


Article

Deep Partial Hedging

Songyan Hou ¹, Thomas Krabichler ²  and Marcus Wunsch ^{3,*}

¹ Department of Mathematics, ETH Zurich, 8092 Zürich, Switzerland; songyan.hou@math.ethz.ch

² Centre for Banking & Finance, Eastern Switzerland University of Applied Sciences, 9001 St. Gallen, Switzerland; thomas.krabichler@ost.ch

³ Institute of Wealth & Asset Management, ZHAW Zurich University of Applied Sciences, 8400 Winterthur, Switzerland

* Correspondence: marcus.wunsch@zhaw.ch

Abstract: Using techniques from deep learning, we show that neural networks can be trained successfully to replicate the modified payoff functions that were first derived in the context of partial hedging by Föllmer and Leukert. Not only does this approach better accommodate the realistic setting of hedging in discrete time, it also allows for the inclusion of transaction costs as well as general market dynamics. It needs to be noted that, without further modifications, the approach works only if the risk aversion is beyond a certain level.

Keywords: machine learning; market frictions; transaction costs; partial hedging; risk management

1. Introduction

In a complete market, the writer of an option can eliminate their risk entirely if they initiate a continuous hedging process with a capital position that equals the Black–Scholes price of the option. If the option writer decides to post strictly less capital to initiate the hedge, they will be exposed to shortfall risk. In this situation, they could try to maximize the probability of replicating the option payoff, a strategy named *quantile hedging* by Föllmer and Leukert (1999). Yet another strategy would be *efficient hedging* (cf. Föllmer and Leukert (2000); Föllmer and Schied (2016)), which has the advantage of taking into account the magnitude of the expected shortfall.

Basic Model Setting

W_t shall denote a standard one-dimensional Brownian motion defined on the complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where $(\mathcal{F}_t)_{t \geq 0}$ is the augmentation of the natural filtration $\mathcal{F}_t^W = \sigma(W_s; 0 \leq s \leq t)$ for all $t \geq 0$. We consider a complete market with a single geometric Brownian motion

$$dX_t = X_t(\mu dt + \sigma dW_t),$$

where the drift $\mu \in \mathbb{R}$ and the volatility $\sigma > 0$ are constant. Given the contingent claim's payoff function $H = H(X_T)$, we look for an admissible hedging strategy (V_0, ζ) with

$$V_t = V_0 + \int_0^t \zeta_s dX_s,$$

where ζ is a predictable process with respect to the Brownian motion W such that either

- $\mathbb{P}[V_T > H]$ becomes maximal in the context of quantile hedging, or;
- $\mathbb{E}[\ell((H - V_T)_+)]$ becomes minimal in the context of efficient hedging.

In this note, we consider loss functions of the form $\ell(x) = x^p/p$, $p \in \mathbb{R}_+$. The limiting case $p \downarrow 0$ is, in fact, identical to the quantile hedging problem; the bigger p , the higher the risk aversion. Föllmer and Leukert showed that quantile/efficient hedging is equivalent



Citation: Hou, Songyan, Thomas Krabichler, and Marcus Wunsch. 2022. Deep Partial Hedging. *Journal of Risk and Financial Management* 15: 223. <https://doi.org/10.3390/jrfm15050223>

Academic Editor: James R. Barth

Received: 8 April 2022

Accepted: 13 May 2022

Published: 19 May 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

to delta-hedging options with certain modified payoffs, cf. (Föllmer and Leukert 2000, Proposition 5.2, and Figure 1).

2. Contribution of this Note

In this research note, we show that deep neural networks can be trained to approximate closely the modified payoffs for efficient hedging with lower partial moments with $p > 1$ derived theoretically by Föllmer and Leukert (2000); Leukert (1999). We stress that no other information is needed for this training but the underlying random environment, the capital amount corresponding to the initial hedge, and the option’s target payoff function. To the best of our knowledge, this is the first algorithmic approach to partial hedging whose optimization takes into account transaction costs.

Remark 1. *There is an important limitation to our approach that we encountered when p lies in the unit interval $[0, 1]$. For risk preferences within this range, Föllmer and Leukert (2000) found that the modified contingent claims to be replicated have a knock-out feature that makes their payoff profiles discontinuous. It appears that deep hedging struggles with detecting these discontinuous profiles. Overcoming the problem is apparently not straightforward. On the other hand, when the modified payoff is continuous (e.g., if $p > 1$), then deep partial hedging does work.*

Pricing and hedging is a classical field of mathematical finance with many ramifications. Tackling partial hedging in realistic financial models is a particularly challenging endeavour. Typical impediments from the analytical viewpoint are model intractability originating from complex underlying dynamics (e.g., stochastic volatility, rough volatility, jumps) market features such as incompleteness and trading frictions, and intricate objective functions. Hence, partial hedging has only been analyzed in comparatively simple contexts; cf. Föllmer and Leukert (2000); Leukert (1999). Deep hedging (cf. Buehler et al. (2019)) offers a fresh and powerful approach to deal with these impediments without further ado. One utilises the backpropagation algorithm to find suitable hedging strategies for a given bundle of scenarios. Introducing market frictions or using any sort of stochastically generated scenarios does not adversely affect the operability of the technique. This research note illustrates both the edge and a possible pitfall when adapting deep hedging to the relevant use case of partial hedging.

3. Numerical Results

In this section, we present the numerical results by using techniques from deep learning. We choose similar parameters as in Föllmer and Leukert (2000); a risk-free rate $r = 0.7$, Black–Scholes dynamics with $\mu = 0.08, \sigma = 0.3$, the initial stock price $S_0 = 100$, and a European call option with strike $\kappa = 110$ and maturity $T = 10$. We choose $N = 100$ time steps for time discretization. This leads us to the discretized value process

$$V_t = V_{t-1} + K_t (S_t - S_{t-1}), \quad t = 1, \dots, T, \tag{1}$$

where $(K_t)_{t=1}^T$ denotes the adapted hedging strategy.

Let us set a bankruptcy bound $B < 0$. In order to incorporate 0-admissibility, we modify the update in the following way: if updating with (1) results in $V_{t+1} < B$, we set

$$V_{t+1} = (V_t - B) \exp \left(\pi_t (\mu dt + \sigma dW_t) - \frac{\pi_t^2 \sigma^2 dt}{2} \right) + B$$

$$\pi_t = \frac{K_t S_t}{V_t}$$

instead. If $V_t = B$, we claim bankruptcy and leave the market, i.e., $K_h = 0$ prevails for all $h \geq t$. With this modification, we make sure all strategies are B -admissible.¹ Then, we numerically minimize the loss consisting of three terms

$$\mathcal{L}(K) = \mathcal{L}_p(K) + \mathcal{L}_{\text{cost}}(K) + \mathcal{L}_{\text{ad}}(K), \tag{2}$$

where

$$\begin{aligned} \mathcal{L}_p(K) &= \mathbb{E} \left[\ell((H - V_T)_+) \right], \\ \mathcal{L}_{\text{cost}}(K) &= c_{\text{cost}} \mathbb{E} \left[\sum_{t=1}^{T-1} |K_{t+1} - K_t| \cdot S_t \right], \\ \mathcal{L}_{\text{ad}}(K) &= c_{\text{ad}} \mathbb{E} \left[(-\min_t V_t)_+ \right], \end{aligned}$$

for the hyperparameters c_{cost} and c_{ad} . The minimization of the first term \mathcal{L}_p is the primary objective of efficient hedging. The second term takes into account proportional transaction costs. The last term penalizes the deviation from 0-admissibility. We consider the so-called deep strategy

$$K_t = \mathcal{N}_t(S_t),$$

where \mathcal{N}_t belongs to a certain class of neural networks \mathcal{NN} . Its structure comprises two hidden layers and each hidden layer consists of 21 nodes. The detailed configuration can be found in the publicly available code². We utilize Adam mini-batch training with a learning rate of 0.01, cf. Kingma and Ba (2015). For different risk aversion levels $p \in \{1, 1.1, 2\}$, Tables 1 and 2 depict the performance line-up for deep hedging and discretized delta hedging respectively.

Table 1. Efficient hedging loss for different risk aversion levels without transaction costs.

	$p = 1$	$p = 1.1$	$p = 2$
deep hedge	18.14	16.75	14.31
delta hedge	19.64	36.69	32.63

Table 2. Efficient hedging loss for different risk aversion levels with proportional transaction costs of 1%.

	$p = 1$	$p = 1.1$	$p = 2$
deep hedge	18.09	16.39	15.58
delta hedge	26.45	44.01	37.84

Finally, we compare the terminal wealth between deep hedging and delta hedging in Figures 1–6, from which we can see that the violation of admissibility by delta hedging is substantial.




Figure 1. Terminal wealth ($p = 1$).




Figure 2. Hedging strategy ($p = 1$).




Figure 3. Terminal wealth ($p = 1.1$).




Figure 4. Hedging strategy ($p = 1.1$).




Figure 5. Terminal wealth ($p = 2$).




Figure 6. Hedging strategy ($p = 2$).

4. Conclusions and Outlook

In this note, we have shown that deep partial hedging can replicate the modified contingent claims as derived theoretically in the context of efficient hedging by Föllmer and Leukert (2000), without any prior knowledge of the latter's payoff profile. Our findings hold for risk-neutral and risk-averse option writers ($p \geq 1$); for risk-taking option writers ($p \in [0, 1)$), it appears to be difficult to capture the modified payoff with two barriers using deep neural networks (cf. Föllmer and Leukert (2000)). We refer further research into these challenges, as well as the investigation of deep partial hedging for more general market dynamics and derivatives, to future work.

Author Contributions: Conceptualization, M.W.; Investigation, S.H., T.K. and M.W.; Software, S.H. and T.K.; Writing, S.H. and T.K.; Visualization, S.H. All authors have read and agreed to the published version of the manuscript.

Funding: M.W. acknowledges partial funding by COST Action 19130: Fintech and Artificial Intelligence in Finance—Towards a transparent financial industry, supported by COST (European Cooperation in Science and Technology). The APC was funded by the ZHAW Zurich University of Applied Sciences.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

Notes

¹ In our numerical experiments, we set $B = -100$.

² <https://github.com/justinhou95/DeepHedging>, accessed on 16 May 2022.

References

- Buehler, Hans, Lukas Gonon, Josef Teichmann, and Ben Wood. 2019. Deep hedging. *Quantitative Finance* 19: 1271–91. [CrossRef]
- Föllmer, Hans, and Peter Leukert. 1999. Quantile hedging. *Finance and Stochastics* 3: 251–73. [CrossRef]
- Föllmer, Hans, and Peter Leukert. 2000. Efficient hedging: Cost versus shortfall risk. *Finance and Stochastics* 4: 117–46. [CrossRef]
- Föllmer, Hans, and Alexander Schied. 2016. *Stochastic Finance*. Berlin and Boston: de Gruyter.
- Kingma, Diederik, and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. Paper presented at the 3rd International Conference on Learning Representations (ICLR 2015), San Diego, CA, USA, May 7–9.
- Leukert, Peter. 1999. Absicherungsstrategien zur Minimierung des Verlustrisikos. Ph.D. thesis, Humboldt Universität zu Berlin, Berlin, Germany.