Annals of Operations Research
https://doi.org/10.1007/s10479-025-06867-3

ORIGINAL RESEARCH

Data-driven hedging with generative models

Rama Cont[1] · Milena Vuletić[1]

Received: 4 June 2025 / Accepted: 18 September 2025

Abstract

We propose a nonparametric data-driven methodology for hedging using generative models. In contrast with model-based hedging approaches relying on sensitivity analysis of model pricing functions, our approach uses a conditional generative model trained on market data to simulate realistic market scenarios given current market conditions and computes hedge ratios which minimize risk across these scenarios. The approach incorporates transaction costs, leads to an optimal selection of hedging instruments, and adapts to market conditions. We illustrate the effectiveness of this methodology for hedging option portfolios using VolGAN, a generative model for implied volatility surfaces. The out-of-sample performance of the method matches and improves over delta and delta-vega hedging, without retraining the model for more than 4 years after the training period.

Mathematics Subject Classification Primary 26A16 · 26A30 · 26A33 · 60L99 · 60H10 · 60H50

1 Introduction

Generative models are a class of machine learning models trained to simulate samples with statistical features similar to a given data set. When properly trained, generative models can be powerful tools for non-parametric, model-free simulation, and have been deployed in this setting for financial applications such as portfolio loss simulation (Cont et al., 2025), forecasting (Vuletić et al., 2024), backtesting of hedging strategies (Vuletić & Cont, 2024), and portfolio construction (Chen et al., 2025). These studies have demonstrated the usefulness of genera-

This paper appears in a special issue of the *Annals of Operations Research* honoring Professor John M. Mulvey (Princeton), on the occasion of his 80th birthday, for his many contributions to large-scale optimization modeling and its applications to support decision-making by businesses and governments.

✉ Rama Cont
Rama.Cont@maths.ox.ac.uk

1 Mathematical Institute, University of Oxford, Oxford, UK

Published online: 16 October 2025

tive models for learning complex features of market data and generating forward-looking scenarios.

We propose a methodology which leverages these features of generative models for the design of hedging strategies and portfolio risk management. By combining a conditional generative model trained on market data with an optimization layer, we introduce a computationally tractable framework for data-driven hedging. In contrast with commonly used model-based approaches for hedging, our method derives a data-driven hedging strategy which exploits information in the training data. Furthermore, it has the ability to perform automatic selection of hedging instruments and incorporate transaction costs and market impact.

1.1 Contribution

We propose a nonparametric data-driven methodology for hedging using generative models. The key idea is to learn the co-movements of potential hedging instruments and the target portfolio from the market data, and then use this information to compute hedging strategies.

In contrast with model-based hedging approaches relying on sensitivity analysis of model-based pricing functions, our approach uses a conditional generative model trained on market data to simulate realistic market scenarios given current market conditions, and computes hedge ratios that minimize risk across these scenarios. This optimization procedure enables the automated selection of hedging instruments, and allows for the incorporation of transaction costs and market impact.

The result is a fully data-driven method for implementing hedging strategies: market data is used to train a generative model, which is then employed to compute hedge ratios. Opting for conditional variance minimization, as opposed to more general risk measures, leads to linear hedge ratios which are easily computable via linear regression. We illustrate the effectiveness of this methodology for hedging option portfolios using VolGAN (Vuletić & Cont, 2024), a customized conditional generative adversarial network (GAN) (Goodfellow et al., 2014; Mirza & Osindero, 2014) for implied volatility surfaces, and compare its performance with delta and delta-vega hedging.

1.2 Relation with previous literature

The dominant paradigm in the design of hedging strategies is *model-based hedging*, which requires full specification of the dynamics of all risk factors affecting a portfolio over the risk horizon. Hedging is then formulated as an intertemporal optimization or stochastic control problem, which is typically solved using dynamic programming or backward induction. Hedging strategies obtained in this manner are exposed to model uncertainty and may exhibit a significant level of 'model risk' (Cont, 2006).

Regression-based hedging is arguably the simplest form of 'data-driven' hedging: the idea is to perform a least squares regression of the daily changes of the target portfolio with respect to the daily changes of the hedging instruments and use the coefficients as hedge ratios (Sercu & Wu, 2000; Cont & Kan, 2011). As shown by Cont and Kan (2011), regression-based hedging can exhibit superior performance when compared to model-based sensitivity hedging, especially in situations where models are arguably miss-specified. However, historical regression-based hedge ratios are backward-looking, assume stationarity of covariance structure in comovements, and may not properly account for changes in market conditions.

Starting with the pioneering work of Hutchinson et al. (1994), who demonstrated that a simple one-layer neural network with four neurons could learn to hedge S& P500 options, neural networks have been used in various ways to calculate hedging strategies (see survey by Ruf and Wang (2020) and references therein). As noted by Ruf and Wang (2020), in most of these studies neural networks are used to learn a pricing function and hedging is either not discussed or it is addressed using a sensitivity-based approach (Cohen et al., 2022; Fan & Sirignano, 2024). In some cases, a neural network is trained to learn the hedge ratios directly from market data (Chen & Sutcliffe, 2012; Carverhill & Cheuk, 2003; Ruf & W. W., 2022).

(Buehler et al., 2019) used deep neural networks to learn optimal hedge ratios within a parametric model via reinforcement learning. This approach is able to incorporate transaction costs, alternative risk measures, and trading constraints but is not data-driven: training is done via deep reinforcement learning using (millions of) simulated scenarios based on a parametric model. Similar (model-based) approaches were explored by Cao et al. (2023); Lütkebohmert et al. (2022); Limmer and Horvath (2024). These approaches are based on *neuro-dynamic programming* (Bertsekas & Tsitsiklis, 1996), i.e., on the parameterization of hedging strategies via a neural network, which is then trained using reinforcement learning. In contrast, our approach uses neural networks for the generation of market scenarios, while the computation of hedge ratios is done through direct optimization, a much simpler approach which results in greater computational efficiency. Also, importantly, these approaches are model-based, not data-driven: training is done using scenarios simulated from a reference model, not on market data.

As explained in Sect. 3, given the one-step ahead scenarios generated by the conditional generative model, we compute hedge ratios by solving a *local risk minimization approach* which is a convex optimization problem. Our optimization criterion combines a conditional variance term with a ℓ_1 penalty for transaction costs. A similar approach was studied by Lamberton et al. (1998), extending work of Föllmer and Schweizer (1988); Schweizer (1995) on quadratic hedging to include transaction costs. However, differently from these references which focus on (not necessarily self-financing) replication strategies computed by backward induction, our hedging strategy is self-financing and computed sequentially using one-step ahead scenarios, provided by the generative model. In particular, our approach does *not* require knowledge of the full price dynamics up to expiry and is thus applicable in sequential setting to conditional generative models based taking current market conditions as input. Our approach is similar in spirit to the model-predictive control approach (Bemporad et al., 2014) and the progressive hedging approach (Rockafellar & Wets, 1991; Rockafellar, 2018), both of which employ scenario-based optimization and proceed forward, not backward, in time. In contrast to multi-stage stochastic programs (Mulvey & Ruszczyński, 1995), we solve a sequence of one-period stochastic programs, leading to a fast and scalable approach applicable to large-scale problems.

Outline. Section 2 reviews commonly used approaches for computing hedging strategies. Section 3 describes our proposed methodology for data-driven hedging using generative models. Section 4 illustrates how our approach can be used to hedge volatility risk with VolGAN, a conditional scenario generator for implied volatility dynamics. Section 5 summarizes our contributions. The appendix contains a coordinate descent algorithm for data-driven hedging with transaction costs, further performance tests, and robustness tests for the VolGAN example.

2 Sensitivity-based hedging vs optimization-based hedging

We consider the problem of hedging a portfolio exposed to a set of risk factors using a set of *hedging instruments*. We distinguish two sets of risk factors:

- market prices $S_t^1, ..., S_t^n$ of tradable primitive/underlying assets;
- other (non-price) risk factors denoted $X_t = (X_t^1, ..., X_t^d)$. Examples of such risk factors may be: implied volatilities, interest rates, credit spreads, etc.

Primitive assets are tradable, while risk factors $X_t^1, ..., X_t^d$ are not directly traded, i.e., do not represent prices of tradable instruments. Examples of such non-price risk factors are yields (for bonds) or implied volatilities for options.

We wish to hedge a portfolio whose value at time *t* is V_t, using a set $(H^i, i \in \mathcal{H})$ of *hedging instruments* whose values H_t^i are sensitive to these risk factors. We assume a rebalancing frequency Δt (for example, one day).

Our risk management framework requires the following ingredients:

1. A *pricing function f* which computes the value of the portfolio as a function of the risk factors
$$V_t = f(t, S_t, X_t) = f(t, S_t^1, ..., S_t^n, X_t^1, ..., X_t^d); \tag{1}$$
2. *Pricing functions* for the hedging instruments:
$$H_t^j = h_j(t, S_t, X_t) = h_j(t, S_t^1, ..., S_t^n, X_t^1, ..., X_t^d); \tag{2}$$
3. A generative model capable of generating one-step ahead (e.g. one-day ahead) scenarios for co-movements of prices S_t and risk factors X_t
$$\text{Inputs}: \ (S_t, X_t, S_{t-\Delta t}, X_{t-\Delta t}, ...), \text{noise term } \omega \quad \xrightarrow{\text{G}} \quad \text{Output}: \ (S_{t+\Delta t}(\omega), X_{t+\Delta t}(\omega)). \tag{3}$$

We consider a self-financing tracking portfolio with positions ϕ_t^i in the hedging instruments $(H_t^i, i \in \mathcal{H})$, and ψ_t in a money market account earning interest at rate r_t. The value Π_t of this tracking portfolio satisfies:

$$\Pi_{t+\Delta t} - \Pi_t = \Delta \Pi_t = \psi_t \ r_t \Delta t + \sum_{i \in \mathcal{H}} \phi_t^i (H_{t+\Delta t}^i - H_t^i), \quad \Pi_0 = V_0. \tag{4}$$

The self-financing condition implies

$$\psi_t = \Pi_t - \sum_{i \in \mathcal{H}} \phi_t^i H_t^i - \sum_{i \in \mathcal{H}} c_t^i |\phi_t^i - \phi_{t-\Delta t}^i|, \tag{5}$$

where c_t^i is the cost of trading one unit of H^i at time *t*. The **tracking error** Z_t is the value of the hedged position:

$$Z_t = V_t - \Pi_t, \quad Z_0 = 0. \tag{6}$$

The two main approaches to determine the hedging strategies ϕ_t are *sensitivity-based hedging* and *conditional risk minimization*, with (local) quadratic hedging being a special case of the latter. We now describe these two approaches in some detail.

2.1 Sensitivity-based hedging

The most common approach for constructing hedging strategies is based on *sensitivities* to risk factors. Given a 'standard' shift ϵ_i for risk factor *i*, we define the sensitivity to a risk factor *i* as the change in portfolio value under a "standardized shift" $X^i \to X^i + \epsilon_i$ to the risk factor $i \in \{1, \ldots, d\}$. Denote the sensitivity to risk factor X^i at time *t* by

$$\zeta_t^i(f) = \frac{f(t, S_t, X_t^0, \ldots, X_t^{i-1}, X_t^i + \epsilon_i, X_t^{i+1}, \ldots, X_t^d) - f(t, S_t, X_t)}{\epsilon_i}. \tag{7}$$

Similarly, denote sensitivities to market prices S^i, for $i \in \{1, \ldots, n\}$ by

$$\Delta_t^i(f) = \frac{f(t, S_t^0, \ldots, , S_t^{i-1}, S_t^i + \epsilon_j, S_t^{i+1}, \ldots, S_t^n, X_t) - f(t, S_t, X_t)}{\epsilon_j}. \tag{8}$$

If the pricing function *f* is differentiable in S_t, X_t, then

$$\lim_{\epsilon_i \to 0} \zeta_t^i(f) = \frac{\partial f}{\partial X^i}(t, S_t, X_t), \qquad \lim_{\epsilon_i \to 0} \Delta_t^i(f) = \frac{\partial f}{\partial S^i}(t, S_t, X_t),$$

Analogously, we define sensitivities of the hedging instruments H^j to risk factors X^i and market prices S^i, which we denote by $\zeta_t^i(h^j)$ and $\Delta_t^i(h^j)$, respectively. In cases where neural networks represent the pricing functions, sensitivities are readily computable using Automatic Differentiation (AD) techniques (Fries, 2019; Capriotti & Giles, 2024).

The sensitivity of the portfolio value Π_t to a risk factor X^i is obtained by summing over positions:

$$\zeta_t^i(h) = \sum_{j \in \mathcal{H}} \phi_t^i \zeta_t^i(h^j),$$

and its sensitivity to S^i is

$$\Delta_t^i(h) = \sum_{j \in \mathcal{H}} \phi_t^i \Delta_t^i(h^j).$$

The idea of sensitivity-based hedging is to choose the positions in the hedging instruments to achieve zero/low overall sensitivity to risk factors (S_t, X_t). The hedged portfolio is immunized to small movements in the risk factor X^i for $i = 1, \ldots, d$ if the sensitivity of the hedged position $Z_t = V_t - \Pi_t$ is zero:

$$\zeta_t^i(f) = \zeta_t^i(h). \tag{9}$$

Similarly, the hedged portfolio is immunized to small changes in market prices S_t^i for $i = 1, \ldots, n$ if

$$\Delta_t^i(f) = \Delta_t^i(h). \tag{10}$$

Examples of commonly used sensitivity-based hedging strategies include delta hedging and delta-vega hedging of option portfolios, and immunization strategies of bond portfolios.

Hedging via risk immunization requires the user to pre-specify both the risk factors and the hedging instruments, a choice that is not necessarily unique. For instance, in delta-vega hedging, volatility risk can be mitigated by holding a position in *any* option on the same underlying asset; sensitivity considerations alone do not guide us in the choice of the hedging instrument.

Additionally, risk immunization typically accounts for only small perturbations in risk factors. Extensions of this approach consider inclusion of higher order sensitivities, for example gamma hedging for options or duration-convexity immunization for bond portfolios (De La Peña et al., 2021). However, even these extensions cannot account for tail events.

Finally, these shifts are applied to individual risk factors in isolation, disregarding their *co-movements*. If risk factors are correlated, this may mean that the scenarios underlying these sensitivity computations are unrealistic.

2.2 Hedging by local risk minimization

Another approach for computing hedging strategies is based on the idea of *local riskminimization* (Föllmer & Schweizer, 1988; Lamberton et al., 1998; Mercurio & Vorst, 1996; Schweizer, 1995). In this approach, one sequentially computes hedge ratios to minimize a one step-ahead risk criterion, based on current information in market variables.

Denote by $\{\mathcal{F}_t\}_{t \geq 0}$ the filtration representing the information contained in the history of market prices and risk factors up to time *t*. For a hedging strategy to be implementable, the hedge ratios ϕ_t have to be $\mathcal{F}_t$-measurable. The idea of local risk minimization is to compute ϕ_t by minimizing a risk measure ρ of the portfolio conditional on current market information

$$\inf_{\phi_t} \ \rho(Z_{t+\Delta t} | \mathcal{F}_t).$$

A tractable class is given by conditional risk measures expressible in the form

$$\rho(Z_{t+\Delta t} | \mathcal{F}_t) = \mathbb{E}\left[L(Z_{t+\Delta t}, Y_t) \middle| \mathcal{F}_t \right]. \tag{11}$$

where *L* is a loss function and Y_t is observable at *t* i.e. $\{\mathcal{F}_t\}$-adapted. In particular, (11) is amenable to computation by simulation.

An important example of conditional risk measure expressible in the form (11) is the conditional variance:

$$\rho(Z_{t+\Delta t}|\mathcal{F}_t) = \mathbb{E}\left[\left(Z_{t+\Delta t} - \mathbb{E}[Z_{t+\Delta t}|\mathcal{F}_t]\right)^2 \middle| \mathcal{F}_t\right], \tag{12}$$

where

$$L(Z,Y) = (Z-Y)^2, \quad Y_t = \mathbb{E}[Z_{t+\Delta t}|\mathcal{F}_t].$$

Remark 2.1 One may of course consider other risk measures, such as Value-at-Risk (VaR), Expected Shortfall etc. However, among these examples the only risk measure expressible in the form (11) is conditional variance.

Minimizing the conditional variance of the hedging error (conditional on market variables $\mathcal{F}_t$) leads to a (local) **regression problem**. The objective is to solve

$$\min_{\phi_t} \operatorname{var}\left(Z_{t+\Delta t}\middle|\mathcal{F}_t\right). \tag{13}$$

Proposition 2.2 *The hedge ratios which minimize the variance of the tracking error $Z_{t+\Delta t}$ conditional on $\mathcal{F}_t$ are given as regression coefficients of ΔV_t on $\{\Delta H_t^i\}_{i\in\mathcal{H}}$:*

$$\min_{\phi_t} \operatorname{var}\left(Z_{t+\Delta t}\middle|\mathcal{F}_t\right) = \min_{A_t,\phi_t} \mathbb{E}\left[\left(V_{t+\Delta t} - V_t - A_t - \sum_{i\in\mathcal{H}} \phi_t^i (H_{t+\Delta t}^i - H_t^i)\right)^2 \middle|\mathcal{F}_t\right]. \tag{14}$$

Proof Since ψ_t, V_t, and ϕ_t^i, H_i^t are $\mathcal{F}_t$-measurable, we have

$$\begin{aligned}
\operatorname{var}\left(Z_{t+\Delta t}\middle|\mathcal{F}_t\right) &= \operatorname{var}\left(Z_{t+\Delta t} - Z_t\middle|\mathcal{F}_t\right) = \operatorname{var}\left(\Delta V_t - \Delta\Pi_t\middle|\mathcal{F}_t\right)\\
&= \operatorname{var}\left(\Delta V_t - \sum_{i\in\mathcal{H}} \phi_t^i \Delta H_t^i - \psi_t r_t \Delta t\middle|\mathcal{F}_t\right)\\
&= \operatorname{var}\left(\Delta V_t - \sum_{i\in\mathcal{H}} \phi_t^i \Delta H_t^i\middle|\mathcal{F}_t\right).
\end{aligned}$$

In fact, for any $\mathcal{F}_t$-measurable A_t, the following holds:

$$\operatorname{var}\left(\Delta V_t - \sum_{i\in\mathcal{H}} \phi_t^i \Delta H_t^i\middle|\mathcal{F}_t\right) = \operatorname{var}\left(\Delta V_t - A_t - \sum_{i\in\mathcal{H}} \phi_t^i \Delta H_t^i\middle|\mathcal{F}_t\right).$$

Hence, we can choose A_t such that

$$\mathbb{E}\left[\Delta V_t - A_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\middle|\mathcal{F}_t\right] = 0, \qquad i.e. \quad A_t = \mathbb{E}\left[\Delta V_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\middle|\mathcal{F}_t\right]. \quad (15)$$

Furthermore, for any $\mathcal{F}_t$-measurable α_t:

$$\begin{aligned}\mathbb{E}\left[\left(\Delta V_t - \alpha_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\right)^2\middle|\mathcal{F}_t\right] &= \mathrm{var}\left(\Delta V_t - \alpha_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\middle|\mathcal{F}_t\right)\\ &\quad + \mathbb{E}\left[\Delta V_t - \alpha_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\middle|\mathcal{F}_t\right]^2\\ &= \mathrm{var}\left(\Delta V_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\middle|\mathcal{F}_t\right)\\ &\quad + \mathbb{E}\left[\Delta V_t - \alpha_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\middle|\mathcal{F}_t\right]^2.\end{aligned} \quad (16)$$

Hence, the value of α_t which minimizes the conditional second moment (16) of $\Delta V_t - \alpha_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i$ is precisely the conditional expectation A_t given by (15):

$$\alpha_t^* = \mathrm{argmin}_\alpha \mathbb{E}\left[\left(\Delta V_t - \alpha - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\right)^2\middle|\mathcal{F}_t\right], \quad \text{i.e.} \quad \alpha_t^* = \mathbb{E}\left[\Delta V_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\middle|\mathcal{F}_t\right] = A_t.$$

This implies that

$$\mathrm{var}\left(Z_{t+\Delta t}\middle|\mathcal{F}_t\right) = \min_{\alpha_t}\mathbb{E}\left[\left(\Delta V_t - \alpha_t - \sum_{i\in\mathcal{H}}\phi_t^i\Delta H_t^i\right)^2\middle|\mathcal{F}_t\right].$$

This is an Ordinary Least Squares (OLS) regression, conditional on the information available at time *t*:

$$\Delta V_t = A_t + \sum_{i\in\mathcal{H}}\phi_t\Delta H_t^i + \epsilon_t,$$

where ϵ_t is uncorrelated with the P &Ls of the hedging instruments and $\mathbb{E}\left[\epsilon_t\middle|\mathcal{F}_t\right] = 0$.

So, finally, we have shown that minimizing the conditional variance of the hedging error (conditional on market variables $\mathcal{F}_t$) leads to a least squares **regression problem**:

$$\min_{\phi_t}\mathrm{var}\left(Z_{t+\Delta t}\middle|\mathcal{F}_t\right) = \min_{A_t,\phi_t}\mathbb{E}\left[\left(V_{t+\Delta t} - V_t - A_t - \sum_{i\in\mathcal{H}}\phi_t^i(H_{t+\Delta t}^i - H_t^i)\right)^2\middle|\mathcal{F}_t\right]. \quad (17)$$

□

Remark 2.3 Equation (14), determines the hedge ratios which minimize the conditional variance of the tracking error. The hedging strategy requires to specify as well the cash/money market component ψ_t. This is determined by the self-financing condition (5). Any other choice for ψ_t would lead to a *non self-financing* hedging strategy. For example (Schweizer, 1995; Föllmer & Schweizer, 1988; Lamberton et al., 1998) choose $\psi_t = A_t$ given by (15), which yields the martingale property for the tracking error but results in the loss of self-financing property.

Remark 2.4 Note that we are not operating under a 'risk-neutral' measure and there is no martingale assumption on price dynamics.

Remark 2.5 The conditional variance minimization problem (13) is a regression of the *change* in portfolio value ΔV_t on the changes ΔH_t^i in values of the hedging instruments, conditional on the information available at time *t*. Using the $\mathcal{F}_t$-measurability of V_t, Π_t, and H_t, we can also express (13) as a regression of values, rather than changes, of these same instruments:

$$\min_{\phi_t} \operatorname{var}\left(Z_{t+\Delta t} \middle| \mathcal{F}_t\right) = \min_{A_t,\phi_t} \mathbb{E}\left[\left(V_{t+\Delta t} - A_t - \sum_{i\in\mathcal{H}} \phi_t^i H_{t+\Delta t}^i\right)^2 \middle| \mathcal{F}_t\right]. \tag{18}$$

2.3 Accounting for transaction costs

Since the transaction cost term $\sum_{i\in\mathcal{H}} c_t^i|\phi_t^i - \phi_{t-\Delta t}^i|$ is $\mathcal{F}_t-$measurable, it contributes to the conditional mean but not to the conditional variance. The optimization problem (14) therefore does not penalize transaction costs. In order to account for transaction costs one may proceed in different ways. A common approach is to minimize the conditional second moment (instead of variance) (Schweizer, 1995; Mercurio & Vorst, 1996, 1997). This leads to an objective function with quadratic penalties for transaction costs, but also cross-terms of the type

$$c_t^i|\phi_t^i - \phi_{t-\Delta t}^i|\ \phi_t^j \Delta H_t^j, \qquad i \neq j,$$

which lacks a financial interpretation and leads to numerical difficulties. Another approach is to use conditional variance, but incorporate transaction costs for the *next* period as in Lamberton et al. (1998). This leads to a variance contribution, but breaks down the analytical tractability of the quadratic problem (14). Furthermore, in this case we lose the interpretation of the objective as variance of the tracking error.

We propose instead a different formulation, which is to include transaction costs as a penalty in the objective function, leading to:

$$\min_{A_t,\phi_t} \mathbb{E}\left[\left(V_{t+\Delta t} - V_t - A_t - \sum_{i\in\mathcal{H}} \phi_t^i (H_{t+\Delta t}^i - H_t^i)\right)^2 \middle| \mathcal{F}_t\right] + \lambda \sum_{i\in\mathcal{H}} c_t^i|\phi_t^i - \phi_{t-\Delta t}^i|, \tag{19}$$

where $\lambda > 0$ is a regularization parameter. A similar approach has been considered by Li et al. (2022) but in a global risk minimization setting. In (19), each term has a clear financial interpretation. Since A_t does not appear in the second term, by Proposition 2.2, the first term is still the conditional variance of the tracking error, while the second term is proportional to the transaction cost. As the transaction cost is naturally expressed as a (weighted) ℓ_1 norm, (19) is in fact a LASSO regression (Tibshirani, 1996). This leads to the desirable property of *sparsity* which, in financial terms, corresponds to selecting a sparse set of hedging instruments.

3 Dynamic data-driven hedging with generative models

An important feature of local risk minimization problem (19) is that it only requires knowledge of the one step-ahead conditional distribution. However, this conditional distribution remains inaccessible, rendering the local risk minimization problem analytically intractable even in simple models. On the other hand, we will demonstrate that this problem may be efficiently solved using conditional generative models. We will now show how to sequentially compute locally risk-minimizing hedging strategies using a conditional generative model.

A conditional generative model G for (S_t, X_t) allows to generate samples $(S_{t+\Delta t}(\omega), X_{t+\Delta t}(\omega))$ whose distribution approximates the conditional distribution of $(S_{t+\Delta t}, X_{t+\Delta t})$ given $\mathcal{F}_t$. These generated samples are forward-looking one step-ahead scenarios, which may be used to estimate the conditional mean and variance of the tracking error via

$$\widehat{A}_t = \frac{1}{N}\sum_{k=1}^{N}\left(\Delta V_t(\omega_k) - \sum_{i\in\mathcal{H}}\phi_t^i \Delta H_t^i(\omega_k)\right), \qquad \widehat{\mathrm{var}(\Delta Z_t|\mathcal{F}_t)} = \frac{1}{N}\sum_{k=1}^{N}\left(\Delta V_t(\omega_k) - \widehat{A}_t - \sum_{i\in\mathcal{H}}\phi_t^i \Delta H_t^i(\omega_k)\right)^2,$$

where $\{\omega_k\}_{k=1,..,N}$ are N i.i.d. scenarios from the generator G. We can then compute an estimator of (19), and compute hedge ratios by optimizing this estimator.

This leads to the following sequential optimization problem: at each step t, given previous hedge ratios $\phi_{t-\Delta t}$ we solve the following LASSO regression problem:

$$\inf_{A_t\in\mathbb{R},\phi_t\in\mathbb{R}^{d+1}} \underbrace{\frac{1}{N}\sum_{k=1}^{N}\left(\Delta V_t(\omega_k) - A_t - \sum_{i\in\mathcal{H}}\phi_t^i \Delta H_t^i(\omega_k)\right)^2}_{\text{Hedging error variance}} + \alpha g_0 \underbrace{\sum_{i\in\mathcal{H}} c_t^i \left|\phi_t^i - \phi_{t-\Delta t}^i\right|}_{\text{Rebalancing cost}}, \tag{20}$$

where

- c_t^i is the cost of trading one unit of H^i at time *t* (we use the half the bid-ask spread for instrument H^i as a proxy);
- g_0 is the initial gross position, and $\alpha > 0$ is a dimensionless regularization parameter.

The solution to (20) can be computed via coordinate descent with soft thresholding (Hastie et al., 2009), details of which are available in the Appendix.

The regularization term in (20) favors lower turnover and hedging instruments H^i with lower transaction costs. Transaction costs lead to an ℓ_1-penalty which is known to induce sparsity in regression: it leads to the "minimal" (sub)set of rebalancing transactions. Sparsity means that not all potential hedging instruments are selected, leading to an **automatic selection of hedging instruments**.

Given that hedging instruments in $\mathcal{H}$ might have highly correlated returns (especially if $|\mathcal{H}| = J$ is large), regularization is useful in (20). Furthermore, other (position, sensitivity) constraints can be easily incorporated in the optimization.

The resulting procedure for data-driven hedging is described in Algorithm 1. Steps to implement efficient coordinate descent for generalized LASSO regression are available in Algorithm 2 in the Appendix.

Algorithm 1 Data-driven hedging with no market impact

Input : Market prices S_5 and risk factors X_t at time t;
Model-free pricing functions: portfolio value $f(t, S_t, X_t)$ and hedging instruments $h_i(t, S_t, X_t)$
A generative model G for simulating $(S_{t+\Delta t}, X_{t+\Delta t})$ given current market information;
Previous hedge ratios $\phi_{t-\Delta t}$;
Regularization parameter α;
Initial gross position g_0

Output: Hedge ratios $(\phi_t^i)_{i\in\mathcal{H}}$ and the cash position ψ_t;
Estimate A_t of $\mathbb{E}\left[\Delta V_t - \sum_{i\in\mathcal{H}} \phi_t^i \Delta H_t^i \middle| \mathcal{F}_t\right]$

Step 1: Simulate forward-looking scenarios
Sample N one-step-ahead scenarios:
$\{(S_{t+\Delta t}(\omega_i), X_{t+\Delta t}(\omega_k))\}_{k=1}^N$ using $G(S_t, X_t, \ldots, \omega_k)$
Compute simulated portfolio values:
$V_{t+\Delta t}(\omega_k) = f(t+\Delta t, S_{t+\Delta t}(\omega_k), X_{t+\Delta t}(\omega_k))$
Compute simulated hedging instrument values:
$H^i_{t+\Delta t}(\omega_k) = h_i(t+\Delta t, S_{t+\Delta t}(\omega_k), X_{t+\Delta t}(\omega_k))$ for $i \in \mathcal{H}$

Step 2: Solve regularized regression (LASSO)
Estimate A_t and hedge ratios $\phi_t = (\phi_t^i)_{i\in\mathcal{H}}$ by solving:

$$\min_{A_t\in\mathbb{R},\, \phi_t\in\mathbb{R}^J} \frac{1}{N}\sum_{k=1}^{N}\left(\Delta V_t(\omega_k) - A_t - \sum_{i\in\mathcal{H}} \phi_t^i \Delta H_t^i(\omega_k)\right)^2 + \alpha g_0 \sum_{i\in\mathcal{H}} c_t^i \left|\phi_t^i - \phi_{t-\Delta t}^i\right|$$

Step 3: Update the cash position

$$\psi_t = \Pi_t - \sum_{i\in\mathcal{H}} \phi_t^i H_t^i - \sum_{i\in\mathcal{H}} c_t^i |\phi_t^i - \phi_{t-\Delta t}^i|$$

return A_t, $(\phi_t^i)_{i\in\mathcal{H}}$, ψ_t

Incorporating market impact

The proportional transaction assumption takes into account the bid-ask spread but not the *market impact* of rebalancing which manifests itself for large portfolios (Webster, 2023). Market impact depends on the size of the rebalacing relative to a measure of market depth, such as average daily traded volume (ADV), volatility of the hedging instrument and possibly other variables such as order flow imbalance. It may be modeled by stating that a trade of size $|\phi_t^i - \phi_{t-\Delta t}^i|$ in instrument H^i gets executed at a price $H_t^i + s_t^i|\phi_t^i - \phi_{t-\Delta t}^i|^\delta$ where $\delta \geq 0$, and s_t^i is a measure, expressed in USD, of market depth for instrument *i*. This leads to a modified expression of profit/loss:

$$\Delta\Pi_t = \sum_{i\in\mathcal{H}} \phi_t^i \Delta H_t^i + r_t \left(\Pi_t - \sum_{i\in\mathcal{H}} \phi_t^i H_t^i - \overbrace{\sum_{i\in\mathcal{H}} |\phi_t^i - \phi_{t-\Delta t}^i|^{1+\delta} s_t^i}^{\text{Rebalancing cost}} \right) \Delta t. \tag{21}$$

For options, impact should be measured in terms of implied volatility. Incorporation of market impact as in (21) leads to an additional term in the optimisation problem:

$$\inf_{A_t\in\mathbb{R},\phi_t\in\mathbb{R}^{d+1}} \frac{1}{N}\sum_{k=1}^{N} \left(\Delta V_t(\omega_k) - A_t - \sum_{i\in\mathcal{H}} \phi_t^i \Delta H_t^i(\omega_k) \right)^2 + \lambda g_0 \sum_{i\in\mathcal{H}} s_t^i \left|\phi_t^i - \phi_{t-\Delta t}^i\right|^{1+\delta}. \tag{22}$$

Empirical evidence supports the assumption of linear impact ($\delta = 1$) across a wide range of trading volumes and frequencies (Cont et al., 2014). In this case (22) leads to a “ridge regression" problem (Hoerl & Kennard, 1981) which has a closed-form solution. However, the values of $\delta > 0$ do not lead to sparsity in (22). Automatic hedging instrument selection can be implemented by first performing a LASSO regression as in (20) at time $t = 0$ in order to determine the hedging instruments, and then using the selected instruments to solve (22). Alternatively, for $\delta > 0$, one could consider an "Elastic Net" regression (Hastie et al., 2009):

$$\begin{aligned} &\inf_{A_t\in\mathbb{R},\phi_t\in\mathbb{R}^{d+1}} \underbrace{\frac{1}{N}\sum_{k=1}^{N} \left(\Delta V_t(\omega_k) - A_t - \sum_{i\in\mathcal{H}} \phi_t^i \Delta H_t^i(\omega_k) \right)^2}_{\text{Hedging error variance}} \\ &+ \underbrace{\lambda g_0 \sum_{i\in\mathcal{H}} s_t^i \left|\phi_t^i - \phi_{t-\Delta t}^i\right|^{1+\delta}}_{\text{Market impact}} + \underbrace{\alpha g_0 \sum_{i\in\mathcal{H}} c_t^i \left|\phi_t^i - \phi_{t-\Delta t}^i\right|}_{\text{Rebalancing cost}}. \end{aligned} \tag{23}$$

4 Example: hedging volatility risk with VolGAN

The above methodology is general, but we will illustrate it with an example: hedging of option portfolios using VolGAN (Vuletić & Cont, 2024), a generative model for arbitrage-free implied volatility surfaces. The model is trained on time series of implied volatility surfaces and underlying prices and is capable of generating realistic scenarios for joint dynamics of the implied volatility surface and the underlying asset. We utilize raw VolGAN outputs, without scenario re-weighting (Cont and Vuletić, 2023).

The portfolio we wish to hedge is a one-month long straddle with strikes $K = m_0 S_0$, for $m_0 \in \{0.75, 0.8, 0.9, 1.1, 1.2, 1.25\}$. The initial portfolio instrument set indexed by $\mathcal{P}$ consists of a *Call*(K, T) and a *Put*(K, T), with $K = m_0 S_0$ and $T = 1/12$.

In addition to data-driven hedging through VolGAN, we also consider delta hedging and delta-vega hedging. The hedging experiment is performed over non-overlapping periods. That is, a long straddle is entered and hedged until expiry, after which the new long straddle is entered and the exercise is repeated. This results in 52 non-overlapping one-month periods. One day is taken to be $\Delta t = 1/252$.

The initial hedging set indexed by $\mathcal{H}_0$ consists of one-month calls and puts whose initial moneyness values are in the set $\{0.9, 0.95, 0.975, 1, 1.025, 1.05, 1.1\}$, where $m < 1$ correspond to puts and $m \geq 1$ to calls. The hedging set indexed by $\mathcal{H}$ consists of all of the options in the initial hedging set that are not in the portfolio we wish to hedge, that is $\mathcal{H} = \mathcal{H}_0 \setminus \mathcal{P}$. The values α used for regularization are obtained through validation on new independent samples from VolGAN, as discussed in the next subsection.

Delta-vega hedging The hedging instruments are the underlying and an option. Denoting by κ_t^i the vega of the option $i \in \mathcal{P}$ in the initial portfolio we wish to hedge (long straddle) at time *t*, the overall vega of the portfolio is

$$\kappa_t^V = \sum_{i \in \mathcal{P}} \psi^i \kappa_t^i.$$

In order to obtain vega-neutrality, it is necessary to include an option H_t^1 with a vega of κ_t^H whose hedge ratio is determined by the ratio of sensitivities:

$$\phi_t^1 = \frac{\kappa_t^V}{\kappa_t^H}.$$

Delta-neutrality is obtained by adjusting the position in the underlying. Denoting by Δ_t^V and Δ_t^H the deltas of the portfolio and the option used for hedging, the hedge ratio ϕ_0 corresponding to the underlying is therefore:

$$\phi_t^0 = \Delta_t^V - \phi_t^1 \Delta_t^H.$$

We opt to hedge with an option initiated at-the-money, i.e. with a strike $K = S_0$. However, the choice of the hedging instrument is not unique. Opting for options with moneyness far away from one, or in general with low vega, may result in unstable hedge ratios ϕ_t^1.

Delta hedging The only hedging instrument is the underlying. The hedge ratio is the portfolio delta:

$$\phi_t^0 = \Delta_t^V.$$

4.1 Data

We use data on SPX options extracted from OptionMetrics for the period Jan 2000-28 Feb 2023. The data from 3rd Jan 2000 to 16th Jun 2018 is used for training the generative model, and the subsequent time period for out-of-sample testing. We construct smooth implied volatility surfaces and bid-ask spread surfaces for calls and puts by applying kernel smoothing to market data (Cont & da Fonseca, 2002; Vuletić & Cont, 2024). More precisely, we use a vega-weighted Nadaraya-Watson kernel estimator with a 2D Gaussian kernel, using the same bandwidth hyperparameters for the implied volatility and for the bid-ask spread surfaces as for VolGAN data pre-processing (Vuletić & Cont, 2024). The grid $(\boldsymbol{m}, \boldsymbol{\tau})$ is the one used for

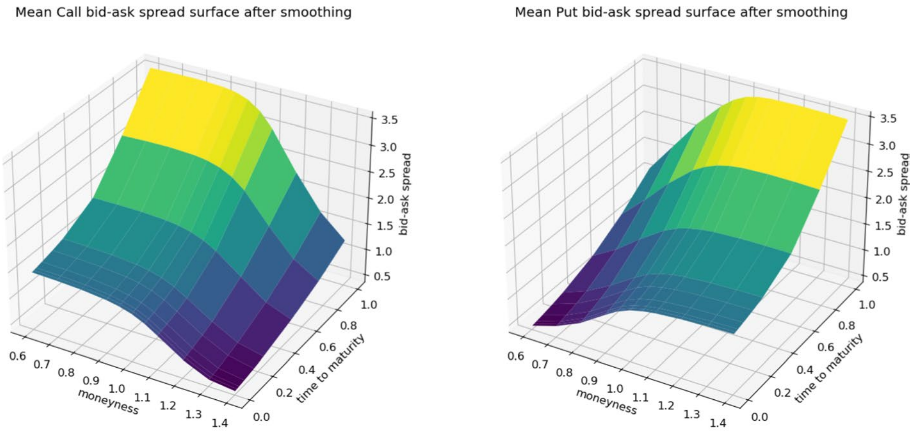

(a) Calls.

(b) Puts.

Fig. 1 Average bid-ask spread for calls and puts

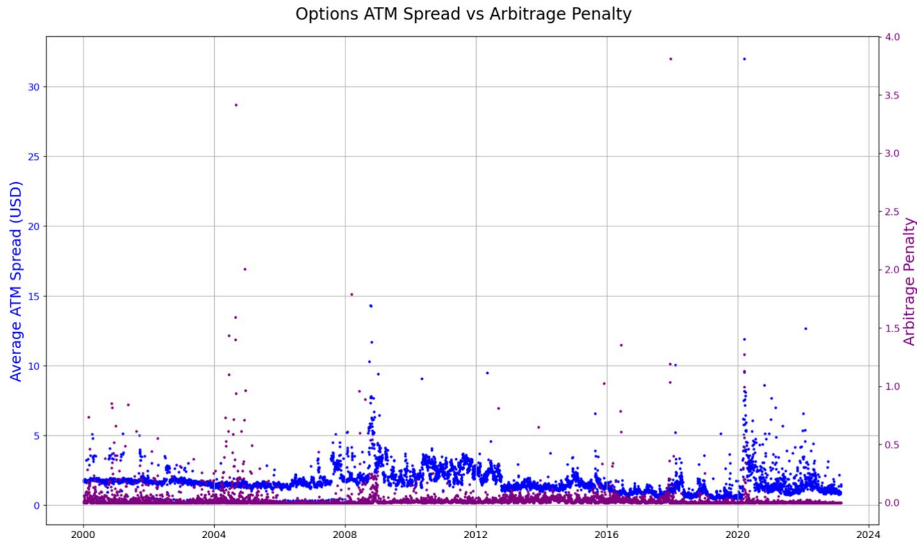

Fig. 2 Average at-the-money bid-ask spread and the arbitrage penalty (Cont and Vuletić, 2023)

VolGAN training, with $0.6 \leq m \leq 1.4$ and times to maturity τ ranging from one day to one year. Interpolation is performed linearly first in moneyness, and then in time to maturity. When extrapolation is necessary, it is linear. The average bid-ask spread surfaces for calls and puts after smoothing are shown in Fig. 1. On average, longer-dated in-the-money options have a higher bid-ask spread than out-of-the-money options with shorter times to expiry. We observe a skew in both the moneyness and time to maturity variables.

We compare the average at-the-money bid-ask spread with the arbitrage penalty (Cont and Vuletić, 2023) in Fig. 2, and note that the instances of non-zero arbitrage penalty coincide with high bid-ask spread values. This observation is consistent with the notion of no-arbitrage, since the arbitrage penalty is calculated using implied volatilities of mid-prices.

4.2 Choosing the regularization parameter

High values of the regularization parameter α would result in low rebalancing due to the high influence of the transaction costs, whereas low values of α would not provide sufficient regularization and would lead to high transaction costs. That is, low values of α might result in overfitting, and high values of α in underfitting. In order to set the appropriate value of α, in each one-month experiment, we perform a search at time $t = 0$ considering the potential α values in $\{0.01, 0.02, \ldots, 0.2\}$, which is a scale similar to that of the transaction-free hedging analysis of Vuletić and Cont (2024). Let $\hat{A}_0(\alpha), (\hat{\phi}_0^i(\alpha))_{i\in\mathcal{H}}$ be the solutions of (20) given α, and let $\{\Delta V_0(\omega_j), (\Delta H_0^i(\omega_j))_{i\in\mathcal{H}}\}_{j=1,\ldots,1000}$ i.i.d. samples from VolGAN for each starting time $t = 0$, used for regression fitting. We choose α minimizing the Akaike Information Criterion (AIC) (Akaike, 1974) for the day on which a new position is entered ($t = 0$), calculated on new $M = 100$ i.i.d. samples from VolGAN $\{\Delta V_0(\omega_j), (\Delta H_0^i(\omega_j))_{i\in\mathcal{H}}\}_{j=N+1,\ldots,N+M}$, independent of the $N = 1000$ simulations on which the regression coefficients were estimated. The (relative) AIC value in the case of linear regression is:

$$AIC(\alpha) = M \log\left(\frac{RSS(\alpha)}{M}\right) + 2\left(1 + \sum_{i\in\mathcal{H}} 1_{\hat{\phi}_0^i(\alpha)\neq 0}\right), \tag{24}$$

where $1 + \sum_{i\in\mathcal{H}} 1_{\hat{\phi}_0^i(\alpha)\neq 0}$ is the number of parameters of the regression fit ($1_{\hat{\phi}_0^i(\alpha)\neq 0}$ is zero if $\hat{\phi}_0^i(\alpha) = 0$ and one otherwise), and the residual sum of squares $RSS(\alpha)$ is estimated on the new M samples from VolGAN:

$$RSS(\alpha) = \sum_{j=N+1}^{N+M} \left(\Delta V_t(\omega_j) - \hat{A}_0(\alpha) - \sum_{i\in\mathcal{H}} \hat{\phi}_0^i(\alpha)\Delta H_t^i(\omega_j)\right)^2.$$

That is, at each starting time $t = 0$ we choose α by minimizing an AIC criterion, where the set of possible α is $\mathcal{A} = \{0.01, 0.02, \ldots, 0.2\}$. Estimating the AIC on independent VolGAN simulations prevents overfitting and provides more accurate estimates of regression fits. Since only VolGAN simulations are used at time $t = 0$, the choice of α is not anticipative, i.e. there is no look-ahead bias. Figure 3 shows that during periods of market turbulence, higher regularization is preferred. The most common choice is $\alpha = 0.01$. Furthermore, we note that the values of m_0 closer to one result in more frequent recalibration of the regularization parameter α. The fact that the highest value of the regularization parameter is selected during periods of market turbulence suggests that a higher value of α would be more suitable than in those instances. However, a reduced search grid for α of $\{0.01, 0.05, 0.1, 0.5, 1\}$ results in a similar performance, as demonstrated in the Appendix. We also note that fixing $\alpha = 0.25$ does not significantly alter the performance either.

The value of V_t is the same as the gross value of the position since V_t is made up of a long call and a long put. Figure 4 shows that V_t is higher for values of m_0 further away from one.

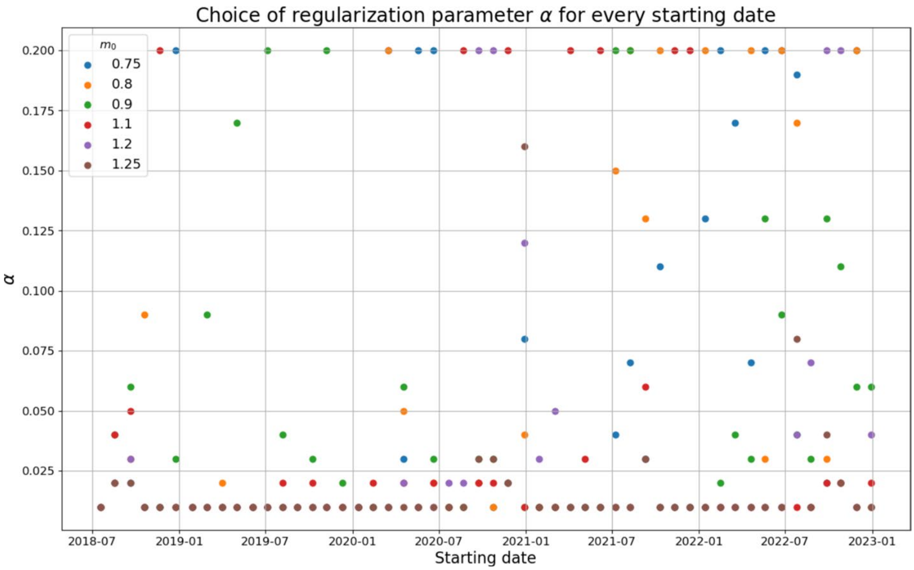

Fig. 3 Choice of regularization parameter α minimizing the AIC (24) for each starting date $t = 0$ and for different values of m_0

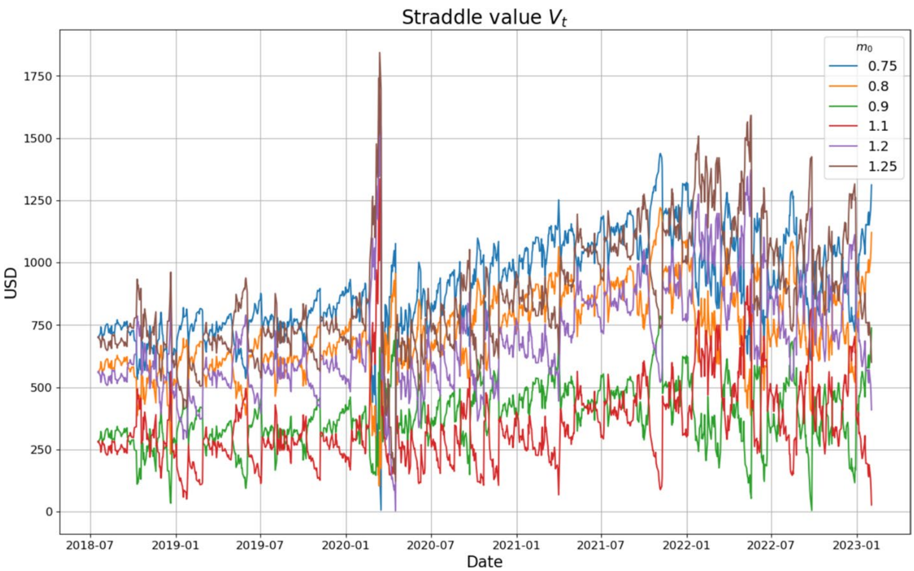

Fig. 4 Value V_t of the long straddle position for different values of m_0

That is, the same value of α would result in a lower value of the regularization term αg_0 in (20) for $m_0 \in \{0.9, 1.1\}$ compared to $m_0 \in \{0.75, 0.8, 1.2, 1.25\}$.

4.3 Number of hedging instruments

Since incorporating transaction cost via an ℓ_1-penalty (20) leads to sparsity, we investigate the number of hedging instruments selected throughout the testing period. Unsurprisingly, the underlying is always selected. A breakdown of the number of times that each possible hedging instrument count was used for each starting value m_0 is available in Table 1. Data-driven hedging with VolGAN predominantly selected only the underlying as the hedging instruments for the straddles with m_0 further away from one. When $m_0 \in \{0.9, 1.1\}$ (the values closest to 1), options were selected alongside the underlying most of the time.

Figure 5 shows the number of hedging instruments used for different starting values m_0. We note a spike in the number of selected hedging instruments at the start of 2019, during the start of the Covid-19 pandemic, and in the second half of 2022.

4.4 Tracking error statistics

We first compare the performance during the entire test set, and then without the initial Covid-19 shock. To exclude the Covid-19 pandemic effect, we remove 5 one-month periods starting from the position entered on the 13th Feb 2020. This results in exclusion of the dates 13th Feb 2020-21st Jul 2020.

Tracking error Z_t statistics under consideration are the mean, median, standard deviation, and Value-at-Risk, defined as

$$VaR_q(Z_t) = -F_{Z_t}^{-1}(q),$$

for $q = 5\%, 2.5\%$, and 1%, where $F_{Z_t}^{-1}$ is the quantile function of Z_t.

The tracking error statistics (for all values of m_0 pooled together) for delta hedging, delta-vega hedging, and VolGAN are given in Table 2. Outside of the Covid-19 pandemic, VolGAN results in the lowest standard deviation and 1% VaR. The corresponding tracking error distributions over the entire test set are shown in Fig. 6. . We also note a significant reduction in standard deviation, and Value-at-risk metrics compared to the unhedged position.

Table 1 Frequency of the number of hedging instruments selected for different m_0 values

	Number of instruments							
m_0	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
0.75	1058	3	11	20	0	0	0	0
0.80	788	156	46	88	3	11	0	0
0.90	193	279	168	277	93	67	15	0
1.10	565	40	72	110	219	69	17	0
1.20	1077	2	7	0	0	1	3	2
1.25	1087	0	0	2	3	0	0	0

The total number of days is $21 \times 52 = 1092$, with an additional day for unwinding. Position expiry dates are not included since the only adjustments in the hedging portfolio happen on days $0, \ldots, 20$

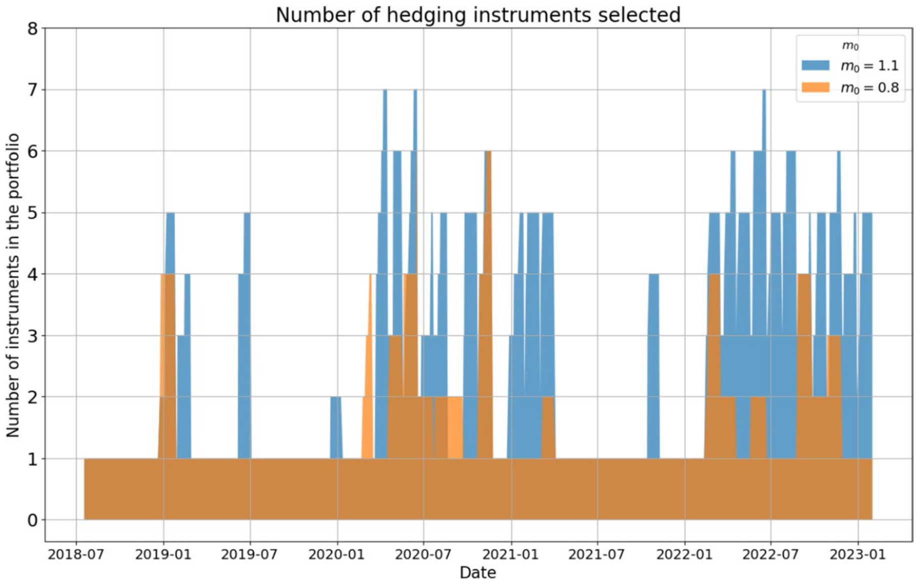

Fig. 5 Number of hedging instruments selected for different values of m_0

Table 2 Tracking error statistics in USD for all values of m_0 over the entire test period

Covid-19 period	Method	Statistics					
		Mean	Median	Std	5% VaR	2.5% VaR	1% VaR
Included	Unhedged position	0.16	0.29	60.58	78.95	102.07	155.29
	Delta hedging	−1.23	−0.44	32.70	19.49	36.33	58.63
	Delta-vega hedging	0.98	0.00	29.70	10.90	19.70	43.72
	Data-driven hedging	0.55	−0.16	32.98	12.79	23.42	50.79
Excluded	Delta hedging	−1.46	−0.36	8.40	13.22	22.84	36.66
	Delta-vega hedging	−0.37	0.00	9.34	9.58	16.67	34.81
	Data-driven hedging	−1.05	−0.18	8.15	10.55	17.32	33.85

Comparing VolGAN hedge with delta hedge in Fig. 7, and with delta-vega hedge in Fig. 8, with the Covid-19 pandemic data excluded for better visualization, we observe that all methods result in similar performance in the bulk of the experiments. In most cases where delta hedging results in a positive PnL, the same holds for VolGAN, and the points in the first quadrant appear to scatter symmetrically around $x = y$. However, there are many instances in which VolGAN results in a positive PnL, while delta hedging does not. In the scenarios in which both approaches result in a negative PnL, the loss is usually more severe for delta-hedging. There is a bit more symmetry when comparing VolGAN with delta-vega hedging in Fig. 8. This observation aligns with the statistics in Table 2, where outside of the Covid-19 pandemic, data-driven hedging leads to tracking errors statistics closer to delta-vega hedging than to delta hedging. As m_0 moves away from 1, differences in hedging approaches are less prominent: the points are tightly concentrated around $x = y$.

A detailed breakdown of performance for each individual value of m_0 is available in the Appendix.

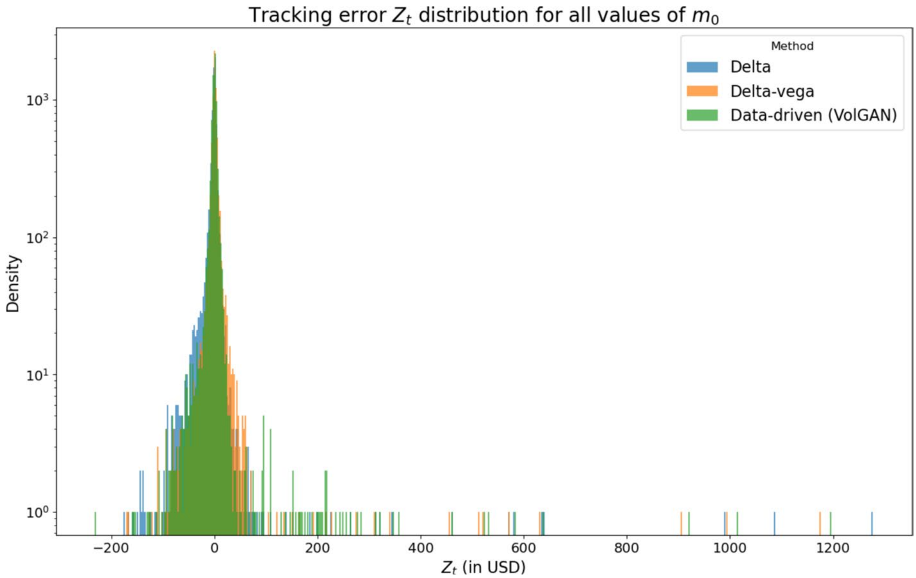

Fig. 6 Distribution of tracking error Z_t for all values of m_0 pooled together, over the entire test set

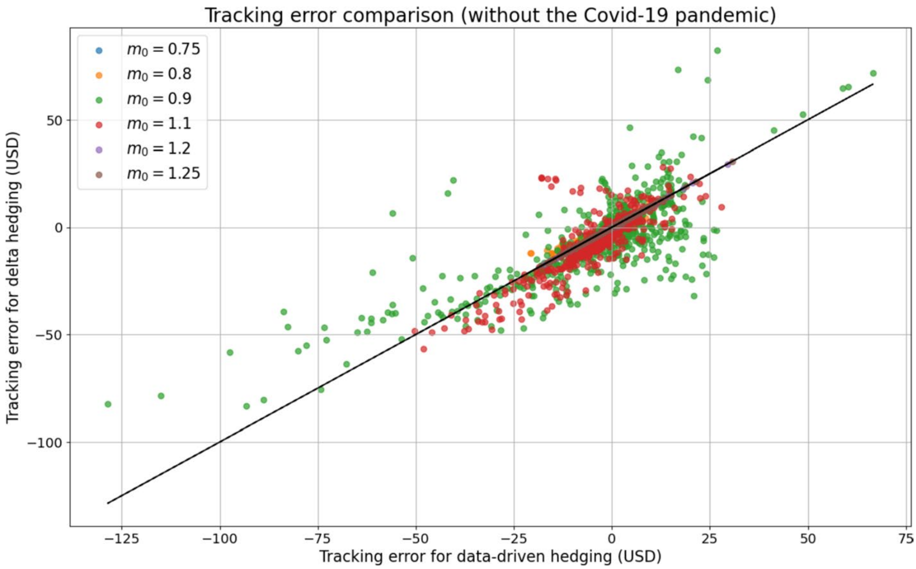

Fig. 7 Tracking error: delta hedge versus VolGAN hedge. Solid black line: $x = y$. Covid-19 data excluded

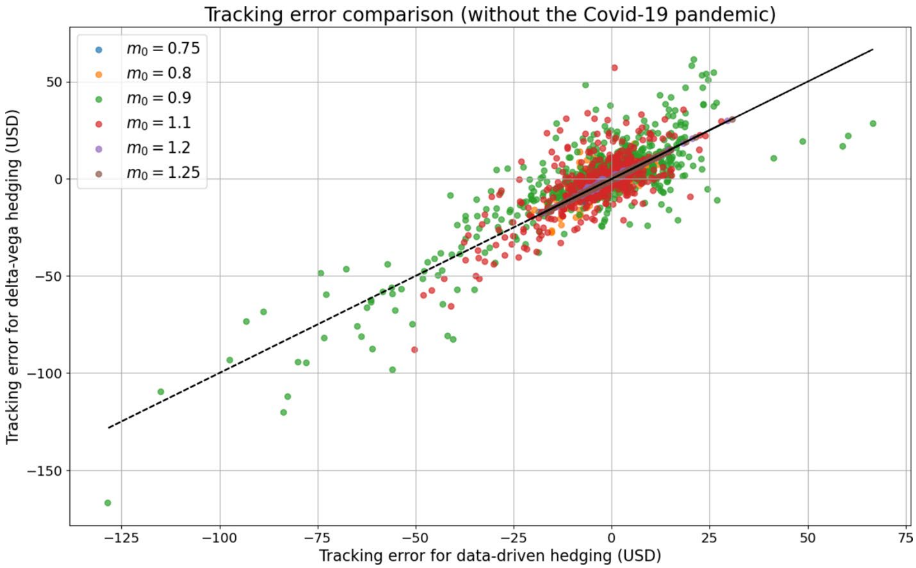

Fig. 8 Tracking error: delta-vega hedge versus VolGAN hedge. Solid black line: $x = y$. Covid-19 data excluded

Figures 9 and 10 compare the performance of the data-driven hedging strategy with delta hedging and delta-vega hedging, with the Covid-19 period included.

4.5 Sensitivity analysis of the hedging strategy

In order to better understand the nature of the hedging strategy resulting from our method (20), we compare the vega and delta of the straddle V_t with the total vega and delta of the hedged position Z_t. As shown in Table 3, the delta exposure of the position is almost completely hedged by data-driven hedging. On the other hand, Table 4 shows that although the total vega of the hedged position Z_t is reduced compared to that of the unhedged position V_t, it is not zero. These observations show that our data-driven hedging approach (20) is not equivalent to delta-vega hedging. So, unlike what has been suggested in the recent literature (Ludkovski, 2023), there is more to data-driven hedging than just "learning the Greeks".

5 Summary

We consider hedging as a conditional risk minimization problem and use a forward-looking data-driven generative model in order to choose hedging instruments and compute hedge ratios for a given portfolio. We simulate joint scenarios for the portfolio and from the hedging instruments, conditional on the current market conditions, via a generative model. In the case of variance minimization, we perform regression of the simulated portfolio increments onto the simulated increments of the hedging instruments. Incorporating transaction costs leads to an ℓ_1-regularized regression problem, which induces sparsity and leads to the minimal subset of rebalancing transactions. We illustrated how our approach can be applied to hedge volatility

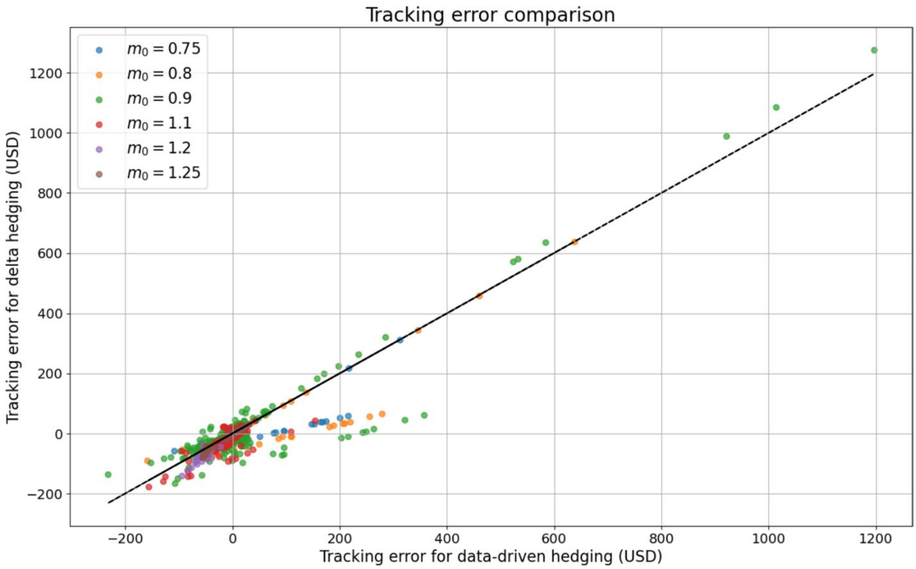

Fig. 9 Tracking error: delta hedge versus VolGAN hedge. Solid black line: $x = y$

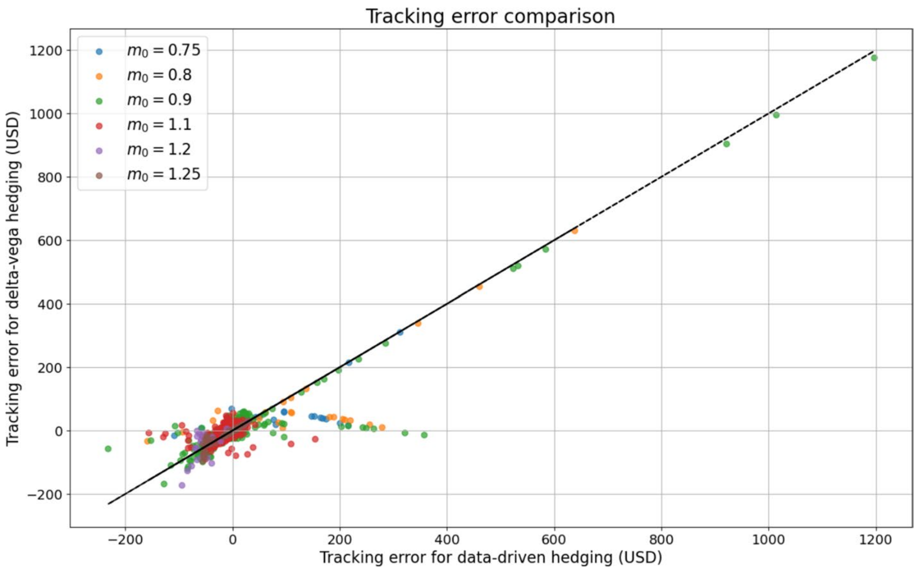

Fig. 10 Tracking error: delta-vega hedge versus VolGAN hedge. Solid black line: $x = y$

risk through hedging long straddle positions using VolGAN (Vuletić & Cont, 2024). However, our approach is not restricted to the use of a particular generative model or architecture, and can be combined with other generative models and applied to other asset classes.

Our methodology combines the benefits of regression-based hedging, such as linearity, interpretability, and incorporation of covariance between different instruments, with the

Table 3 Delta statistics

	Mean		Median		95% Quantile		5% Quantile		Std Dev	
m_0	Z_t	V_t	Z_t	V_t	Z_t	V_t	Z_t	V_t	Z_t	V_t
0.75	0.003	0.994	0.000	1.000	0.005	1.000	−0.001	0.986	0.023	0.067
0.8	0.003	0.986	0.000	1.000	0.009	1.000	−0.016	0.958	0.040	0.097
0.9	−0.011	0.905	−0.011	0.979	0.051	1.000	−0.111	0.618	0.100	0.214
1.1	−0.006	−0.940	0.000	−0.997	0.035	−0.760	−0.049	−1.000	0.066	0.174
1.2	−0.000	−0.996	−0.000	−1.000	0.000	−0.998	−0.001	−1.000	0.016	0.039
1.25	0.000	−0.999	−0.000	−1.000	0.000	−1.000	−0.001	−1.000	0.006	0.009

Table 4 Vega statistics

	Mean		Median		95% Quantile		Std Dev	
m_0	Z_t	V_t	Z_t	V_t	Z_t	V_t	Z_t	V_t
0.75	5.828	6.889	0.078	0.102	32.904	35.288	21.269	22.073
0.8	9.546	17.136	0.170	0.805	66.971	91.581	36.314	40.845
0.9	0.096	111.500	−4.397	37.221	179.155	452.507	89.315	150.315
1.1	26.165	66.529	2.807	7.817	170.733	326.120	68.777	108.969
1.2	3.302	3.637	0.000	0.000	5.399	5.637	23.437	25.145
1.25	1.076	1.079	0.000	0.000	0.239	0.239	10.013	10.022

flexibility of generative models for scenario simulation. We have applied this approach to dynamic hedging, but it may also be deployed for the design of trading strategies, by replacing for instance the variance minimization step with a mean-variance tradeoff involving a target return.

Appendix

Coordinate descent for data-driven hedging with transaction costs

We wish to solve

$$\inf_{A_t\in\mathbb{R},\phi_t\in\mathbb{R}^{d+1}} \underbrace{\frac{1}{N}\sum_{k=1}^{N}\left(\Delta V_t(\omega_k) - A_t - \sum_{i\in\mathcal{H}}\phi_t^i \Delta H_t^i(\omega_k)\right)^2}_{\text{Quadratic hedging error}} + \alpha g_0 \underbrace{\sum_{i\in\mathcal{H}} c_t^i \left|\phi_t^i - \phi_{t-\Delta t}^i\right|}_{\text{Rebalancing cost}}, \tag{25}$$

Equivalently, by decomposing $\phi_t = \phi_{t-\Delta t} + \Delta\phi_{t-\Delta t}$, the problem can be expressed as

$$\inf_{A_t\in\mathbb{R},\Delta\phi_{t-\Delta t}\in\mathbb{R}^{d+1}} \frac{1}{N}\sum_{k=1}^{N}\left(\begin{array}{l}\left(\Delta V_t(\omega_k) - \sum_{i\in\mathcal{H}}\phi_{t-\Delta t}^i \Delta H_t^i(\omega_k)\right) - \\ A_t - \sum_{i\in\mathcal{H}}\Delta\phi_{t-\Delta t}^i \Delta H_t^i(\omega_k)\end{array}\right)^2 + \alpha g_0 \sum_{i\in\mathcal{H}} c_t^i \left|\Delta\phi_{t-\Delta t}^i\right|. \tag{26}$$

With $J = |\mathcal{H}|$, this is a generalized LASSO regression (Tibshirani, 1996) of the form

$$\inf_{\beta \in \mathbb{R}^{J+1}} \frac{1}{N} \sum_{k=1}^{N} \left(y_k - \beta_0 - \sum_{j=1}^{J} \beta_i x_k^i \right)^2 + \sum_{j=1}^{J} \alpha_j \left| \beta_j \right|. \tag{27}$$

To solve it, we apply coordinate descent with soft thresholding, as outlined in Chapter 3.8.6 of Hastie et al. (2009). Denote the loss by

$$J(\beta_0, \beta) = \frac{1}{N} ||\boldsymbol{y} - \beta_0 \mathbf{1} - \boldsymbol{X}\beta||_2^2 + \sum_{j=1}^{J} \alpha_j |\beta_j|,$$

where $\boldsymbol{y} = (y_1, \ldots, y_n)^T$, and $\boldsymbol{X}_{i,j} = x_j^i$. There is no penalty on the intercept, hence the optimal β_0^* will be the solution to $\dfrac{\partial J}{\partial \beta_0} = 0$:

$$\beta_0 = \frac{1}{N} \sum_{i=1}^{n} \left(y_i - \boldsymbol{x}_i^T \beta \right). \tag{28}$$

Since the regularization term in (27) is co-ordinate separable, we can focus on each $j > 0$ separately. Define the partial residual as

$$r^{(-j)} = \boldsymbol{y} - \beta_0 - \sum_{k \neq j} \boldsymbol{X}_k \beta_k, \tag{29}$$

where $\boldsymbol{X}_k$ is the k-th column of the $N \times J$ design matrix $\boldsymbol{X}$. We are now interested in solving the sub-problem for β_j:

$$\min_{\beta_j} \frac{1}{N} ||\boldsymbol{X}_j \beta_j - r^{(-j)}||_2^2 + \alpha_j |\beta_j|. \tag{30}$$

A subgradient at zero is

$$\frac{2}{N} \boldsymbol{X}_j^T \left(\boldsymbol{X}_j \beta_j - r^{(-j)} \right) + \alpha_j \mathrm{sign}(\beta_j). \tag{31}$$

Defining the soft-thresholding map $\mathcal{S}$ as

$$\mathcal{S}(z, \gamma) = \mathrm{sign}(z) \left(|z| - \gamma \right)_+,$$

we derive a closed-form solution for the β_j update:

$$\tilde{\beta}_j = \frac{\mathcal{S}(X_j^T r^{(-j)}, \frac{N\alpha_j}{2})}{\boldsymbol{X}_j^T \boldsymbol{X}_j}. \tag{32}$$

Hence, we may iterate over each $j = 0, \ldots, J$, until convergence, or until the maximum number of iterations, as outlined in Algorithm 2.

At each time step *t*, we use the Algorithm 2 by applying the steps below.

1. Construct y_i as the adjusted portfolio changes: $y_i = \Delta V_t(\omega_i) - \sum_{j\in\mathcal{H}} \phi^j_{t-\Delta t}\Delta H(\omega_i)$.
2. Define $X_{i,j}$ as the instrument price changes: $X_{i,j} = \Delta H^j_t(\omega_i)$.
3. Set α_j proportional to transaction costs: $\alpha_j = \alpha g_0 c^j_t$.
4. Find the optimal β values via the Algorithm 2.
5. Set $\phi_t = \phi_{t-\Delta t} + \beta$ and $A_t = \beta_0$.

Algorithm 2 Coordinate Descent for Generalized Lasso

Input: Design matrix $\boldsymbol{X} \in \mathbb{R}^{N\times J}$ (without intercept column)
Response vector $\boldsymbol{y} \in \mathbb{R}^N$
Regularization parameters $\alpha_j \geq 0$ for $j = 1, \ldots, J$
Maximum iterations max_iter, default 1000
Tolerance tol, default 10^{-6}
Output: Coefficients $\beta_0, \beta_1, \ldots, \beta_J$

Initialization:
$\beta \leftarrow \mathbf{0} \in \mathbb{R}^J$
$\beta_0 \leftarrow 0$
$iter \leftarrow 0$
$converged \leftarrow False$

Precompute:
$\boldsymbol{X}_j^\top \boldsymbol{X}_j \leftarrow \mathrm{diag}(\boldsymbol{X}^\top \boldsymbol{X})$ for $j = 1, \ldots, J$

while $iter < max\_iter$ ***and not*** $converged$ **do**
  $\beta_{old} \leftarrow [\beta_0, \beta]$
  Update intercept:
  $r \leftarrow \boldsymbol{y} - \beta_0 - \boldsymbol{X}\beta$
  $\beta_0 \leftarrow \frac{1}{N}\sum_{k=1}^N r_k$
  Update coefficients:
  for $j = 1$ ***to*** J **do**
    $r \leftarrow \boldsymbol{y} - \beta_0 - \boldsymbol{X}\beta + \beta_j \boldsymbol{X}_j$
    $X_j^T r \leftarrow \boldsymbol{X}_j^\top r$
    Soft-thresholding:
    if $X_j^T r > \frac{N\alpha_j}{2}$ **then**
      $\beta_j \leftarrow \frac{X_j^T r - \frac{N\alpha_j}{2}}{\boldsymbol{X}_j^\top \boldsymbol{X}_j}$
    else if $X_j^T r < -\frac{N\alpha_j}{2}$ **then**
      $\beta_j \leftarrow \frac{X_j^T r + \frac{N\alpha_j}{2}}{\boldsymbol{X}_j^\top \boldsymbol{X}_j}$
    else
      $\beta_j \leftarrow 0$
  Check convergence:
  if $\|[\beta_0, \beta] - \beta_{old}\|_2 < tol$ **then**
    $converged \leftarrow True$
  $iter \leftarrow iter + 1$

return $\beta_0, \beta_1, \ldots, \beta_J$

Hedging with VolGAN: robustness checks

Performance for different values of m_0

Table 5 contains tracking error statistics over the entire dataset, for each value of m_0. Most of the time, for $m_0 < 1$, data-driven hedging results in all tracking error statistics of interest between those corresponding to delta hedging and delta-vega hedging. The higher standard deviation in $m_0 = 0.75$ is due to the options selected as hedging instruments at the beginning of the Covid-19 pandemic, which, as illustrated in Fig. 11, results in a more prominent volatility of the tracking error. As more options are selected for hedging, the performance of data-driven hedging is closer to that of delta-vega hedging than to pure delta hedging.

When $m_0 = 1.1$, hedging with options brings about a very clear reduction in all risk measures being considered. In this instance, delta hedging has the worst performance, delta-vega hedging has the lowest tracking error variance, but data-driven hedging results in the lowest Value-at-Risk. For $m_0 \in \{1.2, 1.25\}$, data-driven hedging produces tracking error of the lowest Variance, with Value-at-Risk statistics that are usually in between those corresponding to the Black-Scholes benchmarks.

Table 5 Performance metrics for different m_0 values over the entire test set (rounded to two decimal places)

m_0	Method	Statistics					
		Mean	Median	Std	5% VaR	2.5% VaR	1% VaR
0.75	Delta hedging	0.17	−0.06	2.81	6.10	7.32	12.37
	Delta-vega hedging	1.38	0.11	12.89	5.07	6.54	9.09
	Data-driven hedging	1.48	0.00	19.23	5.66	7.01	9.04
0.80	Delta hedging	0.39	−0.32	27.20	8.69	12.23	18.39
	Delta-vega hedging	2.43	0.23	26.65	6.31	8.82	17.05
	Data-driven hedging	2.49	−0.16	32.66	7.78	10.31	15.30
0.90	Delta hedging	−0.35	−1.95	70.56	40.24	53.71	81.30
	Delta-vega hedging	4.90	0.55	63.76	26.88	56.35	81.00
	Data-driven hedging	4.02	−0.32	68.66	34.48	55.92	80.78
1.10	Delta hedging	−5.17	−0.98	17.00	33.66	46.66	76.69
	Delta-vega hedging	−1.20	−0.13	11.88	19.59	31.40	51.36
	Data-driven hedging	−2.82	−0.74	13.79	19.17	31.41	47.58
1.20	Delta hedging	−1.41	0.00	11.90	8.59	13.36	73.09
	Delta-vega hedging	−0.87	0.00	10.49	8.07	12.79	31.22
	Data-driven hedging	−1.04	−0.01	9.26	8.45	13.39	56.66
1.25	Delta hedging	−1.02	−0.001	9.33	8.61	13.88	54.51
	Delta-vega hedging	−0.79	0.00	8.47	8.42	13.88	37.12
	Data-driven hedging	−0.82	−0.01	8.04	8.60	13.87	48.52

The initial positions $Z_0 = 0$ are included in the statistics

Table 6 Proportion of observations on which the tracking portfolio Π_t was within 1% of the initial straddle portfolio value V_t for different methods and m_0 values

Method/m_0	0.75	0.80	0.90	1.10	1.20	1.25
Delta-hedging	95.98%	88.11%	43.62%	50.61%	84.53%	89.34%
Delta-vega hedging	94.49%	88.02%	49.04%	51.57%	85.05%	89.42%
VolGAN-based hedging	95.37%	89.25%	51.92%	54.63%	84.53%	89.34%

Table 7 Performance metrics for different m_0 values without the start of the Covid-19 pandemic (rounded to two decimal places)

m_0	Method	Statistics					
		Mean	Median	Std	5% VaR	2.5% VaR	1% VaR
0.75	Delta hedging	−0.23	0.00	2.81	4.99	6.27	7.15
	Delta-vega hedging	0.09	0.04	2.92	4.89	6.47	8.57
	Data-driven hedging	−0.29	0.00	2.87	5.24	6.28	7.56
0.80	Delta hedging	−0.73	−0.17	3.45	6.60	8.01	11.09
	Delta-vega hedging	0.002	0.11	4.34	6.19	8.58	17.35
	Data-driven hedging	−0.71	−0.19	3.83	7.01	9.13	13.22
0.90	Delta hedging	−4.35	−1.77	15.58	34.46	42.10	51.00
	Delta-vega hedging	−0.88	0.40	18.52	27.61	56.87	81.74
	Data-driven hedging	−2.87	−0.39	16.30	32.88	51.28	73.20
1.10	Delta hedging	−3.15	−0.82	10.09	24.34	32.93	43.12
	Delta-vega hedging	−1.19	−0.17	10.47	17.59	27.81	40.52
	Data-driven hedging	−2.13	−0.72	7.84	16.40	24.08	33.99
1.20	Delta hedging	−0.17	−0.02	4.49	7.17	9.48	12.79
	Delta-vega hedging	−0.12	−0.003	4.50	7.22	9.38	12.80
	Data-driven hedging	−0.17	−0.07	4.49	7.14	9.46	12.78
1.25	Delta hedging	−0.14	−0.06	4.67	7.42	9.77	13.34
	Delta-vega hedging	−0.14	−0.05	4.67	7.42	9.78	13.36
	Data-driven hedging	−0.14	−0.07	4.67	7.41	9.78	13.32

The initial positions $Z_0 = 0$ are included in the statistics

The tracking errors Z_t that come from different hedging approaches as functions of time are shown in Fig. 11, for various values of m_0. For better visualization, Fig. 12 displays the tracking values on a scale that is linear in the interval $[-50, 50]$ and logarithmic away from it. All three methods track the initial portfolio satisfactorily well, aside from the Covid-19 pandemic. The three methods are difficult to differentiate from each other during periods of calm. Table 6 highlights that the absolute value of the tracking error is usually less than 1% of the portfolio value, that is, the tracking portfolio and the portfolio we wish to hedge are usually within 1% of each other. Furthermore, Table 6 indicates that all methods perform worse for values of m_0 closer to one. The most volatility is visible during March-June 2020, and during 2022. The corresponding tracking error distributions are available in Fig. 13. The bulk of the tracking error distribution is usually the tightest for data-driven hedging.

Given that we considered all 52 one-month intervals jointly, and that the start of the Covid-19 pandemic resulted in high tracking error variance, as well as that the tracking error was mainly positive during this period (as evidenced in Fig. 11), it is important to repeat the analysis with 13th Feb 2020-21st Jul 2020 excluded. Once again, when $m_0 < 0$ data-driven hedging results in tracking error statistics in between of that corresponding to the Black-Scholes benchmark. As fewer options are selected for hedging, data-driven hedging results in performance more similar to that of delta hedging. All three methods obtain very similar performance to each other for values of $m_0 \in \{0.75, 1.2, 1.25\}$, which have the lowest total vega exposure (Table 4). Interestingly, all three examples result in the underlying being selected by data-driven hedging during the analysis period (Fig. 5). The most significant difference is visible in the $m_0 = 1.1$ example, when data-driven hedging outperforms the benchmarks in both variance and Value-at-Risk statistics. Outside of the Covid-19 pan-

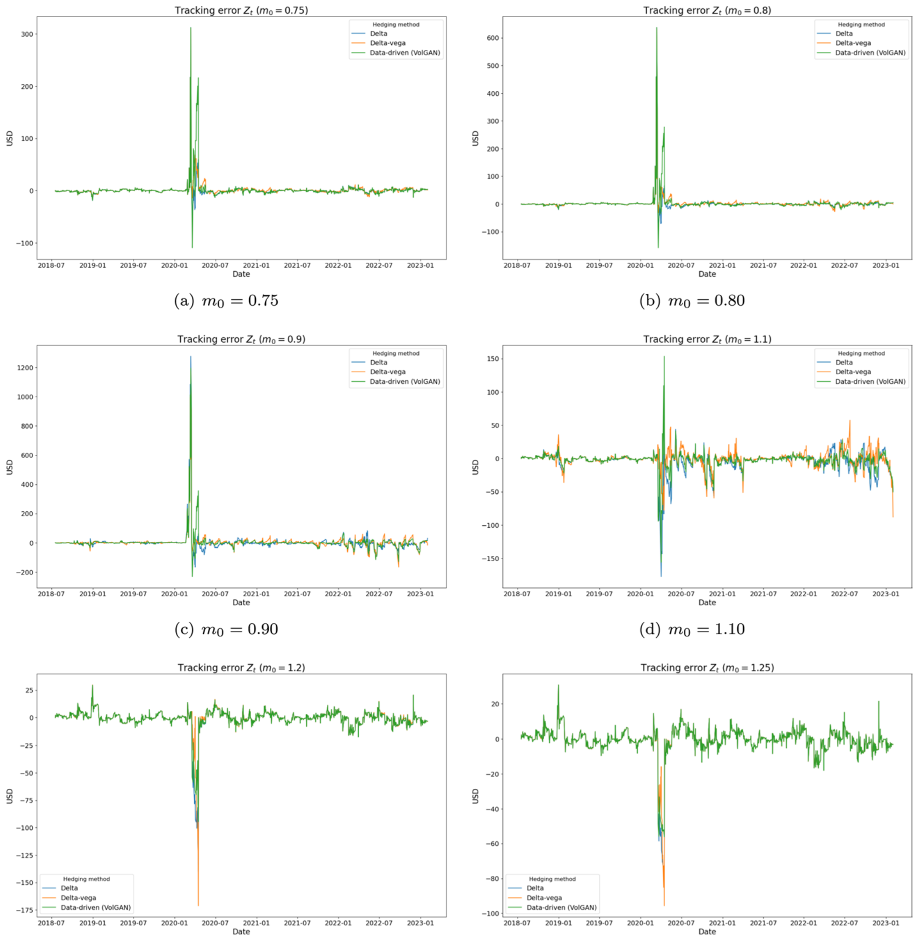

(a) $m_0 = 0.75$ (b) $m_0 = 0.80$ (c) $m_0 = 0.90$ (d) $m_0 = 1.10$ (e) $m_0 = 1.20$ (f) $m_0 = 1.25$

Fig. 11 Tracking error Z_t as a function of time for different values of m_0. Once an initial one-month straddle is expired, a new one is entered

demic, data-driven hedging always results in lower variance and 1% Value-at-Risk than delta-vega hedging, even when more options are selected for hedging, resulting in higher rebalancing costs (Table 7).

Figure 11 indicates that the majority of the tracking error variance after 2020 is due to the behavior in 2022. During this time VolGAN combined with (20) results in much more stable hedging performance compared to the Black-Scholes benchmarks.

Robustness with respect to regularization parameter

We investigate how the results of the data-driven hedging change with the change of the *search grid* $\mathcal{A}$ for the regularization parameter α. Furthermore, we compare the performance for fixed α on three different scales. From Fig. 14, we conclude that for $\alpha = 0.001$

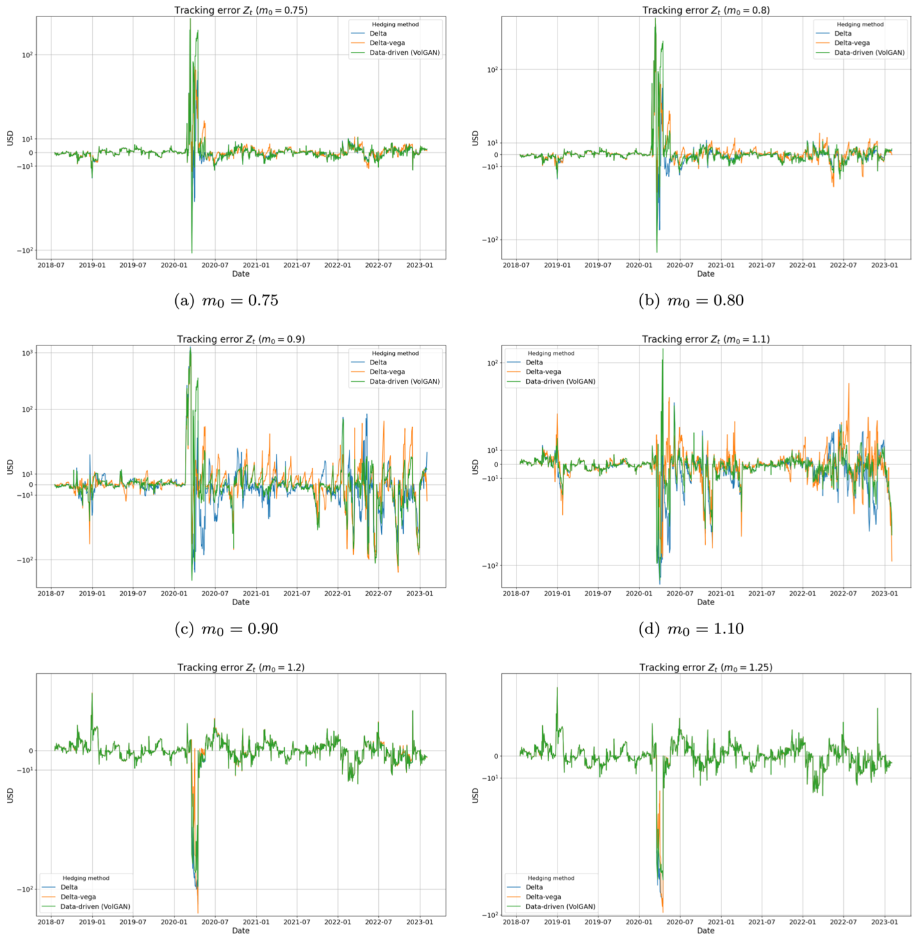

(a) $m_0 = 0.75$

(b) $m_0 = 0.80$

(c) $m_0 = 0.90$

(d) $m_0 = 1.10$

(e) $m_0 = 1.20$

(f) $m_0 = 1.25$

Fig. 12 Tracking error Z_t as a function of time for different values of m_0 on a scale which is linear in $[-50, 50]$, and logarithmic away from this interval. Once an initial one-month straddle is expired, a new one is entered

overfitting occurs, whereas $\alpha = 5$ results in underfitting. However, $\alpha = 0.25$ leads to a similar tracking error to that corresponding to a dynamic choice of α via the AIC criterion.

Table 8 contains the information from Table 2, in addition to the results for data-driven hedging with potential values of the regularization parameter in $\mathcal{A}_1 = \{0.01, 0.05, 0.1, 0.5, 1\}$. The tracking error statistics are very similar for both the search grids $\mathcal{A}_1$ and $\mathcal{A}_2 = \{0.01, 0.02, \ldots, 0.2\}$, indicating robustness to the regularization parameter. However, it is important to note that much lower or higher values of α may result in overfitting and underfitting, as demonstrated with $\alpha = 0.001$ and $\alpha = 5$.

Acknowledgements We thank Peter Bank, Lorenzo Bergomi, Nathan De Carvalho, Fang Rui Lim, Christoph Reisinger, Eric Schaanning, Justin Sirignano, Gabriel Tucci, and seminar participants at QuantMinds

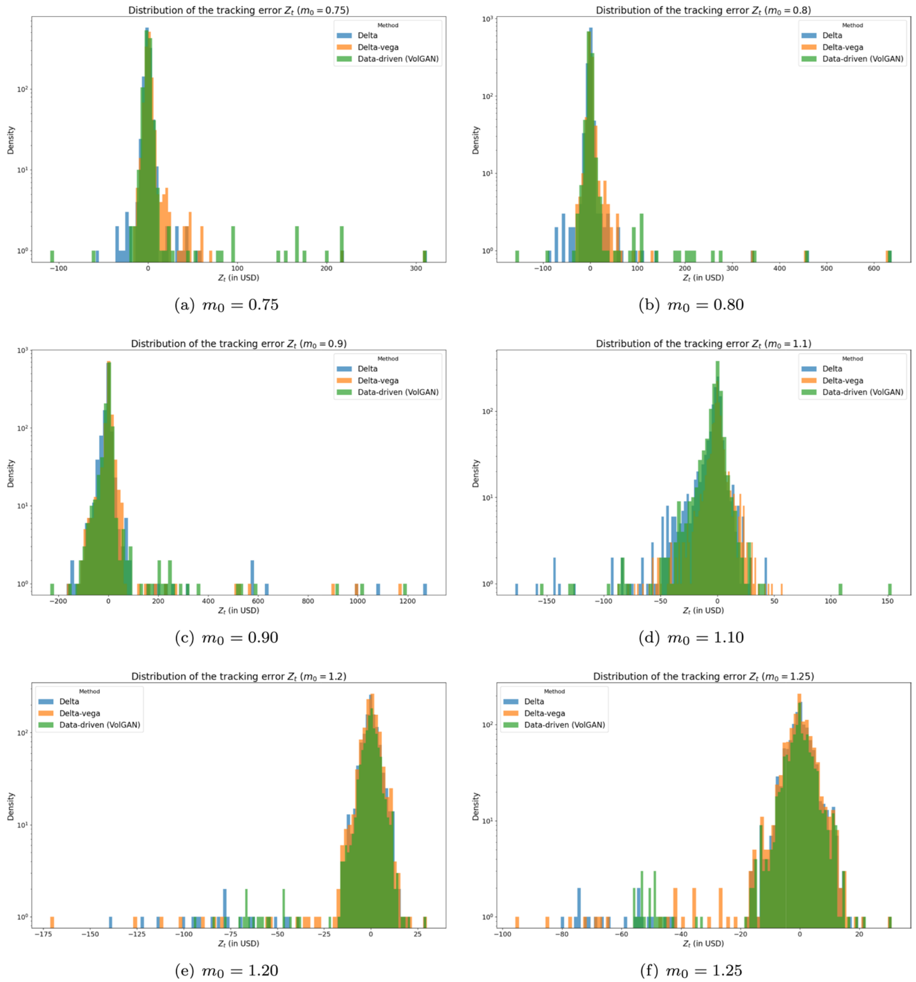

(a) $m_0 = 0.75$ (b) $m_0 = 0.80$

(c) $m_0 = 0.90$ (d) $m_0 = 1.10$

(e) $m_0 = 1.20$ (f) $m_0 = 1.25$

Fig. 13 Histogram of the tracking error Z_t for different values of m_0. Once an initial one-month straddle is expired, a new one is entered

Table 8 Performance metrics for all values of m_0 over the entire test set (rounded to two decimal places)

Covid-19 period	Method	Statistics					
		Mean	Median	Std	5% VaR	2.5% VaR	1% VaR
Included	Delta hedging	-1.23	-0.44	32.70	19.49	36.33	58.63
	Delta-vega hedging	0.98	0.00	29.70	10.90	19.70	43.72
	Data-driven hedging ($\mathcal{A}_1$)	0.58	-0.13	33.01	13.01	24.44	50.45
	Data-driven hedging ($\mathcal{A}_2$)	0.55	-0.16	32.98	12.79	23.42	50.79
Excluded	Delta hedging	-1.46	-0.36	8.40	13.22	22.84	36.66
	Delta-vega hedging	-0.37	0.00	9.34	9.58	16.67	34.81
	Data-driven hedging ($\mathcal{A}_1$)	-1.02	-0.18	8.16	10.80	18.08	35.35
	Data-driven hedging ($\mathcal{A}_2$)	-1.05	-0.18	8.15	10.55	17.32	33.85

The initial positions $Z_0 = 0$ are included in the statistics. Grid searches for α: $\mathcal{A}_1 = \{0.01, 0.05, 0.1, 0.5, 1\}$ and $\mathcal{A}_2 = \{0.01, 0.02, \ldots, 0.2\}$

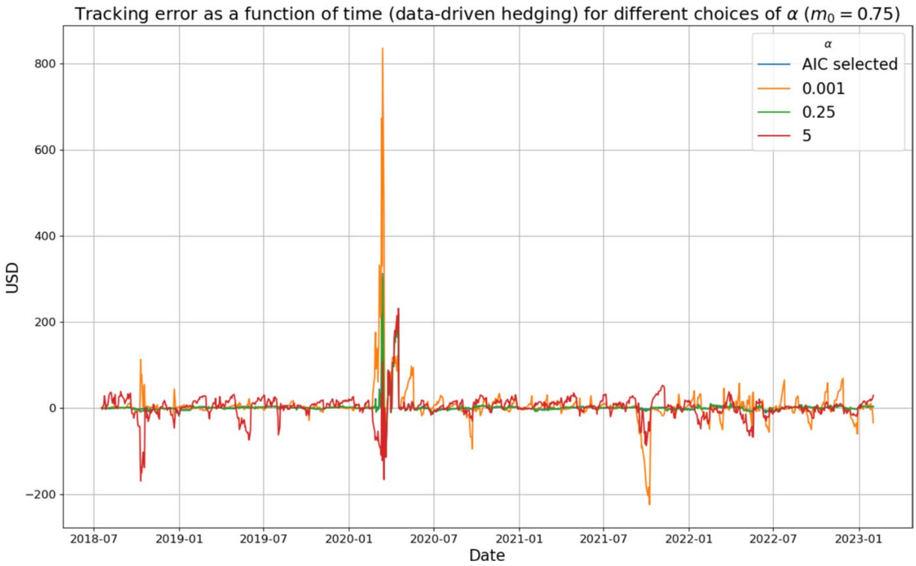

Fig. 14 Tracking error as a function of time, $m_0 = 0.75$. Comparison between different values of α. Opting for very low or very high values of α results in overfitting and underfitting, respectively. Opting for fixed $\alpha = 0.25$ results in a similar performance to performing a grid search over $\mathcal{A}_1$ or $\mathcal{A}_2$. The "AIC selected" refers to searching over $\mathcal{A}_2 = \{0.01, 0.02, \ldots, 0.2\}$

2024, the 18th Financial Risks International Forum, UBS Quant, EL & HOLT Converence 2025, and the 5th Oxford-ETH Workshop on Mathematical & Computational Finance for helpful comments and remarks.

Funding Milena Vuletić's research is supported by BNP Paribas through the *EPSRC Centre for Doctoral Training in Mathematics of Random Systems: Analysis, Modelling and Simulation* (ESPRC Grant EP/S023925/1).

Data availability SPX options data is available from OptionMetrics.

Code Availability VolGAN code is available on GitHub: https://github.com/milenavuletic/VolGAN/.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control, 19*(6), 716–723.

Bemporad, A., Bellucci, L., & Gabbriellini, T. (2014). Dynamic option hedging via stochastic model predictive control based on scenario simulation. *Quantitative Finance, 14*(10), 1739–1751.

Bertsekas, D. P., & Tsitsiklis, J. N. (1996). *Neuro-dynamic programming*. Athena Scientific.

Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2019). Deep hedging. *Quantitative Finance, 19*(8), 1271–1291.

Cao, J., Chen, J., Farghadani, S., Hull, J., Poulos, Z., Wang, Z., & Yuan, J. (2023). Gamma and Vega hedging using deep distributional reinforcement learning. *Frontiers in Artificial Intelligence, 6*, 1129370.

Capriotti, L., & Giles, M. (2024). 15 years of adjoint algorithmic differentiation (AAD) in finance. *Quantitative Finance, 24*(9), 1353–1379.

Carverhill, A. P., & Cheuk, T. H. (2003). Alternative neural network approach for option pricing and hedging. *Available at SSRN 480562*.

Chen, F., & Sutcliffe, C. (2012). Pricing and hedging short sterling options using neural networks. *Intelligent Systems in Accounting, Finance and Management, 19*(2), 128–149.

Chen, M., Xu, R., Xu, Y., & Zhang, R. (2025). Diffusion factor models: Generating high-dimensional returns with factor structure.

Cohen, S. N., Reisinger, C., & Wang, S. (2022). Hedging option books using neural-SDE market models. *Applied Mathematical Finance, 29*(5), 366–401.

Cont, R. (2006). Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance, 16*(3), 519–547.

Cont, R., Cucuringu, M., Xu, R., & Zhang, C. (2025). Tail-GAN: Learning to Simulate Tail Risk Scenarios. *Management Science*.

Cont, R., & da Fonseca, J. (2002). Dynamics of implied volatility surfaces. *Quantitative Finance, 2*(1), 45–60.

Cont, R., & Kan, Y. H. (2011). Dynamic hedging of portfolio credit derivatives. *SIAM Journal on Financial Mathematics, 2*(1), 112–140.

Cont, R., Kukanov, A., & Stoikov, S. (2014). The price impact of order book events. *Journal of Financial Econometrics, 12*(1), 47–88.

Cont, R., & Vuletić, M. (2023). Simulation of arbitrage-free implied volatility surfaces. *Applied Mathematical Finance, 30*, 94–121.

De La Peña, J. I., Iturricastillo, I., Moreno, R., Román, F., & Trigo, E. (2021). Towards an immunization perfect model? *International Journal of Finance & Economics, 26*(1), 1181–1196.

Fan, L., & Sirignano, J. (2024). Machine learning methods for pricing financial derivatives. arXiv preprint arXiv:2406.00459.

Föllmer, H., & Schweizer, M. (1988). Hedging by sequential regression: an introduction to the mathematics of option trading. *ASTIN Bulletin, 18*(2), 147–160.

Fries, C. P. (2019). Stochastic automatic differentiation: Automatic differentiation for Monte-Carlo simulations. *Quantitative Finance, 19*(6), 1043–1059.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (Eds.), *Advances in neural information processing systems.* (Vol. 27). Curran Associates Inc.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.

Hoerl, A. E., & Kennard, R. W. (1981). Ridge regression-1980: advances, algorithms, and applications. *American Journal of Mathematical and Management Sciences, 1*(1), 5–83.

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

Lamberton, D., Pham, H., & Schweizer, M. (1998). Local risk-minimization under transaction costs. *Mathematical Operations Research, 23*(3), 585–612.

Limmer, Y., & Horvath, B. (2024). Robust hedging GANs: Towards automated robustification of hedging strategies. *Applied Mathematical Finance*, 1–38.

Li, X., Uysal, A. S., & Mulvey, J. M. (2022). Multi-period portfolio optimization using model predictive control with mean-variance and risk parity frameworks. *European Journal of Operational Research, 299*(3), 1158–1176.

Lütkebohmert, E., Schmidt, T., & Sester, J. (2022). Robust deep hedging. *Quantitative Finance, 22*(8), 1465–1480.

Ludkovski, M. (2023). Statistical machine learning for quantitative finance. *Annual Review of Statistics and Its Application, 10*(1), 271–295.

Mercurio, F., & Vorst, T. (1996). Option pricing with hedging at fixed trading dates. *Applied Mathematical Finance, 3*(2), 135–158.

Mercurio, F., & Vorst, T. C. F. (1997). Option pricing and hedging in discrete time with transaction costs and incomplete markets. In M. A. H. Dempster & S. R. Pliska (Eds.), *Mathematics of derivative securities* (pp. 190–215). Cambridge, UK: Cambridge University Press.

Mirza, M., & Osindero, S. (2014). Conditional Generative Adversarial Nets.

Mulvey, J. M., & Ruszczyński, A. (1995). A new scenario decomposition method for large-scale stochastic optimization. *Operations Research, 43*(3), 477–490.
Rockafellar, R. T. (2018). Solving stochastic programming problems with risk measures by progressive hedging. *Set-Valued and Variational Analysis, 26*, 759–768.
Rockafellar, R. T., & Wets, R.J.-B. (1991). Scenarios and policy aggregation in optimization under uncertainty. *Mathematical Operations Research, 16*, 119–147.
Ruf, J., & Wang, W. (2020). Neural networks for option pricing and hedging: A literature review. *Journal of Computational Finance, 24*, 1–46.
Ruf, J., & Wang, W. (2022). Hedging with linear regressions and neural networks. *Journal of Business & Economic Statistics, 40*(4), 1442–1454.
Schweizer, M. (1995). Variance-optimal hedging in discrete time. *Mathematical Operations Research, 20*(1), 1–32.
Sercu, P., & Wu, X. (2000). Cross-and delta-hedges: Regression-versus price-based hedge ratios. *Journal of Banking & Finance, 24*(5), 735–757.
Tibshirani, R. (1996). Regression shrinkage and selection via the LASSO . *Journal of the Royal Statistical Society, Series B: Statistical Methodology, 58*(1), 267–288.
Vuletić, M., & Cont, R. (2024). Volgan: A generative model for arbitrage-free implied volatility surfaces. *Applied Mathematical Finance, 31*(4), 203–238.
Vuletić, M., Prenzel, F., & Cucuringu, M. (2024). Fin-GAN: forecasting and classifying financial time series via generative adversarial networks. *Quantitative Finance, 24*(2), 175–199.
Webster, K. T. (2023). *Handbook of price impact modeling*. Chapman and Hall/CRC.