

Seasonal Momentum in Option Returns*

Steven L. Heston[†]

Christopher S. Jones[‡]

Mehdi Khorram[§]

Shuaiqi Li[¶]

Haitao Mo^{||}

June 28, 2022

Abstract

This paper develops a new method to calculate hedged returns on model-free “equity VIX” option portfolios. Our returns are highly correlated with realized variance minus implied variance. Compared to CBOE’s VIX formula, our formulas are more accurate for both simulated and actual prices, and they are more highly correlated with realized variance.

We document a new quarterly cross-sectional continuation pattern in both realized variance and implied variance of individual stocks. Implied variance does not fully anticipate the quarterly pattern of realized variance. Consequently, options that performed well at quarterly lags continue to earn high returns in the future. A long-short portfolio based on seasonal momentum achieves an annualized pre-cost Sharpe ratio of 3.31.

Keywords: options, VIX, seasonal momentum

*We are grateful for comments from seminar participants at LSE, Northeastern University, and the Cancun Derivatives Conference. We thank Bjorn Eraker (discussant) for many helpful suggestions. This paper includes results that previously appeared in two working papers. All errors are our own.

[†]Smith School of Business, University of Maryland, sheston@umd.edu

[‡]Marshall School of Business, University of Southern California, christopher.jones@marshall.usc.edu

[§]Ourso College of Business, Louisiana State University, mkhorr2@lsu.edu

[¶]Systemic Risk Centre, London School of Economics, s.li124@lse.ac.uk

^{||}Ourso College of Business, Louisiana State University, haitaomo@lsu.edu

1 Introduction

The literature studying variance risk premia in the options market is divided between two approaches. One strand of the literature focuses on the differences between implied and realized variances, which it uses to approximate the returns on variance swaps, while another examines average returns on option portfolios. Though the continuous-time equivalence of these two approaches was shown by Madan et al. (1998) and Britten-Jones and Neuberger (2000), a full empirical reconciliation is missing from the literature.

We therefore propose a new approach for computing the integral that represents the model-free implied variance derived in these papers. This “VIX integral,” which we approximate using Simpson’s rule, is both the *exact* price of a traded portfolio of options and the approximate price of an untraded variance swap. While we argue that working with option portfolio returns is preferable to synthetic variance swaps, our new approach nevertheless yields option portfolio returns that are more highly correlated with variance swap returns than would be found using the CBOE’s own VIX formula to compute the integral.

Using our approach, we document a striking new finding of seasonal momentum in individual equity options. Specifically, if a firm’s options performed well in lags that are multiples of three or 12 months, then they are more likely to exhibit high returns in the current month. This effect is distinct from the option momentum documented by Heston et al. (2021) and somewhat stronger in terms of portfolio performance. Using the equivalence between option portfolio and variance swap returns, we show that seasonal momentum in options is the result of substantial periodicity in realized variances that exceeds the observed periodic pattern in implied variances.

In theory, computing model-free risk-neutral variance requires a continuum of strike prices ranging from zero to infinity. Most literature has focused on S&P 500 Index options, which conform reasonably well to this assumption, at least for shorter expiration dates. However,

individual equity options are often liquid for a relatively small number of strikes, which makes choices regarding interpolation and extrapolation more consequential. Some papers, such as Bakshi et al. (2003) and Carr and Wu (2009), attempt to avoid the issue by focusing on a small number of stocks with many liquid strikes. Others, like Driessen et al. (2009), interpolate in the space of implied volatilities and make reasonable but arbitrary assumptions about option values beyond the range of observed prices.

While approximation is a common feature of empirical asset pricing, we argue that the use of approximated model-implied prices in computing risk premia is unnecessary. A cleaner approach is to compute exact rates of return on tradable option portfolios, as was popularized by Coval and Shumway (2001), who examined at-the-money straddles. An issue when analyzing this type of portfolio is that its ability to proxy for pure variance risk may be questionable. It is therefore preferable to work with what we term a “VIX portfolio,” which is any portfolio (not just that implied by the CBOE’s own VIX formula) of tradable options that is specifically constructed to capture the return on a variance claim. The challenge is constructing a portfolio that can be implemented given the data limitations inherent to working with individual equity options data.

Our approach is an application of Simpson’s rule to the problem computing the VIX integral and its implied portfolio weights. Simpson’s rule approximates the integrand with a quadratic polynomial, leading to faster convergence as the number of strike prices increases. As a result, we find that risk-neutral variance can be reasonably approximated with as few as three options, often when the CBOE’s formula and the rectangle rule fail badly. In addition, Simpson’s rule implies option portfolios whose returns are more highly correlated with realized variance than corresponding returns from the CBOE or rectangle rules, particularly when we use corridor adjustments from Bondarenko (2014) and Andersen et al. (2015) to correct for truncation effects resulting from the lack of options with very high or low strike prices.

We apply portfolios constructed using Simpson's rule to the analysis of seasonal momentum, or periodicity, in VIX portfolio returns. We find a periodic pattern that closely matches a quarterly pattern of stock momentum documented by Heston and Sadka (2008). It is also similar in that the persistence of seasonal momentum is long, as it retains strong statistical significance at lags up to five years. Without accounting for transactions costs, the annualized Sharpe ratio of a seasonal momentum strategy based on four quarters of lagged returns is 3.31, which exceeds the Sharpe ratios of all other strategies, including momentum and the volatility differential of Goyal and Saretto (2009).

As seasonal momentum is a type of momentum, it is critical to show that the two phenomena are distinct. We demonstrate that past returns at both periodic lags and non-periodic lags are predictive for future option returns. At horizons ranging from one to five years, both types of lags exhibit significant predictive ability, though the relative importance of seasonal versus non-seasonal momentum increases for more distant formation periods. In addition, both regular and seasonal momentum retain significant explanatory power when included together in Fama-MacBeth regressions. Seasonal momentum also survives controls for other common option return predictors and is not driven by the quarterly earnings cycle.

Building on the close connection between VIX portfolio and variance swap returns, we show that seasonal momentum is the result of strong periodicity in realized variance. When realized variance is regressed on its own lags, periodic lags are of particular importance, and including them separately from non-periodic lags increases forecast accuracy significantly. When the implied variance is also included in the regression, it subsumes almost all of the predictive power of past realized variances, except not the realized variances at periodic lags. Thus, implied variances appear to anticipate future realized variance on average, but they do not appear to fully anticipate the seasonal patterns.

In the next section, we discuss different ways of operationalizing the VIX formula and demonstrate the advantages of using Simpson's rule. Section 3 discusses the data we use

in our empirical work. In Section 4, we examine alternative VIX integration rules in the data, while Section 5 demonstrates seasonal momentum in VIX returns. Finally, Section 6 concludes.

2 Returns on model-free “VIX” portfolios

2.1 Hedged VIX returns

The most prominent published benchmarks for option prices are the Chicago Board of Options Exchange VIX index for S&P 500 Index options (CBOE, 2019¹) and the corresponding equity-VIX indices for options on individual stocks.² These VIX indices are based on portfolios of options, weighted by the squared reciprocals of their strike prices. Carr and Wu (2009) interpolated option prices to measure an idealized continuous VIX portfolio, and then used a continuous-time variance swap approximation to the returns on their portfolio. Although it is not strictly tradable given the approximated swap rate, the variance swap approach has an intuitive advantage of decomposing returns into risk-neutral variance and realized variance. We construct returns on a discrete daily-hedged analog of the continuous variance swap option strategy. This method provides a tradable strategy, while preserving the intuition of the variance swap decomposition.

The CBOE (2019) VIX index is based on the market value of a portfolio at time t comprising options expiring at time T .

$$V(t; T) = 2 \sum_i \frac{O(K_i, t; T) \Delta_i}{K_i^2}, \quad (1)$$

where $O(K, t; T)$ represents time t price of an out-of-the-money call or put option with

¹CBOE White Paper used to construct the VIX index can be found at: www.cboe.com/micro/vix/vixwhite.pdf

²Equity-VIX indices are available for just a few large firms, such as Apple: www.cboe.com/us/indices/dashboard/vxapl

strike price K and expiration T , and Δ_i represents the gap between adjacent strike prices.³ Importantly, VIX portfolios are "model-free" because their construction does not depend on any model parameters. Madan et al. (1998) showed that we can approximate the VIX price with a continuous integral over strike prices.⁴

$$\hat{V}(t; T) = 2 \int_0^\infty \frac{O(K, t; T)}{K^2} dK. \quad (2)$$

Given the spot price $S(T)$ at expiration, the option payoff $O(K, T; T)$ is equal to $\max(S(T) - K, 0)$ for a call option and $\max(K - S(T), 0)$ for a put option. In the absence of intermediate dividends, integrating these option payoffs over strike prices (2) shows the terminal payoff of the idealized VIX portfolio.

$$\hat{V}(T; T) = -2 \log \left(\frac{S(T)}{S(t)(1 + r_f)^{T-t}} \right) + 2 \left(\frac{S(T)}{S(t)(1 + r_f)^{T-t}} - 1 \right), \quad (3)$$

where r_f is the daily risk-free interest rate. The first term in the payoff (3) represents selling two units of the "log-portfolio". The second term represents a costless static hedge that leverages (the present value of) two dollars of stock at time t , and holds this hedge position constant until expiration at time T . The combined payoff is a U-shaped function of the stock price, resembling a squared stock return. Therefore, the price of this portfolio represents the approximate (risk-neutral) variance of return. Since the S&P 500 VIX index and equity-VIX indices on individual stocks represent standard deviation, they are proportional to the square-root of the portfolio value $V(t; T)$.

Due to their U-shaped payoffs, the equity-VIX portfolios have (approximately) zero delta when they are constructed. In other words, they are locally insensitive to movements in the

³The sum uses out-of-the-money options with respect to the forward value of the strike price, $K(1+r_f)^{T-t}$.

⁴See also Demeterfi et al. (1999), Britten-Jones and Neuberger (2000), and Jiang and Tian (2005) for various continuous-time derivations. Breeden and Litzenberger (1978) first expressed the risk-neutral density in terms of the second derivative of the option price with respect to the strike price. Madan et al. (1998) then derived the formula (2) using integration by parts twice.

underlying stock price. But over time, the stock price will drift away from the bottom of the U-shaped payoff, and the equity-VIX portfolios will become sensitive to the stock price. Instead of using a fixed static hedge, we can further reduce risk of the VIX portfolios by dynamic hedging. This replaces the second term of (3) with the gains and losses of a stock position that is rebalanced daily and designed to optimally hedge the log portfolio.

The elasticity of option value with respect to the stock price generally depends on a model. But due to the log-payoff (3), the delta of the idealized continuous-strike VIX portfolio does not. As Bondarenko (2014) shows, the delta-hedge of the log-portfolio buys $1/F(t)$ shares of stock for a price of $S(t)$, and rebalances to maintain a constant hedge exposure of one dollar. So, not only is the value of the VIX portfolio model-free, but its dynamic-hedge is also model-free. This dynamic hedge keeps the delta of the discrete equity-VIX approximately equal to zero, and reduces the volatility of returns (relative to using a fixed static hedge). The dynamically hedged payoff is

$$V_{hedged}(T; T) = -2 \log \left(\frac{S(T)}{S(t)(1+r_f)^{T-t}} \right) + 2 \sum_{u=t+1}^T (r_S(u) - r_f), \quad (4)$$

where $r_S(u)$ represents the stock return on day u . We can replace the stock price in equation (4) to express the $V_{hedged}(T; T)$ payoff in terms of a telescoping series of daily stock returns $r_S(u)$ at times u between t and T :

$$V_{hedged}(T; T) = -2 \sum_{u=t+1}^T \log \left(\frac{1+r_S(u)}{1+r_f} \right) + 2 \sum_{u=t+1}^T (r_S(u) - r_f). \quad (5)$$

When daily returns on the stock and risk-free rate are small, a second-order Taylor series expansion shows that the dynamically hedged option portfolio (5) approximates the payoff

of variance swap contract in Carr and Wu (2009):

$$V_{hedged}(T; T) \approx \sum_{u=t+1}^T (r_S(u) - r_f)^2. \quad (6)$$

The return on the unhedged VIX portfolio from equation (1) is simply the proportional change in its value

$$r_{unhedged}(t; T) = \frac{V(T; T) - V(t; T)}{V(t; T)}. \quad (7)$$

The return on the dynamically hedged VIX portfolio is adjusted by the difference between the static hedge term in (3) and the dynamic risk term in (4).

$$r_{hedged}(t; T) = \frac{V(T; T) - V(t; T) - 2 \left(\frac{S(T)}{S(t)(1+r_f)^{T-t}} - 1 - \sum_{u=t+1}^T (r_S(u) - r_f) \right)}{V(t; T)}. \quad (8)$$

A comparison of the hedged return (8) with the Taylor series approximation (6) shows that the dynamically hedged return on the VIX portfolio is approximately the realized variance relative to the VIX portfolio price.

$$r_{hedged}(t; T) \approx \frac{\sum_{u=t+1}^T (r_S(u) - r_f)^2}{V(t; T)} - 1. \quad (9)$$

Carr and Wu (2009) used the variance swap approximation to analyze variance premiums in the cross-section of option returns, and Bollerslev et al. (2009) used it implicitly when forecasting returns. Dew-Becker et al. (2017) later applied it to multiple asset classes. In unreported diagnostics, we found that the exact return on the underlying S&P 500 Index VIX portfolio is 99% correlated with the variance swap approximation (6). In other words, the dynamically hedged payoff on the index VIX portfolio is very close to the realized variance over the month. But with individual stocks, the correlations of options returns with realized

variance can be lower. Using returns on hedged option portfolios (8) is consistent with previous research that measured delta-hedged returns, while preserving compatibility with the variance swap literature (9).

An additional advantage of our benchmark approach is that it measures option portfolios with all available strike prices. These portfolios maintain consistent sensitivity to volatility because they always include at-the-money options. In a certain sense, a VIX portfolio is always at-the-money. In contrast, Bakshi and Kapadia (2003) approach of delta-hedging a single option will generally change its vega sensitivity when the option drifts away from the money.

2.2 Computing variance indices with discrete strike prices

The calculations above use a finite set of options with discrete strike prices to approximate infinite portfolio with a continuum of strike prices. In practice, the number of strikes may be quite limited for options on individual equities, particularly for longer maturities or if a requirement of positive open interest is imposed.

In this section we derive rectangle, CBOE, and Simpson's methods to compute variance indices. We provide a portfolio interpretation of each index given our aim of computing exact returns.

Let VIX^2 denote a general variance index for horizon T . Then $e^{-rT}VIX^2$ represents the price of an option portfolio, where r is the interest rate. Following Madan et al. (1998), a VIX^2 portfolio should emulate the continuous U-shaped payoff of a hedged log-portfolio:

$$\frac{2}{T} \int_0^\infty \frac{|S_T - K|}{K^2} dK = -2 \log(S_T/F) + 2(S_T/F - 1), \quad (10)$$

where the forward price F equals $e^{rT}S_0$.

We interpret VIX^2 portfolios as linear integration rules. Rectangle rule quadrature

approximates the value of the integral (10) with a summation over option strike prices K_i , or

$$VIX_{rectangle}^2 = \frac{2}{T} e^{rT} \left(\sum_{K_i < F} \frac{Put(K_i)}{K_i^2} \Delta_i + \sum_{K_i > F} \frac{Call(K_i)}{K_i^2} \Delta_i \right), \quad (11)$$

where $\Delta_i = (K_{i+1} - K_{i-1})/2$. The rectangle equation (11) is obviously a portfolio, because it is just a weighted sum of out-of-the-money options (puts on the left of F , calls on the right). If the strike prices have equal spacing Δ and the option functions are smooth, then the rectangle rule converges with order $O(\Delta^2)$.

The CBOE VIX formula is⁵

$$\begin{aligned} VIX_{CBOE}^2 = & \frac{2}{T} e^{rT} \left(\sum_{K_i < K_0} \frac{Put(K_i)}{K_i^2} \Delta_i + \sum_{K_i > K_0} \frac{Call(K_i)}{K_i^2} \Delta_i \right) \\ & + \frac{2}{T} e^{rT} \frac{Put(K_0) + Call(K_0)}{2K_0^2} \Delta_0 - \frac{1}{T} \left(\frac{F}{K_0} - 1 \right)^2, \end{aligned} \quad (12)$$

where K_0 is the closest strike price below-the-money. The CBOE formula (12) uses an average of put and call options with strike price K_0 , and subtracts a quadratic adjustment term. The quadratic adjustment term is the forward value of a payment at expiration equal to $\frac{F-K_0}{TK_0^2}(S_T - K_0)$. It represents the value of a short position of $\frac{F-K_0}{TK_0^2}$ shares, and a long position in a bond paying $\frac{F-K_0}{TK_0^2} K_0$. Therefore, we can interpret the CBOE formula as a portfolio of options, including a small adjustment of stock and bonds.

The CBOE method (12) resembles the trapezoid rule, but it is an inexact application when strike prices are unequally spaced. It treats K_0 and the next higher strike price K_1 asymmetrically, and it gives different answers for currency options expressed in reciprocal numeraire. Finally, it converges at $O(\Delta^2)$, no faster than the rectangle rule.

To improve convergence, we define Simpson's method as an equally weighted average of

⁵See equation (1) of <https://cdn.cboe.com/resources/vix/vixwhite.pdf>.

the rectangle rule, the trapezoid rule centered around K_0 , and the trapezoid rule centered around K_1 . This agrees with the conventional definition of Simpson's composite integration rule in the case of equally spaced strike prices, but it also allows application when strike prices are not equally spaced. This approximation can be adjusted for incorrectly integrating options between K_0 and K_1 by subtracting the extra bowtie-shaped area between K_0 and K_1 ,

$$\frac{1}{T}e^{rT} \int_{K_0}^{K_1} \frac{|Call(K) - Put(K)|}{K^2} dK = \frac{1}{T} \left(\log\left(\frac{K_0}{F}\right) + \frac{F}{K_0} + \log\left(\frac{K_1}{F}\right) + \frac{F}{K_1} - 2 \right). \quad (13)$$

The resulting formula is:

$$\begin{aligned} VIX_{Simpson's}^2 &= \frac{2}{T}e^{rT} \left(\sum_{K_i < K_0} \frac{Put(K_i)}{K_i^2} \Delta_i + \sum_{K_i > K_1} \frac{Call(K_i)}{K_i^2} \Delta_i \right) \\ &+ \frac{2}{T}e^{rT} \times \left(\frac{1}{6} \left(\frac{Call(K_0)}{K_0^2} + \frac{Put(K_1)}{K_1^2} \right) (K_1 - K_0) + \frac{5}{6} \left(\frac{Put(K_0)}{K_0^2} \Delta_0 + \frac{Call(K_1)}{K_1^2} \Delta_1 \right) \right) \\ &- \frac{1}{T} \left(\log\left(\frac{K_0}{F}\right) + \frac{F}{K_0} + \log\left(\frac{K_1}{F}\right) + \frac{F}{K_1} - 2 \right). \end{aligned} \quad (14)$$

Simpson's formula (14) uses one-sixth weight on in-the-money options $Call(K_0)$ and $Put(K_1)$, and five-sixths weight on out-of-the-money options $Put(K_0)$ and $Call(K_1)$. To compute returns, the last line of (14) represents the value of a short position in $\frac{1}{T}(\frac{1}{K_0} - \frac{2}{F} + \frac{1}{K_1})$ shares of stock, and future payment of $\frac{1}{T}(\log(\frac{K_0}{F}) + \log(\frac{K_1}{F}))$ dollars. Simpson's formula gives identical values using currency options from either currency perspective. For smooth option functions (four times differentiable), Simpson's formula converges at $O(\Delta^4)$.

We can accelerate the convergence of Simpson's method by estimating the integration error. The local truncation error of the Simpson's rule applied to a smooth function $f(K)$ is $-\frac{2}{T}e^{rT} \frac{\partial^4 f(K)}{\partial K^4} \frac{\Delta^5}{180}$, where Δ is the strike price interval. If the strike prices are equally spaced, then integrating this truncation error and applying put-call parity gives an error estimate

that is model-independent.

$$-\frac{2\Delta^4}{180T}e^{rT} \left(\int_{-\infty}^F \frac{\partial^4}{\partial K^4} \left(\frac{Put(K)}{K^2} \right) dK + \int_F^{\infty} \frac{\partial^4}{\partial K^4} \left(\frac{Call(K)}{K^2} \right) dK \right) = \frac{\Delta^4}{5TF^4}. \quad (15)$$

The resulting corrected-Simpson's formula is

$$VIX_{corrected-Simpson's}^2 = VIX_{Simpson's}^2 - \frac{\Delta^4}{5TF^4}. \quad (16)$$

This integration correction (15) should be very small. If the strike price spacing Δ is 10% of the forward price F , then the relative magnitude should be only $\frac{.00002}{T}$. After subtracting this error (15), the corrected Simpson's formula (14) converges at an astonishingly fast $O(\Delta^5)$. For example, when the strike price intervals shrink by half, the asymptotic error shrinks by a factor of 32.

Finally, there is a slight discrepancy among the different variance indices, because the rectangle rule uses only out-of-the-money options, but the CBOE formula additionally uses $Call(K_0)$, and Simpson's rule uses both $Call(K_0)$ and $Put(K_1)$. We can eliminate these extra options by using put-call parity. In that case, the CBOE formula (12) simplifies to:

$$VIX_{CBOE}^2 = VIX_{rectangle}^2 + \frac{1}{T} \frac{F - K_0}{K_0^2} (\Delta_0 - F + K_0). \quad (17)$$

The simplified CBOE formula (17) is the forward value of a portfolio of out-of-the-money options, plus the forward value of $\frac{\Delta_0 - F + K_0}{TK_0^2} (S_T - K_0)$.

Put-call parity also simplifies Simpson's formula (14):

$$VIX_{Simpson's}^2 = VIX_{rectangle}^2 + \frac{e^{rT}}{3T} \left(\frac{K_1 - K_0 - \Delta_0}{K_0^2} Put(K_0) + \frac{K_1 - K_0 - \Delta_1}{K_1^2} Call(K_1) \right) + \frac{K_1 - K_0}{3T} \left(\frac{F - K_0}{K_0^2} + \frac{K_1 - F}{K_1^2} \right) + \frac{1}{T} \left(\log\left(\frac{F}{K_0}\right) - \frac{F}{K_0} + \log\left(\frac{F}{K_1}\right) - \frac{F}{K_1} + 2 \right). \quad (18)$$

When the strike prices K_{-1} , K_0 , K_1 , and K_2 are equally spaced, then

$K_1 - K_0 - \Delta_0 = K_1 - K_0 - \Delta_1 = 0$, and the second term on the right side of the simplified formula (18) vanishes. To compute returns, the second line of the simplified Simpson's formula (18) represents the forward value of

$(\frac{K_1 - K_0}{3T}(\frac{1}{K_0^2} - \frac{1}{K_1^2}) + \frac{1}{T}(\frac{2}{F} - \frac{1}{K_0} - \frac{1}{K_1}))S_T + \frac{K_1 - K_0}{3T}(\frac{1}{K_1} - \frac{1}{K_0}) + \frac{1}{T}(\log(F/K_0) + \log(F/K_1))$. We use these simplified formulas (17) and (18) in empirical work because they have lower data requirements and use a common set of out-of-the-money option prices.

2.3 Numerical accuracy

In this section, we perform an initial examination of the accuracy of the different integration schemes using artificial data. Later, we will examine the performance of the different schemes in our actual sample.

The setup we consider here is simple. Suppose that options trade on a grid of strikes that are equidistant and separated by some constant Δ . Options with strikes that are below the current spot price are puts, while options with strikes that are higher than the spot price are calls. Thus, all options are OTM. Let K_0 denote the strike of the put closest to ATM, so that $K_1 = K_0 + \Delta$ is the strike of the call closest to ATM. All option prices are generated using the Black-Scholes formula assuming one period to expiration, a spot price of \$1, and a risk-free rate of zero.

The goal of our analysis is to determine the accuracy of the VIX formulas in Section 2.2 for different values of Δ , K_0 , and volatility. Specifically, how quickly do VIX estimates converge to the true underlying volatility as $\Delta \rightarrow 0$, and how is this convergence affected by different choices of K_0 and volatility?

To understand the importance of K_0 , consider the CBOE formula (17). If K_0 is just below the forward price (which is the same as the spot price in this setting), then $F \approx K_0$

and $VIX_{CBOE}^2 \approx VIX_{rectangle}^2$. In contrast, if K_0 is below the forward price by $\Delta/2$, then

$$VIX_{CBOE}^2 = VIX_{rectangle}^2 + \frac{1}{T} \frac{\Delta^2}{4K_0^2} > VIX_{rectangle}^2. \quad (19)$$

This is just one example of how the choice of K_0 matters, but it illustrates the need to check for sensitivity to its choice.

Figure 1 shows the results of these experiments. Each panel shows, using a logarithmic scale, the absolute proportional error in the implied volatility, or $|\text{Implied}/\text{Actual} - 1|$, for the CBOE, rectangle, Simpson's, and corrected-Simpson's rules. Values along the horizontal axis denote different values for Δ , which are all of the form 2^{-n} , where $n \in \{2, 3, 4, 5, 6\}$. Numbers along the top of each panel represent the number of options that have a non-negligible impact on the VIX calculation, meaning that the value of the option price, divided by the square of the strike price, is greater than 0.0001. Panels differ with respect to the values of K_0 and volatility that are assumed.

Panel A shows the case in which the value of K_0 from the coarsest discretization, in which $\Delta = .25$, is equal to 0.77.⁶ This is approaching the lowest value that K_0 can take in this case. If it were below 0.75, then there would be another OTM put that was closer to ATM, and that put's strike would be chosen as K_0 . The volatility assumed in this panel is 10%. If we interpret each period as one month, then this value corresponds to that of a typical stock.

In this case, we see that the rectangle rule generally outperforms the CBOE's rule. Simpson's rule has smaller errors for all but the coarsest grid of strikes, and the corrected Simpson's rule performs better still. The panel also confirms that the Simpson's rule converges to the true variance at a higher rate than either the rectangle or CBOE rules. The corrected Simpson's rule converges at an even faster rate.

⁶As $\Delta \rightarrow 0$, K_0 and K_1 will both converge to the forward price. The value of K_0 in the title of each panel refers only to the value of K_0 when $\Delta = .25$.

Results are much the same for Panel B, in which K_0 is near its maximum value in the coarsest discretization. The main difference is that Simpson's rule, in this case, outperforms the rectangle rule for even the coarsest discretization. In Panel C, in which the closest-to-ATM call and put are equidistant from the forward price, errors in the CBOE and rectangle formulas are similarly large, but the Simpson's formula is better still. Even when the distance between strikes (Δ) is 25% of the spot price, the error in the Simpson's implied variance is only around 1% of the true variance.

Finally, Panel D, which is otherwise similar to Panel A, shows the effects of an increase in underlying volatility. In this case all methods produce smaller errors, but the rankings are unaffected. The reason for the improved performance is that the larger volatility increases the range of strikes over which the VIX integrand is non-negligible. Effectively, in this setup at least, a higher volatility results in a larger number of options that convey information about volatility. Whether this effect is apparent in actual data, in which options of many strikes do not trade, is less clear.

2.4 The corridor fix

In addition to the assumption of a continuum of strikes, the model-free implied variance (2) assumes the existence of strikes over $(0, \infty)$. In practice, strikes are generally found in a narrower range that brackets the current spot price. The result of this *truncation* is that the integral in (2), when computed over a narrower range, will be downward biased. This bias is not addressed by the use of Simpson's rule, which only improves the interpolation between available strikes.

Truncation does not present a major problem for the calculation of option portfolio returns. It does create some issues, such as the possibility that the portfolio return may not provide a perfect replication of the hypothetical variance swap return. It is also likely that the returns on the ATM options that are available differ from the OTM options that are

not. The option portfolio's return nevertheless represents the rate of appreciation on a set of investable assets.

In contrast, truncation can have a highly deleterious effect on the hypothetical variance swap returns that the literature often studies. The downward bias in the VIX that is due to truncation will cause a corresponding upward bias in variance swap returns. While the issue has long been understood in the literature, many resolutions are unsatisfactory. Some researchers, such as Bakshi et al. (2003) and Carr and Wu (2009) restrict their studies to a small number of stocks with a wide range of strikes. Others, like Driessen et al. (2009), extrapolate beyond the range of available strikes, effectively guessing the prices of options that do not exist.

Fortunately, Andersen et al. (2015) provide an alternative approach that avoids extrapolation and that can be applied even with a narrow range of strikes. While their primary motivation was to separate the impacts of volatility and jump risk, their method is easily applied for the correction of truncation bias. They show that the *corridor implied variance*, defined as

$$\hat{V}(t; T) = 2 \int_{K_L}^{K_H} \frac{O(K, t; T)}{K^2} dK, \quad (20)$$

represents the fair valuation of a contract that pays

$$\sum_{u=t+1}^T x(u), \quad (21)$$

where

$$x(t) = 2 \left(\frac{F(t)}{\bar{F}(t)} \left(\frac{\bar{F}(t) - \bar{F}(t-1)}{\bar{F}(t-1)} \right) - \ln \frac{\bar{F}(t)}{\bar{F}(t-1)} \right) \quad (22)$$

and

$$\bar{F}(t) = \begin{cases} K_L & F(t) < K_L \\ F(t) & K_L \leq F(t) \leq K_H \\ K_H & F(t) > K_H. \end{cases} \quad (23)$$

Andersen et al. (2015) refer to $x(t)$ as the “corridor-truncated version of the squared weighted return,” as its formula closely approximates that description.

In our results below, we use the corridor approach in several ways. First, in addition to using the various VIX calculations in Section 2.2 to compute hypothetical variance swap returns, which will be downward biased as the result of truncation bias, we use the same VIX measures to compute corridor variance swap returns, which use the same denominator but change the numerator from the sum of squared daily returns to (21).

Secondly, following Proposition 4 of Bondarenko (2014), we use a corridor approach for computing the model-free hedge for the VIX portfolio. Specifically, Section 2.1 showed that the model-free hedge for the VIX portfolio was $1/F(t)$. Following Bondarenko (2014), the hedge of the truncated portfolio, with minimum and maximum strike prices of K_L and K_H , respectively, is $1/\bar{F}(t)$. Below, we compare the performance of both hedges with each other and with the Black-Scholes model.

In both cases, the values of K_L is equal to the lowest strike price available minus one half of the distance between that strike price and the one above it. Similarly, K_H is the highest strike plus one half the distance from the next highest. The rationale is that our discrete integrals, whether based on the CBOE, rectangle, or Simpson’s rule, are approximated using rectangles or trapezoids that are centered at each strike price and that therefore include some mass below the lowest strike and above the highest.

2.5 A simulation analysis

In this section, we perform a further examination of the accuracy of different integration schemes based on our actual sample, whose construction is specified in Section 3. The goal is to evaluate the accuracy of the four VIX formulas relative to the true underlying corridor implied variance under realistic market conditions.

We simulate option prices that match the actual sample closely. The underlying stock prices, option strike prices, option maturities, and risk-free rates are the same as those in the data. Following Aschakulporn and Zhang (2022), option prices are generated by Gram-Charlier expansion using time-varying distributional parameters. The volatility for each stock in each month is the average of the implied volatilities of all options in the stock's VIX portfolio. Risk neutral skewness and kurtosis are computed following Bakshi et al. (2003).

The median firm in our simulation has five strikes, matching the sample we describe below. Because these strike prices are the same as those in our actual sample, they are not necessarily equidistant. Furthermore, because the lowest and highest strike prices may not be far enough from the forward price, the values we compute must be interpreted as corridor implied variances.

We compute model-free corridor implied variances using the CBOE, rectangle, Simpson's, and corrected Simpson's rules and compare them to true values, which are calculated by interpolating a 500-point grid of strike prices within the corridor. We measure the accuracy using absolute percentage error:

$$\left| \frac{VIX^2 - VIX_{True}^2}{VIX_{True}^2} \right|.$$

Table 1 reports simulation results: Panel A presents statistics on the percentage errors, while panel B describes the simulation parameters. All statistics are from a sample that

pools all stocks over all dates.

Panel A shows that the Simpson's and corrected Simpson's methods clearly dominate, followed by the rectangle rule and then the CBOE rule. The corrected Simpson's method generates the lowest average error of 1.24%, but the uncorrected Simpson's method is nearly as accurate. In contrast, the mean for the CBOE method is 12.79%, with a standard deviation of 13%. The rectangle rule improves on CBOE substantially with an average error of 4.70% and a standard deviation of 9.07%, but does not match the accuracy of Simpson's.

The medians are below means in all cases, implying that the distribution of errors is right skewed. If we focus on the 90th percentiles, the outperformance of the Simpson's and corrected Simpson's methods is even more obvious, with 90th percentile errors that are less than one tenth of corresponding value for the CBOE rule and one quarter that of the rectangle rule. These differences are economically consequential.

We conclude that the outperformance of the Simpson's rule relative to the CBOE and rectangle rules in realistic data is consistent with our Black-Scholes exercise in Section 2.3. In contrast to that exercise, however, we find relatively little incremental gain from using the corrected Simpson's method. This is likely due to the uncorrected rule already being accurate, and it may also result from the assumption of equally spaced strikes, which was assumed in the derivation of the Simpson's rule correction.⁷ The results below are therefore based on the somewhat simpler uncorrected Simpson's rule.

3 Data

This paper uses data from the OptionMetrics Ivy DB database from January 1996 to November 2020. These data provide daily closing bid and ask quotes for U.S. equity options. We use the T-bill rate of appropriate maturity (interpolated when necessary) from OptionMetrics

⁷When implementing the corrected Simpson's rule following (16), we assume that the Δ , the distance between strike prices, is equal to the difference between the strikes immediately below and above the forward price.

as the risk-free rate. Finally, we obtain information about stock returns, dividends, and firm characteristics from CRSP and Compustat.

We apply a series of filters to our primary sample. Following Driessen et al. (2009), we remove all observations for which the option open interest is equal to zero, in order to eliminate options with no liquidity. We discard options with zero bid prices, and with missing implied volatility or delta (which occurs for options with nonstandard settlement or for options with intrinsic value above the current mid price). We delete all observations whose ask price is lower than the bid price, and eliminate options whose prices violate arbitrage bounds. We also exclude firm-month observations if the underlying stock price is less than \$5 on the formation date or if the stock has a split or pays a dividend during the remaining life of the option. Thus, the early exercise premium is small and the seasonality is not caused by dividends. We require at least three OTM options, including one put and one call, for each stock, and we construct VIX portfolio prices using bid-ask midpoints. We only keep observations with positive VIX prices calculated with option bid prices to avoid very small VIX. Our final sample includes 221,157 firm-month observations with 1,464,062 option contracts. On average, each equity VIX portfolio consists of 6.62 option contracts.

We use the primary sample to compute holding period returns, and the relatively stringent requirements on it serve to ensure that these returns are valid and not overly affected by microstructure biases or other issues. However, requirements such as positive open interest and nonzero bid prices reduce the number of available option contracts, which in a number of cases makes it impossible to compute VIX portfolio returns, thereby reducing our sample size.

We therefore relax the positivity constraints on open interest and bid prices in a secondary sample, which we use solely for computing returns during the formation period. While such returns may not be as meaningful, they are nevertheless valid as signals, and any noise or bias that results from the use of less reliable data should only bias our results against significant

predictability of future returns.

The official CBOE VIX methodology combines options with different expiration dates to achieve a 30-day weighted-average maturity. Our analysis uncovers temporal periodicity in option prices and returns. To measure returns accurately, we must calculate portfolio values without interpolating option prices across different maturities. Therefore, we establish option positions in equity-VIX portfolios on the third Friday of a month. When this date is a holiday, the portfolio formation is one day before. This avoids interpolation by using exact option prices instead of 30-day weighted-averages used in the CBOE's (2019) original VIX methodology.

4 Alternative VIX rules

4.1 Convergence of different integration rules

In Section 2.3, we investigated the accuracy of the different integration rules using several simple Black-Scholes examples. Here, we attempt to assess the convergence of the three VIX formulas in the actual data.

One indication of a problem with the CBOE formula is the fact that it produces negative implied variances, which happens 20 times in our sample. This never occurs for Simpson's rule or the rectangle rule, though it is a theoretical possibility for Simpson's rule. And while the frequency of negative implied variances for the CBOE rule is extremely small, their presence indicates a more fundamental problem with the CBOE formula that leads to inaccuracy in a much larger number of cases, as our results below demonstrate.

Because we are not able to measure the true risk-neutral volatility, it is impossible to say whether any of the methods considered converge to true values. We can nevertheless examine the value of having a larger number of option contracts, and how that value differs across methods, by comparing VIX prices constructed from all available contracts to those

resulting from just a subset.

This is the experiment we report in Table 2. We begin by constructing a sample of firm/months in which there are at least seven OTM options available, with at least three calls and three puts. From these options, we compute a VIX price following each of the three integration rules described above.

Ordering the options from lowest to highest strike price, we then discard the options in even positions, keeping only the odds. The resulting subset therefore contains half or slightly more than half of the original set. We recompute VIX prices from these subsets and compare the values with those of the full set.

Table 2 shows summary statistics on

$$\left| \frac{VIX_{Half}^2 - VIX_{All}^2}{VIX_{All}^2} \right|,$$

where smaller values indicate that VIX prices based on the subset are more similar to those based on the full set. A smaller value is therefore consistent with an integration scheme that converges more quickly. We show statistics for all firms with at least seven strikes and for samples that include firms with exactly seven, 11, or 15 strikes.

The table shows, uniformly, that the Simpson's method dominates the other two. The rectangle rule follows, and the CBOE rule is last. To understand the results, consider the "7 Strikes" row from Panel C. This row reports that the average absolute proportional deviation is just 4.82%. This implies that in options with seven strikes, using just four of them is sufficient to get within a few percentage points of the value obtained with all seven. In contrast, the corresponding mean for the CBOE method is 21.76%, indicating that deviations are more than four times greater, on average, using the CBOE rule. The rectangle rule improves on CBOE substantially, but does not match the performance of Simpson.

The table shows that medians are in all cases below means, implying that the distribution

of absolute deviations is right skewed. Looking at the 90th percentiles, we see that the CBOE rule performs very poorly in some cases. The Simpson's approach is not perfect, the number of substantial failures is comparatively low.

When there are many strikes, Simpson's rule continues to outperform, with a mean less than one third of CBOE and about 60% of the rectangle rule. These differences are large enough to be economically consequential. We conclude that the performance of the Simpson's rule in actual data is consistent with its performance in our Black-Scholes exercise. Even with a small number of options available, our results suggest that VIX prices may be estimated reasonably well.

4.2 Alternative VIX returns

In this section we compare the various ways in which the return on a variance claim can be calculated. There are several dimensions to the differences between approaches. Returns can be computed exactly using the observed prices and payoffs of options, or they can be approximated using theoretical variance swap prices.

Among the approaches that rely on option returns directly, variation arises from different integration schemes, which put different weights on the option contracts included in the VIX portfolio. Option portfolio returns are also affected by whether or not the portfolio is delta-hedged and, if so, whether the a model-free or Black-Scholes hedge ratio is used.

Among variance swap returns, the integration scheme also affects the calculations. Returns will also be affected by whether or not a corridor approach is used to correct for the truncation bias that results from the range of strikes being too narrow to accurately approximate the integral (2).

In fact, as we will see below, the integration scheme is even more important for variance swap returns than it is for option portfolios. With option portfolios, the integration scheme changes the portfolio weights. But regardless of those weights, the portfolio return calculated

is the feasible return on some strategy. In other words, the payoff and the price of the strategy are consistent.

In contrast, integration errors will affect the variance swap price, but the payoff of the variance swap (or the corridor variance swap) is unaffected by these errors. As a result, the payoff and the price are inconsistent, and the variance swap return that is computed does not represent the true return on any tradable strategy.

We begin the analysis by an examination of stock-level returns on different option strategies. Table 3 shows the returns on six different option strategies separately for the CBOE, rectangle, and Simpson methods. The first strategy is an unhedged option portfolio. The next three are dynamically hedged portfolios, where the hedge ratios considered include Black-Scholes, the model-free hedge, and the model-free hedge that incorporates the corridor effect. Finally, we include variance swap returns with and without the corridor adjustment. The simple statistics shown use a pooled sample containing all firms on all dates.

A few results stand out. First, there are relatively minor differences between the returns on different option portfolios, though average returns based on the CBOE integration rule are somewhat less negative than those of the rectangle and Simpson rules. Hedging also makes option portfolio returns slightly more negative on average, which is the result of short-term reversal in stock returns.

While different hedging approaches are similar in terms of their means, the dispersion in hedged returns varies significantly depending on the hedge ratio that is used. In particular, the model-free hedge without a corridor adjustment has a much larger standard deviation than the Black-Scholes hedge or the model-free hedge that makes the corridor adjustment. This indicates that truncation is a significant issue in the data.

While the different integration rules are similar for option portfolio returns, they differ to a much greater degree for variance swap returns. For the CBOE rule, the variance swap return is similar, in terms of means, to a hedged option portfolio, but its standard deviation

is much higher, indicating that these two returns are not as close in practice as they are in theory. Interpreting the VIX as the price of a corridor variance swap, we then replace the numerator of the variance swap return with the payoff of the corridor contract. This results in a standard deviation that is in line with the more accurately hedged option portfolios. However, the mean of the corridor variance swap return is much lower. Altogether, these results suggest that there is no close equivalence between option portfolio and variance swap returns when the CBOE rule is used.

For the rectangle and Simpson's rules, the variance swap returns are on average positive, which is inconsistent with the negative volatility risk premium that is more expected. They are also substantially more volatile than option portfolios hedged using Black-Scholes or the model-free corridor approach. The corridor swap is again lower, consistent with a significant truncation bias, and standard deviations are substantially lower as well. In addition, for both the rectangle and Simpson rules, returns on the corridor variance swap match those of the hedged option portfolios (except for the model-free with no corridor fix) closely, consistent with theory.

For comparison, Panel D of Table 3 includes an alternative integration rule that is outside the framework we have discussed. In this approach, we follow Carr and Wu (2009) and linearly interpolate the implied volatility curve between available strikes. These interpolated values are then converted into a 500-point grid of option prices that are used to perform the integration in (2), except that the integration range is restricted to the lowest and highest available strike.⁸ Since we do not extrapolate beyond the range of available strikes, as do many papers such as Driessen et al. (2009), we interpret the result as a corridor variance swap price and compute the return on that strategy following Andersen et al. (2015).

Because we do not study this interpolation scheme in detail, we view the results in Panel

⁸Because the integration here relies on a fine grid of interpolated implied volatilities, there is no need for K_L and K_H to be lower or higher than the minimum and maximum strikes, as discussed in Section 2.4.

D with caution. Nevertheless, it is interesting to note that they are reasonably close matches for the corridor variance swaps constructed using the rectangle and CBOE rules. They are also similar to several hedged option portfolio returns, including those that use the Black-Scholes and corridor model-free hedges for the rectangle and Simpson rules.

To better establish a strong empirical link between option portfolio returns and variance swap returns, Table 4 analyzes correlations between different returns. In each month of our sample, we compute the correlation between two different return measures, and the table reports statistics on the time series of these correlation measures. As before, we report values for each of the three integration methods.

Overall, correlations are high for all three methods, indicating that option portfolio and variance swap returns largely capture the same risks. However, these correlations are higher for Simpson's rule than they are for the CBOE and rectangle rules. This is particularly true when the option portfolio uses a model-free hedge.

Taking Tables 3 and 4 together, it appears that Simpson's rule implies option portfolio and variance swap returns that are more consistent with each other than are the returns implied by other methods.

5 Seasonal Momentum

5.1 Univariate analysis

Equation (9) shows that the gross return on an equity-VIX portfolio for the i^{th} stock over month t is approximately the realized variance, $RV_i(t)$, divided by the cost of the equity-VIX portfolio $VIX_i^2(t-1)$. In logarithms, this relationship is

$$\log(1 + r_i(t)) \approx \log(RV_i(t)) - \log(VIX_i^2(t-1)). \quad (24)$$

This return decomposition allows for greater visibility into the sources of option return predictability, as we will demonstrate.

Heston et al. (2021) show that option returns display momentum. That is, stocks whose options had relatively high returns in the past 2-36 months tend to have options with high returns going forward. The effect is strong, robust, and unexplained by other predictors or exposures to common factors. It is strong in delta-hedged VIX portfolios or simple at-the-money straddles.

The close relationship between VIX and variance swap returns implies that predictability in equity-VIX returns reflects predictability in realized variance relative to equity-VIX prices. To diagnose the sources of momentum profits, we first run the cross-sectional regression

$$\log(RV_i(t)) = \gamma_{0,t} + \gamma_{k,t} \cdot \log(RV_i(t - k)) + \epsilon_i(t). \quad (25)$$

The coefficient estimate $\gamma_{k,t}$ shows the extent to which the cross-section of realized variance in one month is predicted by the previous cross-section lagged by k -months. The average of $\gamma_{k,t}$ over all months t shows the average relationship. Figure 2 (a) shows that the cross-section of realized variance is persistent, with coefficients exceeding 0.6 for short monthly lags, and declining as lags grow to five years. Figure 2 (a) also displays the corresponding average coefficients for the analogous regression of the cross-section of logarithms of equity-VIX prices $\log(VIX_i^2(t))$ on their own lags. Option prices are even more persistent than realized variance, with average coefficients exceeding 0.8 for short monthly lags, and remaining above 0.5 even for lags of five years.

What is perhaps more striking is that both the realized variance and VIX^2 coefficients exhibit clear periodicity. Coefficients peak at multiples of three lags, indicating a quarterly seasonal pattern across stocks, even at multi-year lags.

The similarity of the seasonal patterns in realized variance and VIX^2 coefficients would

seem to suggest that option prices anticipate the pattern in future variance. Were this true, the log variance swap return $\log(RV_i(t)) - \log(VIX_i^2(t-1))$ might not show any seasonal effects. Figure 2 (b) shows that this is not the case, however. Variance swap coefficients show the same quarterly peaks as their components, suggesting that option prices fail to properly anticipate the periodicity of realized variance. While the persistence of the variance swap return is lower, it remains positive, which is consistent with Heston et al. (2021) finding of momentum.

Finally, Figure 2 (c) examines dynamically hedged returns on the VIX portfolios, where we use model-free hedge ratios with the corridor adjustment. While coefficient estimates are not as clearly downward sloping, quarterly peaks are still evident, which is consistent with the close connection between variance swap and option portfolio returns.

An alternative measurement of seasonal momentum is provided by portfolio sorts, which are examined in Table 5. In Panel A, we analyze formation periods that only include lags that are multiples of three or 12, which respectively capture quarterly and annual seasonal momentum. In some cases, formation periods either include just a single lag, while in other cases they include lags as much as three years before the holding period. All of these formation periods result in positive and highly significant return spreads, though quarterly seasonal momentum is somewhat stronger.

Panel B examines formation periods in which all lags, starting with lag 2, and including lags out to 12 or 36 months.⁹ Return spreads for these strategies are similarly large and significant.

Panel C shows the results of quintile sorts on other variables examined in the empirical options literature. One is the difference between implied and historical volatilities, shown by Goyal and Saretto (2009) to forecast future option returns. Another is the amount of

⁹Similar to the stock momentum literature, we exclude lag 1 because Heston et al. (2021) find some evidence that it is related to short-term reversal rather than momentum.

idiosyncratic volatility in the underlying stock, as defined by Cao and Han (2013). Sorting by market cap of the underlying firm also generates a spread in straddle returns, as demonstrated first by Cao et al. (2021). From Vasquez (2017), the slope of the term structure of at-the-money implied volatilities is the fourth measure. The final measure is the slope of the implied volatility curve (the “smirk”) from one-month options. This is related to the skewness variable examined by Bali and Murray (2013).

Table 6 shows long-short portfolios formed on the basis of different characteristics. The seasonal momentum strategy based on lags 3, 6, 9, and 12 has a monthly Sharpe ratio of 0.956, or 3.31 annualized. While this value does not account for transactions costs, it is nevertheless surprisingly large. The seasonal momentum strategy based on lags $\{3, 6, \dots, 36\}$ also does well, with a monthly Sharpe ratio of 0.739.

Among other strategies, the “2 to 12” momentum portfolios is most competitive, offering a monthly Sharpe ratio of 0.927. The IV-HV difference examined by Goyal and Saretto (2009) has a slightly lower Sharpe ratio, at 0.732. Among these top performers, skewness and maximum drawdown are moderate, at least relative to a strategy that shorts the SPX VIX portfolio or takes an equally weighted short position in equity VIX portfolios. Thus, there is no evidence that the profitability of seasonal momentum can be explained by its exposure to downside risk.

5.2 Seasonal vs. non-seasonal momentum

We perform a number of different analyses to establish that seasonal momentum is distinct from the standard momentum strategy examined by Heston et al. (2021).

Table 7 reports results from portfolio sorts. Each entry in the table is the average return on a high-minus-low portfolio formed on the basis of past returns over some formation period. In all cases, the holding period is one month, so differences between the table are driven entirely by alternative formation periods (and the changes to the sample that result from

using different formation periods).

The first column of the table, labeled “Lags 1 to 12,” shows the results of different formation periods that are subsets of these 12 months. The “All” row includes all 12, while other rows include only months that are in or outside of the same quarterly cycle. The “Annual” row would only include lag 12, while the “Non-annual” row includes lags 1 to 11. Finally, the “Quarterly non-annual” includes the quarterly lags excluding the 12th.

Other columns of the table are similar but are based on longer lags. For example, the “Quarterly” value for the “Lags 25 to 36” column is based on those months, between lags 25 to 36 that are on the same quarterly cycle, namely 27, 30, 33, and 36. The “Annual” value for “Lags 49 to 60” only includes the 60th lag in the formation period.

Incredibly, every formation period shown in the table results in a positive and statistically significant return spread, though results are stronger for more recent formation periods. Throughout the table, the largest spreads are either for the “quarterly” or “quarterly non-annual” strategies. The strong performance of the latter suggests that annual lags (12, 24, ...), while capable of producing return spreads by themselves, are not particularly important relative to other lags included in the quarterly strategy.

The fact that months in and outside of the same quarterly cycle both predict future returns suggests that seasonal and non-seasonal momentum each play a role in predicting future returns. It is possible, however, that the informativeness of one of these measures somehow proxies for that in the other. We examine this possibility using Fama-MacBeth regressions that include seasonal and non-seasonal momentum measures. We also use these regressions to gauge whether annual seasonal patterns offers predictability that is distinct from quarterly seasonality.

Table 8 reports results of regressions with three different explanatory variables. One is quarterly seasonality, one is annual, and one is a standard momentum measure. The panels of the table differ with respect to the length of the formation period. Panel A shows results

in which the formation period includes some or all of the first 12 lags, while Panel B includes the first 36 lags. Panel C uses up to lag 36 to compute seasonality measures but only up to 12 lags to compute momentum.

Panel A shows that each predictor is significant individually. Quarterly seasonality and momentum are also significant in multiple regressions, implying that seasonality and momentum convey distinct information. Annual seasonality remains significant when controlling for quarterly seasonality and/or momentum, but it is clearly the least important of the three regressors. Panel B extends the formation periods as far as lag 36, both for seasonality and momentum measures. While annual seasonality weakens in this case, quarterly seasonality and momentum remain distinct. Finally, Panel C uses the longer formation period for seasonality but the shorter one for momentum. We report these results as a check given the evidence in Heston et al. (2021) and in Table 5 that a shorter momentum formation period may work better.

Overall 8 shows that seasonality and momentum signals contain different information. Between quarterly and annual seasonality, only the former appears to be robust.

Table 9 examines whether the predictive power of seasonality is subsumed by other option predictors. Overall, the table shows that seasonality is distinct from other predictors, retaining its statistical significance when other controls are included. Seasonal momentum is stronger when a longer formation period is used, and including seasonality into a regression with other predictors has little effect on the other coefficients. The exception is momentum, whose importance diminishes when seasonality is included. In one case, momentum becomes insignificant, though this is generally not the case. We conclude that seasonality provides predictive information about future option returns that is distinct from that reflected in momentum or in other common predictors.

5.3 Robustness

In this section, we ask whether results are robust over time and across different types of stocks. We examine the relationship between seasonal momentum and other quarterly cycles in the stock and options markets, such as earnings announcements. Next, we investigate the relationship between seasonal momentum and various proxies for limits to arbitrage. Finally, we ask whether the profitability of seasonal momentum is itself seasonal.

Dubinsky et al. (2019) show that quarterly earnings announcements have large effects on option implied volatility and are associated with a rise in jump risk premia. Gao et al. (2018) also show higher straddle returns in the days around earnings announcements. Since these announcements occur on a regular quarterly cycle, it is natural to think that they might be related to our finding of quarterly seasonal momentum.

We examine this possibility in Table 10. Panel A shows the results of sorts based on seasonal momentum in which the sample consists of firms with earnings announcements in the holding period and in each month of the formation period. Panel B examines a sample in which there are no earnings months in the holding or formation periods. Consistent with the fact that earnings announcements are quarterly, the average number of firms used in Panel A is about half that of Panel B.

Comparing Panels A and B, we see strong evidence of seasonal momentum in both samples, with a moderately larger and more significant high-low spread for firms not making an earnings announcement. This result holds for two different formation periods. Thus, our results cannot be driven by systematic differences between announcement and non-announcement months.

Another quarterly cycle that affects the options market is the cycle of option expiration dates. While there are options expiring in each calendar month, long-dated options are only available with expiration dates in either the first, second, or third month of each quarter.

Stocks differ with respect to which of the three months their long-dated options expire in, which is randomly assigned and persistent over time.

By the time we include them in our sample, all options have one month until expiration. However, those options that are in their firm's expiration cycle have been listed on option exchanges for much longer than options in other months that are not in that cycle. Possibly for this reason, options that are in the firm's quarterly expiration cycle typically have higher volume and open interest.

Panels C and D compare options expiring in the quarterly cycle to those expiring outside of it. In short, the differences in the high-low returns of the two groups are minor. The expiration cycle is not a significant driver of seasonal momentum.

In Table 11, we examine the relationship between seasonal momentum and various proxies for limits to arbitrage. We perform a sequential double sort, first sorting into terciles based on some firm characteristic and then into terciles based on lagged returns over some formation period. The proxies for limits to arbitrage include firm size, stock illiquidity, option illiquidity, and analyst coverage. Size measures the stock's most recent market equity capitalization, while our proxy for stock illiquidity is the average Amihud (2002) measure over the last year. Option illiquidity is computed as the average percentage bid-ask spread of the puts and calls in each VIX portfolio, also also averaged over the most recent 12 months.

In Table 11, each value in the rows labeled "Low," "Medium," and "High" shows the average return of a portfolio that buys options with high seasonal momentum and sells options with low seasonal momentum. The row labels refer to the levels of the four proxies for limits to arbitrage. Finally, the "High - Low" row shows the difference between two different long-short option portfolios.

The table shows that seasonal momentum remains profitable in each tercile based on each limits to arbitrage proxy. This holds for either definition of seasonal momentum. Thus, seasonal momentum is pervasive in the cross-section of VIX portfolios.

The table also shows a consistent pattern of larger return spreads when limits to arbitrage are stronger (smaller size, more illiquidity, less analyst coverage). However, these differences are only statistically significant for three proxies when the shorter formation period is used. Results using the longer formation period are consistent but statistically marginal.

Next, we examine the profitability of seasonal momentum over the 12 months of the year. Figure 3 plots the average high-low returns for each calendar month, where the two panels examine different formation periods. Vertical lines depict 95% confidence intervals.

Overall, it seems that seasonal momentum is robust over the calendar year. With a formation period that includes just the four most recent periodic lags, the average high-low return is significant in each of the 12 months. Using the most recent 12 periodic lags results in some loss of significance in March and August, but there is no obvious pattern that would suggest this is anything other than sampling error.

Our final robustness check is to analyze the profitability of seasonal momentum over different subsamples. Figure 4 plots five-year moving averages of the rate of return on various momentum strategies. Of primary interest is seasonal momentum, though standard and non-seasonal momentum are included for comparison. Panel A examines formation periods based on the first 12 lags of returns, while Panel B considers the first 36 lags.

In both panels, each variety of momentum is positive in each five-year subsample. And while the figure does not show confidence intervals, the seasonal and standard momentum strategy returns are statistically significant, at the 5% level, in each five-year subsample. Non-quarterly momentum is statistically significant in most subsamples, but not all.

Overall, while the profitability of seasonal momentum was slightly lower in the second half of our sample, it is remarkably stable and remains high even at the end of our sample.

6 Seasonal patterns in realized and implied variances

As discussed in Section 5.1, Figure 2 suggests that option prices fail to anticipate the periodicity in realized variance. In this section, we formally test this hypothesis and quantify the size of this failure.

We begin by estimating forecast regressions that are full or restricted versions of the cross sectional regression

$$\ln RCV_i(t) = \alpha + \beta \frac{1}{K/3} \sum_{\substack{k \in \{3, \\ 6, \dots, K\}}} \ln RCV_i(t-k) + \delta \frac{1}{K} \sum_{\substack{k \in \{1, \\ 2, \dots, K\}}} \ln RCV_i(t-k) + \gamma \ln VIX_i(t-1) + \epsilon_i(t), \quad (26)$$

where RCV denotes realized corridor variance and where K is either 12 or 36. The first two regressors represent past average variances during the same month of the quarterly cycle and over all months. The last regressor is the implied variance corresponding to the period over which $RCV_i(t)$ is computed.

Table 12 shows results from these regressions, where Panel A uses a maximum of 12 lags ($K = 12$) and Panel B uses a maximum of 36. The first regressions include only the δ terms and therefore only capture the persistence of volatility over time. Coefficients are moderately below one, indicating reversion to the mean.

The next regression in each panel adds the β terms to measure the strength of quarterly periodicity in realized variances. In both panels, but particularly in panel B, the seasonal average offers significant incremental predictive power relative to the all month average.

The next regression in each panel omits the seasonal average and adds the log implied variance, $\ln VIX_i(t-1)$, observed immediately prior to the period over which the dependent variable $\ln RCV_i(t)$ is calculated. The purpose of these regressions is to show that most of the explanatory power of past realized variances is subsumed by $\ln VIX_i(t-1)$, as evidenced by the fact that the estimate of δ is much smaller than the same coefficient in the first regression

of each panel, which omitted $\ln VIX_i(t - 1)$.

The fourth regression in each panel includes all three explanatory variables. While results in Panel A are weak, Panel B shows clearly that implied variance does not subsume the predictive power of the seasonal average. Option prices therefore do not fully anticipate the periodicity in realized variance.

The fourth regression in each panel provides an alternative way to make this point. These regressions make log implied variance the dependent variable and regress it only on the two realized variance measures. The timing of the dependent variable, $\ln VIX_i(t - 1)$, means that it represents the implied volatility immediately after the period over which the two explanatory variables are computed but immediately prior to the period over which the dependent variable in the earlier regressions was computed.

These regressions show that implied variances are also sensitive to quarterly periodicity in realized variances, but the degree of that sensitivity is too low. In Panel A, for example, actual realized log variance loads on the seasonal variance measure with a coefficient (0.1526) that is almost twice as large as the log implied variance's loading on the same variable (0.0790). We show that these differences are significant in the final regressions in each panel, which examine the difference between log realized and log implied variances.

Thus, there is periodicity in implied variances, as shown in Figure 2 (a), and it is at least partially aligned with the periodicity in realized variance. But the magnitude of the periodicity in implied variance is either too weak, or it is not aligned enough with the periodicity in realized variance to eliminate seasonal momentum in option returns.

7 Conclusions

While the empirical evidence on the variance risk premium in index options is comprehensive and convincing, the methods used in gathering that evidence are difficult to apply for indi-

vidual equities, for which the number of traded options is often limited. Standard methods for computing variance swap rates, which rely on a continuum of strikes, suffer from bias and noise resulting from both interpolation and extrapolation of observed option prices. As a result, the variance swap “return” does not actually represent the rate of return on any portfolio of traded assets, and its validity in estimating variance risk premia is questionable.

In contrast, we propose to estimate variance risk premia from returns on observed option contracts, held in a model-free portfolio that is constructed to replicate the hypothetical return on a variance swap as closely as possible. The portfolio weights are computed using Simpson’s rule, which we show in artificial and actual data to be more accurate than the CBOE’s VIX formula or a simple rectangle rule. Specifically, relative to other methods, Simpson’s rule requires fewer options to obtain an accurate measure of the variance swap rate, and the resulting portfolio returns are more highly correlated with contemporaneous realized variance.

Using this approach, we show that variance risk premia display seasonal momentum at the quarterly frequency. This result is highly robust to the length of the formation period and is apparent in a wide variety of subsamples formed on firm size, liquidity, and analyst coverage. Seasonal momentum is robust to controls for other option predictors (including momentum), is unrelated to the quarterly earnings cycle, and offers returns that are large relative to their variance and tail risk.

Our work highlights some of the advantages of working with options data when attempting to explain the root cause of seasonal momentum. In the equity literature, the lack of a high-frequency cash flow proxy makes it difficult to distinguish between explanations of momentum based on biased expectations and those involving time-varying discount rates. With our options methodology, however, the realized variance of future stock returns is effectively the “cash flow” of our constructed option portfolio. We find a strong periodicity in realized variance that is not fully matched by similar periodicity in implied variance. That

is, the option market appears to suffer from biased expectations by failing to adequately anticipate seasonal volatility patterns into option prices. Understanding the reasons for this failure and connecting them to seasonal patterns in stock markets (Heston and Sadka 2008, Keloharju et al. 2016) represents an interesting area for future work.

References

- Amihud, Yakov (2002), “Illiquidity and stock returns: Cross-section and time-series effects.” *Journal of Financial Markets*, 5, 31–56.
- Andersen, Torben G., Oleg Bondarenko, and Maria T. Gonzalez-Perez (2015), “Exploring return dynamics via corridor implied volatility.” *Review of Financial Studies*, 28, 2902–2945.
- Aschakulporn, Pakorn and Jin E Zhang (2022), “Bakshi, kapadia, and madan (2003) risk-neutral moment estimators: A gram–charlier density approach.” *Review of Derivatives Research*, 1–49.
- Bakshi, Gurdip and Nikunj Kapadia (2003), “Delta-hedged gains and the negative market volatility risk premium.” *Review of Financial Studies*, 16, 527–566.
- Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan (2003), “Stock return characteristics, skew laws, and the differential pricing of individual equity options.” *Review of Financial Studies*, 16, 101–143.
- Bali, Turan G. and Scott Murray (2013), “Does risk-neutral skewness predict the cross-section of equity option portfolio returns?” *Journal of Financial and Quantitative Analysis*, 48, 1145–1171.
- Bollerslev, Tim, George Tauchen, and Hao Zhou (2009), “Expected stock returns and variance risk premia.” *Review of Financial Studies*, 22, 4463–4492.
- Bondarenko, Oleg (2014), “Variance trading and market price of variance risk.” *Journal of Econometrics*, 180, 81–97.
- Breeden, Douglas T. and Robert H. Litzenberger (1978), “Prices of state-contingent claims implicit in option prices.” *Journal of Business*, 621–651.
- Britten-Jones, Mark and Anthony Neuberger (2000), “Option prices, implied price processes, and stochastic volatility.” *Journal of Finance*, 55, 839–866.
- Cao, Jie and Bing Han (2013), “Cross section of option returns and idiosyncratic stock volatility.” *Journal of Financial Economics*, 108, 231–249.
- Cao, Jie, Bing Han, Qing Tong, and Xintong Zhan (2021), “Option return predictability.” *Review of Financial Studies*, forthcoming.
- Carr, Peter and Liuren Wu (2009), “Variance risk premiums.” *Review of Financial Studies*, 22, 1311–1341.
- Coval, Joshua D and Tyler Shumway (2001), “Expected option returns.” *The journal of Finance*, 56, 983–1009.
- Demeterfi, Kresimir, Emanuel Derman, Michael Kamal, and Joseph Zou (1999), “A guide to volatility and variance swaps.” *Journal of Derivatives*, 6, 9–32.

- Dew-Becker, Ian, Stefano Giglio, and Bryan Kelly (2017), “How do investors perceive the risks from macroeconomic and financial uncertainty? Evidence from 19 option markets.” Technical report, Working paper.
- Driessen, Joost, Pascal J. Maenhout, and Grigory Vilkov (2009), “The price of correlation risk: Evidence from equity options.” *Journal of Finance*, 64, 1377–1406.
- Dubinsky, Andrew, Michael Johannes, Andreas Kaeck, and Norman J. Seeger (2019), “Option pricing of earnings announcement risks.” *Review of Financial Studies*, 32, 646–687.
- Gao, Chao, Yuhang Xing, and Xiaoyan Zhang (2018), “Anticipating uncertainty: straddles around earnings announcements.” *Journal of Financial and Quantitative Analysis*, 53, 2587–2617.
- Goyal, Amit and Alessio Saretto (2009), “Cross-section of option returns and volatility.” *Journal of Financial Economics*, 94, 310–326.
- Heston, Steven L., Christopher S. Jones, Mehdi Khorram, Shuaiqi Li, and Haitao Mo (2021), “Option momentum.” *Journal of Finance*, forthcoming.
- Heston, Steven L. and Ronnie Sadka (2008), “Seasonality in the cross-section of stock returns.” *Journal of Financial Economics*, 87, 418–445.
- Jiang, George J. and Yisong S. Tian (2005), “The model-free implied volatility and its information content.” *Review of Financial Studies*, 18, 1305–1342.
- Keloharju, Matti, Juhani T. Linnainmaa, and Peter Nyberg (2016), “Return seasonalities.” *Journal of Finance*, 71, 1557–1590.
- Madan, Dilip B., Peter P. Carr, and Eric C. Chang (1998), “The variance gamma process and option pricing.” *Review of Finance*, 2, 79–105.
- Vasquez, Aurelio (2017), “Equity volatility term structures and the cross section of option returns.” *Journal of Financial and Quantitative Analysis*, 52, 2727–2754.

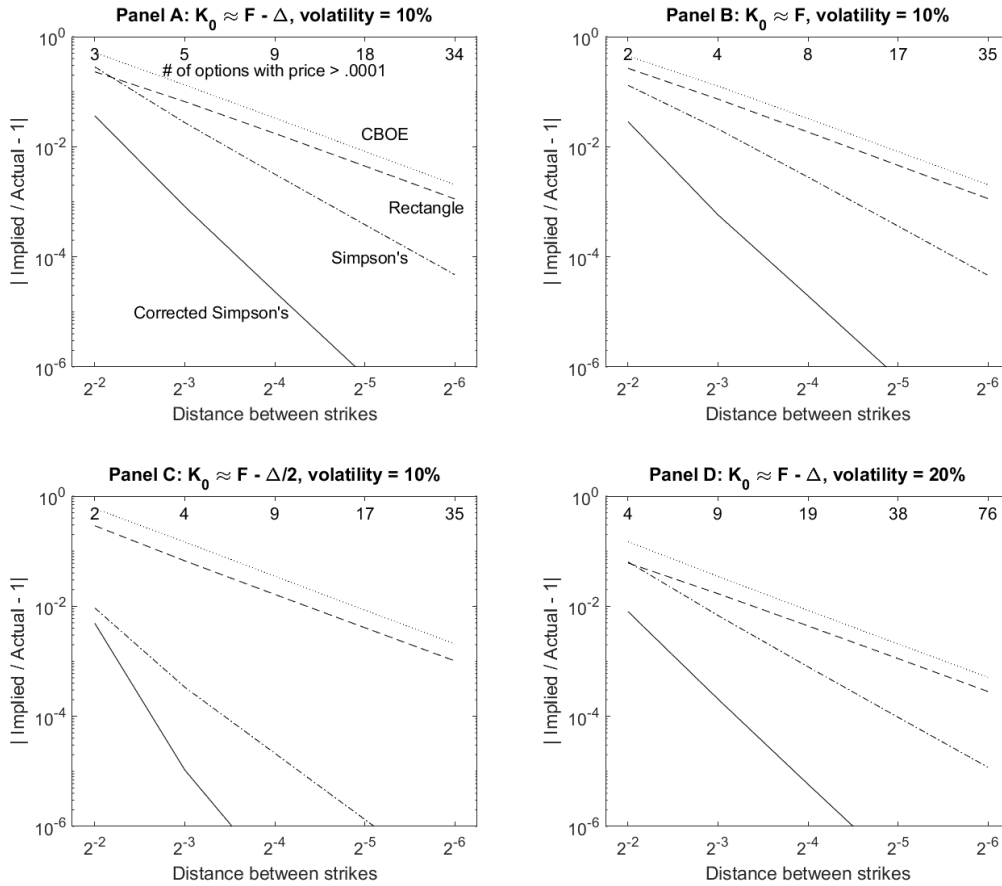


Figure 1.

This figure shows how discretization in strike prices affects VIX prices calculated using the CBOE, rectangle, and Simpson's rules. For each calculation, we assume a set of strike prices separated by Δ , and each panel examines how the choice of Δ affects the accuracy of the VIX calculation. Given strike prices, we compute option prices using the Black-Scholes model using a forward price (F) of \$1 and a volatility that is given in the title of each panel. The panel titles also give the value of K_0 , which is the highest strike price below 1 for the coarsest discretization in which $\Delta = .25$. For each of the three integration methods, the figure shows the absolute error in implied volatility as a proportion of actual volatility.

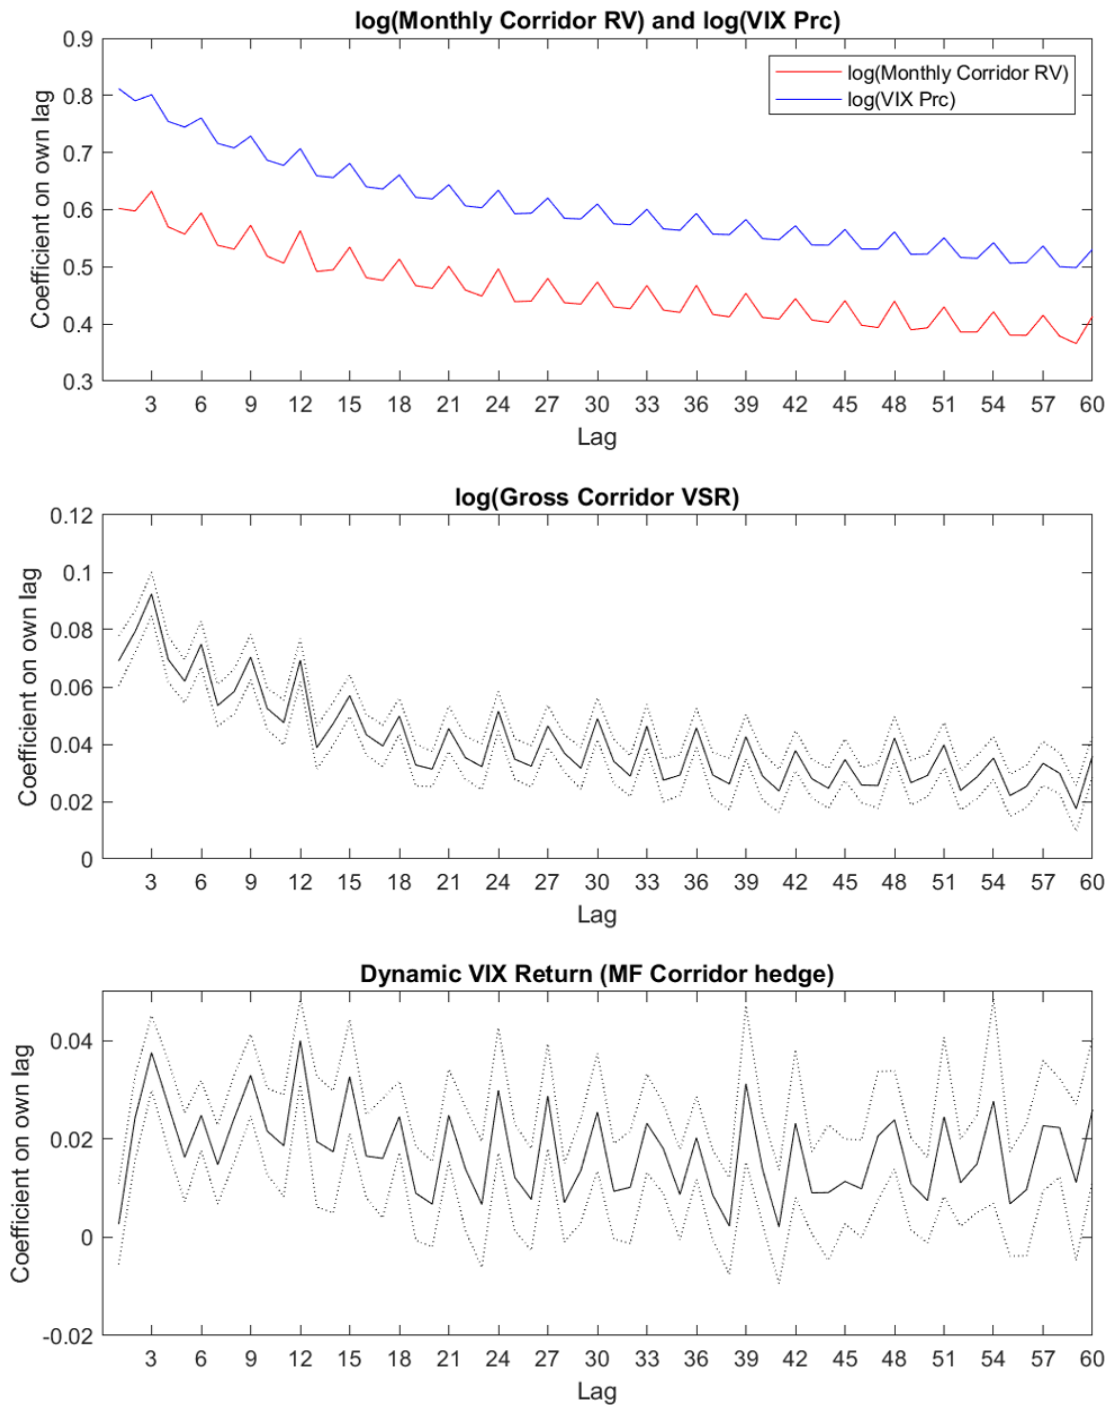


Figure 2. This paper shows cross-sectional regressions in which a variable is regressed on its own k 'th lag, where k is the value on the horizontal axis. In the upper panel, the regression variable is either the log of realized variance or the log of the VIX price, computed using Simpson's rule. In the middle panel, the regression variable is the return on an approximated variance swap. In the lower panel, the regression variable is the return on the VIX portfolio hedged using Black-Scholes.

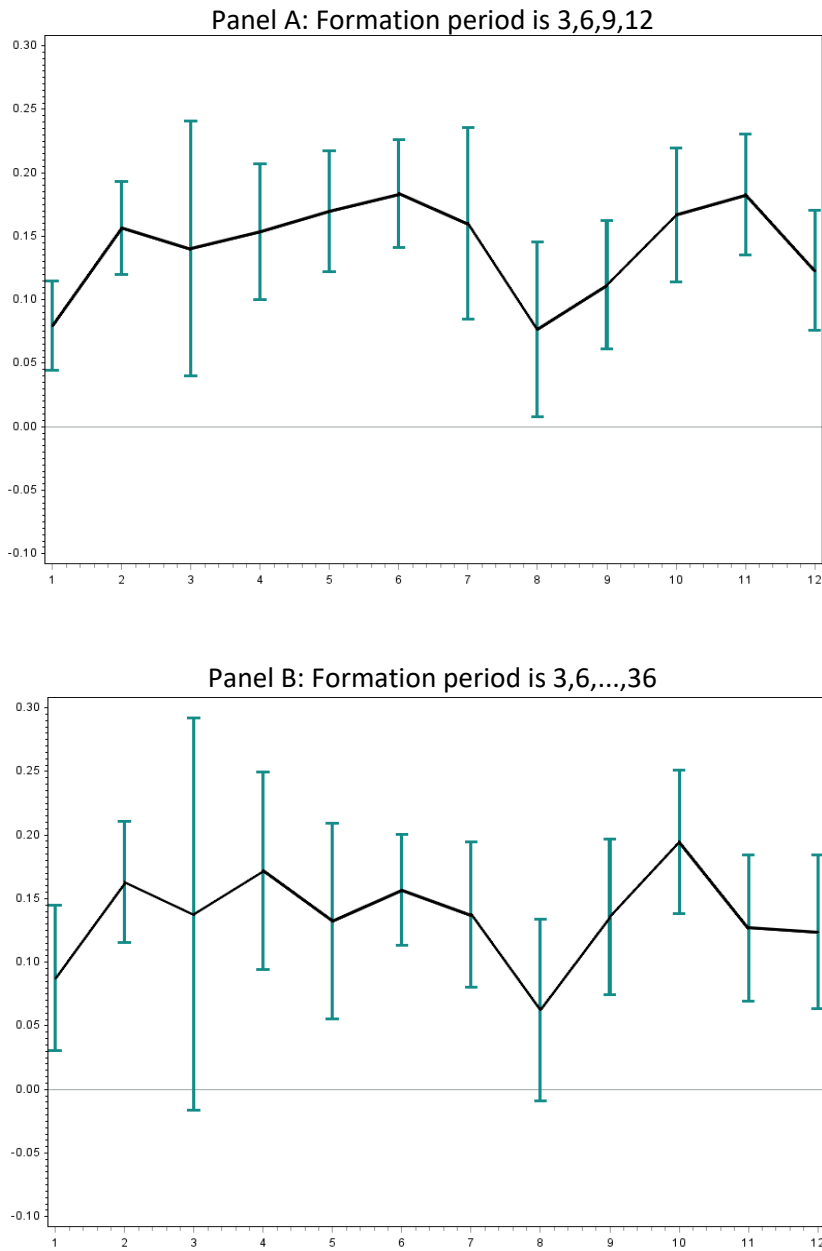


Figure 3.

This figure shows the average returns on two different seasonal momentum strategies in the 12 months of the year. Vertical bars denote 95% confidence intervals.



Figure 4.

This figure shows five-year moving averages of the returns on various momentum strategies. Panel A examines strategies formed on the basis of the past 12 months of returns, while Panel B uses 36 months. In each panel, we show the strategy based on all lags, quarterly lags only, and non-quarterly lags only.

Table 1**Simulation study**

In this table we report simulation results that measure the accuracy of corridor VIX² measures computed using four different integration rules. We simulate a panel of option prices that match the actual sample described in Section 3 in terms of the numbers of options, strike prices, maturities, stock prices, and riskless rates. We assume a time-varying volatility for each stock that is the average of the implied volatilities of all options in each stock's VIX portfolio. Risk neutral skewness and kurtosis also vary across stocks and time and are computed using the approach of Bakshi, Kapadia, and Madan (2003). Following Aschakulporn & Zhang (2019), option prices are generated by Gram-Charlier expansion. We then compute model-free corridor implied variances using three different integration rules and compare them to true values, which are computed using a larger set of options with strike prices over a much finer grid. Panel A of the table reports statistics on the percentage errors, while panel B describes the simulation parameters. All statistics are from a sample that pools all stocks over all dates.

	Mean	StdDev	P10	Median	P90
Panel A: Absolute percentage error in corridor VIX ²					
CBOE	12.79	13.00	2.00	10.07	26.76
Rectangle rule	4.70	9.07	0.39	2.84	11.21
Simpson's rule	1.43	7.81	0.09	0.80	2.47
Corrected Simpson's rule	1.24	7.74	0.08	0.65	2.03
Panel B: Simulation parameters					
# of options	6.62	6.68	3.00	5.00	12.00
# of calls	3.06	3.09	1.00	2.00	6.00
# of puts	3.56	4.06	1.00	2.00	7.00
stock price	51.18	74.00	13.48	36.00	92.94
time to expiration	0.08	0.01	0.08	0.08	0.10
risk-free rate	0.02	0.02	0.00	0.02	0.06
volatility	0.54	0.27	0.26	0.47	0.89
skewness	-0.50	0.68	-1.25	-0.46	0.24
kurtosis	3.84	3.50	1.81	3.02	6.21

Table 2: **Convergence in VIX Prices**

In this table we measure the effects of reducing the number of available strikes on the price of the VIX portfolio. We begin with a sample of firms that have at least seven one-month out-of-the-money options, including at least three calls and three puts, and compute the VIX price (VIX_{All}^2) using each of the three integration rules we consider. Ordering the options from lowest to highest strike price, we then discard all options in even positions, leaving us with half or slightly more than half of the original number. We use this subset to compute a different VIX price (VIX_{Half}^2) and examine the percentage absolute difference between the two, $100 \times \left| \frac{VIX_{Half}^2 - VIX_{All}^2}{VIX_{All}^2} \right|$. The table reports summary statistics of this value, where the sample is pooled across all dates and firms. We also consider subsamples in which the firm has exactly 7, 11, or 15 OTM options. The number of observations is 34,088 when at least 7 options are included, 9,880 when exactly 7 are included, 5,359 when 11 are included, and 2,343 when 15 are included.

	Mean	StdDev	P10	Median	P90
Panel A: CBOE					
At Least 7 Strikes	13.60	10.27	3.98	11.70	24.59
7 Strikes	21.76	11.95	12.55	19.51	32.40
11 Strikes	11.39	7.09	6.16	9.90	17.83
15 Strikes	7.61	4.42	3.83	6.53	12.57
Panel B: Rectangle					
At Least 7 Strikes	5.82	5.32	0.73	4.40	12.80
7 Strikes	8.30	6.42	1.34	7.04	16.80
11 Strikes	5.23	4.41	0.77	4.34	10.56
15 Strikes	3.90	3.47	0.65	3.15	7.85
Panel C: Simpson					
At Least 7 Strikes	3.36	3.27	0.55	2.54	6.87
7 Strikes	4.82	3.96	1.29	3.92	9.20
11 Strikes	3.00	2.71	0.67	2.37	5.86
15 Strikes	2.33	2.31	0.43	1.76	4.55

Table 3: Option and Variance Swap Returns

In this table we report summary statistics on the rates of return on different option strategies, some exact and some approximated. In Panel A to C, VIX portfolio weights are constructed using a different integration rule. For each VIX portfolio, we examine unhedged returns, returns after hedging using Black-Scholes, and returns after hedging using the model-free hedge ratio, both with and without a corridor adjustment. We also examine the approximate returns on variance swaps and corridor variance swaps, both of which use the same VIX price as the swap rate. Panel D reports returns on corridor variance swaps in which the swap rate is approximated using a linear interpolation of implied volatilities. Returns are reported on a monthly-percentage basis, and the number of observations is 221,157.

	Mean	StdDev	P10	Median	P90
Panel A: CBOE					
VIX portfolio	-5.78	165.45	-96.28	-55.73	130.71
VIX portfolio with Black-Scholes hedge	-9.12	76.25	-61.93	-24.60	50.86
VIX portfolio with model-free hedge	-8.53	135.83	-70.88	-22.12	70.53
VIX portfolio with model-free corridor hedge	-8.48	82.33	-70.70	-23.05	64.33
Variance swap	-7.01	155.53	-70.36	-33.75	60.32
Corridor variance swap	-19.39	85.95	-71.39	-36.86	42.02
Panel B: Rectangle					
VIX portfolio	-8.65	188.57	-100.00	-67.24	137.73
VIX portfolio with Black-Scholes hedge	-10.28	84.51	-68.87	-27.35	56.07
VIX portfolio with model-free hedge	-11.60	152.78	-84.07	-28.76	80.58
VIX portfolio with model-free corridor hedge	-11.57	96.06	-83.44	-29.49	74.16
Variance swap	2.62	176.52	-66.83	-26.67	75.63
Corridor variance swap	-11.04	93.40	-68.01	-30.06	56.49
Panel C: Simpson's					
VIX portfolio	-8.35	182.70	-98.77	-63.09	129.63
VIX portfolio with Black-Scholes hedge	-10.22	84.14	-68.29	-27.33	55.85
VIX portfolio with model-free hedge	-11.27	151.68	-69.52	-28.91	64.30
VIX portfolio with model-free corridor hedge	-11.30	84.93	-68.29	-29.45	56.63
Variance swap	2.10	180.58	-66.79	-26.91	74.36
Corridor variance swap	-11.51	92.50	-68.01	-30.30	55.19
Panel D: Implied volatility interpolation					
Corridor variance swap	-10.86	87.07	-68.29	-28.72	57.68

Table 4: Return correlations

In each of the 299 months in our sample, we compute cross-sectional correlations between selected return measures. The table reports means, standard deviations, and quantiles of these correlations for the three different integration rules considered.

	Mean	StdDev	P10	Median	P90
Panel A: CBOE					
Corr(VIX portfolio with Black-Scholes hedge, Variance swap)	0.89	0.08	0.82	0.91	0.95
Corr(VIX portfolio with model-free hedge, Variance swap)	0.71	0.23	0.41	0.80	0.90
Corr(VIX portfolio with Black-Scholes hedge, Corridor variance swap)	0.85	0.22	0.71	0.92	0.97
Corr(VIX portfolio with model-free corridor hedge, Corridor variance swap)	0.77	0.21	0.59	0.82	0.92
Panel B: Rectangle					
Corr(VIX portfolio with Black-Scholes hedge, Variance swap)	0.89	0.08	0.82	0.92	0.96
Corr(VIX portfolio with model-free hedge, Variance swap)	0.69	0.21	0.43	0.76	0.88
Corr(VIX portfolio with Black-Scholes hedge, Corridor variance swap)	0.86	0.21	0.73	0.92	0.97
Corr(VIX portfolio with model-free corridor hedge, Corridor variance swap)	0.77	0.20	0.61	0.82	0.92
Panel C: Simpson's					
Corr(VIX portfolio with Black-Scholes hedge, Variance swap)	0.90	0.08	0.82	0.92	0.96
Corr(VIX portfolio with model-free hedge, Variance swap)	0.80	0.22	0.48	0.89	0.96
Corr(VIX portfolio with Black-Scholes hedge, Corridor variance swap)	0.86	0.21	0.73	0.93	0.97
Corr(VIX portfolio with model-free corridor hedge, Corridor variance swap)	0.92	0.22	0.76	0.99	1.00

Table 5**Univariate sorts of seasonal momentum**

This table reports means and t-statistics from univariate quintile sorts. Dynamically hedged VIX returns with model-free corridor hedge ratios are used. T-statistics, in parentheses, are computed using Newey-West standard errors with three lags. All portfolios are equally weighted, and returns are in excess of the risk-free rate. Observations are included if at least two thirds of all months in the formation period are non-missing. The sample includes common shares (shares with share codes of 10 and 11) and spans between January 1996 to December 2020.

Formation period	Low	2	3	4	High	High - Low	Mean monthly # firms
Panel A: Seasonal momentum							
3	-0.1990 (-13.77)	-0.1594 (-9.98)	-0.1291 (-6.86)	-0.0980 (-5.10)	-0.0815 (-4.49)	0.1175 (13.06)	692.5
3,6,9,12	-0.2167 (-12.99)	-0.1616 (-9.48)	-0.1257 (-6.82)	-0.0975 (-4.73)	-0.0745 (-3.99)	0.1422 (14.80)	666.0
3,6,...,36	-0.2084 (-12.16)	-0.1667 (-9.58)	-0.1221 (-6.23)	-0.0990 (-4.59)	-0.0724 (-3.32)	0.1360 (12.23)	597.9
12	-0.1676 (-9.35)	-0.1493 (-7.66)	-0.1270 (-6.81)	-0.1052 (-5.37)	-0.0830 (-4.47)	0.0846 (9.39)	650.1
12,24,36	-0.1781 (-10.11)	-0.1481 (-7.49)	-0.1204 (-6.30)	-0.1034 (-4.62)	-0.0859 (-4.25)	0.0922 (9.33)	606.5
Panel B: Standard Momentum							
2,3,...,12	-0.2270 (-12.86)	-0.1598 (-10.20)	-0.1298 (-7.46)	-0.0978 (-5.19)	-0.0834 (-4.52)	0.1436 (13.90)	543.2
2,3,...,36	-0.2086 (-11.86)	-0.1515 (-8.86)	-0.1259 (-6.61)	-0.0958 (-4.90)	-0.0903 (-4.26)	0.1183 (9.35)	463.5
Panel C: Other predictors							
Short-term momentum	-0.1921 (-15.04)	-0.1410 (-9.78)	-0.1264 (-7.91)	-0.0944 (-5.24)	-0.1084 (-5.76)	0.0836 (7.38)	595.6
IV-HV	-0.0799 (-4.26)	-0.0697 (-3.03)	-0.1086 (-5.40)	-0.1403 (-8.58)	-0.2174 (-17.46)	-0.1375 (-9.96)	739.8
Idiosyncratic vol	-0.1235 (-4.92)	-0.1102 (-5.60)	-0.1059 (-6.34)	-0.1125 (-7.35)	-0.1638 (-12.15)	-0.0403 (-2.18)	739.8
Market cap	-0.1603 (-14.25)	-0.1158 (-8.01)	-0.1221 (-7.54)	-0.1177 (-6.13)	-0.1000 (-3.34)	0.0603 (2.53)	739.8
IV term spread	-0.1890 (-15.01)	-0.1279 (-8.00)	-0.1095 (-5.92)	-0.0925 (-4.17)	-0.0972 (-4.85)	0.0918 (7.34)	739.8
IV smile slope	-0.1642 (-11.10)	-0.1024 (-5.21)	-0.0999 (-4.97)	-0.1036 (-5.32)	-0.1460 (-10.14)	0.0182 (2.61)	739.8

Table 6**Risk and return for factor portfolios**

This table reports risk and return measures for high-minus-low portfolios based on seasonal momentum and other characteristics. All values are in monthly decimal terms. Means and t-statistics for the long/short factors are identical to those in Table 1 except, in some cases, for the sign. T-statistics, in parentheses, are computed using Newey-West standard errors with three lags. Sample includes common shares (shares with share codes of 10 and 11) and spans the period from January 1996 to December 2020.

	Seasonal momentum (3,6,9,12)	Seasonal momentum (3,6,...,36)	IV - HV	Idiosyncratic vol	Market cap	IV term spread	IV smile slope	Momentum (2- 12)	Short SPX VIX	Short EW stock VIX
Mean	0.1422 (14.80)	0.1360 (12.23)	0.1375 (9.96)	0.0403 (2.18)	0.0603 (2.53)	0.0918 (7.34)	0.0182 (2.61)	0.1436 (13.90)	0.1709 (1.74)	0.1232 (7.20)
Standard deviation	0.149	0.184	0.188	0.307	0.400	0.195	0.112	0.155	1.695	0.280
Sharpe ratio	0.956	0.739	0.732	0.132	0.151	0.471	0.163	0.927	0.101	0.439
Skewness	1.024	2.363	2.621	8.731	9.488	4.794	0.103	0.038	-13.857	-4.267
Kurtosis	7.80	20.45	19.60	116.01	128.10	46.60	1.64	1.06	217.80	34.32
Maximum drawdown	0.373	0.406	0.378	0.982	0.996	0.525	0.861	0.611	> 1	> 1

Table 7**High-minus low spreads for seasonal and non-seasonal momentum**

This table reports average differences between the high and low portfolios formed on the basis of average lagged returns over various formation periods. In all cases, the returns analyzed are for dynamically hedged VIX portfolios, and the holding period is one month. Values in the table differ only as a result of the formation period, which can be determined by taking the intersection of the lags in the row and column headers. For example, the value in the row labeled "Non-quarterly (1,2,4,5,...)" and the column labeled "Lags 25 to 36" is constructed based a formation period that includes all non-quarterly lags in the third year prior to the holding period, namely lags 25, 26, 28, 29, 31, 32, 34, and 35. The sample includes common equities (share codes of 10 and 11) and spans the sample from January 1996 and December 2020.

Formation months	Lags 1 to 12	Lags 13 to 24	Lags 25 to 36	Lags 37 to 48	Lags 49 to 60
All	0.1392 (13.64)	0.0931 (6.77)	0.0648 (5.59)	0.0593 (4.57)	0.0526 (4.45)
Quarterly (3, 6, ...)	0.1422 (14.80)	0.1005 (8.38)	0.0740 (7.50)	0.0669 (5.42)	0.0579 (4.44)
Non-quarterly (1,2,4,5, ...)	0.1179 (8.99)	0.0702 (5.27)	0.0447 (3.71)	0.0365 (2.72)	0.0476 (3.81)
Annual (12, 24, ...)	0.0846 (9.39)	0.0733 (6.55)	0.0473 (6.53)	0.0553 (7.05)	0.0503 (4.42)
Non-annual (1,...,11,13,...)	0.1438 (11.88)	0.0855 (7.01)	0.0688 (5.63)	0.0466 (3.61)	0.0472 (3.69)
Quarterly non-annual (3,6,9,15,...)	0.1366 (14.55)	0.1005 (8.38)	0.0800 (7.34)	0.0613 (4.69)	0.0586 (3.77)

Table 8**Seasonal versus non-seasonal momentum**

This table reports Fama-MacBeth regression results in which model-free VIX portfolio returns are regressed on measures of seasonal and non-seasonal momentum. The sample includes common shares (shares with share codes of 10 and 11) and spans the period from January 1996 to December 2020.

Panel A: Formation periods in last 12 months

	Seasonal momentum (3,6,9,12)	Seasonal momentum (12)	Momentum (2,3,...,12)	Avg. CS R-square
Intercept	0.0850 (13.23)			0.0051
-0.1014 (-5.48)		0.0400 (9.14)		0.0045
-0.1096 (-5.97)			0.1146 (10.99)	0.0057
-0.0978 (-5.39)	0.0800 (9.77)	0.0162 (3.23)		0.0085
-0.0959 (-5.16)	0.0534 (6.31)		0.0677 (5.14)	0.0095
-0.0965 (-5.39)		0.0257 (4.54)	0.0964 (8.36)	0.0094
-0.0938 (-5.12)	0.0503 (5.25)	0.0156 (2.65)	0.0605 (4.43)	0.0128

Panel B: Formation periods in last 36 months

	Seasonal momentum (3,6,...,36)	Seasonal momentum (12,24,36)	Momentum (2,3,...,36)	Avg. CS R-square
Intercept	0.1561 (12.24)			0.0069
-0.0813 (-3.67)		0.0616 (8.24)		0.0047
-0.1020 (-4.92)			0.1734 (8.74)	0.0062
-0.0833 (-3.83)	0.1635 (10.79)	0.0035 (0.44)		0.0106
-0.0758 (-3.43)	0.1171 (8.55)		0.0609 (2.46)	0.0110
-0.0796 (-3.65)		0.0300 (3.86)	0.1551 (7.09)	0.0111
-0.0779 (-3.63)	0.1275 (8.09)	0.0019 (0.22)	0.0621 (2.51)	0.0156

Panel C: Formation periods in last 36 (seasonal momentum) or 12 (momentum) months

	Seasonal momentum (3,6,...,36)	Seasonal momentum (12,24,36)	Momentum (2,3,...,12)	Avg. CS R-square
Intercept	0.1283 (9.65)		0.0427 (3.49)	0.0112
-0.0840 (-4.10)		0.0425 (5.52)	0.0876 (7.65)	0.0101
-0.0910 (-4.61)	0.1366 (8.61)	0.0025 (0.31)	0.0439 (3.54)	0.0153

Table 9**Controlling for other option predictors**

This table reports Fama-MacBeth regression results in which model-free VIX portfolio returns are regressed on measures of seasonal momentum and other option predictors. The sample includes common shares (shares with share codes of 10 and 11) and spans the period from January 1996 to December 2020.

Intercept	Seasonal momentum {3, 6, 9, 12}	Seasonal momentum {3, 6, ..., 36}	Momentum {2,3,...,12}	Momentum {2,3,...,36}	IV - HV	Idiosyncratic volatility	Market capitalization	IV term spread	IV smile slope	Avg. CS R-squared
-0.0550 (-2.47)			0.0842 (7.71)		-0.2050 (-4.80)	-1.8591 (-6.03)	-0.0002 (-1.23)	0.1920 (4.12)	0.1484 (2.98)	0.0324
-0.0534 (-2.42)	0.0498 (6.26)		0.0413 (3.03)		-0.2052 (-4.71)	-1.8877 (-6.02)	-0.0001 (-0.40)	0.1972 (4.17)	0.1421 (2.80)	0.0362
-0.0448 (-1.83)		0.1172 (9.74)	0.0310 (2.32)		-0.1443 (-2.74)	-1.5576 (-4.20)	-0.0002 (-1.40)	0.3267 (5.49)	0.1327 (2.24)	0.0413
-0.0402 (-1.59)				0.1447 (7.73)	-0.1559 (-3.28)	-1.7443 (-4.86)	-0.0002 (-1.11)	0.3161 (5.24)	0.1421 (2.51)	0.0370
-0.0389 (-1.59)	0.0443 (5.09)			0.1079 (5.25)	-0.1343 (-2.70)	-1.7547 (-4.84)	-0.0002 (-1.12)	0.3486 (5.83)	0.1432 (2.42)	0.0416
-0.0367 (-1.46)		0.1169 (9.32)		0.0349 (1.57)	-0.1424 (-3.01)	-1.6462 (-4.56)	-0.0002 (-1.00)	0.3427 (5.65)	0.1425 (2.58)	0.0414

Table 10**Controlling for other quarterly cycles**

This table reports the results of sorting on seasonal momentum in subsets of the full sample. Panel A examines the subsample in which all months in the formation and holding periods contain earnings announcements. Panel B uses the subsample where formation and holding periods include no earnings announcements. Panel C examines the subsample of firms whose options are expiring in a month that is in that firm's quarterly expiration cycle. Panel D uses the subsample of firms whose options are not part of its quarterly expiration cycle.

Formation period	Low	2	3	4	High	High - Low	Mean monthly # firms
Panel A: Earning months during formation and holding periods							
3,6,9,12	-0.1212 (-4.78)	-0.0725 (-3.29)	-0.0448 (-2.08)	-0.0195 (-0.91)	-0.0082 (-0.38)	0.1083 (3.88)	186.4
3,6,...,36	-0.1126 (-4.66)	-0.0721 (-2.99)	-0.0633 (-2.79)	-0.0380 (-1.84)	0.0046 (0.16)	0.1137 (4.41)	156.6
Panel B: Non-earning months during formation and holding periods							
3,6,9,12	-0.2469 (-13.61)	-0.2059 (-10.74)	-0.1463 (-7.24)	-0.1182 (-4.77)	-0.1027 (-4.68)	0.1442 (10.30)	382.8
3,6,...,36	-0.2527 (-12.94)	-0.1986 (-10.50)	-0.1445 (-6.19)	-0.1104 (-4.34)	-0.0932 (-3.58)	0.1571 (9.55)	333.8
Panel C: Options expiring within quarterly cycle							
3,6,9,12	-0.2365 (-13.05)	-0.1868 (-9.76)	-0.1296 (-6.23)	-0.1094 (-4.85)	-0.0937 (-4.48)	0.1428 (10.78)	256.7
3,6,...,36	-0.2310 (-11.75)	-0.1776 (-9.15)	-0.1369 (-6.24)	-0.1148 (-4.99)	-0.0881 (-3.51)	0.1429 (7.45)	227.6
Panel D: Options expiring outside of quarterly cycle							
3,6,9,12	-0.1986 (-11.16)	-0.1553 (-9.61)	-0.1163 (-6.34)	-0.0918 (-4.65)	-0.0666 (-3.65)	0.1320 (11.40)	409.3
3,6,...,36	-0.1948 (-10.66)	-0.1552 (-8.92)	-0.1156 (-5.93)	-0.0880 (-4.10)	-0.0667 (-3.24)	0.1281 (10.28)	370.3

Table 11
Limits to arbitrage

This table reports return means and t-statistics from sequential double sorts on straddles. Every third Friday, we sort straddles into 3 portfolios based on a conditioning variable shown in the column header and then, within each tercile, sort straddles into 3 portfolios based on past average returns over some formation period. Within each tercile of the conditioning variable, we then compute equal-weighted portfolio returns and take long and short positions in the top and bottom terciles. Numbers reported are the resulting high-minus-low return spreads within each tercile of the conditioning variable. T-statistics, in parentheses, are computed using Newey-West standard errors with three lags. We consider four conditioning variables. The first, firm size, is the stock's most recent equity capitalization. Stock illiquidity is proxied by the average Amihud (2002) measure over the most recent 12 months. Option illiquidity is the average the percentage bid-ask spread of the options in each VIX portfolio, averaged over the past 12 months. Analyst coverage is the number of analysts covering the stock, updated monthly. The sample includes common equities (share codes of 10 and 11) and spans the sample from January 1996 and December 2020.

	High-low spreads based on lags 3, 6, 9, & 12				High-low spreads based on lags 3, 6, ..., 36			
	Firm size	Stock illiquidity	Option illiquidity	Analyst coverage	Firm size	Stock illiquidity	Option illiquidity	Analyst coverage
Low	0.1234 (9.82)	0.0887 (11.82)	0.0662 (8.30)	0.1335 (11.23)	0.1152 (7.89)	0.0888 (12.42)	0.0760 (9.16)	0.1103 (7.53)
Medium	0.1185 (10.58)	0.1023 (9.72)	0.0907 (9.18)	0.1174 (13.27)	0.1065 (9.39)	0.0948 (9.21)	0.0850 (8.03)	0.1030 (10.38)
High	0.0904 (12.19)	0.1126 (8.16)	0.1175 (11.05)	0.0789 (9.76)	0.0866 (10.06)	0.1168 (7.40)	0.0995 (7.47)	0.0861 (10.60)
High - Low	-0.0330 (-2.22)	0.0239 (1.53)	0.0513 (3.97)	-0.0546 (-3.66)	-0.0286 (-1.71)	0.0280 (1.75)	0.0235 (1.48)	-0.0242 (-1.58)

Table 12

Seasonal patterns in realized vs. implied variances

Each panel reports regressions in which the variable being forecast is the log of the realized corridor variance (RCV) during month t, the log of the implied variance at the end of month t-1, or the difference between the two. Regressors include averages of lagged log RCVs over all lags or over periodic (multiples of three) lags, as well as lagged log implied variances.

Panels A and B differ with respect to the maximum lag length used for both the periodic and all-month averages of realized variance. The sample includes common shares (shares with share codes of 10 and 11) and spans between January 1996 to December 2020.

Panel A: Longest lag is 12 months

	Intercept	Periodic average of ln RCV(t-3), ln RCV(t-6), ln RCV(t-9), ln RCV(t-12)	All-month average of ln RCV(t-1), ln RCV(t-2), ..., ln RCV(t-12)	ln VIX(t-1)	Avg CS R2
	-0.5015 (-10.94)		0.8957 (83.32)		0.5005
ln RCV(t)	-0.5096 (-11.06)	0.1526 (9.15)	0.7430 (42.52)		0.5059
	-0.3792 (-12.71)		0.2726 (23.90)	0.7047 (61.38)	0.5939
	-0.3786 (-12.63)	0.0906 (10.03)	0.1878 (13.64)	0.6988 (62.49)	0.5961
ln VIX(t-1)	-0.2015 (-5.39)	0.0790 (5.08)	0.8014 (60.28)		0.7311
ln CRV(t) - ln VIX(t-1)	-0.3081 (-10.81)	0.0736 (8.29)	-0.0584 (-5.59)		0.0215

Panel B: Longest lag is 36 months

	Intercept	Periodic average of ln RCV(t-3), ln RCV(t-6), ..., ln RCV(t-36)	All-month average of ln RCV(t-1), ln RCV(t-2), ..., ln RCV(t-36)	ln VIX(t-1)	Avg CS R2
	-0.5434 (-9.75)		0.8919 (62.84)		0.4267
ln RCV(t)	-0.5551 (-9.89)	0.3258 (13.81)	0.5656 (22.44)		0.4354
	-0.3719 (-10.26)		0.1980 (15.62)	0.7894 (65.56)	0.5765
	-0.3686 (-10.12)	0.1602 (11.94)	0.0435 (2.79)	0.7832 (65.27)	0.5792
ln VIX(t-1)	-0.2561 (-6.18)	0.1992 (8.45)	0.6761 (32.82)		0.6447
ln CRV(t) - ln VIX(t-1)	-0.2990 (-8.52)	0.1266 (9.03)	-0.1105 (-7.13)		0.0229