How to Harvest Variance Risk Premiums for the Long-term Investor? ¶

Julian Dörries*, Olaf Korn†, Gabriel J. Power‡

This version: August 2021

JEL Classification: G10, G11, G23

Keywords: Variance Risk Premium; Variance Factor; Trading Strategies; Long-term Investor

¶ We thank Marco Erling, Florian Reibis, Jörg Zimmermann, as well as participants of the 2021 CFR research seminar for helpful comments and suggestions. Vitus Benson provided excellent research assistance.

* Julian Dörries, Chair of Finance, University of Goettingen, Platz der Göttinger Sieben 3, D-37073 Göttingen, Germany; tel. +49 551 39 27264, fax +49 551 39 27665; email julian.doerries@uni-goettingen.de.

† Olaf Korn, Chair of Finance, University of Goettingen and Centre for Financial Research Cologne (CFR), Platz der Göttinger Sieben 3, D-37073 Göttingen, Germany; tel. +49 551 39 27265, fax +49 551 39 27665; email okorn@uni-goettingen.de.

‡ Gabriel J. Power, Université Laval, Department of Finance, Insurance and Real Estate, Faculté des sciences de l'administration, Pavillon Palasis-Prince, 2325 rue de la Terrasse, Québec, Québec G1V 0A6, Canada; tel. +1 418 656-2131; email gabriel.power@fsa.ulaval.ca.

How to Harvest Variance Risk Premiums for the Long-term Investor?

Abstract

Derivative positions that aim to earn variance risk premiums are exposed to sharp price declines during market crises, which calls into question their suitability for the long-term investor. Our paper systematically discusses the problems associated with the design of long-term variance-based investments. Various design elements are proposed to address some of these problems. An empirical study of investment strategies based on these design elements yields significant differences across strategies in terms of risk and return for the S&P 500 index options market. Overall, our results show that variance strategies can be attractive to the long-term investor if properly designed.

JEL Classification: G10, G11, G23

1 Introduction

The variance risk premium is a well-known phenomenon in options markets. A large literature provides empirical evidence for the existence of a negative variance risk premium in equity index options, such as the S&P 500 options (e.g., Carr and Wu 2009; Kozhan, Neuberger, and Schneider 2013). Moreover, the realized premium has a well-documented structure: it exhibits low negative values most of the time and takes on very high positive values during rare extreme events. However, little attention has been paid in the literature to the harvesting of the variance risk premium for the long-term investor. The structure of variance risk with its rare extreme events makes this a difficult—if not impossible—task. The goal of our work is to provide new insights into whether and how investors can use the variance risk premium for long-term capital accumulation. More specifically, this is the first study to systematically compare different variance-based investment strategies in terms of their suitability for long-term investors.

The paper makes theoretical and empirical contributions. First, strategies for earning the variance risk premium are distinguished from other variance strategies. Then, three problems that arise when variance risk premiums are to be exploited for the long-term investor are highlighted: (i) The payoff problem: Which payoff profiles are appropriate? Which instruments should be used to create them?; (ii) The leverage problem: Which risk level should be chosen? How can the ex-ante variance risk of different strategies be measured and compared?; (iii) The finite maturity problem: Which maturities of derivatives should be chosen? When and how often should positions be rolled over? Starting from the three problems, we propose different strategies for earning the variance risk premium that address potential solutions. The analysis of these strategies contributes to a better understanding of the specifics of the variance risk premium in the context of long-term investments.

In an empirical study for the S&P 500, we examine the various strategies in terms of their suitability for capital accumulation for the long-term investor. We use data from January 1996 to December 2020, so that the two most significant stock market crashes in recent decades—the 2008 financial crisis and the Covid pandemic in early 2020—are included in the sample. This is particularly important for variance strategies. For the strategies, we use realistic assumptions about transaction costs and CBOE margin requirements for option positions. We use two variants of strategies to set the risk exposure: First, we use the maximum possible exposure by levering positions until capital is fully tied up. Second, we use an equal ex-ante factor exposure for all strategies. To compare exposure, we propose a model-free measure of nonlinearity (gamma) that can be determined from current option prices alone.

The empirical analysis provides the following main results: (i) Maximum exposure strategies differ greatly in terms of return and risk profile. In some cases, the strategies exhibit excessive risk and there is even a total loss of the invested capital; (ii) Equal exposure strategies still differ from each other, although significantly less than maximum exposure strategies. The remaining differences are due to the different instruments on which the strategies are based, resulting in different payoff profiles and transaction costs; (iii) All variance strategies have a significant positive correlation with the market, but this correlation differs across strategies due to different payoff profiles; (iv) Throughout the whole sample period, variance strategies have continuously earned premiums. This distinguishes them from other factor strategies based on Fama's and French's (2015) five factors or the momentum factor, which have not shown a significant upward trend since the 2008 financial crisis; (v) The variance factor, although correlated with the market, translates into an attractive factor strategy for long-term investors, both as a stand-alone factor and as a complement to a market investment. Overall, our study shows that although variance strategies exhibit extreme distributions with high negative skewness and excess kurtosis, they recover quite quickly from large drawdowns and continue to consistently earn premiums. The results suggest that these strategies, when properly designed, are attractive to long-term investors.

2 Literature Overview

The objective of this paper is to seek effective ways to harvest the variance risk premium for the long-term investor. Formally defined, the variance risk premium is the deviation between the option-implied (risk-neutral) variance and the expected (physical) variance realized over the life of the options. More loosely speaking, the variance risk premium is the expected return on the "variance factor". Although this variance factor is far from uniformly defined in the literature, approaches to obtaining its factor premium have in common that they provide exposure to variance changes. Nevertheless, these approaches can differ greatly from one another. In the following, we briefly review the literature on strategies aimed at earning variance premiums and distinguish our paper from this literature.

A first approach looks at trading strategies that hold options in combination with their underlying, usually a broad market index. Since options have non-linear payoffs, they naturally provide exposure to changes in variance. Whaley (2002) describes and analyzes the Chicago Board Options Exchange (CBOE) BuyWrite Index, which consists of a short-position in an out-of-the money (OTM) S&P 500 index call and a long-position in the S&P 500. Ungar and Moran (2009) outline a put-write strategy, which is the basis for the CBOE S&P 500 PutWrite Index. These strategies have in common that they led to an outperformance of the S&P 500 on a risk-adjusted basis. Other strategies, such as PutWrite and BuyWrite strategies with varying moneyness of the options or with additional caps and floors are assessed by Clark and Dickson (2019). However, all these strategies represent a portfolio approach rather than a pure variance strategy, as they are not delta-hedged but in most cases fully collateralized with a position in the index. This leads to a mixture of market exposure and variance exposure, as outlined in Israelov and Nielsen (2015). The approach of our paper is to study the properties of different pure variance strategies first and then to relate the performance of these strategies to the underlying and to other equity-based factors.

Another strand of literature utilizes variance exposure as an additional element in a broader portfolio context and examines its impact on overall portfolio performance. Brière, Burgues, and Signori (2010) explore long-term equity and bond portfolio strategies and add short-positions in variance swaps to include variance exposure. They find that the additional variance exposure boosted portfolio returns but added little diversification benefits. Fallon, Park, and Yu (2015) employ a similar approach and investigate the role of additional volatility risk in institutional investment portfolios. They find that adding small amounts of exposure to variance risk substantially enhanced long-term returns at the cost of increased short-term tail risk in their sample. Other studies, such as Ge (2016a) and Ge (2017) likewise examine the effect of variance exposure on portfolio returns and portfolio performance in crisis periods. Overall, this literature provides interesting results on the effects of variance exposure in an asset allocation context. However, it does not analyze the trade-offs between alternative design elements of the "variance factor". Moreover, it does not explore the properties of variance strategies themselves but looks only on their impact on specific portfolios.

A different idea to earn premiums associated with variance exposure is to engage into VIX trading strategies. These strategies are typically built on VIX futures (Simon and Campasano 2014) or VIX options (Simon 2017) and aim to roll down the term-structure of option-implied volatility. As the term-structure of implied volatility is typically in contango, investors can enter into a short-position in longer-term futures to roll down the term-structure of VIX futures and, on average, earn the corresponding premium (Ge 2016b). However, in our understanding, this premium is not as much a variance risk premium as it is a VIX term-structure premium. In particular, these VIX strategies earn money via changes in implied (risk-neutral) volatility between two points in time and not via the deviation between implied and expected realized (physical) volatility, i.e., the variance risk premium. Of course, changes in implied volatility might well be correlated with realized volatility. Nevertheless, this type of strategy earns a premium which is conceptually distinct. In our paper, we concentrate on strategies that aim to harvest the variance risk premium.

Lastly, there are few papers that examine strategies aimed to earn pure variance risk premiums. Fallon and Park (2016) analyze a strategy that sells synthetically derived capped one-month S&P 500 variance swaps. Specifically, they employ a stochastic volatility model with jumps (SVJ) to estimate the swap rates. The authors find the corresponding strategy to have very high Sharpe ratios and severe but infrequent crash risk. However, Fallon and Park (2016) do not investigate the long-term performance of such a strategy and the impact of rare crash events on the capital accumulation of the long-term investor. Ge (2016b) comes closest to our approach by examining the long-term performance of, among others, a variance swap investment strategy. He finds that this strategy provides decent returns but a very large maximum drawdown. However, the comparison of different strategies in Ge's study is based on an ex-post standardization of their risks. This is not feasible for real investments. Our paper, in contrast, uses two different ways to determine a meaningful leverage of different strategies on an ex-ante basis. Moreover, our study is the first to include the period of the Covid-19 pandemic. This period and the corresponding stock market shock is potentially very important for the understanding of long-term variance strategies.

3 Strategies, Data, and Study Design

3.1 Three Problems

The design of strategies that seek to harvest variance risk premiums for the long-term investor entails three major problems. The first one is what we call the *payoff problem*. In principle, variance exposure can be generated by selling any convex payoff structure. However, different payoff profiles may be more or less suitable to achieve certain desirable goals. First, a payoff profile should provide sufficient factor exposure that is priced in the market. Second, a payoff profile should produce low correlations with other risk factors to create potential for diversification benefits in a portfolio context. Third, a payoff profile should limit extreme negative returns in order to avoid large drawdowns. Fourth, a payoff profile should be implementable with low transaction costs. This is crucial for long-term investors, as they profit from compound interest throughout the investment period. There are certainly trade-offs between the achievement of these four goals. These trade-offs are at the heart of the payoff problem, i.e. the problem to choose suitable payoff profiles.

The second problem is what we call the *leverage problem*. Variance strategies use derivatives to sell convex payoff profiles. Since some derivative strategies require no initial capital (e.g., swap contracts) and others actually generate capital (e.g., delta-hedged puts), the question of the appropriate amount of leverage arises. The leverage problem can lead to some similar trade-offs as the payoff problem. On the one hand, strategies should provide sufficient exposure to variance risk; on the other hand, they should limit extreme losses that call into question the long-term success of a strategy—all in a cost-effective manner. However, the leverage problem also raises issues that are distinct from the payoff problem. There may be different strategies with similar or even identical payoffs that use different instruments. Therefore, potential leverage may be more constrained for one strategy than another due to differences in initial capital or margin requirements. In addition, differences in leverage between strategies can lead to sharply different factor exposures, making meaningful comparisons difficult.

Third, there is the *finite maturity problem*. Because of the finite maturities of derivatives, long-term investors must periodically roll their positions. This entails further considerations. In particular, investors must decide which maturities to choose, when to roll, and how many instruments to use simultaneously. One of the potential trade-offs is that longer-term contracts may have less factor exposure but need to be traded less frequently, reducing transaction costs. Shorter-term contracts, however, may be more liquid, resulting in lower transaction costs per trade.

3.2 Variance Strategies

We now present strategies that are in principle suited for harvesting variance risk premiums and address some of the trade-offs associated with the three problems. We begin with a short position in an at-the-money (ATM) *straddle* as a first intuitive approach to selling protection against rising variances. The short straddle limits market exposure because it combines a short call (with negative delta) with a short put (with positive delta).[1] Moreover, it consists of only two ATM instruments, which means that transaction costs are limited. The stylized payoff profile of a short straddle is shown in part (a) of Figure 1. It highlights that this payoff is generally consistent with an instrument showing variance exposure. However, large negative or positive price movements of the underlying would result in large negative returns. Even medium price movements already lead to losses.

[Insert Figure 1 about here.]

Alternatively, one could shift the strike prices of the options from ATM to out-of-the-money (OTM) to avoid losses in case of medium price changes of the underlying. In addition, this shift in strike prices reduces the magnitude of losses in case of large price moves of the underlying, as compared to the straddle. A corresponding payoff profile is depicted in part (b) of Figure 1, which shows a short *strangle*. However, the strangle also has some disadvantages. First, OTM options have less variance exposure than ATM options. Second, the maximum payoff is lower than for a corresponding straddle. Finally, large negative returns can still occur since the payoff function has no lower bound.

To counter the problem of extreme losses, one could add a floor to the payoff profile. By adding a long OTM call option and a long OTM put option, the straddle turns into a *butterfly spread*, as shown in part (c) of the Figure 1. Of course, adding a floor is also possible for the strangle. Such a portfolio is called a *condor strangle*. However,

[1] One could even set the strike price of the options as such that the beta of the straddle is exactly zero.

limiting the downside risk also comes with drawbacks. Since two long positions in options enter the portfolio together with short positions, the overall variance exposure of the portfolio is reduced. Moreover, the additional long positions in calls and puts incur costs—both transaction costs and costs of capital for the option premiums.

Another approach to harvesting the variance risk premium is to sell *delta-hedged call or put* options, i.e., puts or calls hedged with positions in the market index. Delta-hedged options have a similar payoff structure as a straddle. Therefore, they have similar advantages, i.e., limited correlation with the market factor and only two instruments making up the portfolio, but also similar disadvantages, i.e., potentially large downside risk. However, since delta-hedged option portfolios contain different instruments than a straddle, the long-term performance of delta-hedged options may differ significantly from the long-term performance of a straddle strategy. This may be due to different potential for leverage or different transaction costs and margin requirements.

Finally, an investor can gain variance exposure through *variance swaps*. The advantage of a variance swap is that it is the most direct way to earn the variance risk premium while limiting correlation with the market factor, since variance swaps are market neutral by construction. However, a variance swap can be viewed as a fairly complex option portfolio with potentially high transaction costs and leverage constraints. Moreover, since variance is calculated as the squared difference from the mean, a variance swap could be prone to extreme losses if the realized variance reaches peaks.

Overall, we have selected seven different but related portfolios: Straddle, strangle, butterfly spread, condor strangle, delta-hedged call, delta-hedged put, and variance swaps. Since all of them offer convex payoff structures, selling these portfolios has the potential to earn variance risk premiums. In addition, differences in their designs are particularly targeted to provide insights on potential trade-offs with respect to variance exposure, downside risk, transaction costs, capital requirements, and margins.

Having introduced different portfolios and their payoff profiles, the next question is what quantities should be held in each portfolio to obtain a good risk-return profile for the long-term investor. We follow two ways to approach this issue in our study. The first is the *maximum exposure* approach. Here, the investor generates the maximum possible exposure for each strategy by levering portfolio positions until her capital is fully invested in long positions and margins for short positions. In essence, the maximum exposure approach is a full invest of the available capital in the variance strategy. It provides information on the impact of different instruments on potential leverage. The second way is the *equal exposure* approach. Here, all strategies are levered until each strategy has the same ex-ante factor exposure. If one strategy requires less capital to achieve this exposure than another, the additional funds are invested in a risk-free account. Equal exposure strategies help us compare different strategies on an ex-ante risk basis.

Equal exposure strategies require an ex-ante measure of factor exposure. The literature suggests different measures associated with the non-linearity of options, in particular vega and gamma (Cremers, Halling, and Weinbaum 2015). However, vega in general depends on the valuation model used, and specifically, the Black-Scholes vega measures variance risk in a setting where no variance risk premium exists. Consequently, using vega creates potential problems in the context of our study. Therefore, we use the gamma of the option as a general measure of non-linearity or convexity. Since we want to avoid model dependence and, in particular, a measure that builds on a model where variance risk premiums do not even exist, we apply a model-free approach. We determine gamma and, for the sake of consistency, also delta in this way. The derivation of such model-free greeks is based on the option-implied risk-neutral return (RND) distribution and is discussed in more detail in the next subsection.

3.3 Data, RNDs and Greeks

In our empirical study, we use European S&P 500 index options data which stems from the OptionMetrics IvyDB U.S. data base. The data base provides historical closing quotes from the CBOE, as well as implied volatilities, interest rates, spot prices of the underlying, and implied dividend yields. Our sample period starts in January 1996 and ends in December 2020, that is, it spans 299 months. We filter our option data set with standard filters from the literature (Goyal and Saretto 2009; Cao and Han 2013). We require best bid and best ask quotes as well as the bid-ask-spread to be non-negative and discard options with special settlement and options that do not have a.m. settlement. Moreover, the implied volatility must be greater than zero and we require options to survive standard no-arbitrage conditions. For every month in our data period, we only use data of the first trading day after the third Friday of that respective month. The latter is the standard expiration date of options traded at the CBOE. Additionally, we require the options to mature in the next month. This retains options with approximately one month to maturity. Our final sample consists of 23,231 call options and 23,697 put options for 299 months with an average time to maturity of approximately 28 days.

Model-free deltas and gammas require an estimate of the risk-neutral distribution of the S&P 500 index. We determine the RND based on the approach outlined in Figlewski (2010). At every point in time we select out-of-the money forward put and call options that survive our filter criteria and convert their midquotes to the corresponding Black and Scholes (1973) implied volatilities. As put and call options might trade at slightly different implied volatilities, we smooth out potential jumps in implied volatilities at the transition point from put options to call options. In an interval of 2.5% around the at-the-money forward price we apply the blending approach suggested in Figlewski (2010) to achieve a smoothed weighted-average implied volatility.[2] Next, we perform a quartic spline to interpolate the implied volatility curve across a broad range of observed strike prices. We then compute a

[2] For this blending approach we have to employ in-the-money (ITM) options that are only used to smooth the implied volatility curve and discarded afterwards.

fine grid of 12,000 equally spaced strike prices on the interval $[0.001, 3 \cdot X_t]$, where X_t is the current index level, and determine the corresponding implied volatilites through a quartic spline. These implied volatilities are then converted back into put option prices, which we use to estimate the risk-neutral density by approximating derivatives of the put price function with respect to the strike price according to Breeden and Litzenberger (1978). This procedure is only valid between the lowest and the highest available strike prices. For strike prices beyond the observed range we estimate a generalized extreme value (GEV) distribution to extend the risk-neutral density to the left and right tail.[3]

Given the RND, we compute model-free deltas and gammas according to a simple comparative static analysis. First, we assume that the entire RND is shifted upward by an infinitesimally small amount ε. Next, we reprice options under this shifted RND and observe the price change. Likewise, we consider a downward shift of the entire RND and perform the same repricing procedure. Lastly, we determine the average of these two price changes, divide it by ε and use this measure as the model-free delta. The same idea is applied to gamma, which is the sensitivity of delta with respect to changes in the price of the underlying. Formally, this analysis can be found in greater detail in Appendix A2. Given the assumptions we make for the derivation, model-free deltas of long call options are simply the probability mass to the right of the strike price, whereas long put options' model-free deltas are the negative value of the probability mass to the left of the strike price. Model-free gammas are simply the probability mass "at" the strike price for both long call and long put options.

3.4 Implementation of Strategies

General Specifications We now turn to the concrete design of the trading strategies, which are based on the portfolios we presented in Subsection 3.2. We use S&P 500 options with a maturity of one month and hold them until expiration, so we

[3] This approach is described in more detail in Figlewski (2010).

have to open new positions every month. Positions are established every Monday after the third Friday of the month. We index each strategy to a level of 100 as of January 1996 and then determine the cumulative wealth through the end of our data period in December 2020.

In terms of transaction costs, we use the ask price for option purchases (long positions) and the bid price for option sales (short positions). However, it is likely that institutional investors can trade at terms better than the quoted spread (Mayhew 2002; De Fontnouvelle, Fishe, and Harris 2003), which is why we choose an effective spread of 25% of the quoted spread as our baseline scenario.[4] This is an even more conservative approach than that used by other papers examining option trading strategies—Goyal and Saretto (2009) and Cao and Han (2013) even assume an effective spread of 0% as the base case, i.e., trading is possible at the midquote. Transaction costs are incurred each time an option or the underlying is bought or sold. However, cash-settled options do not incur additional transaction costs at expiration, which is why we choose to let options expire. We assume that transactions in the underlying, which are necessary for delta hedging, are possible at an quoted bid-ask spread of 3 basis points.

In addition, we require margin accounts to collateralize the positions. For option trades, we apply the CBOE margin rules, which provide detailed guidelines for individual options and option portfolios. Cao et al. (2021), for example, also use the CBOE margin rules for their option trading strategies. An overview of the CBOE margin requirements we use in our study is provided in Appendix A1. We use the margin rules for initial margins and assume that there are no maintenance margins or other adjustments required within the trading month. For short positions in the underlying index, we assume that 150% of the short sale proceeds must be deposited, which is equivalent to the Federal Reserve Board's "Regulation T". Finally, the margin requirements for variance swaps, which require no capital at inception because they are priced so that their initial value is zero, are set in accordance

[4] In practice, institutional investors also use limit order strategies to reduce transaction costs below the quoted spread.

with the “margin requirements for non-centrally cleared derivatives” of the “Basel Committee on Banking Supervision” (BCBS). The BCBS requires that 15% of the variance notional be deposited as initial margin. The variance notional is the notional amount by which the difference between the floating leg and the fixed leg is multiplied. We assume that margin accounts bear interest at the risk-free rate.

For maximum exposure strategies, we calculate the maximum possible exposure by levering the positions until the capital is fully used to buy long positions and provide margins on short positions. For equal exposure strategies, we lever the positions until each strategy has the same factor exposure. In each period, we select the strategy that has the lowest gamma when it is fully levered according to the maximum exposure approach. All other strategies are then scaled down to this exposure and the remaining capital is invested at the risk-free rate.

Straddle and Butterfly Spread For the straddle strategy, we choose both a call option and a put option with strike prices closest to the ATM forward point. A straddle has its own margin rules according to the CBOE. For the butterfly spread, we add two options to the straddle. We choose an additional call option with a model-free delta of 0.05 and an additional put option with a model-free delta of -0.05. Both options enter the straddle as long positions. The choice of deltas is inspired by the CBOE S&P 500 Iron Butterfly Index[5], which is a hypothetical option trading strategy calculated by the CBOE that sells butterfly spreads. The margins for the butterfly spread are determined as the margins of a short call spread plus the margins of a short put spread, as there are no separate rules for butterfly spreads. Both the straddle and the butterfly spread may have a residual index exposure that we delta-hedge with an index position at initiation. At maturity, the option positions are cash-settled and the hedge positions in the index are closed, which is true for all strategies we consider in our study.

[5] `https://www.cboe.com/us/indices/dashboard/BFLY/`.

Strangle and Condor Strangle The way we specify the straddle and condor strangle strategies is very similar to the straddle and butterfly spread strategies. We choose a call option and a put option with expiration next month that have model-free deltas of 0.20 and -0.20, respectively. Selling these two options results in a strangle. For the condor strangle, we again add two options. Consistent with the butterfly spread, we choose an additional call option with a model-free delta of 0.05 and an additional put option with a model-free delta of -0.05. This choice is similar to the CBOE S&P 500 Iron Condor Index[6], which is a hypothetical options trading strategy from the CBOE selling condor strangles. The CBOE provides its own margin rules for the strangle, while for the condor strangle we combine the margins of a short call spread and the margins of a short put spread. Again, any remaining delta of the strategies is hedged via index positions.

Delta-Hedged Call and Delta-Hedged Put For the delta-hedged call and put strategies, we select the call and put options with expiration in the next month that are closest to the ATM forward point. The strategies sell these options at the best bid (taking into account the effective spread) and hedge the resulting delta exposure with a position in the underlying index. Thus, the short call requires a long position in the underlying, while the short put requires a short position in the underlying, for which additional margins must be deposited. To keep transaction costs low, the delta hedge is set up at initiation and is not readjusted until maturity.

Variance Swap The variance swap strategy requires the calculation of the variance swap rate. This swap rate is determined according to Kozhan, Neuberger, and Schneider (2013), which leads to consistent pricing based on the same options for all strategies in our study. We select all OTM forward call and put options to calculate swap rates. Transaction costs are accounted for by using bid quotes instead of midquotes to calculate variance swap rates. The strategy sells variance swaps at the variance swap rate and holds this position until maturity. By design, variance

[6] `https://www.cboe.com/us/indices/dashboard/CNDR/`.

swaps are delta-neutral at inception. At maturity, the realized variance over the last month is taken to determine the variance swap's payoff.

4 Empirical Results

4.1 Maximum Exposure Strategies

The first set of results relates to maximum exposure strategies and aims at a better understanding of the leverage problem. Maximum exposure strategies show whether the capital requirements of an initial investment (if any) plus margins already set reasonable leverage constraints for strategies and lead to desirable risk-return patterns over longer horizons. If not, the question is whether the constraints are already too tight to earn significant premiums or, on the contrary, too loose to protect long-term investors from excessive risk.

[Insert Figure 2 about here.]

Figure 2 provides the accumulated wealth over time as created by different strategies. For each strategy, the initial budget in January 1996 is $100. The figure shows large differences among strategies. One group of extremely risky strategies consists of the straddle and the strangle. These strategies reach very high wealth levels by 2007, but collapse during the financial crisis in October 2008. While the straddle recovers quickly from this shock, the strangle barely survives and does not have enough capital to recover quickly. It is eventually hit by the market turmoil due to the Covid-19 pandemic and is forced into bankruptcy in March 2020. The straddle also takes a massive hit from the pandemic. The butterfly spread and the condor strangle seem to dampen the extremes of the straddle and the strangle somewhat. However, they are still very risky and are also hit hard by the variance shocks of the financial crisis and the Covid-19 pandemic. Another interesting observation is the different behavior of the delta-hedged put and the delta-hedged call. The delta-hedged put fluctuates more and shows a steeper trend. In contrast, the delta-hedged call has a very smooth path, but does not seem to generate enough exposure for a significant upward movement, especially after the financial crisis. Finally, the variance swap combines a very smooth path, except for the most extreme months of the financial crisis and the Covid-19 pandemic, with a clear upward trend.

[Insert Table 1 about here.]

Table 1 characterizes the performance of the different strategies via summary statistics. Panel A provides information on the sample moments of monthly returns and Panel B shows information on downside risk. The first three measures of downside risk (VaR, CVaR, Max Loss) still take a monthly perspective and refer to (potential) losses in the following month. This is sufficient for a monthly investment horizon. However, for the long-term investor, the characteristics of the entire path are also important. In particular, the ability of a strategy to recover from an intermediate downturn is crucial. Therefore, Panel B offers four different drawdown statistics. The maximum drawdown (Max DD) is the maximum percentage loss of a strategy from its current maximum value to a trough; and the average drawdown (Average DD) shows how far, on average over all months of the 25-year period, a strategy is from its previous maximum. Drawdown length indicates how many months it takes to reach a new wealth maximum at a given point in time. For this measure, we report the maximum number of months (Max DD Length) and the average (over all months of the 25-year period) number of months (Average DD Length). For completeness, Panel C of Table 1 repeats the final wealth levels at the end of the data period from Figure 2 and converts them to annual geometric average returns.

Several differences between the strategies are evident in the figures in Table 1. First, when comparing straddle and strangle, straddle returns have a higher mean, a lower standard deviation, are less skewed to the left, and have a lower kurtosis. In terms of downside risk, the straddle is less risky than the strangle by all measures of risk. Therefore, the idea of replacing a straddle with a strangle to achieve risk reduction, especially with respect to large downturns, does not work, at least when taking maximum exposure. Both straddle and strangle strategies are highly speculative, with high mean returns but massive downside risk. Second, the comparison between straddle and butterly spread and between strangle and condor strangle provides insight into the impact of floors imposed via long positions in OTM calls and puts. These floors result in less negative skewness and less kurtosis, but they are also costly in terms of mean returns. Because of the latter effect, it is questionable whether such floors are an effective way to reduce risk. In terms of downside risk, the butterfly spread and the condor strangle are much riskier than the delta-hedged put, the delta-hedged call, and the variance swap. Third, the monthly returns of the variance swap show by far the most negative skewness and the highest kurtosis of all variance strategies. However, these characteristics do not necessarily imply high downside risk. The reason is that the standard deviation and the dynamics of the path are also important. Although the variance swap experiences some heavy losses (the maximum monthly loss is almost 37%), the relatively short drawdown length shows that the strategy can recover quite well from such losses.

In summary, there are clear differences between the various maximum exposure strategies. The delta-hedged call seems to generate relatively low factor exposure, even compared to the delta-hedged put. Other strategies create large exposure and massive risks (straddle, strangle); still others seem to have variance exposure that already offers promising risk-return profiles (delta-hedged put, variance swap) and relatively quick recovery from setbacks. In factor investing, however, investors seek to manage and control exposure. They look for a strategy that most efficiently exploits a given exposure, which can be studied using equal exposure strategies.

4.2 Equal Exposure Strategies

Figure 3 and Table 2 show the wealth accumulation and summary statistics of different equal exposure strategies. As expected, the standardization of the strategies in terms of ex-ante gamma leads to a more homogeneous picture. In terms of differences and similarities among the strategies and the benefits of particular design elements, the following observations are most striking. First, there is no clear evidence that a strangle helps reduce the risk of a straddle. Strangle returns are even more left-skewed and leptokurtic. In terms of downside risk, the different measures point in different directions. The straddle has a lower (less negative) VaR and CVaR, but a higher maximum loss. The maximum drawdown and maximum drawdown length are lower for the strangle, but the average drawdown and average drawdown length are lower for the straddle. Second, the floors introduced by the butterfly spread and the condor strangle improve monthly skewness and kurtosis, but are not effective in reducing downside risk, as a comparison with straddle and strangle shows. Third, delta-hedged put, delta-hedged call, and straddle have very similar paths and distributional properties. This finding implies that the different behavior of delta-hedged calls, delta-hedged puts, and straddles under maximum exposure strategies is due to their different potential for leverage. The delta-hedged call strategy requires a long position in the index to hedge the short call, which consumes a lot of capital. In contrast, the delta-hedged put strategy involves shorting the index. This short position induces margin requirements, but the potential leverage of a delta-hedged put is much greater compared to the call. A (delta-hedged) straddle requires a very small investment (long or short) in the index. In addition, the CBOE's margin requirements for a short straddle are much lower than the sum of the margin requirements for a short call and a short put. Therefore, straddles can achieve much higher gammas than the other strategies. Fourth and finally, the variance swap has the highest return of all the variance strategies while showing the lowest downside risk for the majority of the risk measures.

[Insert Figure 3 about here.]

[Insert Table 2 about here.]

4.3 Why Strategies Differ: Payoff or Costs?

Even with the same gamma exposure, some differences remain between the strategies. Where do these differences come from? A first possible reason is the different payoff structures of the strategies. A second one is potential differences in transaction costs. Since new derivative positions must be set up regularly due to the finite maturity of options, the latter may be substantial.

The similarity of the payoffs should be reflected in high correlations of returns. Table 3 shows such correlations between the monthly returns of the different strategies. Delta-hedged call, delta-hedged put, and straddle have very high correlations of over 0.99, consistent with theoretical considerations.[7] Strangle, condor strangle, and butterfly spread are slightly less correlated, consistent with their modified payoff profiles. However, they still show correlations with the first group above 0.9. The most striking observation from Table 3 is the relatively low correlation (below 0.7) between the variance swap and all other strategies. This indicates the uniqueness of the variance swap.

[Insert Table 3 about here.]

Can the uniqueness of the variance swap really be attributed to its payoff profile? At first glance, this question is difficult to answer. The payoff of a variance swap depends on the realized variance over the price path, while the payoff of the other strategies depends only on the price of the underlying (index level) at the maturity date of the options. However, the initial replicating portfolio of the variance swap can provide some intuition. The payoff function of this options portfolio allows a direct comparison with the payoff profile of the other strategies. As an example, Figure 4 shows the payoff functions of the delta-hedged put and the replicating portfolio of the variance swap for the setup date 24/10/2016. This picture is also quite typical for other setup dates.[8] The payoff function of the delta-hedged put is a piecewise linear function, with a kink at the forward price. This structure implies that losses grow linearly with the distance between the index and the forward price. The variance swap, on the other hand, shows a strongly non-linear payoff profile.

[7] In a world without transaction costs, where put-call parity holds exactly, the payoff profiles of the three strategies are identical for the same gamma exposure.

[8] The payoff function depends on the forward price, the available strike prices and the deltas of the options on the setup date. Therefore, the payoff function varies over time.

There is a wide range around the forward price where the payoff function is almost flat. This feature also leads to a relatively large gap between the two break-even points (-3.62% and 3.36%), compared to the delta-hedged put (-2.01% and 3.4%). The variance swap makes similar profits over a wide range around the ATM point. This is consistent with the strategy's relatively smooth path of cumulated wealth. There is also a strong asymmetry in the payoff of the variance swap. When the index level rises, the losses remain relatively moderate, which also contributes to the smoothness of the path. Only when the index loses massively does the variance swap realize losses that seem to grow exponentially with the index decline. The latter observation is consistent with the high negative skewness and high maximum loss of the variance swap strategy.

[Insert Figure 4 about here.]

Another possible reason for differences between strategies is transaction costs. An intuitive guess is that strategies using more than two option contracts, or OTM and ITM options, are more affected by transaction costs than strategies using only one or two ATM options. The first group consists of the strangle, butterfly spread, condor strangle, and the replicating portfolio of the variance swap, while the second group includes the delta-hedged call, delta-hedged put, and straddle. Figure 5 confirms the conjecture. Part (a) shows the cumulative wealth of the different strategies when the base case effective spread is reduced from 25% to zero; and part (b) shows the cumulative wealth for an effective spread of 50%. Moving from 50% effective spread to zero transaction costs improves the final wealth (after 25 years) of the condor strangle by 30%, the butterfly spread by 24%, and the variance swap by 18%. In contrast, the delta-hedged call, put and straddle only show an increase in wealth of about 13%. However, the variance swap retains the highest final wealth of all strategies even when the effective spread is 50%. Condor strangle, strangle and butterfly spread still deliver the lowest terminal wealth even without transaction costs. So the differences in the performance of strategies are certainly not only caused by different transaction costs.

[Insert Figure 5 about here.]

4.4 Variance Strategies and Market Movements

For a deeper understanding of the differentiating properties of variance strategies, we look more closely at their relationship to market (index) movements. The performance of variance strategies is closely linked to the market for at least two reasons. First, in our study, the index is the underlying of the derivatives positions that implement variance strategies. Therefore, by construction, payoffs are linked to market movements, which can be clearly seen in Figures 1 and 4. The choice between alternative profiles constitutes the payoff problem that we study. Second, previous research has shown an inverse relationship between market movements and variance changes, i.e., variances tend to increase when the market goes down (Black 1976; Christie 1982; Bollerslev et al. 2006). However, different variance strategies might be affected differently by this effect.

Since the additional design elements of strangle, condor strangle, and butterfly spread have not led to improvements in the variance strategies, we focus only on two strategies in the following subsections. The first, the delta-hedged put, represents the group consisting of delta-hedged call and put and the straddle. The second strategy is the variance swap. To give these strategies a chance to exploit their full "variance exposure", we consider their maximum exposure variants.

Figure 6 shows the cumulative wealth of a buy-and-hold index investment in the S&P 500, along with the two variance strategies. An investment in the risk-free instrument is also shown as a reference point. Since 2003, the trend behavior of the market and the two variance strategies has been similar. Setbacks of the index are clearly visible in the delta-hedged put strategy. This is obviously true for the financial crisis in 2008 and the Covid-19 pandemic shock in 2020, but it is also true for the market downturn in 2002 (aftermath of September 11 attacks) and the market drop at the end of 2018. In contrast, the variance swap strategy seems to be mainly affected only by the two very large market shocks. This observation suggests a higher correlation between market returns and delta-hedged put returns compared to variance swap returns. However, exactly the opposite is the case, with monthly return correlations of 0.38 (market and delta-hedged put) and 0.52 (market and variance swap).

[Insert Figure 6 about here.]

To better understand this phenomenon, Figure 7 plots monthly delta-hedged put returns (part (a)) and variance swap returns (part (b)) as a function of market returns. For the delta-hedged put, there is a clear positive relationship between market returns and put returns when the market falls. However, when the market rises, the relationship is negative. Overall, the non-linear relationship between market returns and put returns results in a moderately positive correlation. The overall positive relationship is due to the fact that large negative market movements occur more frequently than large positive market movements. For the variance swap, the relationship is nearly flat over much of the market's return distribution. Only for very large negative market returns is the relationship clearly positive. Since there is no offsetting effect for rising markets, the correlation between the market and variance swap is positive and higher than for the delta-hedged put. These results are well in line with the examples of the payoff profiles in Figure 4. They show that the two variance strategies have significantly different characteristics in terms of their relationship with the market, which is important information for potential investors.

[Insert Figure 7 about here.]

4.5 "Variance" as an Investment Style

So far, we have studied and compared different ways to earn variance risk premiums. In a broader perspective, the strategies under study are just examples of factor investments, because "stock index variance" can be seen as one of (potentially) many equity-based factors. In this section, we relate variance strategies to other factor investments. Specifically, we consider the S&P 500 index (market), long-short portfolios of the remaining four factors of the Fama and French (2015) Five-Factor Model (Size (SMB), Value (HML), Profitability (RMW), Investment (CMA)), plus Momentum (MOM).[9] Factor returns are constructed such that their return periods coincide with the roll-over periods from our option strategies since the original monthly Fama and French (2015) factor returns are based on calendar months. We achieve this by computing geometric returns over the specified investment horizon for every individual portfolio and eventually determine the factor returns according to the definitions from Fama and French (2015). Data on the additional factors is from Kenneth French's website.[10]

[Insert Figure 8 about here.]

Figure 8 shows the cumulative wealth development of the various factor investments. One observation is very striking. After the financial crisis, since mid-2009, none of the SMB, HML, RMW, CMA, and MOM strategies has an upward trend.[11] Only the market and the two variance strategies show significant upward movements during this period. This distinguishes the variance strategies from the other five strategies that try to earn additional premiums in the stock market besides the market risk premium. Also in the period before the financial crisis, the variance strategies together with the momentum strategy show the strongest upward trend.

Table 4 presents return and risk statistics of the different factor investments. In our view, the most relevant reference point for the variance strategies is the market investment, since it is the only one that generates significant premiums after the financial crisis. Looking at the monthly return statistics, market and variance

[9] Since long-short portfolios do not require an initial investment, the initial capital of $100 is invested in an account earning the risk-free rate.

[10] `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

[11] When interpreting the performance of these five factor strategies, one has to keep in mind that they provide an overly optimistic view of the corresponding styles. The reason is that no transaction costs of portfolio revisions are taken into account in the performance calculations, as is the case with the variance strategies.

strategies do not differ strongly in terms of mean returns and standard deviations. However, the variance strategies show more negative skewness and higher excess kurtosis. The variance swap is the most extreme in this respect. These properties seem to support the idea that variance strategies are "picking up nickels in front of a steamroller".

But what do these monthly return moments mean for the long-term investor? Most of the time, the cumulative wealth of variance strategies is above the wealth level of the market investment, and the downside risk statistics provide an additional perspective on this question. As shown in Panel B, the variance swap has lower downside risk than the market according to all measures, except for the maximum loss. Variance strategies do occasionally experience extreme losses, but these losses are of the same magnitude as the extreme losses of a market investment. Moreover, variance strategies are able to recover from previous losses relatively quickly, i.e., they have the ability to generate strong upside movements in a relatively short period of time. This characteristic is evident after the financial crisis and also after the low point of the Covid-19 pandemic.[12] In this sense, variance strategies are able to pick up more than "nickles" in such periods.

[Insert Table 4 about here.]

Finally, we examine the monthly co-movement of variance strategies with other factor portfolios. In doing so, we seek a better understanding of the economic conditions under which variance premiums are small or large. Table 5 shows the results of regressions of the monthly returns of variance strategies on either the market returns or the returns of all six factor portfolios. There is a positive and statistically significant relationship between the variance strategies and the market. As seen in the previous section, the linear approximation does not fully reveal the true non-linear structure, at least for the delta-hedged put. However, an overall positive co-movement is economically intuitive. It is consistent with large market

[12] Unfortunately, our data set ends in December 2020, but some recovery after the Covid-19 shock is already evident since March 2020.

downturns (within a month) being larger than large market upturns, and the notion of increasing variance in falling equity markets.

[Insert Table 5 about here.]

The only other significant loading is a positive one on the size factor. Strategies selling insurance against high market volatility tend to perform poorly when small caps also perform relatively poorly. This is economically plausible. When times become more volatile, it is likely that small caps will struggle more than large caps because they are more vulnerable on average.[13] Taken together, the market and size factors capture 20% (delta-hedged put) or even 33% (variance swap) of the variance strategies' return fluctuations.

Are variance strategies valuable additions of the long-term investor's opportunity set of stock-related factors? Given the results of Table 5, the answer is ambiguous. Looking at the alphas of the delta-hedged put and the variance swap in the full regression model, they are 0.29% and 0.22% per month, respectively, which is economically significant compared with an average market return of 0.9%. However, these alphas are not statistically significant. One reason is the high residual variance caused by only two observations, October 2008 and March 2020. Moreover, due to the non-robustness of the OLS estimator, these two months are also highly influential for the resulting estimates of the factor loadings.[14] We therefore re-estimate the regression models with a more robust method, the least absolute deviation (LAD) estimator, which minimizes the mean absolute residual error.[15] Table 6 presents the results. A first important finding is that LAD regressions lead to much lower factor loadings for the market and the size factor, which become even insignificant for the delta-hedged put. This finding suggests potentially high diversification benefits if a delta-hedged put strategy complements a market investment, given that

[13] This conjecture is supported by, for example, Duffee (1995) and Ang and Chen (2002).

[14] Cook's distances (Cook (1977)) for these two observations reach from 10 to 226 times the average Cook's distance in the respective models, indicating highly influential observations.

[15] See, for example, Hill and Holland (1977). To account for potential heteroskedasticity and autocorrelation in the LAD framework, we determine standard errors with a block bootstrap method. We use a block length of 12, to be consistent with the number of lags applied for the Newey and West (1987) standard errors of Table 5.

such benefits are measured in terms of mean absolute deviations instead of standard deviations.[16] A second important finding is that the LAD regressions attribute a higher proportion of the variance strategies' mean returns to alpha, as compared to the OLS regressions. The estimated alphas of 0.95% and 0.49% per month of the delta-hedged put and variance swap strategies, respectively, are clearly economically and statistically significant.

[Insert Table 6 about here.]

Overall, the results suggest that variance strategies are attractive factor strategies for long-term investors. These strategies can provide alternatives to a market investment that perform similarly overall but have clearly different downside risk characteristics. They can also be useful complements to a market investment to provide diversification benefits and significant alpha.

5 Conclusion

The variance risk premium is a well-documented empirical phenomenon. In this paper, we analyze whether and how investors can exploit this premium for long-term capital accumulation. Our paper identifies three general problems that arise in long-term variance-based investment strategies. Certain design elements that could mitigate some of these problems are suggested, and corresponding trading strategies are proposed. To determine the variance risk to which strategies are exposed ex-ante, we either consider the maximum exposure based on capital requirements or equalize exposure via a model-free measure of convexity (gamma).

In an empirical study for the S&P 500 index options market, we analyze the performance of different strategies. We compare them to each other and to equity-based

[16] Moving from a pure market investment to a portfolio that invests 50% in the market and 50% in the delta-hedged put strategy reduces the monthly standard deviation by about 20 percent but the mean absolute deviation by more than 40 percent. As the results by Goldstein and Taleb (2007) suggest, even finance professionals consider the mean absolute deviation a more intuitive dispersion measure than the standard deviation.

factor investing strategies. The analysis shows that variance strategies differ substantially in some aspects of risk and return, are significantly positively correlated with the market, and consistently earn premiums over the entire study period. The latter distinguishes variance strategies from other factor strategies, which have not generated premiums since the 2008 financial crisis. However, variance strategies can be hit hard by extreme stock market crashes, but also have the potential to recover quickly from these shocks. All in all, the empirical results show that variance strategies can be attractive to the long-term investor—both as an alternative and as a complement to a market investment—if properly designed. Future research could explore other design elements of variance strategies to further improve long-term variance-based investing. For example, comparing the use of options with different maturities or varying the roll-over periods would shed more light on the finite maturity problem.

Appendix

A1 CBOE Margin Requirements

The following appendix provides a summary of relevant CBOE margin requirements that are used to construct the variance strategies. Further information and sample calculations can be found on the homepage of the CBOE.[17]

Short Call Initial Margin Requirement:

- 100% of option proceeds, plus 15% of aggregate underlying index value (number of contracts × index level × $100) less out-of-the-money amount, if any
- minimum requirement is option proceeds plus 10% of the aggregate underlying index value
- proceeds received from sale of call(s) may be applied to the initial margin requirement
- after position is established, ongoing maintenance margin requirement applies, and an increase (or decrease) in the margin required is possible

Short Put Initial Margin Requirement:

- 100% of option proceeds, plus 15% of aggregate underlying index value (number of contracts × index level × $100) less out-of-the-money amount, if any
- minimum requirement is option proceeds plus 10% of the put's aggregate strike price (number of contracts × strike price × $100)
- proceeds received from sale of puts(s) may be applied to the initial margin requirement
- after position is established, ongoing maintenance margin requirement applies, and an increase (or decrease) in the margin required is possible

[17] `https://www.cboe.com/us/options/strategy_based_margin/`.

Short Straddle Initial Margin Requirement:

- short call(s) or short put(s) requirement, whichever is greater, plus the option proceeds of the other side
- proceeds from sale of entire straddle may be applied to initial margin requirement
- after position is established, ongoing maintenance margin requirement applies, and an increase (or decrease) in the margin required is possible

Short Strangle Initial Margin Requirement:

- short call(s) or short put(s) requirement, whichever is greater, plus the option proceeds of the other side
- proceeds from sale of entire strangle may be applied to initial margin requirement
- after position is established, ongoing maintenance margin requirement applies, and an increase (or decrease) in the margin required is possible

Short Call Spread Initial Margin Requirement:

- the amount by which the short call aggregate strike price is below the long call aggregate strike price (aggregate strike price = number of contracts $\times$ strike price $\times$ \$100)
- long call(s) must be paid for in full
- proceeds received from sale of short call(s) may be applied to the initial margin requirement
- the short call(s) may expire before the long call(s) and not affect margin requirement

Short Put Spread Initial Margin Requirement:

- the amount by which the long put aggregate strike price is below the short put aggregate strike price (aggregate strike price = number of contracts $\times$ strike price $\times$ \$100)
- long put(s) must be paid for in full
- proceeds received from sale of short put(s) may be applied to the initial margin requirement
- the short put(s) may expire before the long put(s) and not affect margin requirement

A2 Model-Free Deltas and Gammas

As Black-Scholes greeks depend on the unrealistic assumptions of the Black-Scholes model, we compute deltas and gammas in a model-free fashion. For illustrative purposes, consider the discrete future stock price S_i with future states of the world $i = 1, \ldots, \bar{i}, \ldots, N$, and the corresponding probabilities q_i. Further, assume that the discount rate is $r = 0$. Now, shift the prices S_i by an infinitesimally small amount ε, such that the new prices and corresponding probabilities are

$$S_1 + \varepsilon, \ldots, S_N + \varepsilon,$$
$$q_1, \ldots, q_N.$$

The value of a plain-vanilla call option C^E with $X = S_{\bar{i}}$ prior to the shift is

$$C^E(X) = \sum_{i=\bar{i}+1}^{N} (S_i - X)\, q_i. \tag{A.1}$$

After the shift, the value is

$$\begin{aligned} C^E(X) &= \sum_{i=\bar{i}}^{N} (S_i + \varepsilon - X)\, q_i \\ &= \sum_{i=\bar{i}+1}^{N} (S_i - X)\, q_i + \sum_{i=\bar{i}+1}^{N} \varepsilon q_i + (S_{\bar{i}} + \varepsilon - X)\, q_{\bar{i}} \\ &= \sum_{i=\bar{i}+1}^{N} (S_i - X)\, q_i + \sum_{i=\bar{i}}^{N} \varepsilon q_i. \end{aligned} \tag{A.2}$$

The delta of any option O is defined as the change in value of the option divided by the change in value of the underlying, that is $\Delta = dO/dS$. Equation (A.2) shows that for $dS = \varepsilon$, the change in value of any option O is $dO = \sum_{i=\bar{i}}^{N} \varepsilon q_i$. Thus,

$$\Delta_C^{mf} = \frac{dO}{dS} = \frac{\sum_{i=\bar{i}}^{N} \varepsilon q_i}{\varepsilon} = \sum_{i=\bar{i}}^{N} q_i \tag{A.3}$$

In the case of a plain vanilla call option, this toy example shows that the model-free delta, which can be derived from the whole risk neutral distribution, is equal to the cumulative probability of all future outcomes that are above the strike price. Conversely, it can be shown that the model-free delta for put options is

$$\Delta_P^{mf} = \frac{dO}{dS} = \frac{-\sum_{i=1}^{\bar{i}} \varepsilon q_i}{\varepsilon} = -\sum_{i=1}^{\bar{i}} q_i. \tag{A.4}$$

Going one step further, it is easy to see that a model-free gamma γ^{mf}—the change in the delta of an option induced by a change in the underlying's price—is simply the level of the probability density function $q_{\bar{i}}$ (PDF) for both call options and put options:

$$\gamma_C^{mf} = \gamma_P^{mf} = q_{\bar{i}} \tag{A.5}$$

References

Ang, A. and J. Chen (2002). Asymmetric correlations of equity portfolios. *Journal of Financial Economics 63*(3), 443–494.

Black, F. (1976). Studies of stock market volatility changes. *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, 177–181.

Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy 81*(3), 637–654.

Bollerslev, T., J. Litvinova, and G. Tauchen (2006). Leverage and volatility feedback effects in high-frequency data. *Journal of Financial Econometrics 4*(3), 353–384.

Breeden, D. T. and R. H. Litzenberger (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business 51*(4), 621–651.

Brière, M., A. Burgues, and O. Signori (2010). Volatility exposure for strategic asset allocation. *Journal of Portfolio Management 36*(3), 105–116.

Cao, J. and B. Han (2013). Cross section of option returns and idiosyncratic stock volatility. *Journal of Financial Economics 108*(1), 231–249.

Cao, J., B. Han, X. Zhan, and Q. Tong (2021). Option return predictability. *Review of Financial Studies*, forthcoming.

Carr, P. and L. Wu (2009). Variance risk premiums. *Review of Financial Studies 22*(3), 1311–1341.

Christie, A. A. (1982). The stochastic behavior of common stock variances—Value, leverage and interest rate effects. *Journal of Financial Economics 10*(4), 407–432.

Clark, S. P. and M. Dickson (2019). Performance expectations of basic options strategies may be different than you think. *Journal of Asset Management 20*(2), 91–102.

Cook, R. D. (1977). Detection of influential observation in linear regression. *Technometrics 19*(1), 15–18.

Cremers, M., M. Halling, and D. Weinbaum (2015). Aggregate jump and volatility risk in the cross-section of stock returns. *Journal of Finance 70*(2), 577–614.

De Fontnouvelle, P., R. P. H. Fishe, and J. H. Harris (2003). The behavior of bid-ask spreads and volume in options markets during the competition for listings in 1999. *Journal of Finance 58*(6), 2437–2463.

Duffee, G. R. (1995). Stock returns and volatility—A firm-level analysis. *Journal of Financial Economics 37*(3), 399–420.

Fallon, W. and J. Park (2016). An asset class characterization of the U.S. equity index volatility risk premium. *Journal of Portfolio Management 43*(1), 72–84.

Fallon, W., J. Park, and D. Yu (2015). Asset allocation implications of the global volatility premium. *Financial Analysts Journal 71*(5), 38–56.

Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of Financial Economics 116*(1), 1–22.

Figlewski, S. (2010). Estimating the implied risk-neutral density for the US market portfolio. In T. Bollerslev, J. R. Russell, and M. Watson (Eds.), *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, pp. 323–353. New York: Oxford University Press.

Ge, W. (2016a). Incorporating smart beta into portfolios: A case study with the volatility risk premium. *Journal of Index Investing 7*(1), 17–24.

Ge, W. (2016b). A survey of three derivative-based methods to harvest the volatility premium in equity markets. *Journal of Investing 25*(3), 48–58.

Ge, W. (2017). Stress-testing volatility risk premium harvesting strategies based on S&P 500 index options. *Journal of Index Investing 8*(1), 37–46.

Goldstein, D. G. and N. N. Taleb (2007). We don't quite know what we are talking about. *Journal of Portfolio Management 33*(4), 84–86.

Goyal, A. and A. Saretto (2009). Cross-section of option returns and volatility. *Journal of Financial Economics 94*(2), 310–326.

Hill, R. W. and P. W. Holland (1977). Two robust alternatives to least-squares regression. *Journal of the American Statistical Association 72*(360a), 828–833.

Israelov, R. and L. N. Nielsen (2015). Covered calls uncovered. *Financial Analysts Journal 71*(6), 44–57.

Kozhan, R., A. Neuberger, and P. Schneider (2013). The skew risk premium in the equity index market. *Review of Financial Studies 26*(9), 2174–2203.

Mayhew, S. (2002). Competition, market structure, and bid-ask spreads in stock option markets. *Journal of Finance 57*(2), 931–958.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703.

Simon, D. P. (2017). Trading the VIX futures roll and volatility premiums with VIX options. *Journal of Futures Markets 37*(2), 184–208.

Simon, D. P. and J. Campasano (2014). The VIX futures basis: Evidence and trading strategies. *Journal of Derivatives 21*(3), 54–69.

Ungar, J. and M. T. Moran (2009). The cash-secured PutWrite strategy and performance of related benchmark indexes. *Journal of Alternative Investments 11*(4), 43–56.

Whaley, R. E. (2002). Return and risk of CBOE Buy Write monthly index. *Journal of Derivatives 10*(2), 35–42.

Figure 1: Stylized Payoffs: Straddle, Strangle, and Butterfly Spread

(a) Straddle Short

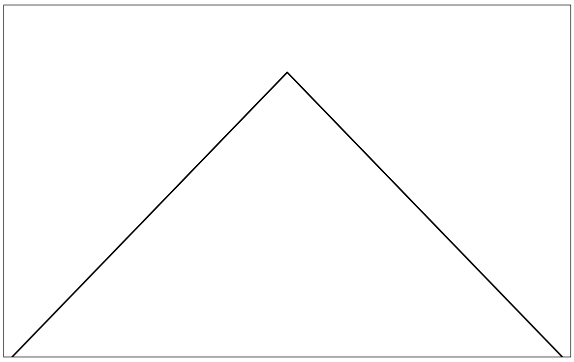

(b) Strangle Short

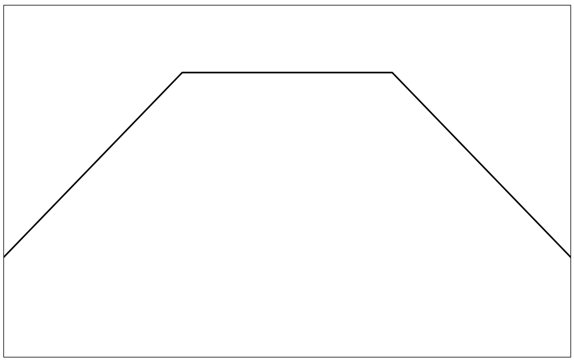

(c) Butterfly Spread Short

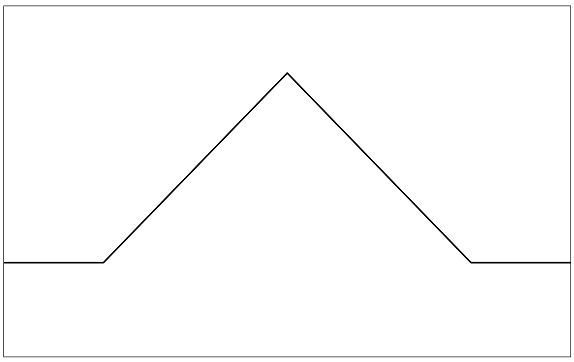

Note: This figure shows stylized payoffs of a short straddle, a short strangle and a short butterfly spread position.

Figure 2: Maximum Exposure Strategies: Cumulative Wealth

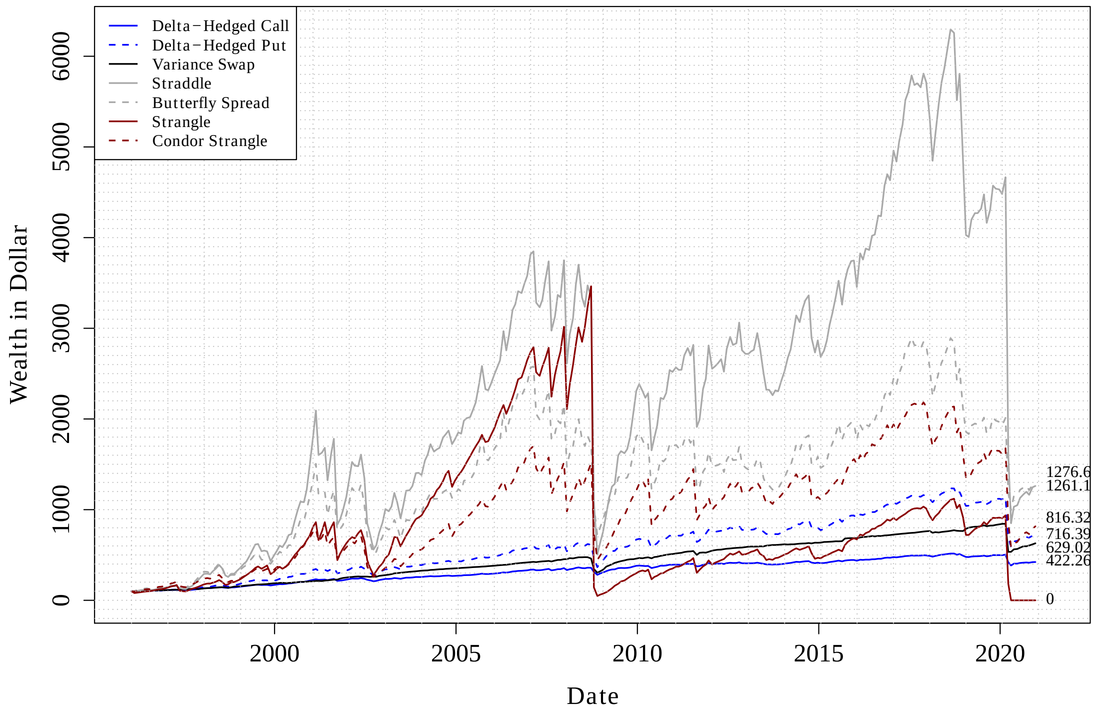

Note: This figure shows the cumulative wealth development of a $100 investment in seven maximum exposure variance risk premium harvesting strategies. The data period covers January 1996 to December 2020.

Figure 3: Equal Exposure Strategies: Cumulative Wealth

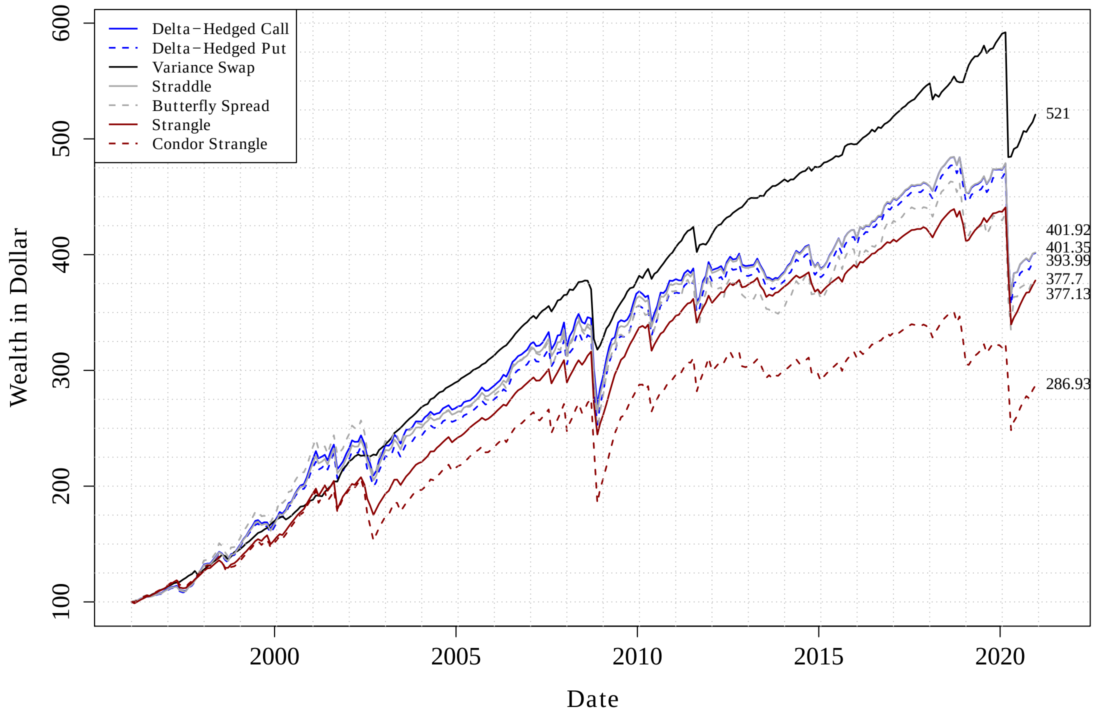

Note: This figure shows the cumulative wealth development of a $100 investment in seven equal exposure variance risk premium harvesting strategies. The data period covers January 1996 to December 2020.

Figure 4: Payoff Plots of Delta-Hedged Put and Variance Swap

(a) Delta-Hedged Put

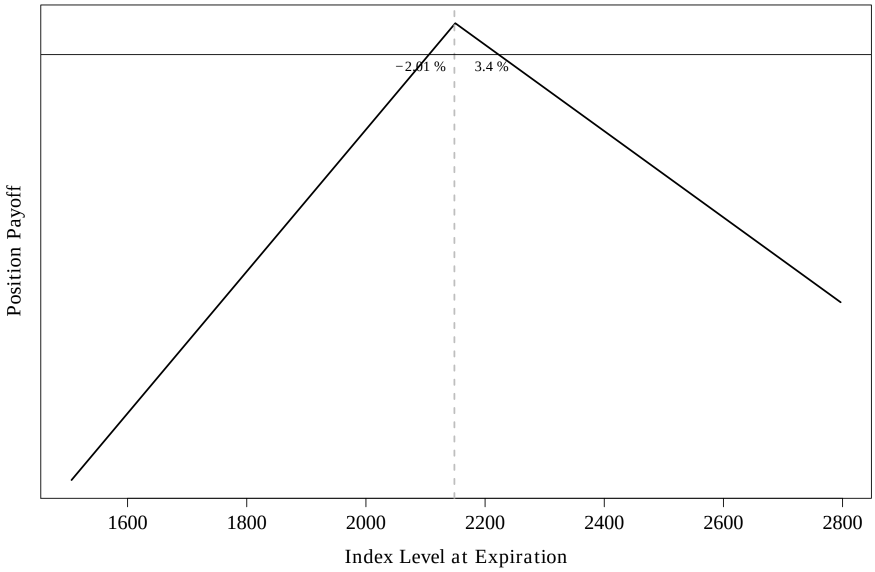

(b) Variance Swap

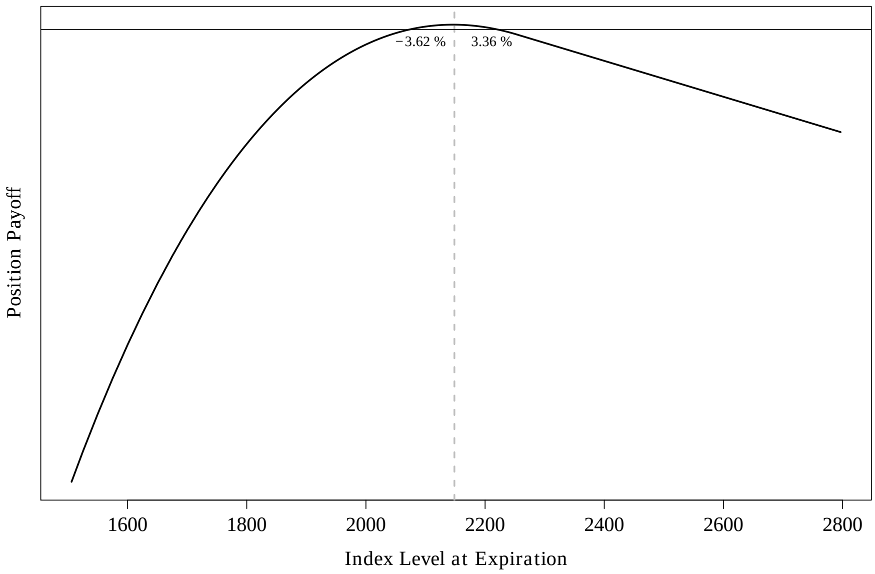

Note: This figure depicts exemplary payoff plots of a delta-hedged put option (part (a)) and a variance swap (part (b)) from October 24, 2016, that is, at initiation of the position. The underlying's index level at expiration is stated on the x-axis and the corresponding position payoff on the y-axis. The horizontal solid black line is the break-even position payoff and the vertical dashed grey line the underlying's index level at initiation.

Figure 5: Equal Exposure Strategies: The Impact of Transaction Costs on Cumulative Wealth

(a) Effective Spread: 0%

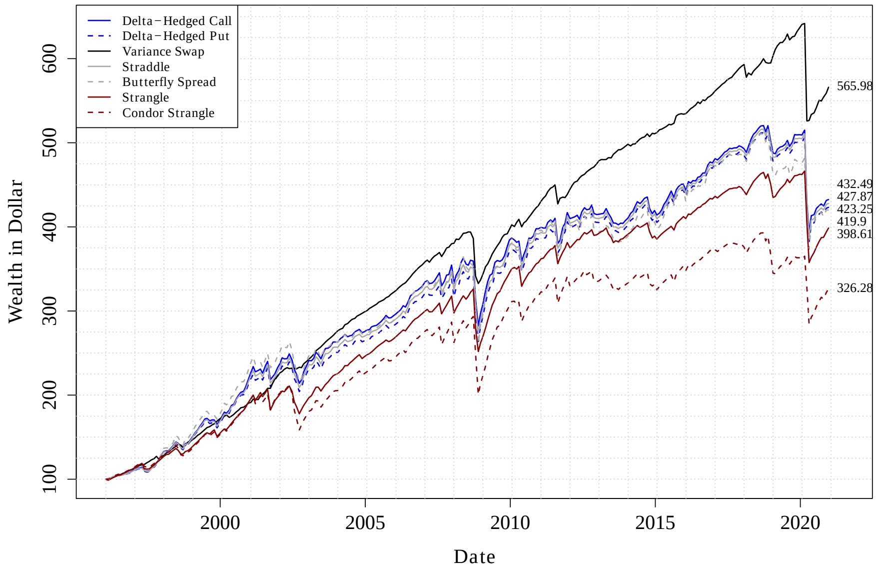

(b) Effective Spread: 50%

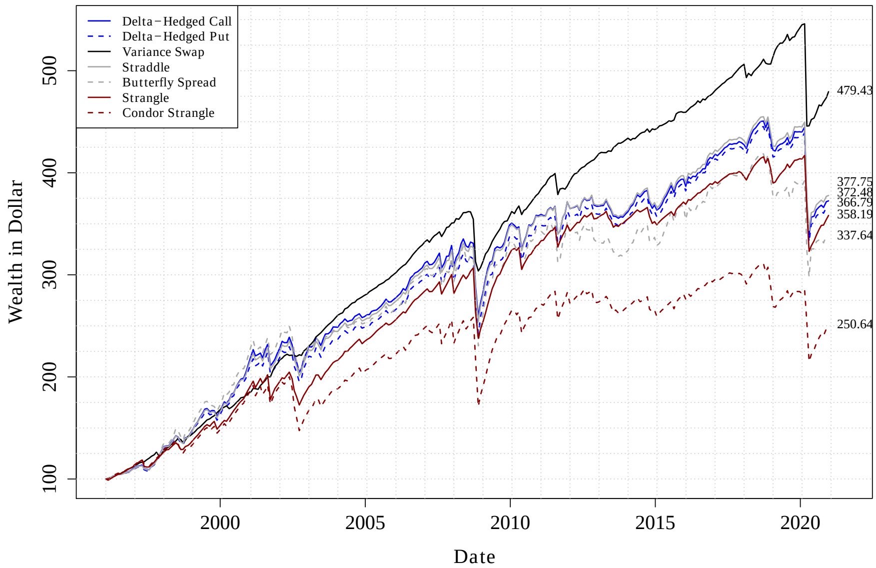

Note: This figure shows the cumulative wealth development of a $100 investment in seven equal exposure variance risk premium harvesting strategies. The data period covers January 1996 to December 2020. Part (a) shows strategies with an effective spread of 0% and part (b) the same strategies with an effective spread of 50%.

Figure 6: Maximum Exposure Strategies vs. Market

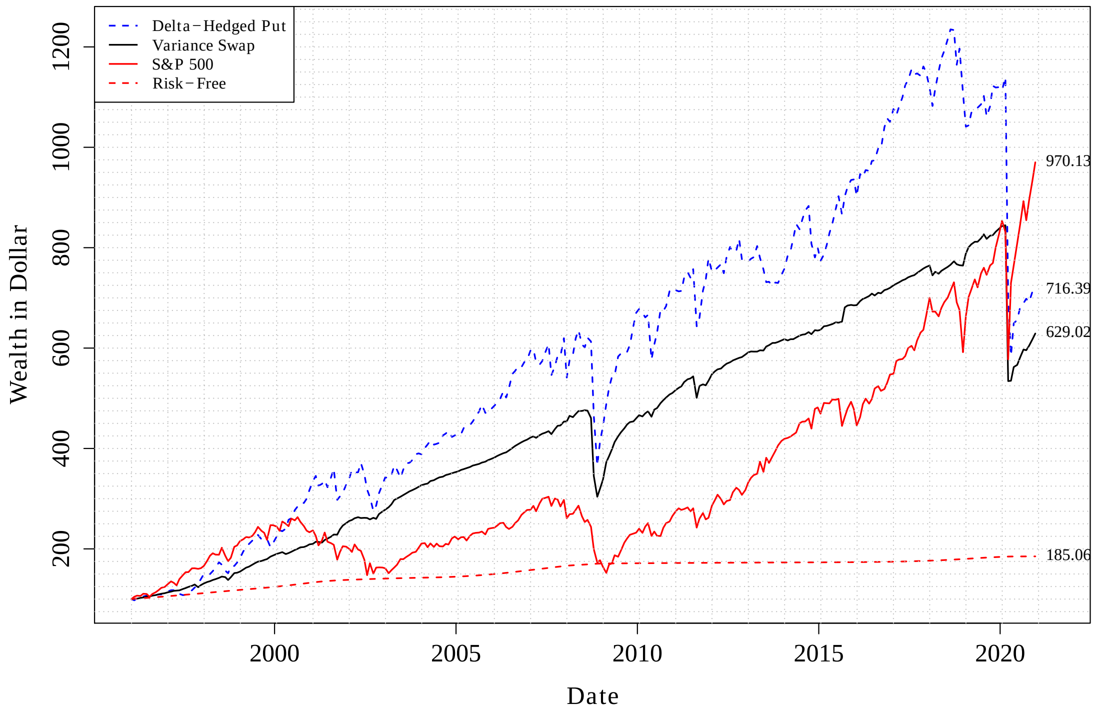

Note: This figure depicts the cumulative wealth development of a $100 investment in four different trading strategies. First, the maximum exposure delta-hedged put (dashed blue line) and variance swap (solid black line) strategies are stated. Second, a buy and hold investment strategy in the S&P 500 total return index is shown (solid red line). Third, an investment in risk-free assets (dashed red line) is shown. The data period covers January 1996 to December 2020.

Figure 7: Return Scatter Plots with Market

(a) Delta-Hedged Put

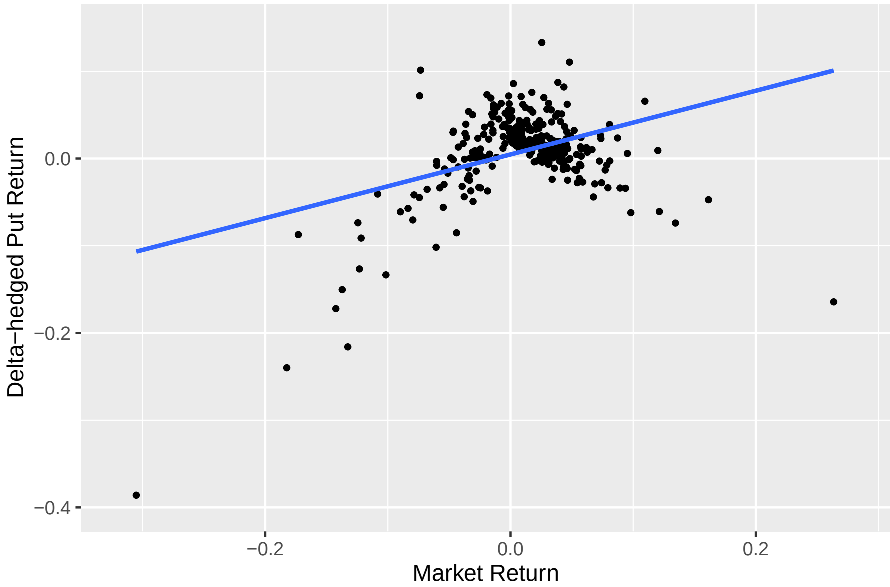

(b) Variance Swap

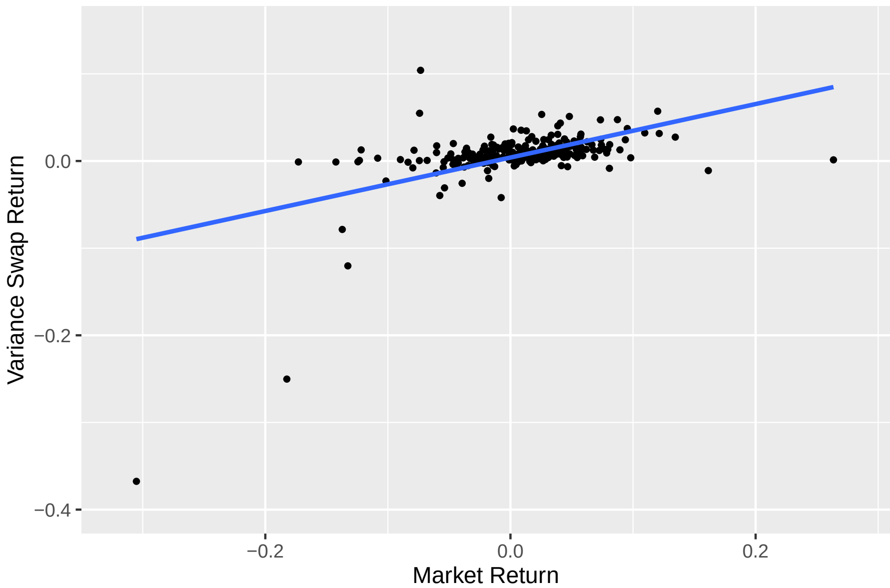

Note: This figure provides return scatter plots that outline the contemporaneous relationship between monthly market returns and monthly delta-hedged put returns (part (a)) and monthly variance swap returns (part (b)), respectively. The blue lines indicate a linear trend estimated with least squares.

Figure 8: Factor Investment Style Performance

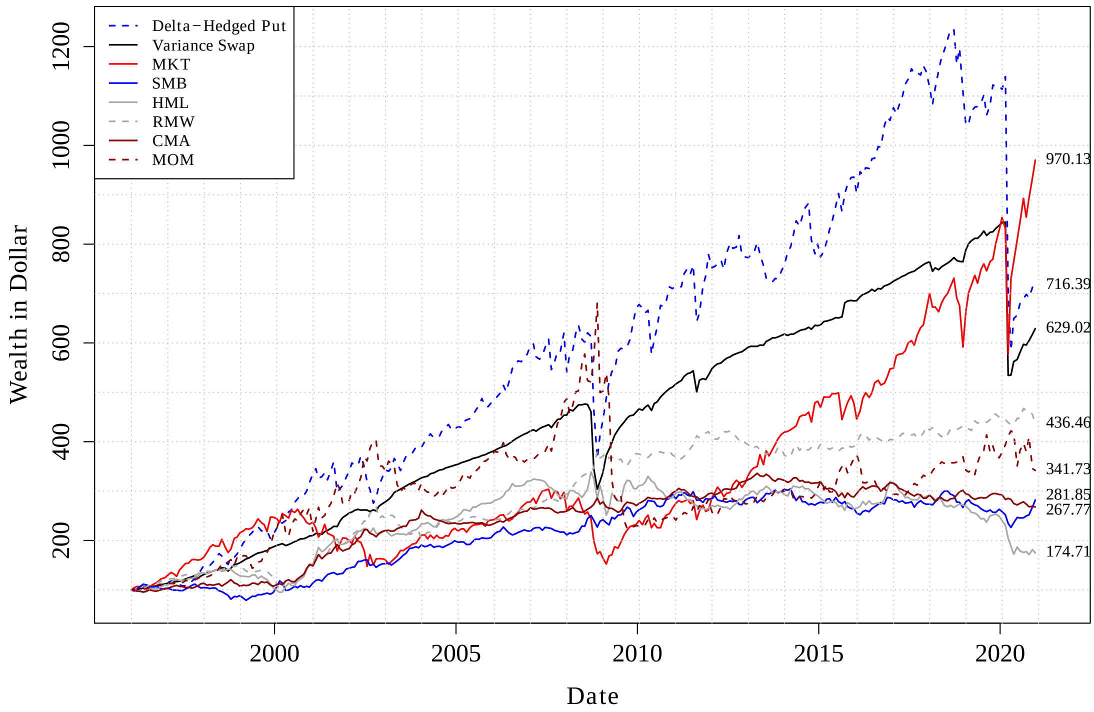

Note: This figure provides a comparison between factor investment style performances and two maximum exposure variance risk premium harvesting strategies (delta-hedged put and variance swap). The data period covers January 1996 to December 2020. Factor returns are calculated according to Fama and French (2015). MKT is the market factor, SMB is the size factor, HML is the value factor, RMW is the profitability factor, CMA is the investment factor, and MOM is the momentum factor.

Table 1: Return- and Risk-Statistics of Maximum Exposure Strategies

Panel A: Basic Monthly Summary Statistics

	Delta-Hedged Call	Delta-Hedged Put	Variance Swap	Straddle	Butterfly Spread	Strangle	Condor Strangle
Mean	0.0051	0.0080	0.0067	0.0182	0.0140	0.0157	0.0129
Standard Dev.	0.0243	0.0494	0.0309	0.1215	0.0986	0.1395	0.0975
Skewness	−2.6440	−2.9729	−8.1252	−2.2880	−1.4947	−4.0907	−2.7520
Exc. Kurtosis	12.9533	17.1776	88.7758	9.4223	3.6283	23.1548	8.6940
Sharpe Ratio	0.1266	0.1195	0.1516	0.1330	0.1208	0.0976	0.1112

Panel B: Downside Risk Statistics

	Delta-Hedged Call	Delta-Hedged Put	Variance Swap	Straddle	Butterfly Spread	Strangle	Condor Strangle
VaR (95%)	−0.0302	−0.0629	−0.0076	−0.1716	−0.1767	−0.1949	−0.1900
CVaR (95%)	−0.0697	−0.1448	−0.0700	−0.3632	−0.2832	−0.4448	−0.3228
Max Loss	−0.1670	−0.3860	−0.3676	−0.6729	−0.4106	−1.0000	−0.5085
Average DD	−0.0372	−0.0755	−0.0433	−0.1934	−0.2295	−0.2733	−0.2006
Max DD	−0.2550	−0.5266	−0.3676	−0.8684	−0.7238	−1.0000	−0.7384
Average DD Length	5.2683	5.7000	4.3750	12.8000	19.0714	18.8462	14.5882
Max DD Length	28	29	23	110	122	148	112

Panel C: Annualized Geometric Mean Return and Terminal Wealth

	Delta-Hedged Call	Delta-Hedged Put	Variance Swap	Straddle	Butterfly Spread	Strangle	Condor Strangle
Ann. Geom. Return [%]	5.95	8.22	7.66	10.71	10.76	−100.00	8.79
Term. Wealth [$]	422.26	716.39	629.02	1261.13	1276.63	0.00	816.32

Note: This table provides return- and risk-statistics of seven maximum exposure variance risk premium harvesting strategies. Panel A depicts basic summary statistics for monthly returns, while Panel B is dedicated to downside risk statistics. In particular, we consider asymmetric risk metrics for monthly returns such as the 95% value-at-risk (VaR), the 95% conditional value-at-risk (CVar), and the maximum loss (max loss). Additionally, we determine path dependent drawdown measures: The average drawdown (Average DD) and the maximum drawdown (Max DD), as well as the average drawdown length (Average DD length) and the maximum drawdown length (Max DD Length). Lastly, Panel C shows the annualized geometric return and the terminal wealth of the strategies.

Table 2: Return- and Risk-Statistics of Equal Exposure Strategies

Panel A: Basic Monthly Summary Statistics

	Delta-Hedged Call	Delta-Hedged Put	Variance Swap	Straddle	Butterfly Spread	Strangle	Condor Strangle
Mean	0.0050	0.0049	0.0057	0.0050	0.0049	0.0047	0.0040
Standard Dev.	0.0238	0.0248	0.0156	0.0242	0.0290	0.0226	0.0296
Skewness	−2.7361	−2.7109	−7.8579	−2.7220	−1.9814	−3.5094	−3.1060
Exc. Kurtosis	13.9543	13.7783	83.5056	13.7999	9.4847	16.0801	13.8736
Sharpe Ratio	0.1213	0.1151	0.2317	0.1202	0.0974	0.1175	0.0653

Panel B: Downside Risk Statistics

	Delta-Hedged Call	Delta-Hedged Put	Variance Swap	Straddle	Butterfly Spread	Strangle	Condor Strangle
VaR (95%)	−0.0294	−0.0316	−0.0067	−0.0303	−0.0397	−0.0323	−0.0536
CVaR (95%)	−0.0691	−0.0724	−0.0374	−0.0702	−0.0812	−0.0738	−0.0966
Max Loss	−0.1670	−0.1705	−0.1821	−0.1673	−0.1697	−0.1417	−0.2064
Average DD	−0.0352	−0.0355	−0.0229	−0.0346	−0.0515	−0.0447	−0.0579
Max DD	−0.2457	−0.2535	−0.1821	−0.2467	−0.2892	−0.2293	−0.3270
Average DD Length	5.1951	5.1220	4.1538	5.0476	6.6765	5.5758	6.8485
Max DD Length	28	28	18	28	33	17	41

Panel C: Annualized Geometric Mean Return and Terminal Wealth

	Delta-Hedged Call	Delta-Hedged Put	Variance Swap	Straddle	Butterfly Spread	Strangle	Condor Strangle
Ann. Geom. Return [%]	5.74	5.66	6.85	5.74	5.47	5.48	4.32
Term. Wealth [$]	401.35	393.99	521.00	401.92	377.13	377.70	286.93

Note: This table provides return- and risk-statistics of seven equal exposure variance risk premium harvesting strategies. Panel A depicts basic summary statistics for monthly returns, while Panel B is dedicated to downside risk statistics. In particular, we consider asymmetric risk metrics for monthly returns such as the 95% value-at-risk (VaR), the 95% conditional value-at-risk (CVar), and the maximum loss (max loss). Additionally, we determine path dependent drawdown measures: The average drawdown (Average DD) and the maximum drawdown (Max DD), as well as the average drawdown length (Average DD length) and the maximum drawdown length (Max DD Length). Lastly, Panel C shows the annualized geometric return and the terminal wealth of the strategies.

Table 3: Equal Exposure Strategies' Return Correlation Matrix

	Delta-Hedged Call	Delta-Hedged Put	Variance Swap	Straddle	Butterfly Spread	Strangle
Delta-Hedged Put	0.9960					
Variance Swap	0.6631	0.6584				
Straddle	0.9990	0.9990	0.6608			
Butterfly Spread	0.9723	0.9751	0.6043	0.9746		
Strangle	0.9306	0.9311	0.6236	0.9319	0.8601	
Condor Strangle	0.9093	0.9156	0.5352	0.9133	0.9000	0.9536

Note: This table depicts the correlation of monthly returns between seven equal exposure variance risk premium harvesting strategies.

Table 4: Return- and Risk-Statistics of Different Factor Investment Styles

Panel A: Basic Monthly Summary Statistics

	Delta-Hedged Put	Variance Swap	MKT	SMB	HML	RMW	CMA	MOM
Mean	0.0080	0.0067	0.0090	0.0039	0.0025	0.0053	0.0035	0.0057
Standard Dev.	0.0494	0.0309	0.0525	0.0298	0.0347	0.0266	0.0217	0.0539
Skewness	−2.9729	−8.1252	−0.9071	−0.0434	0.2812	0.3554	0.5697	−1.6832
Exc. Kurtosis	17.1776	88.7758	6.7505	2.3063	2.6527	6.0900	1.4345	8.4620
Sharpe Ratio	0.1195	0.1516	0.1332	0.0622	0.0116	0.1215	0.0676	0.0670

Panel B: Downside Risk Statistics

	Delta-Hedged Put	Variance Swap	MKT	SMB	HML	RMW	CMA	MOM
VaR (95%)	−0.0629	−0.0076	−0.0747	−0.0399	−0.0432	−0.0302	−0.0298	−0.0900
CVaR (95%)	−0.1448	−0.0700	−0.1323	−0.0609	−0.0728	−0.0566	−0.0396	−0.1491
Max Loss	−0.3860	−0.3676	−0.3052	−0.1341	−0.1278	−0.1078	−0.0548	−0.3401
Average DD	−0.0755	−0.0433	−0.0791	−0.0643	−0.0710	−0.0571	−0.0616	−0.1025
Max DD	−0.5266	−0.3676	−0.4993	−0.2948	−0.4942	−0.3278	−0.2032	−0.6714
Average DD Length	5.7000	4.3750	7.1875	11.2500	13.7368	12.4000	15.4706	13.3500
Max DD Length	29.0000	23.0000	74.0000	82.0000	148.0000	65.0000	93.0000	146.0000

Panel C: Annualized Geometric Mean Return and Terminal Wealth

	Delta-Hedged Put	Variance Swap	MKT	SMB	HML	RMW	CMA	MOM
Ann. Geom. Return [%]	8.22	7.66	9.55	4.25	2.26	6.09	4.03	5.06
Term. Wealth [$]	716.39	629.02	970.17	281.85	174.71	436.46	267.77	341.73

Note: This table provides return- and risk-statistics of two maximum exposure variance risk premium harvesting strategies (delta-hedged put and variance swap) and 6 other equity-based factor investment strategies. Returns of factor investment strategies are calculated according to the Fama and French (2015) Five-Factor Model, such that the return periods coincide with the return periods from the variance risk premium strategies. MKT is the market factor, SMB is the size factor, HML is the value factor, RMW is the profitability factor, CMA is the investment factor, and MOM is the momentum factor. Panel A depicts basic summary statistics for monthly returns, while Panel B is dedicated to downside risk statistics. In particular, we consider asymmetric risk metrics for monthly returns such as the 95% value-at-risk (VaR), the 95% conditional value-at-risk (CVar), and the maximum loss. Additionally, we determine path dependent drawdown measures: The average drawdown (Average DD) and the maximum drawdown (Max DD), as well as the average drawdown length (Average DD length) and the maximum drawdown length (Max DD Length). Lastly, Panel C shows the annualized geometric return and the terminal wealth of the strategies.

Table 5: Maximum Exposure Strategies vs. Factor Investment Styles: OLS Regression

	Delta-Hedged Put	Variance Swap	Delta-Hedged Put	Variance Swap
α	0.0034	0.0025	0.0029	0.0022
	(1.2467)	(1.1412)	(1.0350)	(0.9005)
β_{MKT}	0.3648***	0.3080***	0.3276***	0.2845***
	(3.8692)	(2.9263)	(4.4556)	(2.9022)
β_{SMB}			0.3236**	0.2231**
			(1.9799)	(2.3091)
β_{HML}			0.2907	0.1759
			(1.5783)	(1.3132)
β_{RMW}			0.0973	0.0385
			(0.6987)	(0.3905)
β_{CMA}			−0.1770	−0.1323
			(−0.7123)	(−1.5327)
β_{MOM}			−0.0031	0.0259
			(−0.0388)	(0.4390)
n	299	299	299	299
Adj. R^2	0.1480	0.2718	0.2072	0.3304
F-statistic	52.7617	112.2456	13.9801	25.5060

Note: This table shows the regression results from a regression of two maximum exposure variance risk premium harvesting strategies on other factor investment strategies. Specifically, we consider the S&P 500 index (MKT), long-short portfolios of the remaining four factors of the Fama and French (2015) Five-Factor Model (Size (SMB), Value (HML), Profitability (RMW), Investment (CMA)), plus Momentum (MOM). Newey and West (1987) robust standard errors with 12 lags are reported in parentheses and asterisks indicate significance at the 1% (***), 5% (**), and 10% (*) level, respectively.

Table 6: Maximum Exposure Strategies vs. Factor Investment Styles: LAD Regression

	Delta-Hedged Put	Variance Swap	Delta-Hedged Put	Variance Swap
α	0.0105***	0.0051***	0.0095***	0.0049***
	(0.0027)	(0.0008)	(0.0030)	(0.0009)
β_{MKT}	0.1246	0.0915***	0.0852	0.0826***
	(0.1341)	(0.0239)	(0.1362)	(0.0255)
β_{SMB}			0.1447	0.0736***
			(0.1079)	(0.0275)
β_{HML}			−0.0120	0.0200
			(0.1192)	(0.0323)
β_{RMW}			0.1752	0.0021
			(0.1618)	(0.0329)
β_{CMA}			−0.0844	−0.0016
			(0.2227)	(0.0437)
β_{MOM}			−0.0640	−0.0160
			(0.0656)	(0.0267)
n	299	299	299	299
R^2_{Pseudo}	0.0054	0.0777	0.0216	0.0999

Note: This table shows the least absolute deviation (LAD) regression results from a regression of two maximum exposure variance risk premium harvesting strategies on other factor investment strategies. Specifically, we consider the S&P 500 index (MKT), long-short portfolios of the remaining four factors of the Fama and French (2015) Five-Factor Model (Size (SMB), Value (HML), Profitability (RMW), Investment (CMA)), plus Momentum (MOM). Standard errors are determined using a block bootstrap with a block length of 12 and 20,000 bootstrap samples for each individual regression model. They are reported in parentheses. Asterisks indicate significance at the 1% (***), 5% (**), and 10% (*) level, respectively.