

Deep Local Volatility

Matthew Dixon*, Stéphane Crépey†, and Marc Chataigner‡

Abstract. Deep learning for option pricing has emerged as a novel methodology for fast computations with applications in calibration and computation of Greeks. However, most of these approaches do not enforce any no-arbitrage conditions. In this article, we develop a deep learning approach for no-arbitrage interpolation of European vanilla option prices. In particular, we demonstrate the modification of the standard deep learning methodology to enforce the no-arbitrage constraint and we specify the experimental design parameters that are needed for adequate performance. A novel component is the use of the Dupire formula to enforce bounds on the local volatility associated with (non-arbitrable) option prices, during the network fitting. Numerical results¹ on real datasets of DAX vanilla options demonstrate the numerical error in the price, implied volatility and local volatility surface.

Key words. Option pricing; Deep learning; No-Arbitrage; Local Volatility.

AMS subject classifications. 91G20, 62M45, 91G60.

1. Introduction. A recent approach to option pricing involves reformulating the pricing problem as a surrogate modeling problem. Once reformulated, the problem is amenable to standard supervised machine learning methods such as Neural networks and Gaussian processes. This is particularly suitable in situations with a need for fast computations and a tolerance to in-exact approximation.

In their seminal paper, [Hutchinson et al. \(1994\)](#) use a radial basis function neural network for delta-hedging. Their network is trained to Black-Scholes prices, using the time-to-maturity and the moneyness as input variables, without shape constraints. They use the hedging performance of the ensuing delta-hedging strategy as a performance criterion. [Garcia and Gençay \(2000\)](#) and [Gençay and Qi \(2001\)](#) improve the stability of the previous approach by adding the outputs of two such neural networks, weighted by respective moneyness and time-to-maturity functionals. [Dugas et al. \(2009\)](#) introduce the first neural network architecture guaranteeing arbitrage-free vanilla option pricing. However, [Ackerer et al. \(2019\)](#) show that the corresponding hard constrained networks are very difficult to train in practice. Instead, they advocate the learning of the implied volatility (rather than the prices) by a standard feedforward neural network with soft-constraints, i.e. regularization,

*Department of Applied Mathematics, Illinois Institute of Technology, Chicago; matthew.dixon@iit.edu.

†Department of Mathematics, University of Evry, Paris Saclay; stephane.crepey@univ-evry.fr, marc.chataigner@univ-evry.fr. The PhD thesis of Marc Chataigner is co-funded by the Research Initiative “Modélisation des marchés actions, obligations et dérivés”, financed by HSBC France under the aegis of the Europlace Institute of Finance, and by a public grant as part of investissement d’avenir project, reference ANR-11-LABX-0056-LLH LabEx LMH. The views and opinions expressed in this paper are those of the authors alone and do not necessarily reflect the views or policies of HSBC Investment Bank, its subsidiaries or affiliates.

¹A Python notebook, compatible with Google Colab, and accompanying data are available in <https://github.com/mChataign/DupireNN>. Due to file size constraints, the notebook must be run to reproduce the figures and results in this article.

34 which penalizes calendar spread and butterfly arbitrages. The call and put prices
 35 must also be decreasing and increasing by strike respectively.

36 We propose simple neural network architectures and training methodology which
 37 satisfy these shape constraints. In contrast to [Dugas et al. \(2009\)](#), following [Ackerer
 38 et al. \(2019\)](#), we also explore soft constraints alternatives to hard constraints in
 39 the network configuration, due to the loss of representation power of the latter.
 40 However, our networks are trained to prices, versus implied volatilities in [Ackerer
 41 et al. \(2019\)](#). Moreover, a novel aspect of our approach is to focus on the associated
 42 local volatility surface, considered both for itself and as a penalization device in our
 43 soft constraints approach.

44 **2. Problem Statement.** We consider European vanilla option prices on a stock
 45 or index S . We assume that a deterministic short interest rate term structure
 46 $r(t)$ of the corresponding economy has been bootstrapped from the its zero coupon
 47 curve, and that a term structure of deterministic continuous-dividend-yields $q(t)$ on
 48 S has then been extracted from the prices of the forward contracts on S . Without
 49 restriction given the call-put parity relationship, we only consider put option prices.
 50 We denote by $P^*(T, K)$ the market price of the put option with maturity T and
 51 strike K on S , observed for a finite number of pairs (T, K) at a given day.

Our goal is to construct a nonarbitrable put price surface $P : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$
 in $\mathcal{C}^{1,2}((0, +\infty) \times \mathbb{R}_+^*) \cap \mathcal{C}^0(\mathbb{R}_+ \times \mathbb{R}_+^*)$, interpolating P^* up to some error term. As
 is well known (see e.g. Section 9.2.1 in [Crépey \(2013\)](#)), the corresponding local
 volatility surface, say $\sigma(\cdot, \cdot)$, is given by the [Dupire \(1994\)](#) formula

$$\frac{\sigma^2(T, K)}{2} = \frac{\partial_T P(T, K) + (r - q)\partial_K P(T, K) + qP(T, K)}{K^2 \partial_{K^2}^2 P(T, K)}.$$

52 In terms of $p(T, k) = \exp(\int_0^T q(t)dt)P(T, K)$, where $k = K \exp(\int_0^T (r(t) - q(t))dt)$,
 53 the formula reads

54 (2.1)
$$\frac{\sigma^2(T, K)}{2} = \frac{\partial_T p(T, k)}{k^2 \partial_{k^2}^2 p(T, k)}.$$

55 Given a rectangular domain of interest in maturity and strike, We rescale further
 56 the inputs, $T' = T / (\max(T) - \min(T))$ and $k' = k / (\max(k) - \min(k))$, so that the
 57 domain of the pricing map is $\Omega := [0, 1]^2$. For ease of notation, we shall hereon drop
 58 the superscript. This rescaling avoids any one independent variable dominating over
 59 another during the fitting. We then learn the modified market prices $p^* = p^*(T, k)$.

60 For the Dupire formula (2.1) to be meaningful, its output must be nonnegative.
 61 This holds, in particular, whenever the interpolating map p exhibits nonnegative
 62 derivatives w.r.t. T and second derivative w.r.t. k , which corresponds to (static)
 63 nonarbitrage relationships on the option prices surface. In both networks considered
 64 in the paper, these derivatives are available analytically via the neural network
 65 automatic differentiation features. Hard or soft (regularization) shape constraints
 66 can be used to enforce these shape properties, exactly in the case of hard constraints
 67 and approximately in the case of soft constraints.

3. Shape Preserving Neural Networks. We consider parameterized maps $p =$
 $p_{\mathbf{W}, \mathbf{b}}$

$$(T, k) \ni \mathbb{R}_+^2 \xrightarrow{p} p_{\mathbf{W}, \mathbf{b}}(T, k) \in \mathbb{R}_+,$$

2

68 given as deep neural networks with two hidden layers. As detailed in [Goodfellow](#)
 69 [et al. \(2016\)](#), these take the form of a composition of simpler functions:

$$70 \quad p_{\mathbf{W}, \mathbf{b}}(x) = f_{W^{(3)}, b^{(3)}}^{(3)} \circ f_{W^{(2)}, b^{(2)}}^{(2)} \circ f_{W^{(1)}, b^{(1)}}^{(1)}(x),$$

where

$$\mathbf{W} = (W^{(1)}, W^{(2)}, W^{(3)}) \text{ and } \mathbf{b} = (b^{(1)}, b^{(2)}, b^{(3)})$$

71 are weight matrices and bias vectors, and the $f^{(l)} := \zeta^{(l)}(W^{(l)}x + b^{(l)})$ are semi-
 72 affine, for nondecreasing activation functions $\zeta^{(l)}$ applied to their (vector-valued)
 73 argument componentwise. Any weight matrix $W^{(\ell)} \in \mathbb{R}^{m \times n}$ can be expressed as an
 74 n column $W^{(\ell)} = [\mathbf{w}_1^{(\ell)}, \dots, \mathbf{w}_n^{(\ell)}]$ of m -vectors, for successively chained pairs (n, m)
 75 of dimensions varying with $l = 1, 2, 3$ (starting from $n = 2$, the number of inputs,
 76 for $l = 1$, and ending up with $m = 1$, the number of outputs, for $l = 3$).

77 In the hard constraints case, our network is sparsely connected in the sense that,
 78 with $x = (T, k)$ as above,

$$79 \quad f_{W^{(1)}, b^{(1)}}^{(1)}(x) = [f_{W^{(1,T)}, b^{(1,T)}}^{(1,T)}(T), f_{W^{(1,k)}, b^{(1,k)}}^{(1,k)}(k)],$$

where $W^{(1,T)}, b^{(1,T)}$ and $W^{(1,k)}, b^{(1,k)}$ correspond to parameters of sub-graphs for
 each input, and

$$f^{(1,T)}(T) := \zeta^{(1,T)}(W^{(1,T)}T + b^{(1,T)}), \quad f^{(1,k)}(k) := \zeta^{(1,k)}(W^{(1,k)}k + b^{(1,k)}).$$

To impose the shape constraints relevant for put options, it is then enough
 to restrict ourselves to nonnegative weights, and to convex (and nondecreasing)
 activation functions, namely

$$\text{softplus}(x) := \ln(1 + e^x),$$

80 except for $\zeta^{(1,T)}$, which will be taken as an S-shaped sigmoid $(1 + e^{-x})^{-1}$.

81 Hence, the network is nondecreasing (but not necessarily convex) in T , and
 82 convex and nondecreasing in k , as a composition (restricted to the k variable) of
 83 convex and nondecreasing functions of k . [Figure 1](#) illustrates the shape preserving
 84 feed forward architecture with two hidden layers containing 10 hidden nodes. For
 85 avoidance of doubt, the figure is not representative of the number of hidden neurons
 86 used in our experiments. However, the connectivity is representative. The first
 87 input variable, T , is only connected to the first 5 hidden nodes and the second
 88 input variable, k , is only connect to the last 5 hidden nodes. Effectively, two sub-
 89 networks have been created where no information from the input layer crosses the
 90 sub-networks until the second hidden layer. In other words, each sub-network is a
 91 function of only one input variable. This property is the key to imposing different
 92 hard shape constraints w.r.t. each input variable.

93 **3.1. Approximation error bound.** We recall a result from [Ohn and Kim \(2019\)](#)
 94 which describes how the sparsity in a neural network affects its approximation error.

Let us denote the network parameters $\theta := (\mathbf{W}, \mathbf{b})$. We define the network
 parameter space in terms of the layers width L and depth N , and numbers of inputs
 and outputs,

$$\Theta_{i,o}(L, N) := \{\theta : L(\theta) \leq L, n_{max}(\theta) \leq N, in(\theta) = i, out(\theta) = o\}.$$




Figure 1: A shape preserving (sparse) feed forward architecture with one hidden layer containing 10 hidden nodes. The first input variable, T , is only connected to the 5 left most hidden nodes and the second input variable, k , is only connected to the 5 right most hidden nodes.

We also define the restricted parameter space

$$\Theta_{i,o}(L, N, \Sigma, B) := \{\theta \in \Theta_{i,o}(L, N) : |\theta|_0 \leq \Sigma, |\theta|_\infty \leq B\},$$

where $|\theta|_0$ is the number of nonzero components in θ and $|\theta|_\infty$ is the largest absolute value of elements of θ . Let the activation ς be either piecewise continuous or locally quadratic². For example, softplus and sigmoid functions are locally quadratic. Let the function being approximated, $p \in \mathcal{H}^{\alpha,R}([0, 1]^i)$ be Hölder smooth with parameters $\alpha > 0$ and $R > 0$, where $\mathcal{H}^{\alpha,R}(\Omega) := \{p \in \mathcal{H}^\alpha(\Omega) : \|p\|_{\mathcal{H}^\alpha(\Omega)} \leq R\}$. Then Theorem 4.1 in [Ohn and Kim \(2019\)](#) states the existence of positive constants L_0, N_0, Σ_0, B_0 depending only on i, α, R and ς s.t. for any $\epsilon > 0$, the neural network

$$\theta_\epsilon \in \Theta_{i,1} \left(L_0 \log(1/\epsilon), N_0 \epsilon^{-i/\alpha}, \Sigma_0 \epsilon^{-i/\alpha} \log(1/\epsilon), B_0 \epsilon^{-4(i/\alpha+1)} \right)$$

95 satisfies $\sup_{x \in [0,1]^i} |p(x) - p_{\theta_\epsilon}(x)| \leq \epsilon$. Figure 2 shows the upper bound Σ on the net-
 96 work sparsity, $|\theta|_0 \leq \Sigma$, as a function of the error tolerance ϵ and Hölder smoothness,
 97 α , of the function being approximated. Keeping the number of neurons in each layer
 98 fixed, the graph shows that a denser network, with a higher upper bound, results
 99 in a lower approximation error. Conversely, networks with a low number of non-
 100 zero parameters, due to zero weight edges, exhibit larger approximation error. In
 101 the context of no-arbitrage pricing, the theorem suggests a tradeoff between using

²A function $\varsigma : \mathbb{R} \rightarrow \mathbb{R}$ is locally quadratic if \exists an open interval $(a, b) \subset \mathbb{R}$ over which ς is three times continuously differentiable with bounded derivatives and $\exists t \in (a, b)$ s.t. $\varsigma'(t) \neq 0$ and $\varsigma''(t) \neq 0$.

102 a sparse network to enforce the shape constraints, yet increasing the approxima-
 103 tion error. The adverse effect of using a sparse network should also diminish with
 104 increasing smoothness in the function being approximated.




Figure 2: The upper bound on the network sparsity, $|\theta|_0 \leq \Sigma$, as a function of the error tolerance ϵ and Hölder smoothness, α , of the function being approximated. As α increases, the value of Σ is observed to decrease. The plot is shown for $i = 2$ and $\Sigma_0 = 10$.

105 **3.2. Training.** To fit our fully connected or sparse networks to the grid of
 106 option market prices, we solve a loss minimization problem using observations
 107 $\{x_i = (T_i, k_i), p^*(x_i)\}_{i=1}^n$ of n maturity-strike pairs and the corresponding market
 108 put prices:

$$109 \quad (3.1) \quad \min_{\mathbf{W}, \mathbf{b}} \frac{1}{n} \sum_{i=1}^n \left(|p^*(x_i) - p(x_i)| + \lambda^\top \phi(x_i) \right),$$

110 where $p = p_{\mathbf{W}, \mathbf{b}}$ and $\phi = \phi_{\mathbf{W}, \mathbf{b}}$ is a regularization penalty vector:

$$111 \quad \phi := [(\partial_T p)^-, (\partial_{k^2}^2 p)^-, (\sigma - \bar{\sigma})^+ + (\sigma - \underline{\sigma})^-],$$

112 and σ is related to p through (2.1). The choice to measure the error $p^* - p$ under the
 113 L_1 norm, rather than L_2 norm, in (3.1) is motivated by a need to avoid allocating
 114 too much weight to the deepest in-the-money options. The loss function is non-
 115 convex, possessing many local minima and it is generally difficult to find a global
 116 minimum. The first two elements in the penalty vector enforce the shape constraints
 117 and the third element enforces bounds on the estimated local volatility, σ , where
 118 constants (which could also be functions of time) $0 < \underline{\sigma} < \bar{\sigma}$ respectively denote
 119 desired lower and upper bounds on the surface (at each point in time). Of course,
 120 as soon as penalizations are used, a further difficulty, typically involving grid search,
 121 is the need to determine suitable values of the corresponding ‘‘Lagrange multipliers’’
 122 $\lambda \in \mathbb{R}_+^3$, ensuring the right balance between fit to the market prices and the targeted
 123 constraints.

124 **4. Experimental Design.** In this section we benchmark four different combina-
125 tions of architecture and optimization criteria. In each case the error between the
126 prices of the calibrated model and the market data are evaluated on both the train-
127 ing and an out-of-sample test set. Unless reported otherwise, all numerical results
128 shown below correspond to test sets.

129 Note that only the “hard constraints” approach theoretically guarantees that the
130 associated Dupire formula (2.1) returns a positive function. While soft constraints
131 reduce the risk of statistical arbitrage (in the sense of mismatch between model and
132 market prices), they do not however fully prevent butterfly or calendar spread type
133 arbitrages. In particular, the penalties only control the corresponding derivatives at
134 the training points. Compliance with the arbitrage constraints on the majority of
135 the points in the test set is due only to the regularity of these derivatives.

136 The training and test sets are prepared using daily datasets of DAX index Eu-
137 ropean vanilla options of different available strikes and maturities, listed on the 7th,
138 8th (by default below), and 9th, August 2001, i.e. same datasets as already used in
139 previous work Crépey (2002, 2004), for comparison purposes. The corresponding
140 values of the underlying are $S = 5752.51, 5614.51$ and 5512.28 . Each set of obser-
141 vations has just under 200 options market prices. The associated interest rate and
142 dividend yield curves are constructed from zero-coupon and forward curves, them-
143 selves obtained from quotations of standard fixed income linear instruments and
144 from call/put parity applied to the option market prices.

145 Each network has two hidden layers, each with 200 neurons per hidden layer.
146 Note that Dugas et al. (2009) only uses one hidden layer. Two was found important
147 in practice in our case. All networks are fitted with an ADAM optimizer. The
148 learning rate is divided by 10 whenever no improvement in the error on the training
149 set is observed during 100 consecutive epochs. The total number of epoches is
150 limited to 10,000.

151 After training, a local volatility surface is extracted from the interpolated prices
152 by application of the Dupire formula (2.1), leveraging on the availability of the
153 corresponding exact sensitivities. This local volatility surface is then compared with
154 the one calibrated by a classical nonlinear least squares optimization procedure with
155 Tikhonov regularization (see Crépey (2002)).

156 **5. Numerical Results.** Table 1 shows the pricing RMSEs for four different com-
157 binations of architecture and optimization criteria. For the sparse network with hard
158 constraints, we set $\lambda = \mathbf{0}$. For the sparse and dense networks with soft constraints
159 (i.e. penalization but without the conditions on the weights of Section 3), we set
160 $\lambda = [1.0 \times 10^5, 1.0 \times 10^3, 0]$. For avoidance of doubt, so far we do not impose local
161 volatility bounds.

162 The sparse network with hard constraints is observed to exhibit significant pric-
163 ing error, which suggests that this approach is too limited in practice to approach
164 market prices. This conclusion is consistent with Akerer et al. (2019), who choose
165 a soft-constraints approach in the implied volatility approximation (in contrast to
166 our approach which approximates prices).

167 Figure 3 compares the percentage errors in implied volatilities using the sparse
168 network with hard constraints and the dense network with soft constraints ap-
169 proaches, corresponding to the columns 1 and 3 of Table 1. Errors are capped

	Sparse network		Dense network	
	Hard constraints	Soft constraints	Soft constraints	No constraints
Training dataset	28.13	6.87	2.28	2.56
Testing dataset	28.91	4.09	3.53	3.77
Indicative training times	200s	400s	200s	120s

Table 1: Pricing RMSE (absolute pricing errors) without local volatility constraints.

170 at 10% so as not to overwrite the scale of the figures. This confirms that the error
171 levels of the hard constraints approach are too high to imagine a practical use of this
172 approach: the corresponding model would be immediately arbitrable in relation to
173 the market. Those of the soft constraint approach are much more acceptable, with
174 high errors confined to short maturities or far from the money, i.e. in the region
where prices provide little information on volatility.

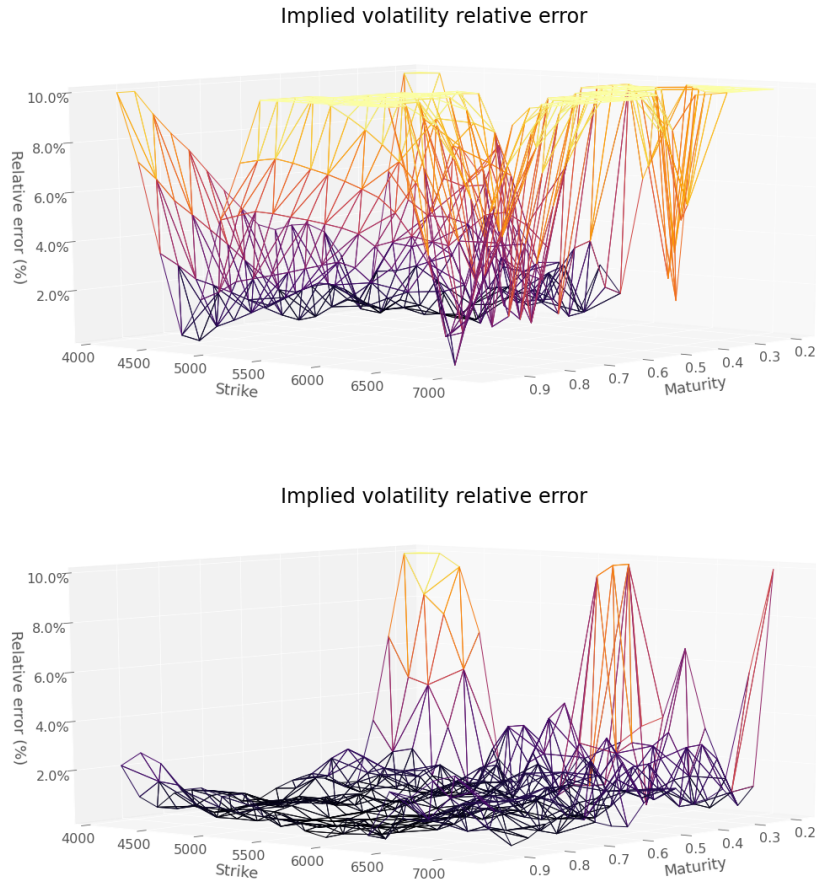


Figure 3: Percentage relative error in the implied volatilities using (a) hard constraints (b) dense networks with soft constraints.

176 **5.1. Local volatility.** We now introduce local volatility bounds into the regular-
177 ization to improve the overall fit in prices and stabilize the local volatility surface.
178 Table 2 shows the RMSEs in absolute pricing resulting from repeating the same set
179 of experiments reported in Table 1, but with the local volatility bounds included in
180 the regularization. For the sparse network with hard constraints, we set $\lambda = [0, 0, 10]$
181 and choose $\underline{\sigma} = 0.05$ and $\bar{\sigma} = 0.4$. For the sparse and dense networks with soft con-
182 straints, we set $\lambda = [1.0 \times 10^5, 1.0 \times 10^3, 10]$. Compared to Table 1, we observe
183 improvement in the test error for the hard and soft constraints approaches when
184 including the additional local volatility penalty term.

	Sparse network		Dense network	
	Hard constraints	Soft constraints	Soft constraints	No constraints
Training dataset	28.04	3.44	2.48	3.48
Testing dataset	27.07	3.33	3.36	4.31
Indicative training times	400s	600s	300s	250s


Table 2: *Price RMSE (absolute pricing errors) with local volatility bounds in the regularization.*

185 Figure 4 shows the comparison of the local volatility surfaces obtained using hard
186 constraints without local volatility bounds, soft constraints without and with local
187 volatility bounds, as well as the Tikhonov regularization approach of Crépey (2002),
188 on data listed on August 7th, 8th, and 9th, 2001. The hard constraint approach yields
189 the smoothest local volatility surface, although the shape is not comparable to those
190 using the Tikhonov regularization approach of Crépey (2002). The soft constraint
191 approach without the local volatility constraints is both spiky (exhibiting outliers
192 on a given day) and unstable (from day to day). In contrast, the soft constraint
193 approach with local volatility bounds yields a reasonable local volatility surface,
194 both at fixed calendar time and in terms of stability across calendar time. The
195 results are then rather comparable to those obtained by Tikhonov regularization
196 (which takes of the order of 30s to run).


197 **6. Conclusion.** Deep learning for option pricing has emerged as a novel com-
198 putational approach for fast option pricing and Greeking. We introduced three
199 variations of deep learning methodology to enforce no-arbitrage interpolation of Eu-
200 ropean option prices: (i) modification of the network architecture to embed shape
201 constraints (hard constraints), (ii) use of shape regularization to enforce the con-
202 straints (soft constraints), and (iii) additional use of local volatility bounds in the
203 regularization via the Dupire formula.

204 Our experimental results suggest that hard constraints reduce the representa-
205 tional power of the network, yet provide the only fail-safe approach to no-arbitrage
206 approximation. Future research would be required to actually prove loss of universal
207 representation in feed-forward networks in the presence of constraints on weights.


208 On the other-hand, soft constraints provide much more accurate prices and
209 implied volatilities. If the Dupire formula is included in the regularization, the
210 corresponding local volatility surface is also reasonably regular (at fixed day) and
211 stable (from day to day). While the neural network training times are longer than




(a) Hard Constraints




(e) Hard Constraints




(i) Hard Constraints




(b) Soft constraints (w/o local vol. constraints)




(f) Soft constraints (w/o local vol. constraints)



(j) Soft constraints (w/o local vol. constraints)




(c) Soft constraints (with local vol. constraints)




(g) Soft constraints (with local vol. constraints)



(k) Soft constraints (with local vol. constraints)



(d) Tikhonov local volatility



(h) Tikhonov local volatility



(l) Tikhonov local volatility

Figure 4: Local volatility over three consecutive days (each row corresponds to a day).

212 the numerical optimization involved in the Tikhonov regularization method, the
213 neural network provides the full surface of prices and local volatilities, as opposed
214 to values at the nodes of a trinomial tree only under the approach of Crépey (2002).

215 **References.**

216 Ackerer, D., N. Tagasovska, and T. Vatter (2019). Deep smoothing of the implied
217 volatility surface. *Available at SSRN 3402942*.

218 Crépey, S. (2002). Calibration of the local volatility in a trinomial tree using
219 Tikhonov regularization. *Inverse Problems* 19(1), 91.

220 Crépey, S. (2004). Delta-hedging vega risk? *Quantitative Finance* 4(5), 559–579.

221 Crépey, S. (2013). *Financial Modeling: A Backward Stochastic Differential Equa-*
222 *tions Perspective*. Springer Finance Textbooks.

223 Dugas, C., Y. Bengio, F. Bédou, C. Nadeau, and R. Garcia (2009). Incorporat-
224 ing functional knowledge in neural networks. *Journal of Machine Learning Re-*
225 *search* 10(Jun), 1239–1262.

226 Dupire, B. (1994). Pricing with a smile. *Risk* 7, 18–20.

227 Garcia, R. and R. Gençay (2000). Pricing and hedging derivative securities with
228 neural networks and a homogeneity hint. *Journal of Econometrics* 94(1-2), 93–
229 115.

230 Gençay, R. and M. Qi (2001). Pricing and hedging derivative securities with neural
231 networks: Bayesian regularization, early stopping, and bagging. *IEEE Transac-*
232 *tions on Neural Networks* 12(4), 726–734.

233 Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning*. MIT Press.

234 Hutchinson, J. M., A. W. Lo, and T. Poggio (1994). A nonparametric approach to
235 pricing and hedging derivative securities via learning networks. *The Journal of*
236 *Finance* 49(3), 851–889.

237 Ohn, I. and Y. Kim (2019, June). Smooth Function Approximation by Deep Neural
238 Networks with General Activation Functions. *Entropy* 21(7), 627.