.5 are displayed in Figure 7a. The marginals from (a) and (b) are omitted as they do not satisfy the conditions from [44, Theorem A.6.], whereas the marginals from (d) are omitted, since the value of lower and upper bound coincide and they therefore do not depend on a pre-specified level of the variance. The neural network was trained on 500000 samples for 1000 epochs with early stopping. We assign 10% of the samples to the test set and 20% of the remaining training samples to the validation set (which is relevant for the early stopping rule). The resulting mean absolute error on the test set is 0.0044, the mean squared error is 0.00003. In Figure 7b we show to which degree the predictions deviate from the target values on the test set implying that the accuracy of the predictions is indeed very high on the test set.

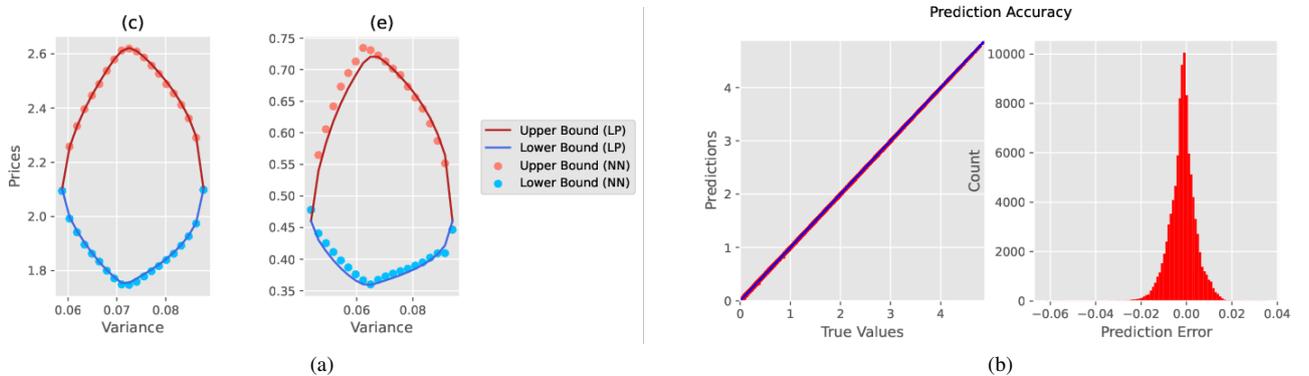


Fig. 7. (a): This figure shows the accuracy of a neural network that was trained with 500000 samples. The accuracy is displayed for the test marginals from (c) and (e) of Example III.5. The points indicate the upper and lower bounds for the prices under the influence of variance information obtained from the trained neural network (NN) in comparison with the precise bounds, computed with a linear programming (LP) approach, that are indicated by the solid lines.

(b): This figure illustrates the accuracy of the predictions of the trained neural network in the MOT setting with variance constraints, evaluated on the test set.

IV. PROOFS

Proof of Theorem II.2 (a). We prove the continuity of $(\mathbf{K}, \pi, \mathbf{S}_{t_0}, \theta) \mapsto \overline{D}_{(\mathbf{K}, \pi, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_{\theta})$. The continuity of $(\mathbf{K}, \pi, \mathbf{S}_{t_0}, \theta) \mapsto \underline{D}_{(\mathbf{K}, \pi, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_{\theta})$ can be obtained analogously.

Let $\varepsilon > 0$ and let $(\mathbf{K}, \pi, \mathbf{S}_{t_0}, \theta) \in \mathbb{K}_1$. For any $\tilde{\delta} > 0$, by the continuity of

$$\begin{aligned} (\Theta, d_p) &\rightarrow \left(C(\mathbb{R}_+^{nd}, \mathbb{R}), d_{\infty, B} \right) \\ \theta &\mapsto \Phi_{\theta}, \end{aligned}$$

we can choose $\delta > 0$ sufficiently small such that for all $\tilde{\theta} \in \Theta$ we have that

$$\|\theta - \tilde{\theta}\|_p < \delta \quad (\text{IV.1})$$

implies that

$$\|\Phi_\theta - \Phi_{\tilde{\theta}}\|_{\infty, B} < \tilde{\delta}. \quad (\text{IV.2})$$

Moreover, we can choose $\tilde{\delta}$ and δ small enough to ensure that

$$\tilde{\delta} + \delta(1 + \kappa)\mathfrak{B} + \delta\mathfrak{B} < \varepsilon/2. \quad (\text{IV.3})$$

We pick some $\delta > 0, \tilde{\delta} > 0$ such that the implication from (IV.1) to (IV.2) is satisfied, and such that (IV.3) holds true and let $(\tilde{K}, \tilde{\pi}, \tilde{S}_{t_0}, \tilde{\theta}) \in \mathbb{K}_1$ satisfy for all i, j, k that

$$\begin{aligned} |K_{ijk}^{\text{call}} - \tilde{K}_{ijk}^{\text{call}}| &< \delta, & |K_{ijk}^{\text{put}} - \tilde{K}_{ijk}^{\text{put}}| &< \delta, \\ \|\theta - \tilde{\theta}\|_p &< \delta, & |S_{t_0}^k - \tilde{S}_{t_0}^k| &< \delta, \\ |\pi_{\text{call}, i, j, k}^+ - \tilde{\pi}_{\text{call}, i, j, k}^+| &< \delta, & |\pi_{\text{call}, i, j, k}^- - \tilde{\pi}_{\text{call}, i, j, k}^-| &< \delta, \\ |\pi_{\text{put}, i, j, k}^+ - \tilde{\pi}_{\text{put}, i, j, k}^+| &< \delta, & |\pi_{\text{put}, i, j, k}^- - \tilde{\pi}_{\text{put}, i, j, k}^-| &< \delta. \end{aligned} \quad (\text{IV.4})$$

First, assume that

$$\overline{D}_{(\tilde{K}, \tilde{\pi}, \tilde{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_{\tilde{\theta}}) \geq \overline{D}_{(K, \pi, S_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta). \quad (\text{IV.5})$$

In this case, consider parameters $\hat{a}, (\hat{c}_{ijk})_{i, j, k}, (\hat{p}_{ijk})_{i, j, k}, (\hat{\Delta}_i^k)$ such that $\Psi_{(\hat{a}, \hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k)}^{(K, S_{t_0})}(\mathbf{s}) \geq \Phi_\theta(\mathbf{s})$ for all $\mathbf{s} \in [0, B]^{nd}$, $\Sigma(\hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k) \leq \mathfrak{B}$, and such that

$$\mathcal{C}\left(\Psi_{(\hat{a}, \hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k)}^{(K, S_{t_0})}, \pi\right) \leq \overline{D}_{(K, \pi, S_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta) + \varepsilon/2, \quad (\text{IV.6})$$

which is possible due to (II.7). Then, we obtain by definition of the semi-static strategies defined in (II.1) and by (IV.4) that

$$\begin{aligned} \left| \Psi_{(\hat{a}, \hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k)}^{(K, S_{t_0})}(\mathbf{s}) - \Psi_{(\tilde{K}, \tilde{S}_{t_0})}^{(\tilde{K}, \tilde{S}_{t_0})}(\mathbf{s}) \right| &= \left| \sum_{i=1}^n \sum_{k=1}^d \sum_{j=1}^{n_{ik}^{\text{opt}}} \left[(\hat{c}_{ijk}^+ - \hat{c}_{ijk}^-) \cdot \left(\max\{s_i^k - K_{ijk}^{\text{call}}, 0\} - \max\{s_i^k - \tilde{K}_{ijk}^{\text{call}}, 0\} \right) \right] \right. \\ &\quad \left. + \sum_{i=1}^n \sum_{k=1}^d \sum_{j=1}^{n_{ik}^{\text{opt}}} \left[(\hat{p}_{ijk}^+ - \hat{p}_{ijk}^-) \cdot \left(\max\{K_{ijk}^{\text{put}} - s_i^k, 0\} - \max\{\tilde{K}_{ijk}^{\text{put}} - s_i^k, 0\} \right) \right] \right. \\ &\quad \left. + \sum_{k=1}^d \left(\hat{\Delta}_0^k (\tilde{S}_{t_0}^k - S_{t_0}^k) + \kappa |\hat{\Delta}_0^k| (|\tilde{S}_{t_0}^k| + |S_{t_0}^k|) \right) \right| \\ &\leq \delta(1 + \kappa) \cdot \Sigma(\hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k) \leq \delta(1 + \kappa)\mathfrak{B} \end{aligned}$$

for all $\mathbf{s} = (s_1^1, \dots, s_1^d, \dots, s_n^1, \dots, s_n^d) \in [0, B]^{nd}$. Thus it holds pointwise on $[0, B]^{nd}$ that

$$\Psi_{(\tilde{\delta} + \delta(1 + \kappa)\mathfrak{B} + \hat{a}, \hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k)}^{(\tilde{K}, \tilde{S}_{t_0})} = \tilde{\delta} + \delta(1 + \kappa)\mathfrak{B} + \Psi_{(\tilde{K}, \tilde{S}_{t_0})}^{(\tilde{K}, \tilde{S}_{t_0})} \geq \tilde{\delta} + \Psi_{(\hat{a}, \hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k)}^{(K, S_{t_0})} \geq \tilde{\delta} + \Phi_\theta \geq \Phi_{\tilde{\theta}}, \quad (\text{IV.7})$$

where the last inequality follows due to (IV.2), since $\|\Phi_\theta - \Phi_{\tilde{\theta}}\|_{\infty, B} < \tilde{\delta}$. This then yields by the definition of the cost function in (II.2), by (IV.2), (IV.3), and by (IV.4) that

$$\begin{aligned} \left| \mathcal{C}\left(\Psi_{(\tilde{\delta} + \delta(1 + \kappa)\mathfrak{B} + \hat{a}, \hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k)}^{(\tilde{K}, \tilde{S}_{t_0})}, \tilde{\pi}\right) - \mathcal{C}\left(\Psi_{(\hat{a}, \hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k)}^{(K, S_{t_0})}, \pi\right) \right| &\leq \tilde{\delta} + \delta(1 + \kappa)\mathfrak{B} + \delta \left(\sum_{i=1}^n \sum_{k=1}^d \sum_{j=1}^{n_{ik}^{\text{opt}}} (\hat{c}_{ijk}^+ + \hat{c}_{ijk}^- + \hat{p}_{ijk}^+ + \hat{p}_{ijk}^-) \right) \\ &\leq \tilde{\delta} + \delta(1 + \kappa)\mathfrak{B} + \delta\mathfrak{B} < \varepsilon/2. \end{aligned} \quad (\text{IV.8})$$

Hence, we obtain that

$$\begin{aligned} &\overline{D}_{(\tilde{K}, \tilde{\pi}, \tilde{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_{\tilde{\theta}}) - \overline{D}_{(K, \pi, S_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta) \\ &= \inf_{\substack{(a, c_{ijk}, p_{ijk}, \Delta_i^k): \\ \Psi_{(a, c_{ijk}, p_{ijk}, \Delta_i^k)}^{(\tilde{K}, \tilde{S}_{t_0})} \geq \Phi_{\tilde{\theta}}, \\ \Sigma(c_{ijk}, p_{ijk}, \Delta_i^k) \leq \mathfrak{B}}} \mathcal{C}\left(\Psi_{(a, c_{ijk}, p_{ijk}, \Delta_i^k)}^{(\tilde{K}, \tilde{S}_{t_0})}, \tilde{\pi}\right) - \inf_{\substack{(a, c_{ijk}, p_{ijk}, \Delta_i^k): \\ \Psi_{(a, c_{ijk}, p_{ijk}, \Delta_i^k)}^{(K, S_{t_0})} \geq \Phi_\theta, \\ \Sigma(c_{ijk}, p_{ijk}, \Delta_i^k) \leq \mathfrak{B}}} \mathcal{C}\left(\Psi_{(a, c_{ijk}, p_{ijk}, \Delta_i^k)}^{(K, S_{t_0})}, \pi\right) \\ &\leq \mathcal{C}\left(\Psi_{(\tilde{\delta} + \delta(1 + \kappa)\mathfrak{B} + \hat{a}, \hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k)}^{(\tilde{K}, \tilde{S}_{t_0})}, \tilde{\pi}\right) - \mathcal{C}\left(\Psi_{(\hat{a}, \hat{c}_{ijk}, \hat{p}_{ijk}, \hat{\Delta}_i^k)}^{(K, S_{t_0})}, \pi\right) + \varepsilon/2 < \varepsilon, \end{aligned}$$

where the last two inequalities are consequences of (IV.3), (IV.5), (IV.6), and (IV.8).

If instead the inequality $\overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta) \geq \overline{D}_{(\widetilde{\mathbf{K}}, \widetilde{\boldsymbol{\pi}}, \widetilde{\mathbf{S}}_{t_0})}^{\mathfrak{B}, B}(\Phi_{\widetilde{\theta}})$ holds, then in this case we choose parameters $\widehat{a}, (\widehat{\mathbf{c}}_{ijk})_{i,j,k}, (\widehat{\mathbf{p}}_{ijk})_{i,j,k}, (\widehat{\Delta}_i^k)$ such that $\Psi_{(\widehat{a}, \widehat{\mathbf{c}}_{ijk}, \widehat{\mathbf{p}}_{ijk}, \widehat{\Delta}_i^k)}^{(\widetilde{\mathbf{K}}, \widetilde{\mathbf{S}}_{t_0})}(s) \geq \Phi_{\widetilde{\theta}}(s)$ for all $s \in [0, B]^{nd}$, $\Sigma(\widehat{\mathbf{c}}_{ijk}, \widehat{\mathbf{p}}_{ijk}, \widehat{\Delta}_i^k) \leq \mathfrak{B}$, and such that

$$\mathcal{C}\left(\Psi_{(\widehat{a}, \widehat{\mathbf{c}}_{ijk}, \widehat{\mathbf{p}}_{ijk}, \widehat{\Delta}_i^k)}^{(\widetilde{\mathbf{K}}, \widetilde{\mathbf{S}}_{t_0})}, \widetilde{\boldsymbol{\pi}}\right) \leq \overline{D}_{(\widetilde{\mathbf{K}}, \widetilde{\boldsymbol{\pi}}, \widetilde{\mathbf{S}}_{t_0})}^{\mathfrak{B}, B}(\Phi_{\widetilde{\theta}}) + \varepsilon/2,$$

and then we repeat the following line of argumentation. This shows part (a). \square

Remark IV.1. For all $\varepsilon > 0$, $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \in \mathbb{K}_1$, there exists some $\delta > 0$ such that if (IV.4) holds for $(\widetilde{\mathbf{K}}, \widetilde{\boldsymbol{\pi}}, \widetilde{\mathbf{S}}_{t_0}, \widetilde{\theta}) \in \mathbb{K}_1$, then for all strategies $\Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})}$ satisfying $\Psi_{(a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\mathbf{K}, \mathbf{S}_{t_0})} \geq \Phi_\theta$ on $[0, B]^{nd}$ and $\Sigma(\mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k) \leq \mathfrak{B}$, there exists some $\tilde{a} \in \mathbb{R}$ with

$$|a - \tilde{a}| < \varepsilon/2 \quad (\text{IV.9})$$

such that $\Psi_{(\tilde{a}, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\widetilde{\mathbf{K}}, \widetilde{\mathbf{S}}_{t_0})} \geq \Phi_{\widetilde{\theta}}$ on $[0, B]^{nd}$. Indeed, let $\varepsilon > 0$, $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \in \mathbb{K}_1$, and choose $\delta, \tilde{\delta} > 0$ analogue as in the proof of Theorem II.2 (a), such that (IV.1), (IV.2), (IV.3), and (IV.4) hold. Then according to (IV.7) we see that the strategy $\Psi_{(\tilde{\delta} + \delta(1+\kappa)\mathfrak{B} + a, \mathbf{c}_{ijk}, \mathbf{p}_{ijk}, \Delta_i^k)}^{(\widetilde{\mathbf{K}}, \widetilde{\mathbf{S}}_{t_0})}$ fulfils (IV.9).

Proof of Theorem II.2 (b). According to Theorem II.2 (a), the map

$$(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \mapsto \left(\overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta), \underline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta)\right)$$

is an element of $C(\mathbb{K}_1, \mathbb{R}^2)$. Hence, we find according to Proposition II.1 a neural network $\mathcal{N}_1 \in \mathfrak{N}_{N_{\text{input}}, 2}$ such that (II.8) holds on \mathbb{K}_1 . \square

Proof of Theorem II.2 (c). Consider a sequence $(\mathbf{K}^{(n)}, \boldsymbol{\pi}^{(n)}, \mathbf{S}_{t_0}^{(n)}, \theta^{(n)})_{n \in \mathbb{N}} \subset \mathbb{K}_2$ with

$$\lim_{n \rightarrow \infty} d_{N_{\text{input}}}((\mathbf{K}^{(n)}, \boldsymbol{\pi}^{(n)}, \mathbf{S}_{t_0}^{(n)}, \theta^{(n)}), (\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)) = 0$$

for some $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \in \mathbb{K}_2$. Since by assumption $\mathfrak{B} < \infty$ and by (II.9), the sequence

$$x^{(n)} := (a^*, (c_{1jk}^*)_{j,k}, (p_{1jk}^*)_{j,k}, (\Delta_0^{k*})_k)(\mathbf{K}^{(n)}, \boldsymbol{\pi}^{(n)}, \mathbf{S}_{t_0}^{(n)}, \theta^{(n)}),$$

$n \in \mathbb{N}$, is bounded. Thus, there exists at least one accumulation point $x \in \mathbb{R}^{1+4M+d}$. Hence, we can find a subsequence $(x^{(n_k)})_{k \in \mathbb{N}}$ with $\lim_{k \rightarrow \infty} d_{1+4M+d}(x^{(n_k)}, x) = 0$. Then, we obtain by the continuity of $\overline{D}^{\mathfrak{B}, B}$, shown in Theorem II.2 (a), and by the continuity of the cost function \mathcal{C} , defined in (II.2), w.r.t. all its arguments, that

$$\overline{D}_{(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0})}^{\mathfrak{B}, B}(\Phi_\theta) = \lim_{k \rightarrow \infty} \overline{D}_{(\mathbf{K}^{(n_k)}, \boldsymbol{\pi}^{(n_k)}, \mathbf{S}_{t_0}^{(n_k)})}^{\mathfrak{B}, B}(\Phi_{\theta^{(n_k)}}) = \lim_{k \rightarrow \infty} \mathcal{C}\left(\Psi_{x^{(n_k)}}^{(\mathbf{K}^{(n_k)}, \mathbf{S}_{t_0}^{(n_k)})}, \boldsymbol{\pi}^{(n_k)}\right) = \mathcal{C}\left(\Psi_x^{(\mathbf{K}, \mathbf{S}_{t_0})}, \boldsymbol{\pi}\right).$$

Using the continuity of Ψ w.r.t. its parameters and the continuity of $\theta \mapsto \Phi_\theta$ as in (II.6), we also obtain that $\Psi_x^{(\mathbf{K}, \mathbf{S}_{t_0})} \geq \Phi_\theta$ on $[0, B]^{nd}$. Moreover, $x := (a, (c_{1jk})_{j,k}, (p_{1jk})_{j,k}, (\Delta_0^k)_k) \in \mathbb{R}^{1+4M+d}$ satisfies $\Sigma(\mathbf{c}_{1jk}, \mathbf{p}_{1jk}, \Delta_0^k) \leq \mathfrak{B}$. Thus, x is indeed a minimal super-replication strategy of Φ_θ for parameters $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)$. So we have shown that any accumulation point x is a minimal super-replication strategy of Φ_θ for parameters $(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta)$. Since we assumed that the minimizer is unique, the accumulation point x is unique and is necessarily the limit of the sequence $(x^{(n)})_{n \in \mathbb{N}}$. Therefore, we have shown that

$$\lim_{n \rightarrow \infty} (a^*, (c_{1jk}^*)_{j,k}, (p_{1jk}^*)_{j,k}, (\Delta_0^{k*})_k)(\mathbf{K}^{(n)}, \boldsymbol{\pi}^{(n)}, \mathbf{S}_{t_0}^{(n)}, \theta^{(n)}) = (a^*, (c_{1jk}^*)_{j,k}, (p_{1jk}^*)_{j,k}, (\Delta_0^{k*})_k)(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta).$$

\square

Proof of Theorem II.2 (d).

According to Theorem II.2 (b), the map

$$(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta) \mapsto \left(a^*, (c_{1jk}^*)_{j,k}, (p_{1jk}^*)_{j,k}, (\Delta_0^{k*})_k\right)(\mathbf{K}, \boldsymbol{\pi}, \mathbf{S}_{t_0}, \theta),$$

is an element of $C(\mathbb{K}_2, \mathbb{R}^{1+4M+d})$. Thus, by Proposition II.1, there exists some $\mathcal{N}_2 \in \mathfrak{N}_{N_{\text{input}}, 1+4M+d}$ such that (II.10) holds on \mathbb{K}_2 . \square

Proof of Theorem III.1. Let $\varepsilon > 0$, $N \in \mathbb{N}$, and pick some compact set $\mathbb{K} \subset \mathbb{R}_+^N$. We observe that the map

$$\begin{aligned} \mathcal{D}^{(N)} : (\mathbb{R}_+^N, d_N) &\rightarrow (\mathcal{P}_1(\mathbb{R}_+), W) \\ (x_1, \dots, x_N) &\mapsto \frac{1}{N} \sum_{i=1}^N \delta_{x_i} \end{aligned}$$

is continuous. Indeed, if

$$\mathbb{R}_+^N \ni \mathbf{x}^m = (x_1^m, \dots, x_N^m) \rightarrow \mathbf{x} = (x_1, \dots, x_N) \in \mathbb{R}_+^N$$

for $m \rightarrow \infty$, then we can consider for all $m \in \mathbb{N}$ the coupling $\pi^m := \frac{1}{N} \sum_{i=1}^N \delta_{(x_i, x_i^m)}$ and obtain that

$$W\left(\mathcal{D}^{(N)}(\mathbf{x}), \mathcal{D}^{(N)}(\mathbf{x}^m)\right) \leq \int_{\mathbb{R}_+^2} |u - v| d\pi^m(u, v) = \frac{1}{N} \sum_{i=1}^N |x_i - x_i^m| \rightarrow 0 \text{ for } m \rightarrow \infty. \quad (\text{IV.10})$$

Since by assumption Φ is upper semi-continuous with $\sup_{x_1, x_2 \in \mathbb{R}_+} \frac{|\Phi(x_1, x_2)|}{1+x_1+x_2} < \infty$, we can apply [56, Theorem 2.9.] which ensures that for any $(\mu_1, \mu_2) \in \mathcal{P}_1(\mathbb{R}_+) \times \mathcal{P}_1(\mathbb{R}_+)$ with $\mu_1 \preceq \mu_2$ and $(\mu_1^m, \mu_2^m) \in \mathcal{P}_1(\mathbb{R}_+) \times \mathcal{P}_1(\mathbb{R}_+)$ with $\mu_1^m \preceq \mu_2^m$, $m \in \mathbb{N}$, satisfying $W(\mu_1^m, \mu_1) \rightarrow 0$, $W(\mu_2^m, \mu_2) \rightarrow 0$ for $m \rightarrow \infty$, it holds

$$\lim_{m \rightarrow \infty} \left| \sup_{\mathbb{Q} \in \mathcal{M}(\mu_1^m, \mu_2^m)} \mathbb{E}_{\mathbb{Q}}[\Phi] - \sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi] \right| = 0. \quad (\text{IV.11})$$

Since for any set of measures $\mathcal{M} \subset \mathcal{P}(\mathbb{R}_+^2)$ we have $\inf_{\mathbb{Q} \in \mathcal{M}} \mathbb{E}_{\mathbb{Q}}[\Phi] = -\sup_{\mathbb{Q} \in \mathcal{M}} \mathbb{E}_{\mathbb{Q}}[-\Phi]$, and since $-\Phi$ is also upper semi-continuous with $\sup_{x_1, x_2 \in \mathbb{R}_+} \frac{|-\Phi(x_1, x_2)|}{1+x_1+x_2} < \infty$, we can apply the same arguments to see that

$$\lim_{m \rightarrow \infty} \left| \inf_{\mathbb{Q} \in \mathcal{M}(\mu_1^m, \mu_2^m)} \mathbb{E}_{\mathbb{Q}}[\Phi] - \inf_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi] \right| = 0. \quad (\text{IV.12})$$

Define the closed set³⁰ $\mathfrak{C}^{(N)} := \left\{ (\mathbf{x}, \mathbf{y}) \in \mathbb{R}_+^N \times \mathbb{R}_+^N : \mathcal{D}^{(N)}(\mathbf{x}) \preceq \mathcal{D}^{(N)}(\mathbf{y}) \right\}$. Then, we obtain by (IV.10), (IV.11), and (IV.12) the continuity of

$$(\mathbf{x}, \mathbf{y}) \mapsto \left(\inf_{\mathbb{Q} \in \mathcal{M}(\mathcal{D}^{(N)}(\mathbf{x}), \mathcal{D}^{(N)}(\mathbf{y}))} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}(\mathcal{D}^{(N)}(\mathbf{x}), \mathcal{D}^{(N)}(\mathbf{y}))} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) \quad \text{on } \mathfrak{C}^{(N)}. \quad (\text{IV.13})$$

Hence, the universal approximation theorem from Proposition II.1 guarantees the existence of a neural network $\mathcal{N} \in \mathfrak{N}_{2N,2}$ such that

$$\sup_{\mathbf{x}, \mathbf{y} \in \mathbb{K} \cap \mathfrak{C}^{(N)}} \left\| \mathcal{N}(\mathbf{x}, \mathbf{y}) - \left(\inf_{\mathbb{Q} \in \mathcal{M}(\mathcal{D}^{(N)}(\mathbf{x}), \mathcal{D}^{(N)}(\mathbf{y}))} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}(\mathcal{D}^{(N)}(\mathbf{x}), \mathcal{D}^{(N)}(\mathbf{y}))} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) \right\|_2 < \varepsilon/2. \quad (\text{IV.14})$$

Now let $(\mu_1, \mu_2) \in \mathcal{P}_1(\mathbb{R}_+) \times \mathcal{P}_1(\mathbb{R}_+)$ with $\mu_1 \preceq \mu_2$. Then [6, Theorem 2.4.11.] ensures that

$$\mathcal{U}^{(N)}(\mu_1) \preceq \mathcal{U}^{(N)}(\mu_2). \quad (\text{IV.15})$$

This and the continuity of the two-marginal MOT problem with respect to its marginals, as stated in (IV.11) and (IV.12), implies that there exists some $\delta > 0$ such that if $W(\mathcal{U}^{(N)}(\mu_1), \mu_1) < \delta$, $W(\mathcal{U}^{(N)}(\mu_2), \mu_2) < \delta$, then

$$\left\| \left(\inf_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) - \left(\inf_{\mathbb{Q} \in \mathcal{M}(\mathcal{U}^{(N)}(\mu_1), \mathcal{U}^{(N)}(\mu_2))} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}(\mathcal{U}^{(N)}(\mu_1), \mathcal{U}^{(N)}(\mu_2))} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) \right\|_2 < \varepsilon/2. \quad (\text{IV.16})$$

By definition of the map $\mathcal{D}^{(N)}$ it holds that $\mathcal{D}^{(N)}(\mathbf{x}^{(N)}(\mu_i)) = \mathcal{U}^{(N)}(\mu_i)$ for $i = 1, 2$. In particular, we have by (IV.15) that $(\mathbf{x}^{(N)}(\mu_1), \mathbf{x}^{(N)}(\mu_2)) \in \mathfrak{C}^{(N)}$. Thus, the triangle inequality and (IV.14) combined with (IV.16) implies that if $W(\mathcal{U}^{(N)}(\mu_1), \mu_1) < \delta$, $W(\mathcal{U}^{(N)}(\mu_2), \mu_2) < \delta$, and $\mathbf{x}^{(N)}(\mu_1), \mathbf{x}^{(N)}(\mu_2) \in \mathbb{K}$, then

$$\begin{aligned} & \left\| \mathcal{N}(\mathbf{x}^{(N)}(\mu_1), \mathbf{x}^{(N)}(\mu_2)) - \left(\inf_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) \right\|_2 \\ & \leq \left\| \mathcal{N}(\mathbf{x}^{(N)}(\mu_1), \mathbf{x}^{(N)}(\mu_2)) - \left(\inf_{\mathbb{Q} \in \mathcal{M}(\mathcal{U}^{(N)}(\mu_1), \mathcal{U}^{(N)}(\mu_2))} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}(\mathcal{U}^{(N)}(\mu_1), \mathcal{U}^{(N)}(\mu_2))} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) \right\|_2 \\ & \quad + \left\| \left(\inf_{\mathbb{Q} \in \mathcal{M}(\mathcal{U}^{(N)}(\mu_1), \mathcal{U}^{(N)}(\mu_2))} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}(\mathcal{U}^{(N)}(\mu_1), \mathcal{U}^{(N)}(\mu_2))} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) - \left(\inf_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi], \sup_{\mathbb{Q} \in \mathcal{M}(\mu_1, \mu_2)} \mathbb{E}_{\mathbb{Q}}[\Phi] \right) \right\|_2 \\ & < \varepsilon, \end{aligned}$$

which shows the assertion. \square

ACKNOWLEDGMENT

Financial support by the Nanyang Assistant Professorship Grant (NAP Grant) *Machine Learning based Algorithms in Finance and Insurance* is gratefully acknowledged.

³⁰We refer to [6, Definition 2.45] and [6, Lemma 2.49] for a characterization of $\mathbf{x}, \mathbf{y} \in \mathbb{R}_+^N$ to satisfy $\mathcal{D}^{(N)}(\mathbf{x}) \preceq \mathcal{D}^{(N)}(\mathbf{y})$.

REFERENCES

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pages 265–283, 2016.
- [2] Beatrice Acciaio, Mathias Beiglböck, Friedrich Penkner, and Walter Schachermayer. A model-free version of the fundamental theorem of asset pricing and the super-replication theorem. *Mathematical Finance*, 26(2):233–251, 2016.
- [3] Jonathan Ansari, Eva Lütkebohmert, Ariel Neufeld, and Julian Sester. Improved robust price bounds for multi-asset derivatives under market-implied dependence information. *arXiv preprint arXiv:2204.01071*, 2022.
- [4] Julio Backhoff-Veraguas and Gudmund Pammer. Stability of martingale optimal transport and weak optimal transport. *The Annals of Applied Probability*, 32(1): 721–752, 2022.
- [5] Michel Baes, Calypso Herrera, Ariel Neufeld, and Pierre Ruysen. Low-rank plus sparse decomposition of covariance matrices using neural network parametrization. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [6] David Baker. Martingales with specified marginals. *Theses, Université Pierre et Marie Curie-Paris VI*, 2012.
- [7] Mathias Beiglböck, Pierre Henry-Labordère, and Friedrich Penkner. Model-independent bounds for option prices: a mass transport approach. *Finance and Stochastics*, 17(3):477–501, 2013.
- [8] Mathias Beiglböck and Nicolas Juillet. On a problem of optimal transport under marginal martingale constraints. *The Annals of Probability*, 44(1):42–106, 2016.
- [9] Yoshua Bengio. *Learning deep architectures for AI*. Now Publishers Inc, 2009.
- [10] Bruno Bouchard and Marcel Nutz. Arbitrage and duality in nondominated discrete-time models. *The Annals of Applied Probability*, 25(2):823–859, 2015.
- [11] Douglas T Breeden and Robert H Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of business*, pages 621–651, 1978.
- [12] Hans Buehler, Lukas Gonon, Josef Teichmann, and Ben Wood. Deep hedging. *Quantitative Finance*, 19(8):1271–1291, 2019.
- [13] Matteo Burzoni, Marco Frittelli, and Marco Maggis. Model-free superhedging duality. *The Annals of Applied Probability*, 27(3):1452–1477, 2017.
- [14] Patrick Cheridito, Matti Kiiski, David J Prömel, and H Mete Soner. Martingale optimal transport duality. *Mathematische Annalen*, pages 1–28, 2020.
- [15] Patrick Cheridito, Michael Kupper, and Ludovic Tangpi. Duality formulas for robust pricing and hedging in discrete time. *SIAM Journal on Financial Mathematics*, 8(1):738–765, 2017.
- [16] Alexander MG Cox and Jan Oblój. Robust pricing and hedging of double no-touch options. *Finance and Stochastics*, 15(3):573–605, 2011.
- [17] Mark Davis, Jan Oblój, and Vimal Raval. Arbitrage bounds for prices of weighted variance swaps. *Mathematical Finance*, 24(4):821–854, 2014.
- [18] Luca De Gennara Aquino and Carole Bernard. Bounds on multi-asset derivatives via neural networks. *International Journal of Theoretical and Applied Finance*, 23(08):2050050, 2020.
- [19] Freddy Delbaen and Walter Schachermayer. A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300(1):463–520, 1994.
- [20] Yan Dolinsky and H Mete Soner. Martingale optimal transport and robust hedging in continuous time. *Probability Theory and Related Fields*, 160(1-2):391–427, 2014.
- [21] Yan Dolinsky and H Mete Soner. Robust hedging with proportional transaction costs. *Finance and Stochastics*, 18(2):327–347, 2014.
- [22] Bruno Dupire. Pricing with a smile. *Risk*, 7(1):18–20, 1994.
- [23] Stephan Eckstein, Gaoyue Guo, Tongseok Lim, and Jan Obloj. Robust pricing and hedging of options on multiple assets and its numerics. *SIAM Journal on Financial Mathematics*, 12(1): 158–188, 2021.
- [24] Stephan Eckstein and Michael Kupper. Computation of optimal transport and related hedging problems via penalization and neural networks. *Applied Mathematics & Optimization*, 83: 639–667, 2019.
- [25] Stephan Eckstein and Michael Kupper. Martingale transport with homogeneous stock movements. *Quantitative Finance*, 21(2):271–280, 2021.
- [26] Stephan Eckstein, Michael Kupper, and Mathias Pohl. Robust risk aggregation with neural networks. *Mathematical Finance*, 30(4):1229–1272, 2020.
- [27] Alfred Galichon, Pierre Henry-Labordère, and Nizar Touzi. A stochastic control approach to no-arbitrage bounds given marginals, with an application to lookback options. *The Annals of Applied Probability*, 24(1):312–336, 2014.
- [28] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.
- [29] Gaoyue Guo and Jan Oblój. Computational methods for martingale optimal transport problems. *The Annals of Applied Probability*, 29(6):3311–3347, 2019.
- [30] Mohamad H Hassoun. *Fundamentals of artificial neural networks*. MIT press, 1995.
- [31] Pierre Henry-Labordère. Automated option pricing: Numerical methods. *International Journal of Theoretical and Applied Finance*, 16(08):1350042, 2013.
- [32] Pierre Henry-Labordère. (martingale) optimal transport and anomaly detection with neural networks: A primal-dual algorithm. *Available at SSRN 3370910*, 2019.
- [33] Steven L Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies*, 6(2):327–343, 1993.
- [34] Mark HA Davis and David Hobson. The range of traded option prices. *Mathematical Finance*, 17(1):1–14, 2007.
- [35] David Hobson. The Skorokhod embedding problem and model-independent bounds for option prices. In *Paris-Princeton Lectures on Mathematical Finance 2010*, pages 267–318. Springer, 2011.
- [36] David Hobson and Anthony Neuberger. Robust bounds for forward start options. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 22(1):31–56, 2012.
- [37] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- [38] Zhaoxu Hou and Jan Oblój. Robust pricing–hedging dualities in continuous time. *Finance and Stochastics*, 22(3):511–567, 2018.
- [39] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015.
- [40] Patrick Kidger and Terry Lyons. Universal approximation with deep narrow networks. In *Conference on Learning Theory*, pages 2306–2327. PMLR, 2020.
- [41] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [42] Frank Knight. *Risk, Uncertainty and Profit*. Houghton Mifflin, 1921.
- [43] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

- [44] Eva Lütkebohmert and Julian Sester. Tightening robust price bounds for exotic derivatives. *Quantitative Finance*, 19(11):1797–1815, 2019.
- [45] Chunsheng Ma. Convex orders for linear combinations of random variables. *Journal of Statistical Planning and Inference*, 84(1-2):11–25, 2000.
- [46] Ariel Neufeld and Marcel Nutz. Superreplication under volatility uncertainty for measurable claims. *Electronic journal of probability*, 18, 2013.
- [47] Ariel Neufeld, Antonis Papapantoleon, and Qikun Xiang. Model-free bounds for multi-asset options using option-implied information and their exact computation. *Management Science*, 2022.
- [48] Ariel Neufeld and Julian Sester. On the stability of the martingale optimal transport problem: A set-valued map approach. *Statistics & Probability Letters*, 176:109–131, 2021.
- [49] Svetlozar T Rachev and Ludger Rüschendorf. *Mass Transportation Problems: Volume I: Theory*, volume 1. Springer Science & Business Media, 1998.
- [50] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- [51] Ludger Rüschendorf. Monge-Kantorovich transportation problem and optimal couplings. *Jahresbericht der DMV*, 3:113–137, 2007.
- [52] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- [53] Wim Schoutens, Erwin Simons, and Jurgen Tistaert. A perfect calibration! now what? *The best of Wilmott*, page 281, 2003.
- [54] Julian Sester. Robust bounds for derivative prices in Markovian models. *International Journal of Theoretical and Applied Finance*, 23(3):2050015, 2020.
- [55] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [56] Johannes Wiesel. Continuity of the martingale optimal transport problem on the real line. *arXiv preprint [arXiv:1905.04574](https://arxiv.org/abs/1905.04574)*, 2019.