It Only Takes a Few Moments to Hedge Options

Andrea Barletta, Paolo Santucci de Magistris, David Sloth

PII: S0165-1889(19)30007-7
DOI: <https://doi.org/10.1016/j.jedc.2018.11.008>
Reference: DYNCON 3653

To appear in: *Journal of Economic Dynamics & Control*

Received date: 30 May 2018
Revised date: 13 September 2018
Accepted date: 20 November 2018

Please cite this article as: Andrea Barletta, Paolo Santucci de Magistris, David Sloth, It Only Takes a Few Moments to Hedge Options, *Journal of Economic Dynamics & Control* (2019), doi: <https://doi.org/10.1016/j.jedc.2018.11.008>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

It Only Takes a Few Moments to Hedge Options

Andrea Barletta* Paolo Santucci de Magistris[†] David Sloth[‡]

January 11, 2019

Abstract

We propose a novel non-structural method for hedging European options, relying on two model-independent results: First, under suitable regularity conditions, an option price can be disentangled into a linear combination of risk-neutral moments. Second, there exists an explicit approximate functional form linking the risk-neutral moments to the futures price of the underlying asset and the related variance swap contracts. We show that S&P 500 call prices are mainly explained by two factors that are related to level and volatility of the underlying index. We empirically compare the performance of two strategies where the vega exposure is adjusted either by a direct position in a variance swap contract or, indirectly, through an at-the-money call. While both strategies ensure effective immunization in periods of market turmoil, taking direct exposure on variance swaps is not optimal during extended periods of subdued volatility.

Keywords: option Greeks; hedging; risk-neutral moments; variance-swap.

JEL Classification: C38; C60; G13.

*Department of Economics and Business Economics, Aarhus University, Denmark. abarletta@econ.au.dk

[†]**Corresponding author:** Department of Economics and Finance, LUISS “Guido Carli” University. Viale Romania 32, 00197 Roma, Italy. CREATES, Aarhus University, Fuglesangs Alle 4, 8210, Aarhus, Denmark. E-mail: sdemagistris@econ.au.dk

[‡]Danske Bank, Denmark. dap@danskebank.dk

1 Introduction

Option contracts are known to facilitate the efficient transmission of risk in the market, enabling traders to run balanced portfolios in terms of profit and loss uncertainty. The uncertainty of an option payoff can be hedged by taking offsetting positions in one or several securities deemed to be driving its risk. In practice, traders compute hedge ratios in terms of the so-called *Greeks*, with the *delta* and *vega* being the most prominent examples. The delta and the vega are the first-order sensitivities of an option price to changes in the underlying asset price and its volatility, respectively. Unfortunately, the Greeks are not observable quantities and the method chosen for their calculation has a crucial impact on the effectiveness of hedging.

To compute the Greeks, traders tend to resort to relatively simple models with a “layer of heuristics” on top. For example, a standard practice to determine the delta of a given vanilla option is to compute the Black-Scholes delta with the implied volatility of the option as volatility parameter. This ad-hoc usage of the Black-Scholes pricing formula is commonly referred to as the “practitioners’ Black-Scholes” (PBS) model, see Christoffersen and Jacobs (2004) and Hull and White (2017). A somewhat similar generalization of the Black-Scholes valuation model is the local volatility model of Dupire (1994), where volatility is a deterministic function of the underlying level and it is uniquely determined by the implied volatility smile. El Karoui et al. (1998) show that the performance of delta hedging through a local volatility model depends on how precisely the “true volatility” is tracked by the model. There is no clear consensus on which specification, between PBS and local volatility, provides the best delta hedging performance, see e.g. Dumas et al. (1998) and Crépey (2004). However, both approaches inherently treat the implied volatility of each option as an individual source of risk. Thus, to adjust the vega exposure of a bucket of options with a fixed maturity, one has to construct a hedging portfolio that is immune to changes in any point of the volatility smile. This, coupled with the fact that the implied volatility itself is not traded, makes vega hedging all but an easy task.

In view of these limitations, the hedging of volatility risk is typically accomplished by amending the standard PBS delta position, for instance by minimizing the variance of the

hedging error, see among others Badescu et al. (2014) and Hull and White (2017). This approach neatly improves the performance of the PBS model in terms of hedging volatility risk, although no direct position on volatility is actually taken. Nowadays, however, traders are offered the opportunity to achieve vega hedging more directly, since volatility has become a fully-fledged asset class thanks to the availability of traded instruments such as variance swaps and VIX derivatives. This is endorsed by several studies suggesting that delta-vega hedging is possible by resorting to only two or three risk factors. Adhering to this principle, stochastic volatility models assume that all the randomness driving the implied volatility surface is embedded in the underlying asset and its instantaneous variance. Many authors have dealt with hedging under stochastic volatility, see for instance Bakshi et al. (1997), Gondzio et al. (2003) and Kaeck (2012). Although the class of stochastic volatility models is widely adopted in the option pricing and hedging literature, it should be noted that any fully parametric pricing model carries a certain amount of *model risk* (Cont, 2006), which may lead to biased computation of hedge ratios, see e.g. Bakshi et al. (2000) and Alexander et al. (2009). Among all sources of model risk, of particular importance is the risk arising from possible model misspecification. Branger et al. (2012) and Kaeck (2012) note that, although the specific functional form of the variance in a stochastic volatility model has little effect on pure delta hedging, the risk of model misspecification may markedly affect the performance of more complex strategies, such as delta-vega hedging.

In the present paper, we propose a new method for computing hedge ratios of European options by only assuming the existence of a risk-neutral density fulfilling mild regularity conditions and avoiding strong assumptions on the dynamics of the underlying process. For this reason, and following the classification of Jondeau et al. (2007, Chapter 11), our methodology belongs to the class of *non-structural* methods. By very definition, a non-structural approach limits the aforementioned risk of model misspecification by a large extent. Over the years, many non-structural hedging techniques have been proposed in the literature. Hutchinson et al. (1994) propose a non-parametric method to estimate delta hedge ratios based on learning networks. Alcock and Gray (2005) de-

rive non-parametric delta hedge formulas based on the canonical valuation approach of Stutzer (1996). Tebaldi (2005) proposes a non-parametric approach to derivative hedging based on least squares regression. Bates (2005) and Alexander and Nogueira (2007) derive model-independent formulas for the delta and gamma of a European option, which can be evaluated through the option sensitivities to the strike price. Finally, Branger and Mahayni (2006) develop robust hedging strategies that do not depend on the exact specification of the stochastic volatility process. A recent comprehensive review of non-structural approaches to hedging is in Davis (2016).

This paper contributes to the existing literature by developing a new non-structural hedging methodology along the following lines: First, we reinterpret a classic non-structural pricing technique as a novel general method to compute hedge ratios of European-style derivatives. Relying on the classic theory of orthogonal polynomials, an option price can be disentangled into a linear combination of risk-neutral moments, see Jarrow and Rudd (1982), Coutant et al. (2001) and Jondeau and Rockinger (2001) for classic references, and Xiu (2014) and Schneider (2015) for recent applications. Second, we derive a model-independent relation between the risk-neutral moments, the underlying futures price, and the associated variance swap. This allows for explicit computation of the risk-neutral moments sensitivities to changes in the risk factors and, therefore, of delta and (variance swap) vega hedge ratios. In our empirical analysis we consider a panel of S&P 500 options and we show that two components associated with futures and variance swaps are the relevant factors to be hedged. Moreover, we empirically compare the performance of two different non-structural hedging strategies, where the vega exposure is adjusted either by taking a direct position on a variance swap contract or, indirectly, through an at-the-money (ATM) call option. We find that both the direct and indirect strategies provide effective immunization during market turmoil, whereas ATM options prove more reliable hedging factors than variance swaps in persistent state of low volatility. Finally, we show that the non-structural delta and vega ratios provide better hedging than the corresponding Greeks based on the PBS model.

This manuscript is organized as follows. Section 2 lays down the general non-structural

hedging methodology constituting the backbone of the paper. Section 3 provides numerical illustrations supporting the validity of the proposed methodology. Section 4 discusses the empirical validity of the assumptions underlying the methodology. Section 5 discusses the performance of two non-structural hedging strategies based on real market data on real market data. Section 6 concludes the paper. The appendix contains further details on the implementation of the methodology and additional empirical results.

2 Non-structural Hedging

Denoting by Z_t the value at time $t \geq 0$ of a generic underlying asset, we work under the standard assumption that all claims on Z are priced through an equivalent risk-neutral measure $\mathbb{Q}$, defined on a filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0})$ satisfying the usual conditions. We further assume that all claims on Z depend on a finite number of traded risky factors $\xi_{t,\tau} = (\xi_{t,\tau}^{(1)}, \dots, \xi_{t,\tau}^{(q)})$. Finally, we make the simplifying assumption of a frictionless market and, for ease of notation and without loss of generality, we set interest rates to zero. Then, at time t , the price $V_{t,\tau}^\Psi$ of a European option on Z with payoff Ψ and expiring at time $t + \tau$ is given by $V_{t,\tau}^\Psi = V_{t,\tau}^\Psi(\xi_{t,\tau}) = E^\mathbb{Q}[\Psi(Z_{t+\tau}) | \mathcal{F}_t]$. From the option seller's perspective, a hedging portfolio takes the form

$$V_{t,\tau}^\Psi = \pi_{t,\tau}^1 \xi_{t,\tau}^{(1)} + \dots + \pi_{t,\tau}^n \xi_{t,\tau}^{(q)} + \pi_{t,\tau}^B B_{t,\tau}, \quad (1)$$

where

$$\pi_{t,\tau}^i = -\frac{\partial V_{t,\tau}^\Psi}{\partial \xi_{t,\tau}^{(i)}}, \quad i = 1, \dots, q, \quad (2)$$

and $\pi_{t,\tau}^B$ is a position in a risk-less asset B , chosen to keep the portfolio self-financing. In the classic Black and Scholes framework, there is only one risk factor (the futures on Z with maturity $t + \tau$) and $\pi_{t,\tau}^1$ is the well-known delta. In stochastic volatility models there is an additional risk factor with $\pi_{t,\tau}^2$ being the so-called vega.

In the following, we outline a general methodology to obtain (2) from a bucket of observed options with a given time-to-maturity. Our approach relies on the inherent

assumption that the marginal (risk-neutral) distribution of future realizations of Z can be deduced from market data. This assumption is prevalent in the theory of model-free pricing and is substantiated by the seminal result of Breeden and Litzenberger (1978), provided that sufficiently many vanilla options are traded. Henceforth, for a fixed n and at any time $t \geq 0$, we assume the existence of a risk-neutral density for $Z_{t+\tau}$ living on a domain $\mathcal{D} \subseteq \mathbb{R}_+$ and possessing finite moments up to order n . Denoting the risk-neutral density by $f_{t,\tau}^{\mathbb{Q}}$, we have

$$V_{t,\tau}^{\Psi} = \int_{\mathbb{R}_+} \Psi(z) f_{t,\tau}^{\mathbb{Q}}(z) dz. \quad (3)$$

The building block of our method to determine hedge ratios is to characterize the risk-neutral density through its first n risk-neutral moments. This is done resorting to the classic theory of orthogonal polynomials, which ensures that under suitable conditions $f_{t,\tau}^{\mathbb{Q}}$ can be approximated as follows

$$f_{t,\tau}^{\mathbb{Q}}(z) \approx \phi(z) \sum_{k=0}^n M_{t,\tau}^{(k)} \left(\sum_{i=k}^n \sum_{j=i}^n w_{i,k} w_{i,j} z^j \right), \quad (4)$$

where $M_{t,\tau}^{(k)}$ denotes the k -th risk-neutral moment, defined as $M_{t,\tau}^{(k)} = E^{\mathbb{Q}}[Z_{t+\tau}^k | \mathcal{F}_t]$ and ϕ is an arbitrarily chosen function possessing finite moments – typically a known probability density with support containing the positive real axis. The term $w_{s,k}$ is the k -th coefficient of the s -th element in the normalized basis of orthogonal polynomials associated with ϕ . For completeness, explicit formulas for $w_{s,k}$ are reported in Appendix A. Replacing $f_{t,\tau}^{\mathbb{Q}}$ with its expansion in (3) yields the following approximate formula for a European-style payoff

$$V_{t,\tau}^{\Psi} \approx H_0^{(\Psi)} + M_{t,\tau}^{(1)} H_1^{(\Psi)} + \dots + M_{t,\tau}^{(n)} H_n^{(\Psi)}, \quad (5)$$

where each coefficient $H_k^{(\Psi)}$ depends on n , Ψ and ϕ , but *not* on the risk-neutral moments, and it is determined by

$$H_k^{(\Psi)} = \sum_{i=k}^n \sum_{j=i}^n w_{i,k} w_{i,j} \int_0^{\infty} z^j \Psi(z) \phi(z) dz.$$

Equation (5) defines a non-structural linkage between option prices and the latent risk-

neutral density in terms of risk-neutral moments. The admissibility of orthogonal polynomial expansions relies on technical conditions relating the tail decay rates of the risk-neutral density and the kernel. This point is further elaborated in Section 2.2. Assuming that a fixed kernel function ϕ is admissible at all times between t and an arbitrarily small increment $t + \Delta$, in view of (5)

$$V_{t+\Delta,\tau}^\Psi - V_{t,\tau}^\Psi \approx \left(M_{t+\Delta,\tau}^{(1)} - M_{t,\tau}^{(1)}\right) H_1^{(\Psi)} + \dots + \left(M_{t+\Delta,\tau}^{(n)} - M_{t,\tau}^{(n)}\right) H_n^{(\Psi)},$$

which implies

$$\frac{\partial V_{t,\tau}^\Psi}{\partial \xi_{t,\tau}} \approx H_1^{(\Psi)} \frac{\partial M_{t,\tau}^{(1)}}{\partial \xi_{t,\tau}} + \dots + H_n^{(\Psi)} \frac{\partial M_{t,\tau}^{(n)}}{\partial \xi_{t,\tau}}. \quad (6)$$

Setting $\Psi(z) = (x - K)^+$ yields a formula for the price $C_{t,\tau}^{(K)}$ of a call option with strike K , that is $C_{t,\tau}^{(K)} \approx H_0^{(K)} + H_1^{(K)} M_{t,\tau}^{(1)} + \dots + H_n^{(K)} M_{t,\tau}^{(n)}$, where, with a slight abuse of notation, we denote $H^{(K)} = H^{(\Psi)}$ as the payoff is now a fixed function of K . More compactly, $C_{t,\tau}^{(K)} \approx H^{(K)} M_{t,\tau}$, where $M_{t,\tau} = [1, M_{t,\tau}^{(1)}, \dots, M_{t,\tau}^{(n)}]'$ and $H^{(K)} = [H_0^{(K)}, \dots, H_n^{(K)}]$. Putting these results together, for a given set of strikes $K_1, \dots, K_p$, the following approximations hold

$$C_{t,\tau} := [C_{t,\tau}^{(K_1)}, \dots, C_{t,\tau}^{(K_p)}]' \approx H M_{t,\tau}, \quad (7)$$

and

$$\frac{\partial C_{t,\tau}}{\partial \xi_{t,\tau}} \approx H \frac{\partial M_{t,\tau}}{\partial \xi_{t,\tau}}, \quad (8)$$

where H is the $p \times (n+1)$ matrix with rows $H^{(K_1)}, \dots, H^{(K_p)}$, while $\frac{\partial M_{t,\tau}}{\partial \xi_{t,\tau}}$ is the $(n+1) \times q$ matrix formed by the gradients of each component $1, M_{t,\tau}^{(1)}, \dots, M_{t,\tau}^{(n)}$ or, in other words, by the Greeks of each risk-neutral moment. Equations (6)-(8) offer a representation of the hedge ratios $\pi^1, \dots, \pi^q$ where the dependence on the risk factors, ξ , and that on the payoff function, Ψ , are explicitly decoupled. In particular, the dependence on ξ is fully described by the moments sensitivities to the risk factors, whereas the dependence on Ψ is embedded in H . Using the moments sensitivities as proxies has the advantage of dealing with simple payoff functions, which are easy to relate to the risk factors. From a mathematical point of view, an investor could directly exploit the expansion (5)

and take a direct position on the risk-neutral moments with the purpose of hedging the risks associated with a call option. However, it is practically unfeasible to take a direct position on the first n risk-neutral moments since the latter are non-traded quantities. In the following section, we characterize $\xi_{t,\tau}$ in terms of traded quantities and derive a (non-parametric) functional link between $M_{t,\tau}$ and $\xi_{t,\tau}$.

2.1 Delta and Variance Swap Vega

In this section, we consider two “natural candidates” for the role of risk-factors. The first factor is the futures on the underlying asset, denoted as $F_{t,\tau} = M_{t,\tau}^{(1)}$. Being exactly the risk-neutral mean of $Z_{t+\tau}$, the futures carries information about the location of the risk-neutral density and for this reason the associated “delta risk” is the classic object of interest by portfolio hedgers. We complement the futures with the variance swap, defined as $VS_{t,\tau} = -\frac{2}{\tau}E^{\mathbb{Q}}\left[\log\left(\frac{Z_{t+\tau}}{F_{t,\tau}}\right) \mid \mathcal{F}_t\right]$, see Carr and Madan (1998). There are two main reasons supporting this choice. First, variance swap contracts are traded securities allowing for the direct inclusion of other features of the risk-neutral density (e.g. the volatility) in the hedging portfolio. Second, being proportional to the risk-neutral mean of the log-returns, the variance swap is defined by a smooth payoff function, i.e., the logarithm. Thanks to this property, an approximate relation between $VS_{t,\tau}$ and the first k risk-neutral moments can be obtained by looking at the k -th order Taylor-expansion of $\log(x)$ around $x = 1$, i.e. $\log\left(\frac{Z_{t+\tau}}{F_{t,\tau}}\right) \approx \sum_{i=0}^k a_i^{(k)} \left(\frac{Z_{t+\tau}}{F_{t,\tau}}\right)^i$. Hence, by taking position on $VS_{t,\tau}$, we allow for the higher-order risk-neutral moments to influence – and possibly improve – the quality of the hedging strategy. Therefore, the Greeks of interest are the delta and the *variance swap vega*, henceforth simply vega, which are denoted as $\Delta_{t,\tau}^{(K)} := \frac{\partial C_{t,\tau}^{(K)}}{\partial F_{t,\tau}}$ and $\nu_{t,\tau}^{(K)} := \frac{\partial C_{t,\tau}^{(K)}}{\partial VS_{t,\tau}}$. The hedging portfolio (1) becomes

$$C_{t,\tau}^{(K)} - \Delta_{t,\tau}^{(K)} F_{t,\tau} - \nu_{t,\tau}^{(K)} VS_{t,\tau} + \pi_{t,\tau}^B B_{t,\tau}. \quad (9)$$

To compute $\Delta_{t,\tau}^{(K)}$ and $\nu_{t,\tau}^{(K)}$ through formula (8), we are required to determine moments sensitivities to the underlying futures and variance swaps, i.e., $\frac{\partial M_{t,\tau}^{(k)}}{\partial F_{t,\tau}}$ and $\frac{\partial M_{t,\tau}^{(k)}}{\partial VS_{t,\tau}}$, $k =$

$2, \dots, n$. A model-free heuristic argument outlined in Appendix B leads to the following approximate relation between the k -th moment, the futures and the variance swap

$$M_{t,\tau}^{(k)} \approx F_{t,\tau}^k e^{\beta_k \text{VS}_{t,\tau}}, \quad (10)$$

where β_k measures the dependence of the k -th risk-neutral moment on the variance swap. The empirical analysis on S&P 500 options reported in Sections 4.1-4.2 not only strongly supports the evidence that $F_{t,\tau}$ and $\text{VS}_{t,\tau}$ are the main risk factors driving options, but it also confirms the validity of the functional relation in (10). By plugging (10) into equations (7)-(8) we get

$$\Delta_{t,\tau} := [\Delta_{t,\tau}^{(K_1)}, \dots, \Delta_{t,\tau}^{(K_p)}]' \approx HD, \quad \nu_{t,\tau} := [\nu_{t,\tau}^{(K_1)}, \dots, \nu_{t,\tau}^{(K_p)}]' \approx HW, \quad (11)$$

where $D = [0, 1, 2F_{t,\tau}e^{\beta_2 \text{VS}_{t,\tau}}, \dots, nF_{t,\tau}^{n-1}e^{\beta_n \text{VS}_{t,\tau}}]'$, $W = [0, 0, \beta_2 F_{t,\tau}^2 e^{\beta_2 \text{VS}_{t,\tau}}, \dots, \beta_n F_{t,\tau}^n e^{\beta_n \text{VS}_{t,\tau}}]'$.

Finally, each coefficient β_k , $k = 2, \dots, n$, can be obtained non-parametrically by inverting (10), i.e.,

$$\beta_k = \frac{1}{\text{VS}_{t,\tau}} \log \left(\frac{M_{t,\tau}^{(k)}}{F_{t,\tau}^k} \right), \quad k = 2, \dots, n, \quad (12)$$

so that the Greeks in (11) can be computed very efficiently given $F_{t,\tau}$, $\text{VS}_{t,\tau}$ and $M_{t,\tau}^{(k)}$.

The accuracy of any hedging strategy based on the delta and the vega hedge ratios defined in (11) relies on two main assumptions. First, that $F_{t,\tau}$ and $\text{VS}_{t,\tau}$ are the main risk factors driving options. Second, that the approximate relation in (10) holds to a high degree of precision. As we show in Sections 4.1-4.2, the validity of both assumptions is strongly supported by empirical evidence in the context of S&P 500 options. Furthermore, although the relation in (10) is an approximation, no modeling assumptions are made to derive it. Thus, the equations in (11) preserve the non-structural nature of the proposed approach for the determination of the option Greeks. The entire work-flow necessary to compute $\Delta_{t,\tau}^{(K)}$ and $\nu_{t,\tau}^{(K)}$ can be reproduced with minimal computational effort by means of the MATLAB App *rndfittool* publicly available on GitHub, see Barletta and Santucci de Magistris (2018) for a tutorial. In the remainder of this section, we discuss details on the

practical implementation of the methodology, such as the choice of the kernel and the estimation of the risk-neutral moments.

2.2 Choice of the Kernel Function

As anticipated in Section 2, in general the function ϕ cannot be totally independent of the true (but unknown in practice) underlying model. Based on the theoretical results outlined in Filipović et al. (2013) and Barletta and Nicolato (2017), the choice of ϕ is dictated by the rate of tail decay and the support of the true density. In other words, the expansion (4) is non-structural since it does not depend on the (parametric) form of the risk-neutral density, but rather on general properties such as its tail decay and support. For instance, a kernel featuring exponential tail decay can be adopted to approximate any positively supported risk-neutral density decaying at least exponentially, regardless of the underlying true model from which it originates. As for the support, a common approach followed in the literature is to consider the risk-neutral density associated with the standardized log-returns. In this case, a kernel with support on the entire real axis (such as a Gaussian density, see e.g. Aït-Sahalia, 2002 and Li et al., 2013) is to be used. However, the argument used for the derivation of the Greeks in (11) does not allow for working with a transformed variable. This restricts the set of admissible kernels to those whose support is contained in the positive real axis.

Based on these considerations, we propose the following “double-beta” density

$$\phi(x) \sim w \left(\frac{x}{\lambda_1} \right)^{a_1} \left(1 - \frac{x}{\lambda_1} \right)^{b_1} \mathbb{1}_{0 \leq x \leq a_1} + (1 - w) \left(\frac{x}{\lambda_2} \right)^{a_2} \left(1 - \frac{x}{\lambda_2} \right)^{b_2} \mathbb{1}_{0 \leq x \leq a_2}, \quad (13)$$

with $a_1, b_1, a_2, b_2, \lambda_1, \lambda_2 > 0$ and $0 \leq w \leq 1$, and where these parameters are chosen so to minimize the distance between the observed and the kernel-implied option prices. Compared to other kernels, such as the lognormal, the double-beta provides the best fitting of (7) to the observed call prices even for a moderately low expansion order n . Given that we are interested in hedging options that are observed only for a finite set of strikes, it is not unreasonable to choose a kernel with bounded support such as the

double-beta. This leads to several advantages. First, any condition on the asymptotic tail decay of ϕ can be disregarded. Second, and perhaps more importantly, requiring the existence of the risk-neutral moments is no longer a binding constraint. This means that the risk-neutral moments appearing in formula (5) are replaced with their “corridor” counterparts. Using the corridor moments is economically advisable because they can be replicated using traded vanilla options more precisely than the “regular” moments, see e.g. Carr and Madan (1998) and Bakshi and Madan (2000). Furthermore, working with a bounded domain makes our approach robust to possible moment explosions typical of processes characterized by tails decaying slower than those of the lognormal density, see e.g. Andersen and Piterbarg (2007).

In Section 3 we carry out numerical exercises to confirm the validity of our method under several parametric specifications for the underlying asset. Our illustrations demonstrate that, as a matter of fact, working with a (arbitrarily large) bounded domain has very little impact on the accuracy of the approximation (5) and the related Greeks, even when the true risk-neutral density has unbounded support and potentially explosive moments. In Section 2.3, we briefly present the method adopted to compute the risk-neutral moments up to order n from an observed set of options for fixed t and τ .

2.3 Computation of the Risk-Neutral Moments

Given that the risk-neutral moments $M_{t,\tau}^{(k)}$ necessary to compute (11) are not directly observed, they have to be retrieved from the available option prices. For given t and τ , we obtain the risk-neutral moments up to order n by minimizing the square pricing errors between (7) and the observed option prices. We denote these prices by $V_{t,\tau}^{\Psi_1}, \dots, V_{t,\tau}^{\Psi_m}$ where m denotes the total number of available contracts and $\Psi_1, \dots, \Psi_m$ denote the related payoff functions. To avoid scale effects produced by working with high order moments, we rearrange the terms in (5) and obtain $V_{t,\tau}^{\Psi} \approx A_0^{(\Psi)} + c_1 A_1^{(\Psi)} + \dots + c_n A_n^{(\Psi)}$, where $A_k^{(\Psi)} = \sum_{j=0}^k w_{k,j} \int_{\mathcal{D}} x^j \Psi(x) \phi(x) dx$ and the coefficients

$$c_{t,\tau}^{(k)} = \sum_{j=1}^k w_{k,j} M_{t,\tau}^{(j)}, \quad (14)$$

are the unknown terms to be estimated. The estimates $\tilde{c}_{t,\tau}^{(1)}, \dots, \tilde{c}_{t,\tau}^{(n)}$ are obtained by solving the following least-squares problem

$$\left[\tilde{c}_{t,\tau}^{(1)}, \dots, \tilde{c}_{t,\tau}^{(n)} \right] = \arg \min_{c_1, \dots, c_n} \sum_{i=1}^m \left| A_0^{(\Psi_i)} + c_1 A_1^{(\Psi_i)} + \dots + c_n A_n^{(\Psi_i)} - V_{t,\tau}^{\Psi_i} \right|^2.$$

Finally, estimates of $M_{t,\tau}^{(1)}, \dots, M_{t,\tau}^{(n)}$ are obtained by inverting the linear mapping (14). In all the applications of this paper, we choose $n = 5$. Given the estimates of $M_{t,\tau}^{(1)}, \dots, M_{t,\tau}^{(n)}$, the variance swap is computed through (5), with $\Psi(z) = -\frac{2}{\tau} \log(z/F_{t,\tau})$. Alternatively, one could use the spanning formulas of Carr and Madan (1998) and Bakshi and Madan (2000) to estimate both the risk-neutral moments and the variance swap, but we find that (10) and (11) are generally less accurate when $M_{t,\tau}^{(1)}, \dots, M_{t,\tau}^{(n)}$ and $VS_{t,\tau}$ are computed in this way. To verify the accuracy of this approach, we look at the pricing error on a sample of monthly S&P 500 options collected in the period 2006–2016. Specifically we consider mid quotes relative to all available calls and puts with positive bid price. The goodness-of-fit to the observed option prices emerges by looking at the time series of the relative error on the Black and Scholes implied volatilities displayed in Figure 1 (solid blue line). Indeed, the pricing error amounts to 2.6% on average over both time and moneyness with a standard deviation of 1.3%. Figure 1 also shows that the pricing error seems unrelated to the level of volatility of the underlying asset, as measured by the variance swap (dashed red line). Although the pricing errors display some mild degree of

Figure 1: Pricing errors and variance swap, $\tau = 1$ month. The solid blue line is the time series of the mean (across strikes) absolute relative distance (expressed in percentage) between the Black-Scholes implied volatilities associated with observed and estimated option prices. The dashed red line depicts the time series of the 1-month variance swap.

(negative) correlation (-11.5%), this is not significant (the p -value is 19.6%). Thus, we can safely conclude that no clear relation exists between the pricing error and the level of market volatility. This further proves the accuracy of the estimates of the risk-neutral moments in terms of pricing residuals and demonstrates the consistency of the estimated moments with the information provided by the market.

3 Numerical Illustrations

In this section, we assess the validity of the approximating formulas in (11) by means of numerical experiments under different parametric setups. In particular, we illustrate the practical validity of the proposed methodology by assessing the flexibility of the hedging formulas in (11) when the true values of $\Delta_{t,\tau}^{(K)}$ and $\nu_{t,\tau}^{(K)}$ are known. To this purpose, we consider option prices generated from several parametric specifications, for which values of $F_{t,\tau}$, $VS_{t,\tau}$, $\Delta_{t,\tau}^{(K)}$ and $\nu_{t,\tau}^{(K)}$ can be computed exactly. These values are compared with those obtained with our methodology. More specifically, we rely on three different parametric models for the underlying process Z , namely a CEV-type variance as, e.g., in Aït-Sahalia and Kimmel (2007), a stochastic vol-of-vol extension of the Heston model (SVV) and the affine jump-diffusion model of Duffie et al. (2000) with Gaussian jumps in returns and exponential jumps in variance (SVJJ). The reason behind the choice of these models is twofold: First, they generate a number of features of the underlying process, such as jumps, heavy tails in volatility and stochastic vol-of-vol; Second, for all the three models, the increment on $VS_{t,\tau}$ (produced by perturbing the instantaneous variance) can be measured exactly, thus simplifying the computation of the true values of $\nu_{t,\tau}^{(K)}$. The dynamics of the three models can be embedded in the following general specification

$$\begin{aligned} d \log Z_t &= -\frac{1}{2}v_t dt + \sqrt{v_t} dW_t^1 + dJ_t^X, \\ dv_t &= \kappa(\bar{v} - v_t)dt + \epsilon_t v_t^\delta dW_t^2 + dJ_t^v, \\ d\epsilon_t &= \ell(\bar{\epsilon} - \epsilon_t)dt + \eta \sqrt{\epsilon_t} dW_t^3, \end{aligned}$$

where $W = (W^{(1)}, W^{(2)}, W^{(3)})$ is a correlated Brownian motion with correlation matrix $\rho = (\rho_{ij})_{i,j=1,2,3}$, $J = (J^X, J^v)$ is a compound Poisson process with intensity λ . The jump size characteristic function is given by

$$E^{\mathbb{Q}} \left[e^{\xi_1 \Delta J_{t+\tau}^X + \xi_2 \Delta J_{t+\tau}^v} \mid \mathcal{F}_t \right] = e^{\xi_1 \mu_X + \frac{1}{2} \xi_1^2 \sigma_X^2} \frac{1}{1 - \mu_v (\rho_J \xi_1 + \xi_2)}.$$

Under this general model, the variance swap can be computed explicitly as an affine function of the instantaneous variance. Specifically

$$\text{VS}_{t,\tau} = \frac{1 - e^{-\kappa\tau}}{\kappa\tau} v_t + \left(\frac{\lambda \mu_v}{\kappa} + \bar{v} \right) \left(1 - \frac{1 - e^{-\kappa\tau}}{\kappa\tau} \right) + 2\lambda \left(\frac{e^{\mu_X + \frac{1}{2} \sigma_X^2}}{1 - \mu_v \rho_J} - 1 - \mu_X \right). \quad (15)$$

The parameters used in the numerical illustrations are reported in Panel (a) of Table 1. The CEV model is obtained as a limit case for $\ell = \eta = 0$ and $\lambda = 0$, i.e., when $\epsilon_t = \epsilon$ is constant and the jump components disappear in both the log-price and the variance processes. The SVV model is obtained for $\delta = \frac{1}{2}$ and $\lambda = 0$. The SVJJ model is obtained for $\delta = \frac{1}{2}$ and $\epsilon_t = \epsilon$.

In the numerical analysis, we evaluate option prices associated with $p = 50$ strikes through Monte Carlo simulations for each model specification. For these options, true model-implied values of $\Delta_{t,\tau}^{(K)}$ and $\nu_{t,\tau}^{(K)}$ are compared with the approximate values retrieved by applying our methodology to the discrete set of model-implied options. Panel (b) of Table 1 shows that the values of the variance swap approximated through (5), with $\Psi(z) = -\frac{2}{\tau} \log(z/F_{t,\tau})$, closely replicate the model-implied ones, computed by (15). The estimated values of $\Delta_{t,\tau}^{(K)}$ and $\nu_{t,\tau}^{(K)}$ for different strike values are reported in Figure 2, along with the true theoretical values. The associated relative errors are reported in Panel (c) of Table 1. These illustrations indicate that we can effectively and precisely replicate the true delta and vega for the entire range of moneyness considered. In particular, the approximation of the delta proves extremely accurate irrespective of the model specification and the chosen maturity, with a relative error well below 1% in all cases. The approximation of the vega is also very reliable, with a relative error ranging from 2% to 8%.

Figure 2: Call options Greeks for different model specifications and times-to-maturity. Each plot reports true model-implied values (asterisks) versus approximate values (solid line) retrieved from model-implied prices, as functions of the option moneyness (x-axis).

	CEV		SVV		SVJJ	
(a) Model parameters						
$\sqrt{v_0}$	0.221	$\sqrt{v_0}$	0.307	$\sqrt{v_0}$	0.178	
κ	4.103	κ	5.070	κ	3.250	
$\bar{v}$	0.045	$\bar{v}$	0.046	$\bar{v}$	0.179	
ϵ	0.858	ϵ_0	0.480	ϵ	0.290	
δ	0.655	ℓ	1.000	μ_X	1.072	
ρ	-0.760	$\bar{\epsilon}$	0.480	σ_X	-0.100	
		η	1.000	μ_V	0.100	
		ρ_{12}	-0.767	σ_V	0.062	
		ρ_{13}	-0.500	ρ	-0.700	
		ρ_{23}	0.800	ρ_J	0	
(b) Variance swaps square-root percentage values						
	True	Approximate	True	Approximate	True	Approximate
$\tau = 1$ month	22.00%	22.04%	29.26%	29.27%	26.94%	26.42%
$\tau = 3$ months	21.80%	21.77%	27.10%	26.93%	32.36%	32.31%
$\tau = 6$ months	21.62%	21.62%	25.20%	25.03%	36.96%	36.92%
(c) Relative errors on Greeks						
	Delta	Vega	Delta	Vega	Delta	Vega
$\tau = 1$ month	0.16%	6.17%	0.04%	1.57%	0.60%	8.22%
$\tau = 3$ months	0.49%	5.87%	0.21%	1.98%	0.08%	4.82%
$\tau = 6$ months	0.69%	5.48%	0.42%	3.78%	0.18%	2.02%

Table 1: Numerical validation. The parameters for the SVJJ model are inspired by Sepp (2008), while parameters for the CEV and the SVV models are inspired by Aït-Sahalia and Kimmel (2007). We set $S_0 = 2000$ for all models.

Given these comforting numerical results, the rest of the paper is dedicated to the assessment of the empirical soundness of the assumptions underlying the non-structural hedging outlined in Section 2 and to the analysis of the performance of delta-vega hedging strategies when employed on real market data.

4 Empirical Analysis

In this section, we provide empirical foundation to the main economic assumptions underlying our methodology. More specifically, we show that when the underlying asset is the S&P 500 index the optimal number of hedging factors is 2 and that the approximate relation in (10) has a high degree of accuracy. In our analysis, we consider a panel of S&P 500 options for different times and maturities. The dataset of option prices is obtained from the OptionMetrics database. To deal with irregular ranges of available call strikes we take average option prices over fixed moneyness intervals. For each time-to-maturity

and for each moneyness interval, the sample consists of a total of 126 call mid prices collected at monthly frequency in the period 2006 - 2016. Figure 3 displays the times series of S&P 500 1-month call options over different levels of moneyness. The plot reports marked differences between the dynamic behaviors of in-the-money options (ITM; top-left) and out-the-money options (OTM; bottom-right). ITM options are characterized by a pronounced (non-stationary) stochastic trend, whereas OTM options exhibit a clear stationary behavior (see Appendix C for further analysis).

The remainder of this section proceeds as follows. First, in Section 4.1 we assess the main risk drivers of index options by analyzing the commonalities in the S&P 500 call prices. The purpose is to prove that two main risk-factors of the option prices are indeed closely related to the futures and to the variance swap, thus supporting the definition of the hedging portfolio in (9). Second, in Section 4.2 we provide evidence to support the empirical consistency of the approximate relation in (10) between the risk-neutral moments, the underlying futures price, and the variance swap prices.

4.1 Common Factors in S&P 500 Options

We adopt principal component analysis (PCA) with the purpose of studying the commonalities in the S&P 500 options. A number of papers have proposed PCA to identify the common factors explaining the dynamic evolution of the implied volatility surface, see e.g. Cont and Da Fonseca (2002), Fengler et al. (2003) and Christoffersen et al. (2009, 2018). The common finding of these papers is that the structure of implied volatilities can be described as a random surface driven by a small number of factors. While these papers are mostly focused on explaining the term-structure of the implied volatility, we keep the time-to-maturity fixed and concentrate our analysis on moneyness effects. Furthermore, we perform our analysis on option prices rather than on implied volatilities. This enables us to find a direct link between the common factors and traded assets through which we can hedge options. In particular, the non-stationary behavior of ITM options can be related to that of the underlying S&P 500 index, since the price of a call option must converge to the value of the underlying futures as the strike price goes to zero. A similar

Figure 3: Call prices time series. Each plot reports times series of average call prices belonging to the moneyness range indicated by I .

consideration can be drawn for deep OTM options, whose stationary behavior can be associated with that of volatility. Figure 3 further suggests a “smooth transition” from a non-stationary to a stationary dynamic behavior of call options as the moneyness runs from ITM to OTM. Therefore, we could intuitively argue that S&P 500 call options are driven by a combination of two factors that can be interpreted as the “level” and the “volatility” of the underlying index.

The PCA confirms this intuition since the first two principal components have clear-cut financial interpretation as illustrated in Figure 4. Specifically, Panel (a) of Figure 4 highlights the close link between the first principal component and the S&P 500 futures price. Analogously, Figure 4-(b) shows that the second component evolves similarly to the square of the VIX, i.e. the 1-month variance swap. Table 2 reveals that the first

Figure 4: Principal components of X , $\tau = 1$ month. The left panel displays the SPX futures (solid line) and the first principal component (dashed line) of X . The right panel displays the square of the VIX index (solid line), i.e., the 1-month variance swap scaled by $1e4$, and the second principal component (dashed line) of X .

two principal components explain more than 97% of the variability of our panel of options.¹ Looking at the loading coefficients on the two principal components, we note that most ITM options load positively – and with largest magnitude – on the first component, which explains the non-stationary stochastic trend common to this subset of options.

¹Table 2 reports results for $\tau = 1$ and analogous results are obtained for maturities up to 6 months. For longer maturities (e.g. 1 year), a third principal component might be necessary to explain the same level of variability, due to the so-called “rho risk” arising from changes in interest rates, see Bakshi et al. (2000) and Lioui and Poncet (2000). However, this is of marginal importance here as options with very long maturities are not considered in our analysis.

Conversely, the OTM options load positively on the second component. Furthermore, the OTM options load negatively into the first component, i.e. the futures. Having established that OTM options are mostly driven by a volatility term and since the principal components are orthogonal by construction, the negative weights associated with OTM options on the first component reflect the leverage effect. Interestingly, the bucket of ATM options with moneyness $0.990 - 1.005$ exhibits very little sensitivity to changes in the futures, as it loads almost exclusively on the second component. This result suggests that ATM options are less sensitive to the leverage effect than ITM and OTM options.²

Component	Eigenvalue	Proportion	Cumulative	Moneyness	W_{PCA1}	W_{PCA2}	R^2 (PCA)	R^2 (OBS)
1	10.490	0.6232	0.6171	0.780 – 0.840	0.308	-0.006	0.9938	0.9934
2	6.0300	0.3573	0.9717	0.840 – 0.855	0.308	-0.002	0.9953	0.9980
3	0.4200	0.0251	0.9967	0.855 – 0.870	0.308	0.001	0.9955	0.9984
4	0.0240	0.0015	0.9981	0.870 – 0.885	0.308	0.010	0.9953	0.9972
5	0.0089	0.0005	0.9986	0.885 – 0.900	0.308	0.023	0.9961	0.9966
6	0.0046	0.0003	0.9989	0.900 – 0.915	0.307	0.038	0.9959	0.9954
7	0.0045	0.0003	0.9992	0.915 – 0.930	0.304	0.062	0.9955	0.9914
8	0.0031	0.0002	0.9994	0.930 – 0.945	0.299	0.098	0.9968	0.9809
9	0.0025	0.0001	0.9995	0.945 – 0.960	0.285	0.156	0.9969	0.9493
10	0.0021	0.0001	0.9996	0.960 – 0.975	0.243	0.248	0.9907	0.8662
11	0.0018	0.0001	0.9997	0.975 – 0.990	0.146	0.351	0.9690	0.7183
12	0.0014	0.0001	0.9998	0.990 – 1.005	0.011	0.396	0.9442	0.6336
13	0.0011	0.0001	0.9999	1.005 – 1.020	-0.088	0.382	0.9610	0.7029
14	0.0008	0.0000	0.9999	1.002 – 1.035	-0.139	0.362	0.9915	0.8210
15	0.0006	0.0000	1.0000	1.035 – 1.050	-0.153	0.351	0.9877	0.9081
16	0.0004	0.0000	1.0000	1.050 – 1.065	-0.149	0.342	0.9396	0.9589
17	0.0003	0.0000	1.0000	1.060 – 1.130	-0.125	0.318	0.7737	0.9671

Table 2: Principal components analysis. The table reports the sorted eigenvalues of the correlation matrix of panel of option prices, the proportion of explained variance by each component and the cumulative explained variance. The table also reports the weights associated to the first two components for each moneyness range. The last two columns report, for each moneyness, the R -square of the OLS regression on the two principal components, denoted as R -square (PCA), and on the two observed factors, denoted as R -square (OBS).

We also perform time series regressions of the 1-month options, for each moneyness level, on the S&P 500 futures (with 1-month maturity) and the square of the VIX, which are the *observed* risk factors. The values of the R -square for each regression are reported in the last column of Table 2. Not surprisingly, the R -square is close to one for the ITM options, which are characterized by the $I(1)$ trend, as well as for the OTM options,

²Appendix C also reports the results of a factor analysis based on a state-space model. The results are qualitatively very similar to those obtained with the PCA analysis.

which are used to determine the VIX. In contrast, the R-square is around 70% for the options falling in the ATM region. This reduced explanatory power of the two observed risk factors for the ATM options can be justified by looking again at Panel (b) of Figure 4. Here, we observe that the second principal component and the square of the VIX significantly deviate from each other in the last part of the sample, approximately starting from 2013. As a matter of fact, the VIX has been close to its historical low values for a long period after 2013. Section 5 investigates this phenomenon from a hedging perspective and shows that, in presence of a persistent state of low volatility, alternative “indirect” forms of vega hedging, less sensitive to the leverage effect, should be preferred over taking a direct position on volatility.

4.2 Risk-neutral moments and variance swap

We now turn our attention to the empirical consistency of formula (10), which is at the core of (11). In all the illustrations reported in this section, we fix $n = 5$ and we estimate the values of the risk-neutral moments and the variance swaps, required in (11), by following the procedure described in Section 2.3. For each $k = 2, \dots, 5$, we regress the dependent variable, $M^{(k)}$, on the two independent variables, F and VS. In particular, we consider the following time series regression

$$Y_t^{(k)} = \beta_k X_t + \epsilon_t^{(k)}, \quad k = 2, \dots, 5, \quad (16)$$

where

$$Y_t^{(k)} = \log \left(M_{t,\tau}^{(k)} \right) - k \log (F_{t,\tau}), \quad X_t = VS_{t,\tau}, \quad \tau = 1 \text{ month},$$

and where the coefficient β_k is estimated by OLS. As it emerges from Figure 5, Equation (10) provides an excellent approximation of the risk-neutral moments. This is confirmed by the R-square coefficients, which are above 99% for all $k = 2, \dots, 5$. As $M_{t,\tau}^{(k)}$ is not directly observed, we repeat regression (16) using the lagged values of $Y_t^{(k)}$ as an instrumental variable, to check the robustness of the β_k estimate to the presence of measurement error in $M_{t,\tau}^{(k)}$, see Hansen and Lunde (2014). The estimates of β_k and the

R-square coefficients are basically the same as those obtained under the OLS regression, which confirms the negligible role of the error implied by measuring $M_{t,\tau}^{(k)}$ through options. These results endorse the empirical validity of (10) and, in particular, the existence of a deterministic relation between the risk-neutral moments and the risk factors. An important consequence is that each coefficient β_k – and the corresponding delta and vega hedge ratios – can be determined through (12), i.e., solely based on present values of $F_{t,\tau}$, $VS_{t,\tau}$ and $M_{t,\tau}^{(k)}$.

Figure 5: Risk-neutral moments, $\tau = 1$ month. The figure displays time series of $M_{t,\tau}^{(k)} / F_{t,\tau}^k$ (solid-blue line) and $e^{\hat{\beta}_k VS_{t,\tau}}$ (red-dashed line). The left and the right panels display the results for $k = 2$ and $k = 5$, respectively.

5 Hedging S&P 500 Options

The goal of this section is to empirically illustrate the effectiveness of the proposed methodology for the determination of the Greeks. To this purpose, we assess the “real-world performance” of non-structural delta-vega hedging strategies for S&P 500 call options under different market scenarios. In particular, the vega risk is adjusted either by a direct position on a variance swap contract or indirectly through an ATM call option. We implement replicating portfolios for different buckets of S&P 500 call options and we evaluate their hedging performance at daily and weekly re-balancing frequencies. For fixed strike K , issuing day t_0 , time-to-maturity $\tau > 0$ and re-balancing times $t_1 = t_0 + \tau/N, \dots, t_N = t_0 + \tau$, we consider the hedging portfolio (9) with hedge vector $\xi_{t,\tau} = (F_{t,\tau}, VS_{t,\tau})$ and portfolio weights computed by (11). For brevity, throughout the

section we denote

$$C_i^{(K)} = C_{t_i, \tau - t_i}^{(K)}, \quad F_i = F_{t_i, \tau - t_i}, \quad VS_i = VS_{t_i, \tau - t_i}, \quad \Delta_i^{(K)} = \Delta_{t_i, \tau - t_i}^{(K)}, \quad \nu_i^{(K)} = \nu_{t_i, \tau - t_i}^{(K)},$$

so that, at the i -th re-balancing time, we consider the following portfolio

$$C_i^{(K)} + \Delta_i^{(K)} F_i + \nu_i^{(K)} VS_i, \quad (17)$$

which is associated with the replication squared error (RSE) for the i -th rebalancing time and strike K

$$\text{RSE}_i^{(K)} = \left[\Delta_{i-1}^{(K)} (F_i - F_{i-1}) + \nu_{i-1}^{(K)} (VS_i - VS_{i-1}) - (C_i^{(K)} - C_{i-1}^{(K)}) \right]^2. \quad (18)$$

We normalize (18) by the replication error generated by the practitioners Black-Scholes (PBS) delta hedging approach, so to have a scale-free error metric. Therefore, for a bucket of call options with strikes $K_1, \dots, K_p$, we define the replication “average gain” as follows

$$G = \frac{1}{M} \sum_{j=1}^M \left(1 - \frac{\sum_{i=1}^N \text{RSE}_i^{(K_j)}}{\sum_{i=1}^N \text{BRSE}_i^{(K_j)}} \right). \quad (19)$$

where

$$\text{BRSE}_i^{(K)} = \left[\Delta_{i-1}^{(\text{BS}, K)} (F_i - F_{i-1}) - (C_i^{(K)} - C_{i-1}^{(K)}) \right]^2,$$

and where $\Delta_{i-1}^{(\text{BS}, K)}$ is the Black-Scholes delta computed with volatility parameter set to the implied volatility of $C_{i-1}^{(K)}$.

Panels (a), (d) and (g) of Table (3) report average gain values resulting from hedging buckets of options with maturity 1 and 6 months during their entire lifetime. We choose a total of seven hedging periods corresponding to different levels of market stress. Specifically, the selected periods cover the 2008 financial crisis, the 2010 Flash Crash, the Greek debt crisis and the more recent low volatility environment. For every period, the hedging is repeated for different option maturities (1 and 6 months) and re-balancing frequencies (daily and weekly). Except for a very few special cases, non-structural delta-vega hedg-

$\tau = 1$ month, daily rebalancing									
Hedging period	(a) DDVH			(b) IDVH			(c) PBSDVH		
	ITM	OTM	All strikes	ITM	OTM	All strikes	ITM	OTM	All strikes
Jun 08–Dec 08	0.33 (47)	0.70 (56)	0.54 (103)	0.61 (47)	0.80 (56)	0.72 (103)	0.45 (47)	0.53 (56)	0.50 (103)
Mar 10–Sep 10	0.21 (71)	0.50 (27)	0.30 (99)	0.14 (71)	0.68 (27)	0.29 (99)	0.29 (71)	0.26 (27)	0.27 (99)
Sep 11–Mar 12	0.30 (87)	0.41 (32)	0.35 (119)	0.36 (87)	0.54 (32)	0.39 (119)	0.23 (87)	0.28 (32)	0.20 (119)
Jan 13–Jun 13	0.15 (51)	0.23 (19)	0.17 (71)	0.18 (51)	0.67 (19)	0.32 (71)	0.30 (51)	0.42 (19)	0.34 (71)
Jul 13–Dec 13	0.13 (55)	0.00 (28)	0.09 (83)	0.20 (55)	0.60 (28)	0.34 (83)	0.38 (55)	0.44 (28)	0.39 (83)
Jan 14–Jun 14	-0.01 (67)	0.20 (28)	0.05 (95)	0.13 (67)	0.70 (28)	0.29 (95)	0.27 (67)	0.56 (28)	0.35 (95)
Mar 15–Sep 15	0.30 (97)	0.40 (36)	0.33 (132)	0.32 (97)	0.69 (36)	0.42 (132)	0.29 (97)	0.36 (36)	0.30 (132)
Average	0.21 (475)	0.40 (226)	0.28 (702)	0.27 (475)	0.68 (226)	0.40 (702)	0.30 (475)	0.42 (226)	0.33 (702)
$\tau = 6$ months, daily rebalancing									
Hedging period	(d) DDVH			(e) IDVH			(f) PBSDVH		
	ITM	OTM	All strikes	ITM	OTM	All strikes	ITM	OTM	All strikes
Jun 08–Dec 08	0.80 (18)	0.92 (29)	0.87 (47)	0.79 (18)	0.98 (29)	0.90 (47)	0.86 (18)	0.84 (29)	0.84 (47)
Mar 10–Sep 10	0.51 (22)	0.75 (15)	0.61 (37)	0.46 (22)	0.91 (15)	0.64 (37)	0.50 (22)	0.71 (15)	0.59 (37)
Sep 11–Mar 12	0.52 (23)	0.55 (18)	0.54 (41)	0.47 (23)	0.76 (18)	0.60 (41)	0.55 (23)	0.48 (18)	0.52 (41)
Jan 13–Jun 13	0.14 (25)	0.06 (13)	0.11 (38)	0.19 (25)	0.59 (13)	0.33 (38)	0.30 (25)	0.49 (13)	0.36 (38)
Jul 13–Dec 13	0.34 (24)	0.38 (16)	0.36 (40)	0.35 (24)	0.45 (16)	0.39 (40)	0.40 (24)	0.33 (16)	0.38 (40)
Jan 14–Jun 14	0.26 (26)	0.10 (12)	0.21 (38)	0.24 (26)	0.87 (12)	0.44 (38)	0.26 (26)	0.58 (12)	0.36 (38)
Mar 15–Sep 15	0.32 (33)	0.24 (17)	0.29 (50)	0.27 (33)	0.82 (17)	0.46 (50)	0.22 (33)	0.59 (17)	0.35 (50)
Average	0.39 (171)	0.50 (120)	0.43 (291)	0.37 (171)	0.79 (120)	0.54 (291)	0.41 (171)	0.60 (120)	0.49 (291)
$\tau = 6$ months, weekly rebalancing									
Hedging period	(g) DDVH			(h) IDVH			(i) PBSDVH		
	ITM	OTM	All strikes	ITM	OTM	All strikes	ITM	OTM	All strikes
Jun 08–Dec 08	0.42 (18)	0.80 (29)	0.65 (47)	0.71 (18)	0.90 (29)	0.82 (47)	0.79 (18)	0.74 (29)	0.76 (47)
Mar 10–Sep 10	0.44 (22)	0.74 (15)	0.56 (37)	0.54 (22)	0.92 (15)	0.69 (37)	0.61 (22)	0.81 (15)	0.69 (37)
Sep 11–Mar 12	0.41 (23)	0.50 (18)	0.45 (41)	0.56 (23)	0.82 (18)	0.67 (41)	0.67 (23)	-0.48 (18)	0.16 (41)
Jan 13–Jun 13	0.58 (25)	0.61 (13)	0.59 (38)	0.63 (25)	0.79 (13)	0.69 (38)	0.62 (25)	0.86 (13)	0.70 (38)
Jul 13–Dec 13	0.55 (24)	-0.18 (16)	0.26 (40)	0.60 (24)	0.34 (16)	0.50 (40)	0.49 (24)	0.33 (16)	0.42 (40)
Jan 14–Jun 14	-0.09 (26)	0.53 (12)	0.10 (38)	0.26 (26)	0.92 (12)	0.47 (38)	0.19 (26)	0.47 (12)	0.28 (38)
Mar 15–Sep 15	0.47 (33)	0.78 (17)	0.58 (50)	0.64 (33)	0.89 (17)	0.72 (50)	0.74 (33)	0.56 (17)	0.68 (50)
Average	0.40 (171)	0.57 (120)	0.47 (291)	0.56 (171)	0.80 (120)	0.66 (291)	0.58 (171)	0.47 (120)	0.54 (291)

Table 3: Hedging average gains. Each block corresponds to a different combination of length of the hedging period, τ , and re-balancing frequency. Each entry in the table reports the average gain G , as defined in Equation (19), resulting from hedging a fixed maturity bucket (in-the-money, out-the-money, all available strikes) of call options over their entire lifespan. Panels (a), (d), and (g) refer to the direct delta-vega hedging portfolio (DDVH) in (17); Panels (b), (e), and (h) refer to the indirect delta-vega hedging portfolio (IDVH) in (20); Panels (c), (f), and (i) refer to the PBS delta-vega hedging (PBSDVH) in (22). The last row of each panel displays column-wise average values weighted by the number (reported in parentheses) of options contained in each bucket.

ing always provides positive gain compared to delta hedging based on the benchmark PBS model, which can be quantified in a 30 – 50% reduction in the replication error. On average, results are consistent across different re-balancing frequencies but display differences across both moneyness and maturity. Specifically, we notice better hedging performance for OTM than for ITM options, which is in accordance with the empirical analysis carried out in Section 4.1: the volatility risk is more relevant for options in the OTM range of moneyness. As for maturity effects, we notice slightly higher gains for longer maturities, reflecting the positive effect of time on prices sensitivity to volatility. As we discuss below, an interesting clue emerges when relating the hedging performance with the observed level of market volatility.

5.1 Hedging in Low Volatility Markets

Options are instruments used to limit unfavorable but uncertain outcomes and, as such, derive their price from uncertainty. However, in an environment with unprecedented accommodating monetary policy in both Europe and the United States, zero yields and a consistent rally in asset prices, the option market has reacted by producing unusual features in volatility. More precisely, in 2017 the VIX entered single-digit territory, while indicators of economic risk such as the policy uncertainty index of Baker et al. (2016) have signaled high risk during the same period. This phenomenon is known as *low VIX puzzle* see Pastor and Veronesi (2017). In Minsky’s boom-bust theory, periods of subdued volatility lulls investors into a false perception of risk. As a result, they over-leverage themselves increasing the risk in the financial system until the system finally crashes. At a more granular level, when the financial markets experience a big volatility short trade, volatility itself is being put under a heavy supply pressure. The trade comes in various forms: Directly, through inverse-VIX ETFs or by selling volatility in the option markets in an attempt to harvest the implied-realized volatility premium. Indirectly, through investment strategies such as targeted volatility portfolios, where low volatility amplifies leverage. A recent empirical analysis by Sornette et al. (2018) shows that, although volatility cannot be consistently used to forecast crashes, “in two-thirds of the studied

bubbles, the crash follows a period of lower volatility, reminiscent of the idiom of a lull before the storm”.

Our empirical analysis unveils the reduced role of the variance swap as a risk driver in the most recent years. In particular, Panels (a), (d) and (g) of Table 3 reveal higher gains under market turmoil and less pronounced gains in periods characterized by subdued volatility. At first glance, these results comply with the intuition that hedging vega risk is most needed in a turbulent market, whereas it is of lesser importance under tranquil market conditions. In other words, in a low volatility environment, the delta position is expected to hedge a larger portion of the overall risk of an option, making the benchmark replication error smaller. However, while this intuition is in principle correct, the results of Section 4.1 and the discussion presented in this section offer a more troubling conclusion. Using variance swaps as the only measure of uncertainty can lead to dangerous or faulty investment decisions in a low volatility regime. To address this issue an alternative hedging strategy may be devised by adopting an ATM option as a proxy to hedge a portion of market uncertainty. For instance, an “indirect” form of vega hedging can be obtained through an ATM option, denoted as $C_i^{(K^*)}$, which, as shown in Section 4.1, displays negligible dependence on the first risk factor (the futures). This modification is in line with the intuition that a correction of the hedging strategy is needed when dealing with correlated risk-factors, see for instance Daluiso and Morini (2017). In particular, at i -th re-balancing time, the portfolio in (17) can be replaced by

$$C_i^{(K)} + \pi_i^{(K,F)} F_i + \pi_i^{(K,K^*)} C_i^{(K^*)}, \quad (20)$$

where the hedge ratios

$$\pi_i^{(K,F)} = -\Delta_i^{(K)} - \pi_i^{(K,K^*)} \Delta_i^{(K^*)}, \quad \pi_i^{(K,K^*)} = -\frac{\nu_i^{(K)}}{\nu_i^{(K^*)}},$$

can be directly computed through (11). The RSE associated with (20) is

$$\text{RSE}_i^{(K)} = \left[\pi_{i-1}^{(K,F)} (F_i - F_{i-1}) + \pi_{i-1}^{(K,K^*)} (C_i^{(K^*)} - C_{i-1}^{(K^*)}) + (C_i^{(K)} - C_{i-1}^{(K)}) \right]^2. \quad (21)$$

Panels (b), (e), and (h) of Table (3) report average gain values produced by the indirect delta-vega hedging strategy (20). These figures are obtained by replacing (18) with (21) in the gains formula (19). For most periods in our sample, the direct and indirect hedging schemes offer comparable results. But especially in the low volatility regime, the indirect vega hedging appears superior to the direct scheme. The analysis carried out in Section 4.1 identifies the leverage effect as the main culprit for this and unveils the necessity of adjusting the classic notion of volatility hedging, for instance through the indirect approach suggested here. Either way, it is important to take into account the current volatility regime when devising hedging strategies for options. Summarizing, in market conditions characterized by extremely low volatility, it should be preferred to adopt orthogonal hedging factors (such as $F_{t,\tau}$ and ATM options), rather than factors linked by the leverage effect ($F_{t,\tau}$ and $VS_{t,\tau}$), not optimally spanning the full extent of the market risk.

5.2 Vega hedging under PBS

We conclude the empirical analysis by assessing the relative advantage of the adopting a non-structural vega hedging strategy compared to a structural approach based on the Greeks computed through the PBS model. In particular, we augment the PBS delta hedging strategy with the inclusion of the sensitivity to the implied volatility (vega hedging). Unfortunately, a direct vega hedging strategy cannot be implemented since the vega hedge ratios computed through the Black-Scholes model are the sensitivities to the individual implied volatilities for each strike. This makes it not trivial to reconcile these sensitivities with market volatility, measured as the square root of the price of a one-month variance swap contract. Therefore, similarly to the “indirect” form of vega hedging outlined above, we consider the following portfolio at the i -th re-balancing time,

$$C_i^{(K)} + \pi_i^{(\text{BS},K,F)} F_i + \pi_i^{(\text{BS},K,K^*)} C_i^{(K^*)}, \quad (22)$$

with

$$\pi_i^{(\text{BS},K,F)} = -\Delta_i^{(\text{BS},K)} - \pi_i^{(\text{BS},K,K^*)} \Delta_i^{(\text{BS},K^*)}, \quad \pi_i^{(\text{BS},K,K^*)} = -\frac{\nu_i^{(\text{BS},K)}}{\nu_i^{(\text{BS},K^*)}},$$

where $\Delta_i^{(\text{BS},K)}$ and $\nu_i^{(\text{BS},K)}$ denote the Black-Scholes delta and vega, respectively, computed with volatility parameter set to the implied volatility of $C_i^{(K)}$. The RSE associated with (22) is

$$\text{RSE}_i^{(K)} = \left[\pi_{i-1}^{(\text{BS},K,F)} (F_i - F_{i-1}) + \pi_{i-1}^{(\text{BS},K,K^*)} (C_i^{(K^*)} - C_{i-1}^{(K^*)}) + (C_i^{(K)} - C_{i-1}^{(K)}) \right]^2.$$

The RSEs associated with this strategy are summarized in Panels (e), (f), and (i) of Table 3. Not surprisingly, the inclusion of the new indirect vega hedging leads to significant gains over the baseline PBS strategy. These gains are consistent over periods and maturity. However, comparing the performance of the augmented PBS strategy with that of the non-structural indirect delta-vega hedging reveals the relative advantage of avoiding a model-based computation of the hedge ratios. Indeed, the empirical evidence suggests that the indirect delta-vega hedging strategy based on the non-structural hedge ratios leads to higher gains than the corresponding strategy based on PBS Greeks. The economic gain is milder (approx. 5–7% on average) when the re-balancing occurs daily and it is more substantial (approx. 12% on average) when the hedging portfolio is rebalanced less frequently (weekly). This outcome is particularly meaningful to practical applications, where transaction costs are involved and do not allow for an arbitrarily high rebalancing frequency.

6 Conclusions

We propose a non-structural method to compute delta and vega hedges of European options, by only assuming the existence of a risk-neutral density possessing a finite number of moments. Pivotal to the technique is a model-independent relation between the risk-neutral moments, the underlying futures price, and the related variance swaps, which is heuristically derived and shown to empirically hold. Numerical experiments substan-

tiate the validity of the technique as the non-structural hedge ratios provide accurate approximations of the true Greeks implied by a number of different model specifications. Evaluating the performance of delta-vega hedging strategies applied to S&P 500 call options reveals that accounting for vega risk adds an important layer of immunization, especially in periods of market turmoil. However, taking a direct position in variance swaps might not be optimal during an extended period of subdued market volatility. In such periods, alternative “indirect” forms of vega hedging should be preferred. This is in line with the results of our empirical analysis, in which we found that although the underlying volatility is the predominant risk factor in periods of market turmoil, there might be other sources of uncertainty priced in the options, whose importance emerges in periods of low volatility.

Acknowledgments

We are grateful to Leopoldo Catania, Wolfgang Härdle, Elisa Nicolato, Roberto Renò, Paul Schneider, Michael Schatz, Viktor Todorov, Lorenzo Torricelli and Dacheng Xiu for useful suggestions that improved the quality of this work. The authors would also like to thank the participants of the XIX Quantitative Finance Workshop (Rome 2018) and the CENTRAL Workshop: Machine Learning in Economics (Berlin 2018) and the seminar participants at University of Verona, Aarhus University and Humboldt University for helpful comments.

Funding

Andrea Barletta acknowledges support by the People Programme (Marie Curie Actions) of the European Union’s Seventh Framework Programme FP7/2007-2013/ under REA grant agreement #289032 and by the Danish Council for Independent Research under grant agreement DFF – #6109-00056. Paolo Santucci de Magistris acknowledges the research support of CREATES, funded by the Danish National Research Foundation (DNRF78).

Disclosure statement

Declarations of interest: none.

References

- Aït-Sahalia, Y. (2002). Maximum likelihood estimation of discretely sampled diffusions: a closed-form approximation approach. *Econometrica* 70(1), 223–262.
- Aït-Sahalia, Y. and R. Kimmel (2007). Maximum likelihood estimation of stochastic volatility models. *Journal of Financial Economics* 83(2), 413–452.
- Alcock, J. and P. Gray (2005). Dynamic, nonparametric hedging of European style contingent claims using canonical valuation. *Finance Research Letters* 2(1), 41–50.
- Alexander, C., A. Kaeck, and L. M. Nogueira (2009). Model risk adjusted hedge ratios. *Journal of Futures Markets* 29(11), 1021–1049.
- Alexander, C. and L. M. Nogueira (2007). Model-free hedge ratios and scale-invariant models. *Journal of Banking & Finance* 31(6), 1839–1861.
- Andersen, L. B. and V. V. Piterbarg (2007). Moment explosions in stochastic volatility models. *Finance and Stochastics* 11(1), 29–50.
- Badescu, A., R. J. Elliott, and J.-P. Ortega (2014). Quadratic hedging schemes for non-Gaussian GARCH models. *Journal of Economic Dynamics and Control* 42, 13–32.
- Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics* 131, 1593–1636.
- Bakshi, G., C. Cao, and Z. Chen (1997). Empirical performance of alternative option pricing models. *The Journal of Finance* 52(5), 2003–2049.
- Bakshi, G., C. Cao, and Z. Chen (2000). Pricing and hedging long-term options. *Journal of Econometrics* 94(1), 277–318.

- Bakshi, G. and D. Madan (2000). Spanning and derivative-security valuation. *Journal of Financial Economics* 55(2), 205–238.
- Barletta, A. and E. Nicolato (2017). Orthogonal expansions for VIX options under affine jump diffusions. *Quantitative Finance*, 1–18.
- Barletta, A. and P. Santucci de Magistris (2018). Analyzing the risks embedded in option prices with rndfittool. *Risks* 6(2).
- Bates, D. S. (2005). Hedging the smirk. *Finance Research Letters* 2(4), 195–200.
- Branger, N., E. Krautheim, C. Schlag, and N. Seeger (2012). Hedging under model misspecification: All risk factors are equal, but some are more equal than others. . . . *Journal of Futures Markets* 32(5), 397–430.
- Branger, N. and A. Mahayni (2006). Tractable hedging: An implementation of robust hedging strategies. *Journal of Economic Dynamics and Control* 30(11), 1937–1962.
- Breeden, D. T. and R. H. Litzenberger (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621–651.
- Carr, P. and D. Madan (1998). Towards a theory of volatility trading. *Volatility: New estimation techniques for pricing derivatives* (29), 417–427.
- Christoffersen, P., M. Fournier, and K. Jacobs (2018). The factor structure in equity options. *The Review of Financial Studies* 31(2), 595–637.
- Christoffersen, P., S. Heston, and K. Jacobs (2009). The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. *Management Science* 55(12), 1914–1932.
- Christoffersen, P. and K. Jacobs (2004). The importance of the loss function in option valuation. *Journal of Financial Economics* 72(2), 291–318.
- Cont, R. (2006). Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance* 16(3), 519–547.

- Cont, R. and J. Da Fonseca (2002). Dynamics of implied volatility surfaces. *Quantitative Finance* 2(1), 45–60.
- Coutant, S., E. Jondeau, and M. Rockinger (2001). Reading PIBOR futures options smiles: The 1997 snap election. *Journal of Banking & Finance* 25(11), 1957–1987.
- Crépey, S. (2004). Delta-hedging vega risk? *Quantitative Finance* 4(5), 559–579.
- Daluiso, R. and M. Morini (2017). Hedging efficiently under correlation. *Quantitative Finance* 17(10), 1535–1547.
- Davis, M. H. (2016). Model-free methods in valuation and hedging of derivative securities. In *The Handbook of Post Crisis Financial Modeling*, pp. 168–189. Palgrave Macmillan UK.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68(6), 1343–1376.
- Dumas, B., J. Fleming, and R. E. Whaley (1998). Implied volatility functions: Empirical tests. *The Journal of Finance* 53(6), 2059–2106.
- Dupire, B. (1994). Pricing and hedging with smiles. *Risk* 7, 18–20.
- El Karoui, N., M. Jeanblanc-Picquè, and S. E. Shreve (1998). Robustness of the Black and Scholes formula. *Mathematical Finance* 8(2), 93–126.
- Fengler, M. R., W. K. Härdle, and C. Villa (2003). The dynamics of implied volatilities: A common principal components approach. *Review of Derivatives Research* 6(3), 179–202.
- Filipović, D., E. Mayerhofer, and P. Schneider (2013). Density approximations for multivariate affine jump-diffusion processes. *Journal of Econometrics* 176(2), 93–111.
- Gondzio, J., R. Kouwenberg, and T. Vorst (2003). Hedging options under transaction costs and stochastic volatility. *Journal of Economic Dynamics and Control* 27(6), 1045–1068.

- Hansen, P. R. and A. Lunde (2014). Estimating the persistence and the autocorrelation function of a time series that is measured with error. *Econometric Theory* 30(1), 60–93.
- Harris, D. (1997). Principal components analysis of cointegrated time series. *Econometric Theory* 13(4), 529–557.
- Hull, J. and A. White (2017). Optimal delta hedging for options. *Journal of Banking & Finance* 82, 180–190.
- Hutchinson, J. M., A. W. Lo, and T. Poggio (1994). A nonparametric approach to pricing and hedging derivative securities via learning networks. *The Journal of Finance* 49(3), 851–889.
- Jarrow, R. and A. Rudd (1982). Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics* 10(3), 347–369.
- Johansen, S. (1995). *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford University Press, Oxford.
- Jondeau, E., S.-H. Poon, and M. Rockinger (2007). *Financial modeling under non-Gaussian distributions*. Springer Science & Business Media.
- Jondeau, E. and M. Rockinger (2001). Gram–Charlier densities. *Journal of Economic Dynamics and Control* 25(10), 1457–1483.
- Kaeck, A. (2012). Hedging surprises, jumps, and model misspecification: A risk management perspective on hedging S&P 500 options. *Review of Finance* 17(4), 1535–1569.
- Li, C. et al. (2013). Maximum-likelihood estimation for diffusion processes via closed-form density expansions. *The Annals of Statistics* 41(3), 1350–1380.
- Lioui, A. and P. Poncet (2000). The minimum variance hedge ratio under stochastic interest rates. *Management Science* 46(5), 658–668.
- Pastor, L. and P. Veronesi (2017). Explaining the puzzle of high policy uncertainty and low market volatility. *VOX, CEPR portal*, available [here](#).

- Schneider, P. (2015). Generalized risk premia. *Journal of Financial Economics* 116(3), 487 – 504.
- Sepp, A. (2008). VIX option pricing in a jump-diffusion model. *Risk*, 84–89.
- Sornette, D., P. Cauwels, and G. Smilyanov (2018). Can we use volatility to diagnose financial bubbles? Lessons from 40 historical bubbles. *Quantitative Finance and Economics* 3, 1–109.
- Stutzer, M. (1996). A simple nonparametric approach to derivative security valuation. *The Journal of Finance* 51(5), 1633–1652.
- Tebaldi, C. (2005). Hedging using simulation: a least squares approach. *Journal of Economic Dynamics and Control* 29(8), 1287–1312.
- Xiu, D. (2014). Hermite polynomial based expansion of european option prices. *Journal of Econometrics* 179(2), 158–177.

A Recursive Formulas for Orthogonal Polynomials

Throughout the section we denote the s -th moment of ϕ by μ_s , i.e.

$$\mu_s = \int_{\mathcal{D}} z^s \phi(z) dx.$$

The coefficients of lower triangular matrix $w = (w_{s,k})_{s,k=1,\dots,n}$ are uniquely determined (up to a sign in each of its rows) as solution of the following quadratic system

$$\begin{cases} w_{s,k} = 0, & \text{if } k > s, \\ \sum_{k=0}^s \sum_{j=0}^i w_{s,k} w_{i,j} \int_0^\infty z^{k+j} \phi(z) dx = 1, & \text{if } k \leq s \text{ and } s = i, \\ \sum_{k=0}^s \sum_{j=0}^i w_{s,k} w_{i,j} \int_0^\infty z^{k+j} \phi(z) dx = 0, & \text{if } k \leq s \text{ and } s \neq i. \end{cases}$$

In practice, the rows of w can be computed efficiently based on well-known recursive relations between orthogonal polynomials as follows. The first two rows are given explicitly

$$w_{0,0} = 1, \quad w_{1,0} + w_{1,1}z = -\frac{\mu_1}{\sqrt{\mu_2 - \mu_1^2}} + \frac{1}{\sqrt{\mu_2 - \mu_1^2}}z. \quad (\text{A.1})$$

The s -th row, for $s \geq 2$, is determined by the following relation

$$w_{s,0} + w_{s,1}z + \dots + w_{s,s}z^s = \frac{1}{C_s} \left[(x - a_s) \sum_{k=0}^{s-1} w_{s-1,k}z^k - b_s \sum_{k=0}^{s-2} w_{s-2,k}z^k \right],$$

where

$$a_s = \sum_{k=0}^{s-1} \sum_{q=0}^{s-1} \int_{\mathcal{D}} w_{s-1,k} w_{s-1,q} z^{k+q+1} \phi(z) dx, \quad b_s = \sum_{k=0}^{s-1} \sum_{q=0}^{s-2} \int_{\mathcal{D}} w_{s-1,k} w_{s-2,q} z^{k+q+1} \phi(z) dx,$$

and C_s is a normalization constant chosen so that $[\sum_{k=0}^s w_{s,k}z^k]^2$ integrates to 1. At step s , we end up with explicit formulas as follows

$$w_{s,k} = \frac{1}{C_s} \tilde{w}_{s,k}, \quad s = 2, \dots, n, k = 0, \dots, n, \quad (\text{A.2})$$

where

$$\tilde{w}_{i,j} = \begin{cases} -a_s w_{s-1,0} - b_s w_{s-2,0} & \text{if } j = 0 \\ w_{s-1,j-1} - a_s w_{s-1,j} - b_s w_{s-2,j} & \text{if } j = 1, \dots, s-2 \\ w_{s-1,s-2} - a_s w_{s-1,s-1} & \text{if } j = s-1 \\ w_{s-1,s-1} & \text{if } j = s \\ 0 & \text{if } j > s \end{cases}$$

and

$$a_s = \sum_{k=0}^{s-1} \sum_{q=0}^{s-1} w_{s-1,k} w_{s-1,q} \mu_{k+q+1}, \quad b_s = \sum_{k=0}^{s-1} \sum_{q=0}^{s-2} w_{s-1,k} w_{s-2,q} \mu_{k+q+1},$$

$$C_s = \pm \left(\sum_{k=0}^s \sum_{q=0}^s \tilde{w}_{s,k} \tilde{w}_{s,q} \mu_{k+q} \right)^{\frac{1}{2}}.$$

B Relation between Moments and Risk Factors

For a fixed k , the Taylor-expansion of $\log(x)$ around $x = 1$ leads to the following approximation

$$\log\left(\frac{Z_{t+\tau}}{F_{t,\tau}}\right) \approx \sum_{i=0}^k a_i^{(k)} \left(\frac{Z_{t+\tau}}{F_{t,\tau}}\right)^i, \quad (\text{B.3})$$

where the specific form of the coefficients $a_0^{(k)}, \dots, a_k^{(k)}$ is not relevant here, but they are such that

$$a_0^{(k)} + \dots + a_k^{(k)} = \log(1) = 0. \quad (\text{B.4})$$

By taking conditional expectations in (B.3), we get

$$\text{VS}_{t,\tau} = -\frac{2}{\tau} E^{\mathbb{Q}} \left[\log\left(\frac{Z_{t+\tau}}{F_{t,\tau}}\right) \mid \mathcal{F}_t \right] \approx -\frac{2}{\tau} \sum_{i=0}^k a_i^{(k)} \frac{M_{t,\tau}^{(i)}}{F_{t,\tau}}.$$

Now, for $k = 2$,

$$M_{t,\tau}^{(2)} \approx F_{t,\tau}^2 \left(1 + \frac{\tau}{2a_2^{(2)}} \text{VS}_{t,\tau} \right) = F_{t,\tau}^2 (1 + \beta_2 \text{VS}_{t,\tau}). \quad (\text{B.5})$$

Similarly, for $k = 3$,

$$M_{t,\tau}^{(3)} \approx -\frac{2 F_{t,\tau}^3}{\tau a_3^{(3)}} \left[\frac{\tau}{2} \left(a_0^{(3)} + a_1^{(3)} + a_2^{(3)} \frac{M_{t,\tau}^{(2)}}{F_{t,\tau}^2} \right) + \text{VS}_{t,\tau} \right],$$

and, in view of (B.5) and (B.3), we get

$$M_{t,\tau}^{(3)} \approx -\frac{2 F_{t,\tau}^3}{\tau a_3^{(3)}} \left[\frac{\tau}{2} \left(-a_0^{(3)} + a_2^{(3)} \beta_2 \text{VS}_{t,\tau} \right) + \text{VS}_{t,\tau} \right] = F_{t,\tau}^3 (1 + \beta_3 \text{VS}_{t,\tau}).$$

By iterating this argument, we eventually have

$$M_{t,\tau}^{(k)} \approx -\frac{2 F_{t,\tau}^k}{\tau a_k^{(k)}} \left[\frac{\tau}{2} \sum_{i=0}^{k-1} a_i^{(k)} \frac{M_{t,\tau}^{(i)}}{F_{t,\tau}^i} + \text{VS}_{t,\tau} \right] \approx -\frac{2 F_{t,\tau}^k}{\tau a_k^{(k)}} \left[\frac{\tau}{2} \sum_{i=0}^{k-1} a_i^{(k)} (1 + \beta_i \text{VS}_{t,\tau}) + \text{VS}_{t,\tau} \right],$$

and, since $\frac{1}{a_k^{(k)}} \sum_{i=0}^{k-1} a_i^{(k)} = -1$ in view of (B.4), then for some $\beta_k = \beta_k(\tau)$ we have

$$M_{t,\tau}^{(k)} \approx F_{t,\tau}^k (1 + \beta_k \text{VS}_{t,\tau}). \quad (\text{B.6})$$

Finally, we retrieve (10) by using $1 + x \approx e^x$. While the latter approximation might seem unnecessary as one may rely on (B.6) to compute the Greeks, we find (10) a more convenient choice as it proves exact when the underlying process is driven, e.g., by a geometric Brownian motion.

C Factor analysis

In Section 4.1, we applied PCA to assess whether the panel of SPX options is driven by two common components, interpreted as level and volatility factors. We point out that the time series of ITM options display a pervasive non-stationary component. This is confirmed by the Dickey-Fuller test statistic, reported in Panel (a) of Table 4, which cannot reject the null hypothesis of unit root for ten out of seventeen time series of options. In particular, the non-stationary option prices are those in the ITM region with moneyness between 0.78% and 0.975%. For these options, we additionally perform the Johansen cointegration test, finding a strong evidence for a unique non-stationary common stochastic trend since the cointegration rank is $r = 9$, see Panel (b) in Table 4. Following Harris (1997), the PCA delivers consistent estimates of the cointegrating relations among the non-stationary series.

As a further robustness check for the analysis in Section 4.1, we specify a model that explicitly characterizes the dynamic behavior of the latent factors. In particular, we define the following factor model

$$\begin{aligned}\tilde{C}_t &= \Lambda f_t + \varepsilon_t, & t = 1, 2, \dots, T \\ f_{1,t} &= f_{1,t-1} + \eta_{1,t} \\ f_{2,t} &= \psi_1 f_{2,t-1} + \eta_{2,t},\end{aligned}$$

where $\tilde{C}_t$ is a $p \times 1$ vector of standardized call option prices at time t , Λ is a $p \times 2$ vector of loadings and f_t is a 2×1 vector of latent factors at time t . In particular, the first factor is modeled as a random walk, while the second factor is a simple stationary AR(1) process governed by the coefficient ψ_1 .

(a) Unit-root tests				(b) Cointegration tests					
Moneyness	Lags	DF Test	p-value	Rank	Eigenvalue	Trace	p-value	Lmax	p-value
0.780 – 0.840	1	-0.131	0.944	0	0.699	872.12	0.00	152.38	0.00
0.840 – 0.855	0	-0.144	0.941	1	0.661	719.75	0.00	137.58	0.00
0.855 – 0.870	0	-0.144	0.941	2	0.615	582.17	0.00	121.27	0.00
0.870 – 0.885	0	0.037	0.960	3	0.579	460.90	0.00	110.00	0.00
0.885 – 0.900	0	0.068	0.962	4	0.519	350.91	0.00	92.916	0.00
0.900 – 0.915	0	-0.003	0.956	5	0.474	257.99	0.00	81.553	0.00
0.915 – 0.930	0	-0.006	0.956	6	0.425	176.44	0.00	70.314	0.00
0.930 – 0.945	0	-0.330	0.916	7	0.348	106.12	0.00	54.391	0.00
0.945 – 0.960	2	-0.680	0.850	8	0.334	51.733	0.00	51.681	0.00
0.960 – 0.975	1	-2.009	0.283	9	0.001	0.0531	0.82	0.0530	0.82
0.975 – 0.990	0	-4.130	0.001						
0.990 – 1.005	0	-5.027	0.000						
1.005 – 1.020	0	-4.798	0.000						
1.020 – 1.035	0	-4.369	0.000						
1.035 – 1.050	0	-4.071	0.002						
1.050 – 1.065	0	-3.975	0.002						
1.065 – 1.130	2	-3.602	0.006						

Table 4: Unit root and cointegration tests. Panel (a) reports the Augmented Dickey-Fuller test for non-stationarity for the times series of call option prices for 17 different moneyness intervals from 0.78 – 0.84 to 1.065 – 1.13. The optimal number of lags is chosen with the BIC criterion. Panel (b) reports the results of the cointegration rank tests (trace and maximum eigenvalue) of Johansen (1995).

The innovations of the latent factors $\eta = [\eta_{1,t}, \eta_{2,t}]'$ are assumed to be Gaussian with mean 0 and covariance matrix Q , which allows for a non-zero correlation between $f_{1,t}$ and $f_{2,t}$. The model can be cast in state-space form and the Kalman filter routine used to retrieve the log-likelihood function. Hence, the parameters can be estimated by maximum likelihood.

Table 5 reports the estimation results. The estimated loadings have a qualitatively similar interpretation as the weights associated with the PCA. Notably the loadings associated with the OTM options on the first (non-stationary) factor are negative but non significant. Similarly, the loadings for the most ITM options on the second (stationary) factor are non-negative. The estimated AR coefficient is $\psi_1 = 0.7885$, thus signaling a rather persistent but stationary second factor. Notably, the ATM options (0.990 – 1.005) load positively and significantly on the second factor only. The estimated correlation of the innovations to the factors ρ_{η_1, η_2} is strongly negative (-88%), thus directly accommodating the leverage effect. Hence, when looking at the extracted factor dynamics in

Moneyiness	Λ_1	s.e.	Λ_2	s.e.	R^2
0.780 – 0.840	0.1017	0.0513	0.0375	0.0203	0.9538
0.840 – 0.855	0.1022	0.0516	0.0385	0.0208	0.9633
0.855 – 0.870	0.1025	0.0517	0.0406	0.0209	0.9679
0.870 – 0.885	0.1029	0.0519	0.0455	0.0212	0.9722
0.885 – 0.900	0.1031	0.0520	0.0505	0.0223	0.9695
0.900 – 0.915	0.1036	0.0523	0.0581	0.0241	0.9681
0.915 – 0.930	0.1040	0.0524	0.0707	0.0283	0.9659
0.930 – 0.945	0.1041	0.0524	0.0883	0.0363	0.9607
0.945 – 0.960	0.1021	0.0514	0.1167	0.0520	0.9303
0.960 – 0.975	0.0927	0.0468	0.1592	0.0789	0.8275
0.975 – 0.990	0.0656	0.0339	0.2000	0.1092	0.6172
0.990 – 1.005	0.0228	0.0159	0.2068	0.1233	0.4659
1.005 – 1.020	-0.0102	0.0131	0.1896	0.1193	0.4969
1.020 – 1.035	-0.0275	0.0182	0.1758	0.1109	0.5661
1.035 – 1.050	-0.0320	0.0197	0.1739	0.1070	0.6011
1.050 – 1.065	-0.0303	0.0188	0.1761	0.1036	0.5996
1.065 – 1.130	-0.0219	0.0149	0.1718	0.0915	0.5134
Covariance of η_t					
$\sigma_{\eta_1}^2$	4.0737	$\sigma_{\eta_2}^2$	7.5942	ρ_{η_1, η_2}	-0.8858

Table 5: Estimates of the factor model coefficients.

Figure 6: Latent Factors. The left panel displays the SPX futures (solid-red line) and the first (non-stationary) factor (solid-blue line) extracted with the Kalman filter. The right panel displays the square of the VIX index (solid-red line), i.e., the 1-month variance swap, and the second (stationary) factor (solid-blue line) extracted with the Kalman filter.

Figure (6), we notice that the second factor does not deviate from the square-VIX in the second part of the sample. This is the consequence of the fact that the leverage effect is explicitly modeled.

ACCEPTED MANUSCRIPT