The Shape of the Pricing Kernel and Expected Option Returns

Christian Schlag*

Tobias Sichert†

This version: November 4, 2021

Abstract

A growing literature analyzes the cross-section of single stock option returns, virtually always under the (implicit or explicit) assumption of a monotonically decreasing pricing kernel. Using option returns, we non-parametrically provide significant and robust evidence that the pricing kernel as a function of single stock returns is indeed U-shaped. This shape of the pricing kernel has strong implications for the impact of volatility on expected options returns. For example, we show both theoretically and empirically that higher volatility can increase or decrease expected call option returns, depending on moneyness. Furthermore, on the basis of a U-shaped pricing kernel, we shed new light on some recent findings from the literature on expected option returns, such as anomalies related to ex-ante option return skewness and to lottery characteristics of the underlying stock.

Keywords: Pricing kernel, equity options, option returns, volatility, expected option returns, cross-section of option returns

JEL: G0, G12, G13

* Goethe University Frankfurt and Leibniz Institute for Financial Research SAFE, Theodor-W.-Adorno-Platz 3, 60323 Frankfurt, Germany. E-mail: schlag@finance.uni-frankfurt.de.

† Stockholm School of Economics and Swedish House of Finance, Drottninggatan 98, Stockholm SE-111 60, Sweden. E-mail: tobias.sichert@hhs.se, Tel.: +46 8 736 91 00. Corresponding author.

The authors would like to thank Magnus Dahlquist, Burton Hollifield, Duane Seppi, Paolo Sodini, Per Strömberg, Chris Telmer, Julian Thimme, Gyuri Venter, and seminar participants at Goethe University, Carnegie Mellon (Tepper), NFN Young Scholars Finance Webinar Series, Swedish House of Finance, DGF and the Virtual Derivatives Workshop for valuable comments and suggestions.

1 Introduction

A growing body of research analyzes the returns of single stock equity options, and the richness of option data makes it possible to study a large variety of important questions. On the one hand, there are classical issues from the area of asset pricing, like the identification of risk factors which can explain the observed option returns, and along similar lines, the relation of option returns to the returns of their underlying securities. On the other hand, based on these data, one can address even more fundamental economic issues, such as, e.g., investor preferences, the trading behavior of market participants, or the transmission of information between the stock and option market.

Starting with Coval and Shumway (2001), virtually all studies of expected option returns have relied on the paradigm of a pricing kernel that is monotonically decreasing in the return of the underlying. Here, the term "pricing kernel" (PK) stands for the projection of the stochastic discount factor (SDF) on the return of the underlying asset.

In this paper, we investigate the shape of the PK in a non-parametric fashion using returns on single-stock options, and we find strong evidence for a U-shape. This is interesting and relevant, since most asset pricing models imply that the PK is a monotonically decreasing function of the return on the underlying wealth portfolio (empirically frequently approximated by a broad stock market index). Prominent examples are classics like the CAPM or the approach suggested by Rubinstein (1976), but also more recent theories, such as, e.g., the long-run risk model developed by Bansal and Yaron (2004), share this property. While many models exclusively study the aggregate portfolio (i.e., the index) and remain silent about single stocks, one can transfer the monotonicity argument to the typical stock, since all stocks taken together must ultimately aggregate to the index. From a different perspective, also workhorse factor models used for single stocks as e.g., the three-factor model proposed by Fama and French (1993) and its successors, imply that the PK is monotonically decreasing in the return of the typical single stock.

There is, however, growing evidence that the PK as a function of the returns of the index is non-monotonic and actually U-shaped (see, e.g., Bakshi et al. (2010), Christoffersen et al. (2013), and Babaoğlu et al. (2017)). The typical finding is that after initially decreasing from very high values for large negative returns, the PK is increasing again for large positive returns. Since the index return is a weighted average of individual stock returns, this result raises the question whether the projection of the SDF on the returns of single stocks also exhibits this U-shaped pattern.

The answer to this question is important for at least two reasons. On the one hand, it is informative about the relative pricing of options on the index versus options on individual stocks. For example, Chaudhuri and Schroder (2015) argue that the two types of options are priced by different PKs, and, while the PK for the index is non-monotonic, the analogous quantity for the typical stock is monotonically decreasing. This results seems rather puzzling, since it either implies pricing inconsistencies between index and single stock options, or, since the index is aggregated from individual stocks, it would require rather counter-intuitive joint dynamics for the stocks, first of all negatively correlated returns, for which one would have a hard time finding empirical support. On the other hand, we know next to nothing about the implications of a U-shaped PK for the cross-section of single stock option returns, where the term "cross-section" refers to different underlying securities.[1]

Our paper offers three main contributions, where the overarching link is the relation between the shape of the PK and expected option returns. First, via a simple test on single stock option returns, we show that the PK as a function of the returns on a typical stock is indeed U-shaped. Second, this finding of a U-shaped PK has far reaching implications for the impact of volatility on expected option returns. For example, in contrast to the case of a decreasing PK, an increase in volatility can either lead to either increasing or decreasing expected call returns, depending on moneyness. Third, we show that a U-shaped PK offers a risk-based explanation for a number of option return patterns considered anomalous in the literature.

[1] When we refer to the properties of option returns across different strike prices, but for the same underlying asset, we state this explicitly.

Concerning the first result, our approach is fully non-parametric and we find that our result is robust with respect to a large number of variations in the empirical methodology, and it is furthermore not limited to periods of either high or low volatility, to a specific return horizon, or to options with specific characteristics (such as low liquidity). This finding of a U-shaped PK is important in its own right, since the PK is the key object of interest in modern asset pricing, and its shape provides fundamental insights into the relevance of different sources of risks and the associated risk premia, as well as into investor preferences. Furthermore, such robust and non-parametric evidence is important in general, since the shape of the PK is still vividly debated (see e.g. Linn et al. (2018) and Barone-Adesi et al. (2020) for recent examples).

The use of single stock option returns comes with (at least) two advantages. First, the empirical result of a U-shaped PK becomes much more robust and consequently also more significant in a statistical sense, since there are many more securities available compared to an investigation based on index options. For example, for S&P 500 index options Bakshi et al. (2010) find results which are qualitatively very similar to ours, but not statistically significant.[2] Second, since we find the PK for single stocks to be also U-shaped, consistent with the index, the two option markets do not appear to be segmented in a particularly strong fashion, as was recently suggested by, among others, Lemmon and Ni (2014).

In our analysis we follow Bakshi et al. (2010) and rely on claims that pay off in the positive range of the return distribution, such as calls and bull spreads. A monotonically decreasing PK implies that expected call returns are monotonically increasing in the strike price and always higher than the expected return of the underlying. In contrast to this, a U-shaped PK implies expected call returns will only increase up to a certain strike and decrease again afterwards.

To test the hypothesis of a U-shaped PK, we sort call options into bins with respect to their moneyness X (defined as the ratio of strike price to current stock price) and find exactly the pattern of average returns consistent with this shape. In particular, average call returns are found to be increasing in moneyness until X is between 1.025 and 1.05, but decreasing for X greater than 1.05. In addition, for high X, the average return even turns negative and

[2] This is also the case for the more recent sample ranging until Dec 2017 (see Online Appendix).

is also lower than the average return of the underlying. These results strongly contradict the main implications of a monotonically decreasing PK, namely that expected call returns are increasing in moneyness and always greater than the expected return of the underlying.

The theoretical explanation for the empirical pattern for the single stock PK is likely to be similar to that proposed for the index. For example, Chabi-Yo (2012) and Christoffersen et al. (2013) both suggest a PK which is decreasing in the return of the underlying and increasing in the conditional return variance. The existence of a variance risk premium in combination with a U-shaped relationship between returns and variance then induces an U-shaped projection of the SDF onto the return of the underlying. A similar argument is brought forward by Bakshi et al. (2019). [3]

The paper closest to us with respect to testing PK montonicity using option returns is Chaudhuri and Schroder (2015), who argue in favor of a decreasing PK. The main source for this divergence between their result and ours, are differences in the fundamental empirical approach. Chaudhuri and Schroder (2015) use intermediate option returns with a one week holding period, while we study monthly, hold-until-maturity returns which is the standard in the literature on option returns, and is furthermore widely considered as more robust in the presence of noise in the data (Broadie et al., 2009; Duarte et al., 2019). Furthermore, Chaudhuri and Schroder (2015) pool options with widely varying times to maturity, which have rather heterogeneous return properties. In addition to this, they only qualitatively take the options' moneyness into account, while moneyness is the key variable in our analysis in a quantitative sense. We perform extensive robustness checks and show that our results are robust with respect to a large set of variations in the empirical approach, including those suggested by Chaudhuri and Schroder (2015).

Our second main contribution is to show both theoretically and empirically that the shape

[3] A U-shaped PK can also be generated in a model with rank-dependent utility and probability weighting (Polkovnichenko and Zhao, 2013) or cumulative prospect theory and probability weighting (Baele et al., 2019). In contrast to these papers, we base our analysis on the aforementioned fully rational explanation, which, in addition to matching the characteristics of index options, also captures the properties of the cross-section of single stock option returns.

of the PK is an important determinant for the impact of volatility on expected option returns. Volatility is arguably the most important variable in the context of options, and its relation to option returns has recently received increased attention. For example, Hu and Jacobs (2020) study the effect of total volatility on expected option returns. Aretz et al. (2018) refine the analysis and show that idiosyncratic and systematic volatility have qualitatively different effects on expected option returns, and a similar analysis can be found in Chaudhury (2017). However, all these papers derive their theoretical implications in a framework with a monotonically decreasing PK. We consider the impact of a U-shaped PK, and thus contribute to this literature by offering comprehensive and more general theoretical and empirical analyses of the effects of volatility on the cross-section of call option returns.

To this end, we propose a new model featuring a U-shaped PK and a cross-section of stocks with both idiosyncratic and systematic volatility components. In the presence of a U-shaped PK, the model shows that an increase in the idiosyncratic volatility of a stock can either increase or decrease expected call option returns, depending on moneyness. This is in contrast to the case of a monotonically decreasing PK, where an increase in idiosyncratic volatility always decreases expected call returns. The intuition behind the new result is as follows: In many option pricing models, with increasing idiosyncratic volatility, the expected return of any call option with a positive strike converges to the expected return of the underlying. As stated above, however, with a U-shaped PK expected call returns are positive (and higher than that of the underlying) for low moneyness and negative (and lower than that of the underlying) for high moneyness. Together with the previous result, this implies that with higher idiosyncratic volatility expected call option returns will decrease for low and increase for high moneyness, since for low moneyness, they will approach the expected return of the underlying from above, while the opposite is true when moneyness is high. We test this prediction empirically, and find it strongly supported by the data.

Since total volatility is highly correlated with idiosyncratic volatility in the cross-section of stocks (see, e.g., Ang et al. (2006)), it is not surprising that, for total volatility, we find a pattern that is very similar to the one described for idiosyncratic volatility. This is is a very important finding, as it implies that the results presented by Hu and Jacobs (2020) regarding the behavior of expected returns on at-the-money calls with respect to total volatility cannot simply be generalized to a wider range of strikes. While the pattern of results is pretty clear for idiosyncratic and total volatility, the case is not so clear-cut for systematic volatility, where expected returns for calls of a given moneyness can be increasing or decreasing, with the respective moneyness ranges depending on the parameters of the data-generating process.

Finally, as our third main contribution, we revisit several anomalies documented in the literature on option returns. Recent papers like Ni (2008), Boyer and Vorkink (2014), Byun and Kim (2016), Blau et al. (2016), and Filippou et al. (2018, 2019) argue that investor preferences for lottery-like (right-skewed) payoffs affect asset prices and lead to negative expected option returns.[4] These studies suggest that investors have strong preferences for lottery-like payoffs, i.e., they seem to accept large losses on average from investing in out-of-the-money call options, in exchange for pronounced positive skewness. Given the competitive nature of option markets it seems somewhat unlikely, however, that it is demand pressure by retail investors that causes higher option prices and thus lower expected returns, and not additional risk compensation required by market participants selling these contracts.

The lottery-like character of an option has been identified in the literature mostly in two ways, either via the ex ante skewness measure proposed by Boyer and Vorkink (2014) applied to the options themselves or via the maximum daily past return of the underlying stock as suggested by Bali et al. (2011). The paper then provides evidence that the expected returns on options where these characteristics are high tend to be low.

Our findings regarding the shape of the PK also shed new light on this issue. In the presence of a U-shaped PK, we show that after properly controlling for the effect of an option's moneyness, the impact of ex ante skewness on expected option returns disappears almost completely. Regarding the use of the maximum daily return, we confirm the finding that at-

[4] Examples for papers studying other aspects of returns on single stock option are Driessen et al. (2009), Goyal and Saretto (2009), Muravyev (2016), and Vasquez (2017). In addition, there are several papers that study delta-hedged option returns, which we consider an entirely different object than raw returns.

the-money expected call option returns decrease in this variable. However, the direction of monotonicity reverses for out-of-the-money calls. This result substantially challenges the above gambling argument, since, if anything, the effect should be stronger for out-of-the-money calls which exhibit even more pronounced skewness. In addition, the maximum daily return is highly correlated with idiosyncratic and total volatility, so that the results we find for this variable are not surprising, given the behavior of expected call returns with increasing idiosyncratic and total volatility discussed above. In sum, our analysis suggests that the fact that options with higher skewness have lower returns, is perfectly consistent with a risk-based explanation, which we offer via the U-shaped PK.

The remainder of the paper is organized as follows. Section 2 presents the setup and result of the test for a U-shaped PK. Section 3 analyzes the effect of the U-shaped PK on the cross-section of single stock option returns at the example of the underlying's stock volatility, and Section 4 examines these effects in existing findings in the literature. In Section 5 we present robustness analyses, and Section 6 concludes. The appendix contains detailed additional information on our analysis. In addition, we provide an online appendix.[5]

2 Pricing Kernel for the Typical Stock

2.1 Theoretical considerations

The fundamental research question in this paper is whether the PK for the typical stock, i.e., the projection of the SDF on single stock returns, is U-shaped or monotonically decreasing. Bakshi et al. (2010) point out that only contingent claims with a non-zero payoff in the range of (large) positive returns on the underlying can be used as an instrument to distinguish between the two potential shapes of the PK. This makes sense intuitively, since a monotonically decreasing and a U-shaped PK share the region with high positive values for (large) negative returns of the underlying, but differ in the right tail of the return distribution. Examples for such claims are

[5] https://sites.google.com/site/tsichert/.

standard calls and digital calls (also sometimes called cash-or-nothing or binary options), where the latter pays 1$ if the underlying at maturity is greater than or equal to the strike, and zero otherwise. Analogously, put options are not informative in this context, and this is why we exclusively focus our analysis on calls.

To formalize the intuition a little bit more, we rely on the properties of expected option returns under the two alternative PK scenarios. As shown by Coval and Shumway (2001), under a monotonically decreasing PK, (i) expected call returns are monotonically increasing in the strike, (ii) expected digital option returns are monotonically increasing in their strike, and (iii) expected digital option returns are always below call returns for the same strike.

The situation is very different under a U-shaped pricing kernel, as shown by Bakshi et al. (2010). There, (i) expected call returns are increasing in the strike price until a certain point K^*, and then monotonically decreasing for $K > K^*$, (ii) expected digital call returns are increasing in the strike until a certain point $K^{**} > K^*$, and then monotonically decreasing, and (iii) expected digital call returns are below expected call returns for $K < K^*$, and higher afterwards.

Figure 1 illustrates the relationships for the two cases graphically, and one can see that the structure of expected call returns across strikes is basically inverse to that of the PK. This in turn implies that the pattern of expected option returns across strikes is a useful device to distinguish between the competing hypotheses of a monotonically decreasing and a U-shaped PK in the data.

2.2 Empirical analysis

2.2.1 Data

For the empirical analysis we use all available single stock options from OptionMetrics IvyDB US. The data contain information on the entire exchange-traded U.S. single stock options market and include the daily closing bid and ask quote as well as implied volatility, Greeks, trading volume, and open interest for the period from from January 1996 until December 2017. We obtain return data for the underlying stocks from CRSP, and accounting data from Compustat. We apply a series of standard data filters to the option data, which are described in detail in Appendix A. Most importantly, since we are interested in reliable returns of deep out-of-the-money options, we exclude options that have a best bid of zero or a mid-price below 1/8$. In addition, we exclude options which violate standard no-arbitrage bounds, have zero open interest[6] or with a bid-ask spread lower than the minimum tick size. Since single stock options are of the American type, we exclude options that pay a dividend during the holding period to mitigate effects from the early exercise premium.

On average, the final sample contains 1,566 unique underlying stocks per month, and in a typical month the aggregate market value of these stocks represents around 62% of the total CRSP market capitalization (which contains the market value of equity of many stocks that are not optionable).

Given the theoretical considerations outlined above, we limit our analysis to call options. In the context of our paper, it is interesting to note that for single stocks, in contrast to major stock market indices, calls are more actively traded than puts in terms of volume. They also feature a larger number of available traded strikes and higher open interest.

2.2.2 Benchmark analysis

In our benchmark analysis, we use non-overlapping (roughly) one-month option returns to maturity, computed from the Monday after the third Friday in the given month to the third Friday in the next month (which is the expiration date following the standard monthly expiration cycle).[7] We focus on these options, since they usually represent the most liquid contracts. The

[6] All return patterns documented in the following are also found when options with zero volume are excluded, but the findings are not always statistically significant, as the number of option contracts drops significantly. Nevertheless, we believe that positive open interest is a sufficient criterion, since the prices are always closing quotes, and a few trades during the course of the day need not to have any implications on the closing quote.

[7] Using hold-until-maturity returns mitigates the concerns regarding microstructure noise in daily option returns pointed out by Duarte et al. (2019).

options are bought on the first trading day each third Friday to avoid expiration day effects. We use returns until maturity, since the above theoretical results only apply to maturity returns, where the payoff is purely a function of the return of the underlying and state variables like volatility are irrelevant.

We sort call options into 11 moneyness bins spanning a range for $X \equiv K/S_0$ from 0.9 to 2, where the cutoff points for the bins are chosen such that the intervals are roughly equally populated.[8] In case there is more than one option for a given stock in a given moneyness bin in a given month, we only retain the option with the strike price closest to the midpoint of the bin.

We prefer this larger number of moneyness bins to provide a more detailed picture of the results. This comes at the potential cost of a loss of a statistical power, since with more bins each bin will necessarily contain fewer observations, but we provide results for alternative cuts of the data in the section on robustness checks.

To give an idea about the data quality behind our analyses, Table 1 reports averages of monthly trading volume, open interest and number of options for each moneyness bin. One might be concerned that out-of-the-money options are substantially less liquid and less actively traded than at-the-money contracts. This is not the case, however, as the table clearly shows. At least until Bin 8 there is a sizable number of call each month. Furthermore, out-of-the-money calls ($X > 1.05$) exhibit even slightly larger volume and open interest on average. This shows that out-of-the-money calls with higher strikes are definitely actively traded and sufficiently liquid assets. A usual concern about option prices are the relatively large bid-ask spreads. In a robustness analysis where we account for transaction costs by buying the options at the ask price we find even stronger results. This is due to the fact that relative bid-ask spreads tend to increase with moneyness, so that average returns are even lower for out-of-the-money options.

We calculate the average option return in a given moneyness bin over the full sample by

[8] The cutoff points are 0.9, 0.95, 0.975, 1, 1.025, 1.05, 1.1, 1.2, 1.3, 1.4, 1.6, and 2. We include deep-out-of-the-money calls in our analysis, as some papers in the literature (in particular those we refer to in Sections 4 and 5.1.2) use the full moneyness spectrum in their analysis. Our results, however, also hold up when we restrict the analysis to calls with a moneyness up to 1.3.

first calculating equally weighted returns for each month and then averaging over all months. This approach mimics the standard approach in the asset pricing literature on stock returns, and allows to interpret the results as returns to an implementable trading strategy. In addition, to test for monotonicity, we calculate the differences in returns between adjacent moneyness bins. Due to sometimes smaller sample sizes and the pronounced skewness of option returns, one may have concerns with respect to the use of standard statistical tests. That is why we base our inference on bootstraps using 100,000 i.i.d. draws for time series analyses and i.i.d. pairwise draws for differences.[9]

The limited number of available strike prices for many individual stocks poses a challenge for our approach. In an ideal setting, one would perform a monotonicity test for each stock separately, in the same way as Bakshi et al. (2010) do it for the S&P 500 index. However, only few stocks offer a sufficient number of actively traded strikes to perform that kind of analysis.

There are, of course, different approaches to this issue. Chaudhuri and Schroder (2015), e.g., sort options based on moneyness and then assign them to portfolios based on the *ordinal* moneyness value. This means, with at-the-money options in Bin 3 according to the convention used in Chaudhuri and Schroder (2015), the option with the next higher strike would be assigned to Bin 4 *irrespective of the numerical value of* K, then the next one to Bin 5, and so on. This approach has the advantage that, for a given stock, there will be options in consecutive bins by definition, so that one could basically perform monotonicity tests at the stock level. However, one can directly see that this will lead to options with potentially very different moneyness values being assigned to the same bin. As we will show below, this approach, in combination with the actual non-monotonicity in average call option returns as a function of moneyness, can in fact be problematic.

Our benchmark approach of sorting options into bins based on their actual (cardinal) moneyness values seems the most natural one to us. Of course, we then have to deal with the

[9] We choose i.i.d. draws, since the Ljung-Box test for up to 30 lags cannot reject the null of zero autocorrelation in the time series of option returns. We also calculated t-statistics based on a Newey and West (1987) adjustment and found that they are very similar.

problem that adjacent bins will not contain options on exactly the same stocks. However, as long as the availability of strike prices is not correlated with the expected option return, this will not be a major problem in the context of our analysis.[10]

Table 1 shows the results of our benchmark analysis, which fully support the hypothesis of a U-shaped PK. First, one can see from the numbers in Panel A that average call option returns increase in moneyness until X between 1.025 and 1.05, and then decrease again to eventually approach zero. As stated above, such a pattern is one of the key implications of a U-shaped PK. The inflection point is in a moneyness area where there are many traded options, with high turnover and open interest. The difference in returns is negative already between Bin 4 and 5, and statistically significantly so between Bins 5 and 6.

The second implication of a U-shaped PK relates to the behavior of expected returns on digital call options across the moneyness spectrum. Since digital options are not traded directly, we approximate a digital call option by a portfolio consisting of one call long with strike K_1 and one call short with the next available strike $K_2 > K_1$. We then take the average strike of the two calls to represent the strike of the digital call.

The results in Panel B of Table 1 show that, in line with the implications of a U-shaped PK, the average returns of the synthetic digital calls exhibit an inflection point in the moneyness dimension, and that this point is located further to the right on the moneyness axis than that for calls.

The third key implication of U-shaped pricing kernel is that the difference in expected returns between calls and digital calls is positive for low moneyness values and then changes sign as X increases. Panel C of Table 1 shows that this is exactly what we find in the data, since average call returns are greater than average digital call returns around the money, i.e., for $X \approx 1.05$, but lower for values of X greater than 1.05.

Since returns of call options written on different stocks are averaged in our analysis, the interpretation is that the PK is U-shaped as a function of the return of the "typical" stock.

[10] We provide extensive robustness checks with respect to our approach of assigning options to moneyness categories in Section 5.1 below.

There is, however, strong evidence in the data against the hypotheses that this result might be driven by certain special types of stocks. In Section 5 below we repeat the analysis for various subsamples with respect to stock characteristics such as stock volatility, stock liquidity, or stock size, and the pattern of average call option returns across the moneyness spectrum obtained in our benchmark analysis always prevails.[11] In sum, one can conclude that the finding of a U-shaped PK for the typical stock is very robust with respect to a large set of variations of the empirical approach.

3 U-shaped PK: Impact of Volatility on Expected Option Returns

3.1 Motivation

The existence of a U-shaped PK has fundamental implications for the cross-section of option returns, i.e., for the structure of expected option returns across underlying stocks which differ with respect to important characteristics. In the context of option pricing, the most important of these characteristics is arguably the return volatility of the underlying stock. In the following we therefore investigate the impact of idiosyncratic, systematic, and total volatility on expected option returns.

In a recent study, Hu and Jacobs (2020) provide empirical evidence for expected returns of at-the-money call options being monotonically decreasing in total (i.e., systematic plus idiosyncratic) volatility. They motivate their analysis by showing that in both the BSM model (Black and Scholes, 1973; Merton, 1973) and the Heston (1993) model, this pattern is implied theoretically. However, both Chaudhury (2017) and Aretz et al. (2018) point out that the Hu and Jacobs (2020) result crucially depends on the assumption that higher volatility is *not* com-

[11] As further robustness checks, we conduct double sorts on moneyness and other stock or option characteristics, investigate various subsamples with respect to time, and consider alternative return horizons and sorting approaches. The above findings turn out to be extremely robust with respect to these variations in our empirical approach. Details can be found in Section 5.4.

pensated by higher expected stock returns, which seems somewhat counterintuitive. In the basic versions of the two models analyzed in Hu and Jacobs (2020), there is no distinction between systematic and idiosyncratic volatility, and in addition, the model-implied PK is monotonically decreasing in the return of the underlying. These issues motivate us to separately study the impact of the different types of volatility on expected call returns, with a particular focus on the impact of a U-shaped PK.

3.2 Model

To guide our empirical analysis and to gain deeper insights regarding the link between the different types of volatility and expected call option returns, we build on the model proposed by Christoffersen et al. (2013) (CHJ, henceforth). We choose this model, since it can generate a U-shaped PK with reasonable parameter values. Next, we extend the model to accommodate idiosyncratic volatility as well as different levels of systematic volatility.

In the model, daily log returns are conditionally normally distributed, with a time-varying volatility represented via GARCH processes. The dynamics of the model are given as follows:

$$\ln\left(\frac{S_{t,i}}{S_{t-1,i}}\right) = r + \left(b_i\,\mu - \frac{b_i^2}{2}\right) h_{z,t} - \frac{c_i^2}{2} h_{y,t} + b_i\sqrt{h_{z,t}}z_t + c_i\sqrt{h_{y,t}}y_t, \tag{1}$$

$$h_{z,t} = \omega_z + \beta_z h_{z,t-1} + \alpha_z\left(z_{t-1} - \gamma_z\sqrt{h_{z,t-1}}\right)^2, \tag{2}$$

$$h_{y,t} = \omega_y + \beta_y h_{y,t-1} + \alpha_y\left(y_{t-1} - \gamma_y\sqrt{h_{y,t-1}}\right)^2, \tag{3}$$

$$y_t, z_t \sim N(0,1). \tag{4}$$

Here, $S_{t,i}$ is the price of stock i at time t and r is the daily continuously compounded interest rate. The shocks z_t and y_t are assumed to be independent, and they are the systematic (priced) and idiosyncratic (unpriced) return innovations, respectively. $h_{z,t}$ and $h_{y,t}$ represent the systematic and the idiosyncratic conditional variance, respectively, while b_i and c_i control the exposure of the stock to the different sources of risk. We model different exposures to systematic risk in a

"CAPM style", i.e., different stocks have different exposures b_i to the market factor $\sqrt{h_z}z$. The stock's return volatility has two components, a systematic one represented by the exposure to z, and an unpriced, idiosyncratic one given by the exposure to y. Furthermore, note that an increase to the market factor always implies an increase in the expected stock return.

We follow CHJ and employ an SDF with the following functional form:

$$M_t = M_0 \left(\frac{S_{t,I}}{S_{0,I}} \right)^{\phi} \exp \left[\delta t + \eta \sum_{s=1}^{t} h_{z,s} + \xi \left(h_{z,t+1} - h_{z,1} \right) \right], \tag{5}$$

where I denotes the index, i.e., an asset with $b_I = 1$ and $c_I = 0$ in (1). This SDF is monotonically decreasing in the return of the asset and monotonically increasing in variance for $\xi > 0$. However, the PK, i.e., the projection of this SDF on the asset's return, is U-shaped. The reason is that volatility is not only high for large negative returns, but also for large positive returns. For these large positive returns, the negative variance risk premium (represented by the free parameter $\xi > 0$, as shown in CHJ) dominates the positive equity premium, so that the projection is increasing again. For a more detailed description of the model, we refer to the original paper. The solution for the parameters ϕ, δ, and η are provided in Appendix C.1.

The original CHJ model for the stock market index is nested in the above specification with $b_I = 1$ and $c_I = 0$. Since we use the model for qualitative predictions only, we use the parameters provided by the CHJ estimation for the index. For simplicity, we use the index' GARCH parameters $\omega_z, \alpha_z, \beta_z, \gamma_z$ also for the idiosyncratic volatility GARCH parameters $\omega_y, \alpha_y, \beta_y, \gamma_y$. To ensure the robustness of our results, we perform an extensive simulation study and many different set of parameter value (see Appendix C.2).

To simulate a cross-section of expected option returns written on stocks with different levels of idiosyncratic volatility ($IVOL$) and systematic volatility ($SVOL$), we need to specify values for the coefficients b_i and c_i. For the individual stocks, we choose levels of $IVOL$ and $SVOL$ close to the 10th, 50th and 90th percentiles of the distribution of these quantities in our sample. The details on the parameterization and simulation are provided in Appendix C.2.

As a benchmark example for a monotonically decreasing PK we use the BSM model. This model has the advantages that all quantities of interest are available in closed form. We slightly modify the model in order to be able to study cross-sectional effects.[12] The stock price dynamics are as follows:

$$\frac{dS_{t,i}}{S_{t,i}} = \mu_i dt + \sigma_i dB_t, \tag{6}$$

$$\mu_i = r + b_i(\mu - r), \tag{7}$$

$$\sigma_i = \sqrt{b_i^2 \sigma_z^2 + c_i^2 \sigma_y^2}. \tag{8}$$

The volatility specification is qualitatively equivalent to the one shown above in (1).

3.3 Idiosyncratic volatility

First, to study the impact of a pure variation in idiosyncratic risk ($IVOL$), we keep systematic risk constant. We then simulate our model and study expected one-month hold-until-maturity option returns for different levels of moneyness and $IVOL$, by fixing $b_i = 1$ and varying c_i in $\{0, 1, 2, 4\}$.

The results are shown in Figure 2, with the solid line representing the expected return of the underlying stock. First observation is that with increasing $IVOL$, the expected call returns approach the expected stock return, as can be seen from the high $IVOL$ line approaching the solid black line. This observation is maybe not very surprising, as it is for example a well-known property of the BSM model. The second observation is more striking: since out-of-the money expected call returns are negative, the expected returns of these option approach the expected stock return from below, i.e., they are *increasing* in $IVOL$. Hence, in the presence

[12]Note that for $\xi = 0$ the GARCH based model also implies a monotonically decreasing PK. In fact, for $\xi = 0$ the model nests the BSM model in a one period setting. Therefore, when restricting the model to $\xi = 0$, we find that the implication of our model on expected option returns are quantitatively very similar to those of the BSM model. However, we prefer to use the latter model, as is allows for closed form solutions, especially of the sensitivity to parameter changes.

of a U-shaped PK, the relationship between expected call returns and $IVOL$ depends on the options' moneyness. In contrast, in models with a monotonically decreasing PK, such as the BSM model, the relationship between $IVOL$ and expected call returns is negative for all levels of moneyness (see Hu and Jacobs, 2020), as illustrated in the first plot in Figure 2. For the interested reader we provide a more detailed analysis of the relationship between expected call returns and $IVOL$ in both models in Appendix D.

These insights lead to our main hypothesis concerning the impact of $IVOL$ on expected call option returns. When the expected return of any call option approaches the expected return of the underlying for increasing $IVOL$, it must do so from above for calls with higher expected returns than the underlying, and from below for those calls which exhibit lower expected returns than the underlying.

In the context of a U-shaped PK, the first group of options will be those with low strike prices, while those with high strikes will be in the second group. Our hypothesis thus is that for calls with low moneyness, we should find expected returns to decrease with increasing idiosyncratic volatility, while we should obtain the opposite result for calls with high strikes.

Of course, different stocks will have different expected returns, so that expected call returns will not tend towards a uniform constant for all stocks. We do not consider this a major issue in the context of our analysis though, since option returns are typically orders of magnitudes larger than stock returns, so that cross-sectional differences in expected stock returns should be negligible.

Following Ang et al. (2006), we measure idiosyncratic volatility for stock i at day t as

$$IVOL_{i,t} = \sqrt{\sum_{j=0}^{251} e_{i,t-j}^2}, \tag{9}$$

where $e_{i,t}$ are the residuals from a regression of the daily returns on stock i on the three Fama-French factors MKT, HML, and SMB.[13] We use one year (252 trading days) of data here as well

[13]Using the latest five Fama-French factors gives very similar results.

as for all other variables as a compromise between having enough observations to estimate the variables reliably, and being flexible enough to pick up structural changes in the relationship between returns and the factors.[14]

We then perform independent double sorts for the options in our sample with respect to moneyness and the most recent underlying stock $IVOL$.[15] More precisely, at a given date we form quintiles with respect to underlying stocks $IVOL$ and within each quintile assign options to moneyness bins like in our benchmark analysis. Again we use at most one option per stock and per bin each month, namely the one that is closest to the midpoint of the respective bin. Finally, we calculate average call returns for all moneyness-$IVOL$ portfolios, in exactly the same way as in the benchmark approach. We therefore can interpret the results as returns to an implementable trading strategy.

The first graph in Figure 3 confirms our fundamental hypothesis that expected call returns should be decreasing in $IVOL$ for low to medium moneyness levels, but increasing in this variable for higher X. The dashed red line for the low $IVOL$ portfolio is above the dashed yellow line for moneyness up to 1.12, and, when crossing into negative territory, the ordering switches. In addition, we also plot average S&P 500 index call returns sorted by moneyness as a proxy for stocks without any $IVOL$. As predicted by our model, they have the highest returns for in-the-money and at-the-money calls, and the highest curvature, i.e., steepest decreasing line for out-of-the money calls. The second graph of Figure 3, that focuses on the return difference between the high and the low $IVOL$ portfolio in a given moneyness bin, confirms that the differences are also statistically significant. Furthermore, the average return spread decreases monotonically until the option return-moneyness inflection point ($X \approx 1.05$), where it is most negative, and subsequently monotonically increasing in X.

Of course, due to the pronounced volatility of option returns, we cannot expect these differences to be uniformly significant. Still, for X up to a value of around 1.05, we obtain

[14] Our results remain valid when we use alternative return windows of 21, 42, 63, 126, 504 and 1260 days.

[15] We exclude the deepest out-of-the-money bin $[1.6, 2]$, since it is too thinly populated for meaningful double sorts.

significance even at the 1% level, and the same is true for X greater than roughly 1.35. In sum, the results in the data are very much in line with the above considerations. A more detailed analysis is provided Section 3.5.

3.4 Systematic volatility

We next turn to the analysis of systematic volatility ($SVOL$). In order to get an idea of the relationship between expected call option returns and $SVOL$, we again perform a simulation analyses of our model, but now vary both c_i and b_i in (1) and (6). To study the impact of a pure variation in $SVOL$, we fix a level of idiosyncratic and vary only the systematic risk component. More specifically, we fix a c_i from the set $\{0, 1, 2, 4\}$, and then choose a b_i from $\{0.5, 1, 2\}$. Further details regarding the simulation study are provided in Appendix C.

Figure 4 shows the results. For very low strikes, average call returns must basically increase in $SVOL$, since in the limiting case of a zero strike, a call option simply corresponds to the underlying stock. This is exactly what we see in all plots of Figure 4: For all specifications, average call returns are increasing in $SVOL$ until a moneyness level of at least 0.9. However, for higher moneyness levels the relationship is not always monotonic, and in addition depends on the level of $IVOL$.

With respect to this, the figure shows an important difference between a model with a U-shaped PK relative to one where the PK is monotonically decreasing. When comparing the two models for a given level of $IVOL$, one can see that for the relevant moneyness range, i.e., for moneyness between roughly 0.9 and (depending on $IVOL$) 1.02 to 1.8, $SVOL$ has opposite effects on expected call returns for the two types of PKs. The reason for this is solely the negative variance-risk premium (see Appendix D.3 for more details). To the right of the moneyness levels where the curves intersect a second time monotonicity again reverses. However, in our model, the probability that the terminal stock return exceeds those levels of moneyness is very low, such that these scenarios are of only minor relevance. For example, in the case with medium $IVOL$, only (at most) 1% of the terminal returns are in that moneyness range. In the other cases, these numbers are comparable.[16]

For our empirical analysis, we define the systematic volatility of stock i at day t, $SVOL_{i,t}$, as

$$SVOL_{i,t} = \sqrt{TVOL_{i,t}^2 - IVOL_{i,t}^2},$$

where $TVOL$ denotes total volatility measured as

$$TVOL_{i,t} = \sqrt{\sum_{j=0}^{251} r_{i,t-j}^2 - \left(\sum_{j=0}^{251} r_{i,t-j}\right)^2},$$

and $IVOL$ is given by (9). A natural starting point for an empirical test of the effect of $SVOL$ on expected call returns would be an independent double support with respect to $IVOL$ and $SVOL$. However, disentangling the effects of these two variables is complicated by the relatively high correlation. The time-series average of the cross-sectional correlation between $IVOL$ and $SVOL$ is about 54%. The results for this double sort are shown in the online Appendix and qualitatively support our predictions, but are not statistically significant. We therefore turn to a regression analysis in the next section.

3.5 Fama-MacBeth regressions

We perform Fama-MacBeth (1973) regressions to disentangle the effects of $SVOL$ and $IVOL$, while controlling for other variables that are known to affect expected call returns. We follow Hu and Jacobs (2020) in their selection of risk factors and characteristics, but omit their options-based variables such as delta, vega and gamma, since they are all highly correlated with moneyness.

[16] All these patterns are observed consistently across a wide range of different parameter vectors, and it is also not simply a representation of the $IVOL$ effect discussed above. We provide further details regarding these rich patterns in Appendix D.

Every month, we run the following cross-sectional regression:

$$R_{t,i}^C = \gamma_{0,t} + \gamma_{1,t}\ IVOL_{t-1,i} + \gamma_{2,t}\ SVOL_{t-1,i} + \gamma_{3,t}' Z_{t-1,i} + \epsilon_{t,i},$$

where $R_{t,i}^C$ is the hold-until-maturity return of a call option on stock i from month $t-1$ to month t, and $Z_{t-1,i}$ is a vector of controls including momentum, reversal, the variance risk premium, and illiquidity. We do not separately include the CAPM-β, size and book-to-market as controls, as they are already in a certain sense included in $SVOL$. The exact definitions of these variables can be found in Appendix B.

To make our results comparable to the literature we choose a linear baseline specification. Our model, however, suggests a decreasing marginal relationship between $IVOL$ and expected call returns for all moneyness levels. We therefore also include a version where we use $\ln(IVOL)$. For consistency, we also do so for $SVOL$, although the model's implications from Figure 4 regarding the marginal relationships are ambiguous.

We run the analysis separately for three broad moneyness categories which subsume our finer moneyness grid used above. We focus our discussion on the results for out-of-the-money call returns with a moneyness X between 1.1 and 1.6.[17], as here the U-shaped PK is most important. Results for the other moneyness categories are presented in Appendix E. For each underlying stock, we keep the call with moneyness X closest to 1.25. Table 2 reports the time-series averages of the γ coefficient estimates, together with t-statistics based on Newey and West (1987) standard errors.

The results in Columns (2) - (9) show that $IVOL$ always has a positive and statistically significant effect on expected out-of-the-money call returns. The effect is even stronger for the log-transformed variable, which is consistent with our model. For $SVOL$, we also find a positive and statistically significant effect.[18] Here the log-transformation does not lead to a

[17] The results remain quantitatively unchanged when using 1.4 or 1.5 as upper moneyness limit.

[18] We also run the regression separately for the $IVOL$ portfolios $1-5$ from Section 3.3. We find that the coefficients for both $IVOL$ and $SVOL$, as well as $\ln(IVOL)$ and $\ln(SVOL)$, decrease from portfolio 1 to portfolio 5, which is in line with the model predictions. However, these decreases are not statistically significant.

stronger effect. The inclusion of $SVOL$ in the regression leads to a smaller effect of $IVOL$, due to the pronounced correlated of the two variables discussed above. Still, both $SVOL$ and $IVOL$ have separate effects on expected call returns, and these effects also remain significant when we include several well-known cross-sectional determinants of stock and option returns in the regression (see Columns (8) and (9)).

Comparing Columns (1) and (2) of Table 2 shows that is not enough to condition the analysis on a certain moneyness category, but moneyness X itself must also be included as a control. As soon as X is included as a regressor the coefficient for $IVOL$ becomes positive and highly statistically significant. The results in Column (3) show that the exact way how X is included is of minor relevance for the significance of $IVOL$. Controlling for moneyness thus identifies the autonomous effect of $IVOL$ much more clearly. The reason is that the further out-of-the money a call, the higher is the likelihood that it is written on a stock with high (idiosyncratic) volatility.

We conclude that the results in Table 2 are fully consistent with our model predictions from the previous two sections. Regarding $SVOL$, the results that out-of-the-moeny call returns are increasing in this variable are in line with the bottom left plot of Figure 4.

We perform the analogous Fama-MacBeth regressions for in-the-money ($0.8 \leq X < 0.95$) and at-the-money calls ($0.95 \leq X < 1.05$), where the mid-point is used as the reference moneyness level to select calls. The results are reported Tables F.1 and F.2 in the appendix. We find that in both cases $IVOL$ has a strongly negative and statistically significant effect on expected call returns, and this effect is also present after including control variables. The effect of $SVOL$ is positive for in-the-money calls, which is in line with the model predictions, but not statistically significant. For at-the-money calls, the effect of $SVOL$ is statistically insignificant and rather ambiguous, as it switches sign depending on the specification. These findings are also consistent with our model's prediction regarding turning points and non-monotonicities around $X = 1$.

3.6 Total volatility

To ensure robustness, we want to make sure that the effect of $IVOL$ on expected call returns is not driven by $TVOL$. First, when performing an independent double sort on moneyness and stocks' $TVOL$, we find almost the same pattern as for the analogous $IVOL$ analysis in Figure 3. However, when we perform a dependent double sorts on first $TVOL$ and then $IVOL$ we find no effect of $TVOL$ on average call returns.[19] When we switch the order and first sort on $IVOL$ and then on $TVOL$, we find a statistically significant effect of $IVOL$ on average call returns as in Section 3.3 above, at least for in-the-money calls. The results are displayed in Figures F.1-F.3 in the Appendix.

4 U-shaped Pricing Kernel: Option Return Anomalies

We now take a closer look on what our findings regarding the shape of the PK imply for option return anomalies documented in the literature. Our main result here is that the assumption of a monotonically decreasing PK, which implies expected call returns were increasing with moneyness, can obstruct a clear view of the phenomena actually present in the data and thus lead to questionable conclusions and interpretations.

We illustrate this point based on the two most common measures for lottery characteristics, namely the *ex ante skewness* measure suggested by Boyer and Vorkink (2014) and the MAX measure of Bali et al. (2011), used for options in Byun and Kim (2016). In both of these papers, option returns are used as a basis for the claim that investors' preferences for

[19] The major challenge in disentangling $TVOL$ and $IVOL$ empirically is the high correlation of the two variables in the cross-section of stocks (see, e.g., Ang et al. (2006)). Indeed, in our sample, the time series average of this cross-sectional correlation exceeds 98%. Hence, independent double sorts are unlikely to be a promising empirical strategy. Therefore, we follow the approach proposed by Bali et al. (2011) for such a situation and use dependent double sorts.

For this, we first sort all underlying stocks into quintiles based on $IVOL$. Then, within $IVOL$ quintile j, we sort stocks into quintiles based on $TVOL$. Next, for each moneyness bin, we compute the average call return across all $IVOL$ quintiles for a given $TVOL$ quintile. For example, for moneyness bin 1 and $TVOL$ quintile 1, we compute the average across $IVOL$ quintiles 1 to 5. The key quantity of interest for a given moneyness bin then is the difference of these averages between $TVOL$ quintiles 5 and 1, i.e., the average return on a high-minus-low portfolio. When we sort on $IVOL$ and control for $TVOL$, we switch the role of the two variables.

lottery-like assets do affect asset prices. The basic argument is this type of investors is willing to accept on average large losses on investments in (overpriced) call options, when they are compensated with large positive return skewness.

Options are indeed suitable instruments to study such preferences, since they offer highly skewed return distributions. As we will show in our analysis below, however, the patterns in the data which the authors of the two studies cited above take as evidence supporting certain behavioral biases, are perfectly in line with a U-shaped PK and thus with a risk-based explanation. Given that there are hardly any short-selling restrictions in the options market, professional investors could arbitrage away the potential mispricing rather easily, which makes this risk-based interpretation of the data appear more appealing than a behavioral story. Furthermore, if limits to arbitrage were to be the driver of these results, the effects should mostly come from stocks with pronounced trading frictions such as small or illiquid stocks. However, in our robustness analysis below, we find no such patterns.

Boyer and Vorkink (2014) construct a measure for the ex ante skewness of option returns and use this measure to argue that there is a negative relation between the lottery-like characteristics of options (measured by this ex ante skewness) and subsequent option returns. The measure they propose relies on the assumption of lognormal stock returns, and it uses as inputs the mean and the variance of stock returns based on daily data over the preceding six months. The ex ante skewness of the option return is then given in closed form as a function of only the expected stock return, stock return volatility, the time to maturity of the option, and its moneyness.

We start with a replication of the main result in Boyer and Vorkink (2014) and, like them, we now also consider all call options without any restrictions on moneyness.

Table 3 contains the average call returns of five portfolios sorted on the ex ante skewness measure proposed by the authors ($BVSkew$ henceforth). The results exactly confirm the pattern documented in the original paper, namely that returns exhibit a hump-shaped pattern across the portfolios and are highly negative (although not significantly so) for the portfolio with high

$BVSkew$.

There is, however, another very important point to be taken from Table 3. As one can see, $BVSkew$ seems to be very highly positively correlated with moneyness. In fact, the measure is a highly convex function of moneyness and more sensitive to this than to any other of the variables, as also indicated by Figure 1 in Boyer and Vorkink (2014). In the context of a U-shaped PK, this is a crucial observation. Given such a PK, the expected returns on calls with high moneyness will be negative only due to the specific properties of the PK, irrespective of any behavioral biases on the part of investors.

To test if it is indeed just moneyness which is responsible for the observed return pattern, we perform a dependent double sort on moneyness and $BVSkew$ by first sorting on moneyness, and then within each moneyness bin, generating skewness portfolios of equal size. We prefer dependent double sorts here, since the strong correlation between $BVSkew$ and moneyness would produce many poorly populated portfolios for an independent double sort.

Figure 5 shows that when controlling for moneyness via the double sort, average returns are not decreasing in $BVSkew$ anymore. As one can see, the return differential between the high and the low $BVSkew$ portfolio is not significantly different from zero for any level of moneyness. This is a first strong indication that $BVSkew$ does not offer explanatory power for average call returns beyond moneyness.

Boyer and Vorkink (2014) also perform Fama-MacBeth regressions. When we perform the analogous analyses, the results again show in a very strong fashion that, once moneyness is properly controlled for, the effect of the $BVSkew$ disappears. We use the same set of controls as the original paper, and in addition add illiquidity and the variance risk premium, as they are known to be strong predictors of option returns.

Table 4 presents the results. We perform the regression for subsamples with in-the-money and at-the-money options in one group ($X \leq 1.02$) and out-of-the-money options in the other ($X > 1.02$). The choice of the cutoff point between the two groups is motivated by the results shown in Figure 5.

First, the simple regression for the full sample indeed confirms a significantly negative link between $BVSkew$ and call returns. The sample split, however, casts serious doubts on whether $BVSkew$ is the true return driver. For the subsample with $X \leq 1.02$, the associated coefficient is insignificant in the specification without controls, but even significantly *positive* in the regression with stock and option controls. Moneyness on the other hand has a positive impact for this subsample, which is again in line with the notion of a U-shaped PK. For the subsample with out-of-the-money calls, we find $BVSkew$ to be insignificant as soon as we control at least for moneyness, and the picture does not change when we consider the larger set of control variables. Note that X here has a significantly negative coefficient, which again supports our fundamental hypothesis of a U-shaped PK.

The second option return anomaly in the context of lottery-like assets is documented in Byun and Kim (2016). They find evidence for average call returns to be lower when the underlying stock exhibits more pronounced lottery characteristics. The proxy for lottery characteristics used by the authors is the MAX measure introduced by Bali et al. (2011), i.e., the average of the ten highest daily stock returns over the past quarter. The authors conclude from their analyses (based on at-the-money and slightly out-of-the-money options) that there is a strongly negative relationship between a stock's lottery characteristics and average call option returns.

In the context of U-shaped PK we can shed new light on this link by considering also options with higher moneyness than those used in the original paper. Figure 6 graphically presents the return difference between the high and the low MAX portfolio across all our moneyness bins together with the bootstrapped 99% confidence intervals.

Indeed, for low values for X the spread in average call returns between the high and the low MAX portfolio is negative, and significantly so. For X around 1.1, however, the effect starts to go the other way, and for deep-out-of-the-money calls, we finally find the spread to have the opposite sign relative to the predictions in the Byun and Kim (2016) paper.

Overall, one can see from our analyses of option return anomalies related to $BVSkew$ and MAX that in the light of a U-shaped PK, the observed patters for average call returns are either not anomalous anymore, as in the case of $BVSkew$, or they even even change direction, as in the case of MAX.[20]

5 Robustness

5.1 Discussion of alternative moneyness sorts

5.1.1 Pre-defined moneyness bins (cardinal sort)

In our benchmark approach we use all call options satisfying our filtering criteria and sort them into bins according to their numerical, i.e., cardinal, moneyness value. As indicated above in Section 2.2.2, the advantage of this approach is that it relies on a large amount of data, but on the other hand it is very likely that the sets of stocks in two adjacent bins are not the same. In this section we provide a number of robustness checks with respect to the procedure applied in our main analysis.

A stricter approach would be to compute the return differences between options on the same underlying in adjacent bins, since the pattern of these (average) differences across moneyness bins would represent a potentially more direct test of the hypothesis regarding the monotonicity the PK. In the extreme case, one would actually require literally the same stocks to be in all bins. The key problem here is that for many individual stocks, in contrast to the index, only few strikes are available. Table 5 illustrates this problem. It shows that, on average in our final sample, there is only one call option for 39% of the stocks, and only two for another 30%.[21]

[20] Byun and Kim (2016) also use Fama-MacBeth regressions to support their finding. We replicate their approach, and confirm their results for at-the-money call returns. However, for out-of-the-money calls we again find that the relation between MAX and call returns switches sign. The results are shown in the Tables F.3 and F.4 in the appendix.

[21] Since we allow only one call in any given bin per underlying at each formation date, these numbers potentially even overstate the number of actually available strikes, due to the fact that more than one strike of a given stock may fall in a given moneyness bin. Furthermore, we only consider standard calls in the following robustness analyses, since due to the approximation of digital calls via a portfolio of two standard calls, the data requirements in terms of strike availability would even be tougher for digital calls.

Instead of computing average call returns in each moneyness bin and then testing for the significance of the difference between adjacent bins as in our benchmark analysis, we now directly focus on the differences in returns between two call options on the same underlying in adjacent bins. This means that the average return difference between the two bins is computed on the basis of calls on only those underlying stocks which are contained in both bins. Nevertheless, the set of such stocks may vary depending on which bins are considered. Due to this stricter data requirement, we have to combine two adjacent bins from the original set into one, with the special case of an at-the-money bin containing the strike prices previously assigned to Bins 3, 4 and 5.

Table 6 presents the results, which strongly confirm the findings from the benchmark analysis. For lower values of X, average call returns are increasing in moneyness, as indicated by a significantly positive average return difference of 4.4% per month between the bin with $X \in [0.975, 1.050]$ and that with $X \in [0.900, 0.975]$. For higher moneyness levels, the differences become significantly negative, which shows that for out-of-the-money calls, average returns are decreasing in moneyness. The two results together provide strong support for the hypothesis of a U-shaped PK for the typical stock.

In the extreme case, one might require that, at a given date, the differences across bins are computed based on calls on only those underlying stocks, which are represented in *all* of the bins. This even stricter data requirement only allows for two meaningfully populated bins in the end, one for at-the-money calls with $X \in [0.9, 1.1]$ and one for out-of-the-money options $X \in [1.1, 1.5]$. As before, we only include the call closest to the midpoint of the respective bin. This leaves us with on average 511 observations per month, which still represents more than 50% of the on average 950 stocks for which one could at least potentially compute a difference in option returns between moneyness bins (see Table 5).

The results in Table 7 again clearly support the notion of a U-shaped PK with a strongly significant average return difference between at-the-money and out-of-the-money calls of -17% per month. Furthermore, these results lends strong credibility to our benchmark approach, which features a finer moneyness grid.

5.1.2 Ordinal moneyness sort

As pointed out above, Chaudhuri and Schroder (2015) propose a different procedure for assigning options to moneyness bins. In their portfolio formation, they treat moneyness as an ordinal variable. In detail, they first identify the option with a strike closest to the current stock price and assign it to a reference bin (Bin 3 in their case).[22] The next higher strikes for each stock are assigned to Bins 4 and 5, the next lower ones to Bins 2 and 1. It is important to note that this assignment takes place *irrespective of the numerical value of moneyness*. This approach has the advantage that even with the requirement that return differences are computed from calls on only those underlyings which are present in two adjacent bins, the analysis can be based on a fairly large sample, since for many stocks at least a few strikes are available.

This approach offers a number of degrees of freedom, e.g., the choice of the reference point for the start of the sort. Chaudhuri and Schroder (2015) start their procedure with the at-the-money bin, but one might as well pick any other reference point on the moneyness axis.

We no investigate the impact of this aspect of the methodology and present the first results for such an "ordinal sort" in Table 8. Here, we use the full sample of all calls satisfying the basic data filters described in Section 2. The reference point for the start of the assignment of options to bin is at-the-money, and we consider only bins with increasing strikes, i.e., we move from at-the-money to out-of-the-money calls.

Again, the results are strongly in favor of a U-shaped PK. The average return difference between any pair of adjacent bins is negative, and between Bins 2 and 1, and Bins 3 and 2, this difference is also strongly statistically significant.

Table 8 also highlights a potential problem with the ordinal sorting approach. The reported moneyness statistics clearly show that this procedure of assigning options to bins can have rather undesirable side effects, since the range of X in the different bins can be huge. Already in Bin 2, there can be options with values of X which would be considered extremely far

[22] If there is no option for a given stock in this pre-defined at-the-money range, that stock is discarded on that formation date.

out-of-the-money, while the minimum moneyness would usually be considered almost exactly at-the-money.

We next perform the ordinal sort and require options in all bins to have exactly the same underlying stocks. Tables 9 and 10 present the result when we use four and three bins, respectively. In both cases the reference point is at-the-money. Both variations of the approach confirm our main finding that out-of-the-money, average call returns decrease in moneyness, and this decrease is also often highly statistically significant.

There is no theoretical reason to use an at-the-money reference point to start the sort. In the following, we use the deepest out-of-the-money call for each stock as this reference point, i.e., for each stock we identify the call with the highest strike and put it into Bin 1. The call option with the next lower strike is then assigned to Bin 2, and so on. Since we are mainly interested in the properties of out-of-the-money calls, we require that all options have a moneyness greater than one. Finally, we impose the requirement that the same stock must be in all bins.

Table 11 presents the results for two and three out-of-the-money bins. The numbers again confirm our main finding. The average returns on out-of-money calls significantly decrease with moneyness, which is one of the main implications of a U-shaped pricing kernel.

The choice of an alternative reference point for the assignment of options to bins is also important with respect other explanations for negative expected call returns suggested in the literature. Chan and Ni (2018) argue that negative returns for out-of-the-money calls could also be due to measurement error. Some option prices may be too small to be quoted, and hence they are not included in the data.

However, the “reverse” sorting with an out-of-the-money reference point shows that this explanation is not very likely. Even relatively far out-of-the-money, call returns decrease on average in moneyness (and do so in a highly significant fashion), and these are, of course, the cheapest options available. In addition, the average number of options in the sample each month is fairly large.

5.1.3 Alternative definition of moneyness

The same monthly stock return has a different probability to occur for stocks with different volatilities. To account for this, we perform a robustness analysis where we replace our original definition $X = K/S_0$ by $\widehat{X} = (K/S_0 - 1)/\sigma$, i.e., we standardize the deviation from $X = 1$ by volatility. For example, a $\widehat{X} = 3$ mean that the strike is three return standard deviations away from the current stock price. We use three alternative measures for σ, namely historical return volatility measured from daily returns over the previous month and over the previous year, and the implied volatility of the respective option. We scale each volatility measure with time to maturity.

We then continue as in our benchmark approach, just with standardized moneyness. For brevity, we again focus on return differences, and present the results in Table 12. The numbers again confirm the hump-shaped return pattern implied by the U-shaped PK.

5.2 Alternative return horizons

As an alternative to our monthly returns, we now consider two week and six week maturity returns. When we stick to our portfolio formation on Mondays as in the benchmark analysis, two and six weeks effectively correspond to 18 (17 if Monday is not available) and 46 (45) days, respectively. Furthermore, given the above considerations, we present the results for an analysis analogous to Table 1 (using all available options) and the version where we require the same stocks top be in all bins corresponding to Table 7.

Table 13 presents the results for these alternative return horizons when we we use all available options. The results represent very strong support for the hypothesis of a U-shaped PK. As in the benchmark case, average call returns are significantly increasing in moneyness up to a level of around 1.05, and then start to decrease just as significantly for X beyond that point. In total, we observe the same hump shape is in the benchmark case.[23]

[23] In the case of six week returns, the returns are overlapping. We consider this in the statistics and use block bootstrap with a size of five in this case.

In Table 14 we show the results for the case when all bins (in this case again only two) contain exactly the same stocks at a given date. The approach is the same as described above for the standard monthly returns, and the structure of the results is analogous to that in Table 7. The results are again strongly statistically significant. Average out-of-the-money call returns are highly negative, and the difference between the at-the-money and the out-of-the-money returns is negative, large in absolute terms, and strongly significant.

5.3 Weekly intermediate option returns

5.3.1 Basic considerations

Up to now, our analysis has exclusively focused on maturity option returns, i.e., options were bought and then held until maturity. In contrast to this, Chaudhuri and Schroder (2015) consider weekly intermediate returns of options with three to six weeks maturity, calculated from Tuesday to Tuesday. When we investigate their approach in detail, we find that their results are not really robust to variations of the empirical setup, and that overall, also weekly returns are consistent with a U-shaped PK.[24]

The use of weekly intermediate option returns comes with several data-related as well as theoretical problems. First, for most of the sample there is just one option expiry day per month, so there are several weekly returns per month. Due to their special way of computing weekly returns, Chaudhuri and Schroder (2015) actually have options with maturities of 3.5, 4.5 and 5.5 weeks in their sample, over the returns of which they compute averages. The characteristics of intermediate weekly returns, however, vary with maturity, so that potentially very different types of returns are averaged. Since the maturity choice by Chaudhuri and Schroder (2015) seems somewhat ad hoc, we also add weekly intermediate returns on options with maturities of 2.5 and 6.5 week. On the one hand, these additional maturities correspond to those for which

[24] Chaudhuri and Schroder (2015) use fewer option price filters than is commonly done in the literature. In particular, they do not require a minimum option price, positive open interest or volume, and do not impose any bounds the options' moneyness. We nevertheless follow their approach to allow for the best possible comparability between their findings and ours.

maturity returns are discussed in Section 5.2 above, and on the other hand, they will help to show that intermediate weekly returns for options with different maturities actually exhibit very different properties.[25]

Second, to compute a return, not only the initial, but also the terminal option price is required, while for maturity returns it is enough to have the stock price at the maturity date. We find that in more than 15% of the cases, the option price in the following week is missing. More importantly, we find that option prices are not missing randomly, but that the fact that a price is not there is correlated with returns, which obviously biases the results. A deeper analysis shows that the problem is most severe for out-the-money calls, where prices are frequently missing when the stock price has moved so much that the option is now very deep out-of-the-money, and hence it is not quoted or traded anymore. The opposite is found for in-the-money calls. This should bias the observed out-of-the-money call returns upwards, and the observed in-the-money call returns downwards.

Finally, there is also a substantial conceptual problem with the approach of using intermediate instead of maturity returns. The theoretical relationships between call returns and the shape of the PK discussed above in Section 2.1 only apply to maturity returns, since these returns are just a function of the underlying price at maturity, exactly like the PK, whereas an option price before maturity is a function of the underlying price and a (potentially large) vector of state variables. Hence, a test for the shape of the PK based on intermediate returns is very likely to be inconclusive.

In a strict sense, the link between intermediate option returns and the shape of the PK can only be established in a fully parameterized model. Chaudhuri and Schroder (2015) choose the Merton jump-diffusion model, but this is a problematic choice, since the model features a monotonically decreasing PK, which clearly favors this hypothesis over the alternative of a U-shaped PK.

[25] This has also been pointed out by Broadie et al. (2009).

5.3.2 Cardinal sort

We first present all weekly returns sorted by moneyness in Figure 7. Other than working with a different return concept we completely follow our benchmark approach. As indicated above, Chaudhuri and Schroder (2015) average over weekly call returns with 3.5, 4.5 and 5.5 weeks to maturity, which is represented in the graph by the line labeled "3w-6w".

One can see from the lines in the graph that the return patterns for different types of average intermediate returns are indeed very different. Nevertheless, average returns for all maturities exhibit a hump shaped pattern, albeit with different locations, and the peak of average returns is attenuated with increasing maturity, and the moneyness level associated with the peak shifts to the right.

There are also several striking structural differences between the intermediate returns presented here and the maturity returns discussed above. Most importantly, now the average returns of out-of-the-money calls are always positive. Furthermore, average in-the-money call returns are negative, which is puzzling. The average weekly stock return in our sample is positive, so that this is likely to be a consequence of the data problem described above, which biases observed returns on in-the-money calls downwards.

Table 15 presents the average call return differences between adjacent moneyness bins for maturities from 2.5 to 6.5 weeks. The general tendency is that average returns on out-of-the-money calls are decreasing in moneyness, with very strong overall significance for the two short maturities, and a number of significantly negative differences between Bins k and $k+1$ also for longer maturities.

To complement the analysis, Table 16 shows the analogous results, but with the additional requirement that return differences are only computed from options on those underlying which are represented in both of two adjacent bins. The tendency of the results using all available options is confirmed here, albeit with in part less pronounced statistical significance.

They also confirm the general finding, but the decrease in returns is not statistically significant anymore for the 4.5w and 5.5w maturities. In sum, one can conclude that average intermediate weekly option returns are clearly hump-shaped for options with shorter maturities. For longer maturities, we observe decreasing average returns to a lesser degree and, if so, mostly for deep-out-of-the-money calls. We would like to emphasize again here that especially these latter findings are to be interpreted cautiously due to data issues discussed above.

5.3.3 Ordinal sort

We also replicate the ordinal sorting approach used in Chaudhuri and Schroder (2015) and find results similar to those for the cardinal sort. We present the results for the cases of three, four, five, and six bins in in the Online Appendix.

For short maturities, the average intermediate call returns exhibit a hump-shaped pattern across moneyness. For longer maturities, the fact that the approach combines structurally very different options in the same bin, as discussed above in Section 5.1.2, makes the results hard to interpret.

5.4 Further double sorts

To check whether the hump shaped relationship between moneyness and option returns is driven by any obvious subsample, we perform independent double sorts on various stock characteristics and moneyness. The approach is analogous to the analysis of the impact of $IVOL$ in Section 3.3 above. We do this for the stock's beta, book-to-market ratio, size, the Amihud (2002) illiquidity measure, turnover, momentum, short term reversal and the difference between historical and at-the-money implied volatility (variance risk premium) studied in Goyal and Saretto (2009). In sum, we do not find any double sort where the hump-shaped relationship between moneyness and expected call returns cannot be observed for any subsample. The results are included in the online Appendix.

5.5 Subsamples in the time dimension

The theoretically implied return patters (as e.g. in Figure 1) are conditional statements. However, our non-parametric approach using *average* option returns only allows to study an *unconditional* relationship. To get closer to a conditional analysis, we repeat our main analysis using meaningful subsamples.

5.5.1 High and low volatility periods

A typical finding in the asset pricing literature is that patterns differ between good and bad times, between times of distress and normal times, between recessions and expansions.

We divide our base sample into two subsamples by conditioning on the level of the volatility index at the portfolio formation date. When the VIX at formation date is below its sample median of 18, we assign that month to the low volatility regimes, and vice versa. Figure 8 shows that there is no structural difference between the two volatility regimes when it comes to average call returns as a function of moneyness. The two lines for average call returns move largely in parallel across the moneyness axis, and the 95% confidence bands around the sample averages overlap substantially.

5.5.2 Calendar-time subsamples

Over the years, the option market underwent some changes, as for example a growing number of underlying stocks, and increased automation. We therefore split our sample in four parts, all of equal duration of 5.5 years. The numbers in Table 17 again confirm our main hypothesis of a U-shaped PK.

5.6 S&P 500 constituents

Finally, we repeat our main analysis using only those underlying stocks, that where constituents of the S&P 500 index at portfolio formation date. The S&P 500 contains the largest US stocks and accounts on average for about 80% of the total CRSP stock market value. The results in the last panel of Table 17 again corroborate our main findings.

6 Conclusion

We provide strong empirical evidence in favor of a U-shaped pricing kernel for the typical stock. Average call returns are a hump-shaped function of moneyness, with the maximum being in the region of 2.5% to 5.0% out-of-the money and average returns on out-of-the-money options being negative. These findings are robust to a wide set of variations of the empirical methodology. Among other things, we carefully investigate the impact of the actual procedure of assigning options to moneyness bins and of computing option returns.

The insight that the PK for the typical stock is U-shaped has important implications for the increasingly popular analysis of the cross-section of option returns, i.e., of option returns across different underlying stocks. For example, we show that higher idiosyncratic volatility will tend to decrease expected returns on in-the-money calls, while the opposite is likely to happen for out-of-the-money calls. This means that the impact of idiosyncratic volatility on expected call returns does not only go in one direction, but it rather depends on the option's moneyness. Furthermore, our new insights concerning the shape of the PK also provide a starting point for risk-based instead of behavioral explanations for seemingly anomalous patterns in expected option returns.

Table 1: Returns of Call Options and Digital Call Options Across Moneyness Bins

	1	2	3	4	5	6	7	8	9	10	11
Panel A: Calls											
$X \in$	(0.9, 0.95]	(0.95, 0.975]	(0.975, 1]	(1, 1.025]	(1.025, 1.05]	(1.05, 1.1]	(1.1, 1.2]	(1.2, 1.3]	(1.3, 1.4]	(1.4, 1.6]	(1.6, 2]
Avg. X	0.926	0.962	0.988	1.013	1.037	1.073	1.140	1.240	1.342	1.473	1.734
Return (%)	5.7**	7.4***	9.7***	11.9***	11.5***	7.9*	-1.7	-21.2***	-29.7***	-46.6***	-71.6***
(t-stat)	(2.22)	(2.47)	(2.84)	(3.02)	(2.58)	(1.63)	(-0.35)	(-4.03)	(-3.52)	(-5.58)	(-12.71)
99% conf. int.	[-0.2, 11.8]	[0.5, 14.4]	[1.8, 17.6]	[2.8, 21.2]	[1.2, 22.0]	[-3.1, 19.5]	[-13.0, 10.0]	[-32.9, -8.5]	[-47.4, -8.1]	[-63.9, -25.3]	[-83.7, -56.8]
Avg. return difference to previous bin (%)		1.7***	2.3***	2.2***	-0.5	-3.6***	-9.6***	-19.5***	-8.6	-16.8*	-25.0***
(t-stat)		(2.72)	(2.92)	(2.31)	(-0.41)	(-2.85)	(-4.83)	(-6.54)	(-1.21)	(-1.73)	(-2.98)
99% conf. int.		[0.3, 3.1]	[0.5, 4.1]	[0.0, 4.6]	[-3.0, 2.1]	[-6.5, -0.6]	[-14.2, -5.0]	[-26.3, -12.5]	[-23.5, 9.3]	[-40.2, 5.3]	[-48.7, -6.7]
Avg. no. of calls	615	367	384	370	370	557	502	185	77	50	21
Option price	3.95	2.92	2.14	1.60	1.16	0.78	0.47	0.35	0.31	0.27	0.25
Volume	915	1218	1444	1663	1654	1471	1266	1273	1364	1585	2077
Open interest	71	138	231	346	328	262	185	158	174	174	259
Panel B: Digital Calls											
Avg. X	0.926	0.963	0.988	1.013	1.037	1.073	1.139	1.240	1.340	1.471	1.731
Return (%)	3.2**	4.3**	5.5***	7.8***	9.3***	12.8***	18.5***	11.0*	0.8	-29.4***	-52.2***
(t-stat)	(2.24)	(2.34)	(2.52)	(2.97)	(2.93)	(3.28)	(3.69)	(1.65)	(0.09)	(-3.04)	(-4.81)
99% conf. int.	[-0.2, 6.3]	[0.0, 8.4]	[0.4, 10.5]	[1.7, 13.8]	[1.9, 16.7]	[3.9, 22.0]	[7.1, 30.3]	[-3.9, 27.0]	[-18.3, 22.2]	[-50.2, -5.1]	[-74.4, -19.9]
Avg. return difference to previous bin (%)		1.1**	1.2**	2.3***	1.5*	3.5***	5.7***	-7.5	-10.2	-30.2***	-22.8
(t-stat)		(2.01)	(2.22)	(3.22)	(1.76)	(2.97)	(2.54)	(-1.58)	(-0.95)	(-2.90)	(-1.71)
99% conf. int.		[-0.2, 2.4]	[0.0, 2.6]	[0.6, 3.9]	[-0.5, 3.5]	[0.8, 6.3]	[0.8, 11.2]	[-17.9, 4.1]	[-27.5, 12.2]	[-55.7, -5.9]	[-56.7, 13.7]
Avg. no. of digitals	473	305	321	302	270	354	276	101	45	29	14
Panel C: digitals vs. calls											
$R^C - R^D$	2.6*	3.2**	4.2***	4.2***	2.2	-4.9***	-20.2***	-32.2***	-33.7***	-16.4*	-18.5*
(t-stat)	(1.75)	(2.11)	(2.55)	(2.29)	(1.15)	(-2.56)	(-9.45)	(-6.20)	(-3.78)	(-1.85)	(-1.67)
99% conf. int.	[-0.7, 6.1]	[-0.3, 6.7]	[0.5, 8.1]	[0.1, 8.5]	[-2.1, 6.7]	[-9.2, -0.3]	[-25.2, -15.2]	[-44.8, -20.7]	[-55.3, -13.2]	[-37.1, 4.3]	[-51.2, 4.9]

The table reports statistics of one month hold-until-maturity call option returns. All option positions are bought on the first trading day after the third Friday of each month and hold until maturity. For each moneyness bin and date, only the option closest to the midpoint of the moneyness interval is considered for each underlying stock. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 independent draws (pairwise draws for differences). The t-stats are additionally reported in parenthesis. Moneyness X is defined as K/S_0. Digital calls are approximated as a bull spread, and the midpoint between the two call strikes is used to calculate the moneyness level of the digital. The sample is from Jan 1996 to Dec 2017.

Table 2: Fama-MacBeth Regressions of Out-Of-The-Money Call Returns with $IVOL$ and $SVOL$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$IVOL$	0.018	0.289***	0.287***			0.202**		0.227**	
	(0.21)	(3.33)	(3.27)			(2.21)		(2.23)	
$\ln(IVOL)$				0.220***			0.178***		0.200***
				(3.86)			(3.06)		(3.31)
$SVOL$					0.725**	0.606**		0.495*	
					(2.50)	(2.01)		(1.72)	
$\ln(SVOL)$							0.054*		0.035
							(1.79)		(1.22)
X		−2.250***	2.347	0.635	0.382	0.751	−0.385	0.620	−0.664
		(−8.39)	(0.42)	(0.11)	(0.07)	(0.14)	(−0.07)	(0.11)	(−0.11)
X^2			−1.928	−1.289	−1.073	−1.314	−0.893	−1.307	−0.834
			(−0.85)	(−0.55)	(−0.49)	(−0.61)	(−0.40)	(−0.56)	(−0.34)
MOM								0.036	0.040
								(0.74)	(0.83)
REV								−0.134	−0.133
								(−1.20)	(−1.18)
VRP								−0.016	−0.018
								(−0.26)	(−0.30)
$ILLIQ$								−19.850	−26.213*
								(−1.29)	(−1.66)
R^2(%)	0.4	0.8	1.0	1.1	1.2	1.5	1.7	2.6	2.8

This table shows the results for Fama-MacBeth regressions of one month hold-until-maturity call returns on stock and option characteristics plus additional controls. For each month, we estimate the cross-sectional regression

$$R^C_{t,i} = \gamma_{0,t} + \gamma_{1,t}\ IVOL_{t-1,i} + \gamma_{2,t}\ SVOL_{t-1,i} + \gamma'_{3,t} Z_{t-1,i} + \epsilon_{t,i},$$

where $R^C_{t,i}$ is the return for call option i at time t, and Z is a vector of controls. The definitions for the different variables are provided in Appendix B. For all out-of-the-money calls with a moneyness $X = K/S_0$ between 1.1 and 1.6, for each underlying stock, we keep only the call which has an X closest to 1.25. The table reports the time-series averages of the γ-coefficients and the associated Newey and West (1987) adjusted t-statistics (with five lags) in parentheses. The last row gives the averages of the adjusted-R^2 from the regressions. ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively. The total number of observations is $151,713$ in all regression specifications.

Table 3: Call Returns for Portfolios Sorted on $BVSkew$

PF	Avg. return (%)	99% conf. int.	avg. $BVSkew$	avg. X
L	3.87^{***}	[0.3, 7.4]	1.29	0.74
2	5.91^{***}	[0.0, 11.8]	1.68	0.88
3	7.03^{**}	[-0.9, 15.3]	2.15	0.96
4	5.70	[-4.2, 16.3]	2.95	1.05
H	-6.69	[-17.5, 4.6]	16.60	1.18
L-H	10.56^{***}	[1.3, 19.2]		

This table presents a replication of the results in Table IV of Boyer and Vorkink (2014). It shows average one month hold-until-maturity call returns for portfolios sorted on $BVSkew$. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 independent draws (pairwise draws for differences). In addition, the table reports the average moneyness X and average $BVSkew$ of the portfolios.

Table 4: Fama-MacBeth Regressions with $BVSkew$

	All			$X \leq 1.02$			$X > 1.02$		
$BVSkew$	-0.007^{*}	-0.006^{**}	-0.002	0.015	-0.022	0.044^{***}	-0.008^{***}	0.003	0.004
	(-1.72)	(-2.04)	(-0.66)	(0.85)	(-1.28)	(2.62)	(-3.20)	(1.03)	(0.81)
X		-0.109	-0.307^{***}		0.268^{***}	-0.120		-1.363^{***}	-1.187^{***}
		(-1.31)	(-3.22)		(4.61)	(-1.40)		(-8.29)	(-6.38)
$IVOL$			-0.380^{***}			-0.260^{***}			-0.299^{**}
			(-6.24)			(-4.71)			(-2.57)
$\ln(B/M)$			0.009			0.006			0.017
			(1.03)			(1.07)			(1.06)
MOM			-0.003			0.019			0.004
			(-0.14)			(1.41)			(0.09)
REV			-0.172^{***}			-0.062^{**}			-0.261^{**}
			(-2.97)			(-2.40)			(-2.53)
$SIZE$			-0.017^{**}			-0.006			-0.018
			(-2.25)			(-1.49)			(-1.04)
CAPM-β			0.007			-0.006			0.020
			(0.30)			(-0.39)			(0.54)
$ILLIQ$			-0.771			-1.887			-6.714
			(-0.14)			(-0.54)			(-0.51)
VRP			0.254^{***}			0.101^{***}			0.362^{***}
			(6.96)			(4.87)			(4.23)
$R^2(\%)$	0.4	1.9	4.5	2.2	3.5	8.2	0.2	0.8	3.3

This table is analogous to Table VIII in Boyer and Vorkink (2014) and shows the results for Fama-MacBeth regressions of one month hold-until-maturity call returns on stock and option characteristics plus additional controls. For each month, we estimate the cross-sectional regression

$$R^C_{t,i} = \gamma_{0,t} + \gamma_{1,t}\ BVSkew_{t-1,i} + \gamma'_{2,t} Z_{t-1,i} + \epsilon_{t,i},$$

where $R^C_{t,i}$ is the return for call option i at time t, and Z is a vector of controls. The definitions for the different variables are provided in Appendix B. The regression is run separately for options with initial moneyness $X = K/S_0$ below and above 1.02. The table reports the time-series averages of the γ-coefficients and the associated Newey and West (1987) adjusted t-statistics (with five lags) in parentheses. The last row gives the averages of the adjusted-R^2 from the regressions. ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively.

Table 5: Available Number of Strike Prices

Number of strikes	1	2	3	4	5	≥ 6
Average number of stock-months	616	471	209	101	59	110
in % of all	39.3	30.1	13.3	6.5	3.8	7.0

This table presents the average number of call option strike prices per stock and month in our sample, where moneyness $X = K/S_0$ is restricted to be between 0.9 and 2. The average is taken across the stocks which have exactly n calls with different strikes.

Table 6: Call Return Differences Across Bins (I)

Bin	2	3	4	5
$X \in$	(0.975, 1.050]	(1.050, 1.200]	(1.200, 1.400]	(1.400, 2.000]
Avg. ret. diff.				
to prev. bin (%)	4.4***	-5.4***	-22.7***	-11.6**
(t-stat)	(2.55)	(-2.45)	(-7.92)	(-2.29)
99% conf. int.	[0.5, 8.5]	[-10.3, 0.0]	[-29.2, -15.8]	[-22.4, 1.2]
Avg. no. of calls	442	463	165	43

The table shows average call return differences across moneyness bins. Call returns are one month hold-until-maturity returns. The average return difference between the higher and the lower bin is computed on the basis of calls on only those underlying stocks which are represented in both bins. The set of such stocks will vary with the bins under consideration. For the Bin 1 we use only calls with $X \in [0.9, 0.975]$. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 pairwise draws of differences. The t-stats are additionally reported in parenthesis.

Table 7: Call Return Differences Across Bins (II)

	$X \in [0.9, 1.1]$	$X \in (1.1, 1.5]$	Difference
Avg. return (%)	7.0*	-10.0*	-17.0***
(t-stat)	(1.92)	(-1.80)	(-6.62)
99% conf. int.	[-1.9, 16.0]	[-21.5, 2.6]	[-22.5, -10.7]
Avg. X	1.01	1.20	
Avg. no. of calls	511		

The table shows average call return differences across two moneyness bins. Call returns are one month hold-until-maturity returns. The average return difference between the bins is computed on the basis of calls on only those underlying stocks which are represented in both bins. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 independent draws (pairwise draws for differences). The t-stats are additionally reported in parenthesis.

Table 8: Call Return Differences: Ordinal Sort, Reference Point At-The-Money

Bin	1	2	3	4
Avg. return difference				
to previous bin (%)		-5.3***	-11.2***	-2.5
(t-stat)		(-3.05)	(-4.93)	(-0.63)
99% conf. int.		[-9.2, -1.2]	[-16.4, -5.8]	[-11.5, 6.7]
Avg. no. of calls	1,470	790	325	155
Avg. X	1.00	1.11	1.18	1.23
Std(X)	0.04	0.07	0.12	0.16
Min(X)	0.90	0.92	1.00	1.01
Max(X)	1.10	2.63	2.92	3.60

The table presents differences in average call returns between adjacent bins with ordinal sorting and the reference point at the money. The at-the-money reference bin (Bin 1) is restricted to include only calls with $X \in [0.9, 1.1]$. The values for std(X), min(X), and max(X) are calculated from all observations, i.e., without first averaging within each month. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 pairwise draws of differences. The t-stats are additionally reported in parenthesis.

Table 9: Call Return Differences: Ordinal Sort, Reference Point At-The-Money, Same Stocks in All Bins (I)

Bin	1	2	3	4
Avg. return difference				
to previous bin (%)		-1.5	-5.9***	-8.4***
(t-stat)		(-1.28)	(-3.44)	(-3.39)
99% conf. int.		[-4.1, 1.3]	[-9.7, -1.8]	[-13.9, -2.3]
Avg. no. of calls			155	
Avg. X	0.99	1.07	1.15	1.23
Std(X)	0.02	0.05	0.10	0.16
Min(X)	0.90	0.92	1.00	1.01
Max(X)	1.10	1.80	2.50	3.60

The table presents differences in average call returns between adjacent bins with ordinal sorting, reference point at the money, and options on the same underlying stocks in all bins. The at-the-money bin (Bin 1) is restricted to include only calls with $X \in [0.9, 1.1]$. The values for std(X), min(X), and max(X) are calculated from all observations, i.e., without first averaging within each month. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 pairwise draws of differences.

Table 10: Call Return Differences: Ordinal Sort, Reference Point At-The-Money, Same Stocks in All Bins (II)

Bin	1	2	3
Avg. return difference			
to previous bin (%)		-1.2	-13.8***
(t-stat)		(-0.82)	(-7.50)
99% conf. int.		[-4.4, 2.4]	[-18.0, -9.5]
Avg. no. of calls		325	
Avg. X	0.99	1.08	1.18
Std(X)	0.03	0.06	0.12
Min(X)	0.90	0.92	1.00
Max(X)	1.10	1.94	2.92

The table presents differences in average call returns between adjacent bins with ordinal sorting, reference point at the money, and options on the same underlying stocks in all bins. The at-the-money bin (Bin 1) is restricted to include only calls with $X \in [0.9, 1.1]$. The values for std(X), min(X), and max(X) are calculated from all observations, i.e., without first averaging within each month. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 pairwise draws of differences. The t-stats are additionally reported in parenthesis.

Table 11: Call Return Differences: Ordinal Sort, Reference Point Out-Of-The-Money, Same Stocks in All Bins

	3 bins			2 bins	
Bin	3	2	1	2	1
Avg. return difference to previous bin (%)		-3.41^{*}	-9.71^{***}		-12.7^{***}
(t-stat)		(-1.81)	(-3.29)		(-7.13)
99% conf. int.		[-7.6, 1.1]	[-16.2, -2.5]		[-16.7, -8.4]
Avg. no. of calls		225			521
Avg. X	1.09	1.18	1.28	1.10	1.21
Std(X)	0.15	0.18	0.24	0.14	0.19
Min(X)	1.00	1.01	1.03	1.00	1.01
Max(X)	3.77	4.46	4.75	4.46	4.75

The table presents differences in average call returns between adjacent bins with ordinal sorting, reference point out of the money, and options on the same underlying stocks in all bins. The reference out-of-the-money bin (Bin 1) is defined by the call option with the highest available strike. We only consider options with $X > 1$. The values for std(X), min(X), and max(X) are calculated from all observations, i.e., without first averaging within each month. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 pairwise draws of differences. The t-stats are additionally reported in parenthesis.

Table 12: Call Return Differences: Standardized Moneyness

$\widehat{X}$ bin	$(-3,-2]$	$(-2,-1.5]$	$(-1.5,-1]$	$(-1,-0.5]$	$(-0.5,0]$	$(0,0.5]$	$(0.5,1]$	$(1,1.5]$	$(1.5,2]$	$(2,3]$	$(3,4]$
					σ measured using past 21 day returns						
Ret. 1^{st} differences	0.9**	0.4	1.1**	0.2	2.1***	0.8	-1.1	-3.9***	-2.9*	-7.9***	-10.7***
(t-stat)	(2.05)	(0.76)	(2.18)	(0.34)	(2.83)	(0.88)	(-1.02)	(-2.85)	(-1.81)	(-3.40)	(-2.43)
99% conf. int.	[-0.1, 1.9]	[-0.7, 1.5]	[-0.1, 2.2]	[-1.1, 1.5]	[0.4, 3.9]	[-1.3, 3.1]	[-3.5, 1.4]	[-7.1, -0.7]	[-6.6, 0.8]	[-13.3, -2.4]	[-20.5, -0.1]
					σ measured using past 252 day returns						
Ret. 1^{st} differences	0.6	0.6	1.1**	0.9	2.0***	2.1**	-1.8*	-1.9	-8.9***	-15.1***	-22.4***
(t-stat)	(1.39)	(1.38)	(2.17)	(1.54)	(2.62)	(2.20)	(-1.69)	(-1.23)	(-4.17)	(-5.11)	(-3.15)
99% conf. int.	[-0.4, 1.6]	[-0.4, 1.7]	[-0.1, 2.2]	[-0.4, 2.2]	[0.3, 3.8]	[-0.1, 4.4]	[-4.2, 0.7]	[-5.4, 1.7]	[-13.9, -4.0]	[-21.9, -8.2]	[-37.9, -4.8]
					σ measured as implied volatility						
Ret. 1^{st} differences	-0.5	-0.5	0.8	1.1*	1.4*	1.1	-1.8*	-1.9	-7.6***	-12.0***	-9.9
(t-stat)	(-0.68)	(-1.29)	(1.43)	(1.69)	(1.87)	(1.22)	(-1.85)	(-1.53)	(-4.36)	(-3.62)	(-0.56)
99% conf. int.	[-2.3, 1.3]	[-1.5, 0.5]	[-0.5, 2.2]	[-0.4, 2.6]	[-0.3, 3.2]	[-0.9, 3.2]	[-4.0, 0.5]	[-4.7, 1.0]	[-11.7, -3.6]	[-19.5, -4.2]	[-44.0, 37.0]

The table presents differences in average call returns between adjacent bins. The options are sorted according to their standardized moneyness, i.e. $\widehat{X} = (K/S_0 - 1)/\sigma$. For example, $\widehat{X} = 3$ means that the strike is three return standard deviations away from the current stock price. We use three alternative measures for σ, namely historical return volatility measured from daily returns over the previous month and over the previous year, and the implied volatility of the respective option. We scale each volatility measure with time to maturity. The first bin (Bin 1) is restricted to include only calls with $\widehat{X} \in [-4, -3]$. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 pairwise draws of differences. The t-stats are additionally reported in parenthesis.

Table 13: Two and Six Week Call Returns by Moneyness (All Available Options)

Bin	1	2	3	4	5	6	7	8	9	10	11
$X \in$	(0.9, 0.95]	(0.95, 0.975]	(0.975, 1]	(1, 1.025]	(1.025, 1.05]	(1.05, 1.1]	(1.1, 1.2]	(1.2, 1.3]	(1.3, 1.4]	(1.4, 1.6]	(1.6, 2]
Panel A: Return horizon two weeks											
Avg. X	0.926	0.963	0.988	1.013	1.037	1.073	1.141	1.242	1.345	1.483	1.756
Avg. return (%)	4.7**	7.3***	9.0***	11.3***	11.3***	1.3	-23.4***	-53.0***	-69.1***	-81.4***	-90.2***
(t-stat)	(1.99)	(2.53)	(2.67)	(2.76)	(2.41)	(0.27)	(-5.55)	(-12.72)	(-16.52)	(-22.17)	(-20.70)
99% conf. int.	[-0.8, 10.2]	[0.7, 13.9]	[1.2, 17.0]	[1.9, 20.9]	[0.7, 22.5]	[-9.5, 12.7]	[-32.9, -13.3]	[-62.1, -42.7]	[-78.0, -58.7]	[-88.8, -72.0]	[-97.2, -77.7]
Avg. return difference to previous bin (%)		2.6***	1.8**	2.2**	0.1	-10.0***	-24.7***	-29.6***	-16.1***	-12.3***	-8.8*
(t-stat)		(3.59)	(2.20)	(1.97)	(0.05)	(-6.66)	(-10.00)	(-11.60)	(-5.91)	(-3.59)	(-1.88)
99% conf. int.		[0.9, 4.3]	[-0.1, 3.7]	[-0.4, 4.9]	[-2.8, 3.1]	[-13.6, -6.6]	[-30.6, -19.0]	[-35.5, -23.7]	[-22.5, -9.8]	[-20.4, -4.4]	[-18.8, 3.3]
Avg. no. of calls	682	433	461	441	423	582	528	239	129	117	77
Panel B: Return horizon six weeks											
Avg. X	0.926	0.963	0.988	1.013	1.037	1.074	1.144	1.243	1.344	1.481	1.755
Avg. return (%)	5.5	7.0*	7.5*	10.8**	10.2*	8.3	-0.9	-16.4***	-35.2***	-50.1***	-64.8***
(t-stat)	(1.87)	(2.12)	(2.12)	(2.64)	(2.35)	(1.78)	(-0.18)	(-3.09)	(-6.42)	(-9.46)	(-10.01)
99% conf. int.	[-2.8, 13.5]	[-2.5, 15.8]	[-2.5, 17.0]	[-0.5, 21.6]	[-2.0, 21.5]	[-4.5, 20.4]	[-13.8, 11.8]	[-30.3, -1.4]	[-49.1, -18.3]	[-63.8, -33.0]	[-77.7, -48.2]
Avg. return difference to previous bin (%)		1.5*	0.6	3.2***	-0.6	-1.9*	-9.2***	-15.5***	-18.8***	-14.8***	-14.8***
(t-stat)		(2.19)	(0.72)	(3.25)	(-0.61)	(-1.64)	(-6.71)	(-7.10)	(-7.20)	(-4.74)	(-2.71)
99% conf. int.		[-0.5, 3.2]	[-1.2, 2.4]	[1.0, 5.8]	[-3.2, 1.6]	[-4.5, 0.7]	[-13.0, -5.3]	[-21.1, -9.4]	[-24.7, -12.2]	[-21.7, -7.8]	[-26.8, -1.6]
Avg. no of calls	474	293	318	315	322	539	634	330	172	142	84

The table presents statistics on hold-until-maturity call returns. In Panel A the call maturity is two weeks, and in Panel B six weeks. All option positions are bought on the first trading (usually Monday) day after 21 (49) days before the third Friday of each month and hold until maturity. For each moneyness bin and date, only the option closest to the midpoint of the moneyness interval is considered for each underlying stock. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 independent draws (pairwise draws for differences). For the statistics of the six week maturity options in Panel B we use block bootstraps for both levels and differences, with a block size of five. The t-stats are additionally reported in parenthesis. The sample is from Jan 1996 to Dec 2017.

Table 14: Two and Six Week Call Returns (Same Stocks in All Bins)

	Two weeks			Six weeks		
	$X \in [0.9, 1.1]$	$X \in (1.1, 1.5]$	Diff.	$X \in [0.9, 1.1]$	$X \in (1.1, 1.5]$	Diff.
Avg. return (%)	5.9	-42.0***	-48.0***	5.5	-13.9**	-19.4***
(t-stat)	(1.49)	(-10.66)	(-17.42)	(1.38)	(-2.65)	(-8.41)
99% conf. int.	[-3.2, 15.4]	[-50.8, -32.4]	[-54.4, -41.5]	[-5.1, 16.3]	[-27.6, 2.0]	[-25.7, -11.5]
Avg. X	1.01	1.22		1.02	1.24	
Avg. no. of calls		574			676	

The table presents average two week and six week hold-until-maturity call returns when we require the two bins to contain exactly the same underlying stocks. For each stock at each date, we only include the calls with moneyness X closest to 1.0 and 1.3, respectively. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 independent draws (pairwise draws for differences). For the statistics of the six week maturity options we use block bootstraps for both levels and differences, with a block size of five.

Table 15: Differences of Intermediate Call Option Returns across Moneyness and Time to Maturity

Bin	2	3	4	5	6	7	8	9	10	11	12	13
Range of x	(0.975, 1]	(1, 1.025]	(1.025, 1.05]	(1.05, 1.075]	(1.075, 1.1]	(1.1, 1.2]	(1.2, 1.3]	(1.3, 1.4]	(1.4, 1.6]	(1.6, 1.8]	(1.8, 2]	(2, 2.2]
Panel A: 2.5 weeks												
Avg. ret. diff. to prev. bin (%)	2.3***	3.7***	5.4***	5.5***	3.7***	0.0	-3.6***	-3.0***	-4.2***	-3.6***	-2.0***	-2.6***
(t-stat)	(5.52)	(7.07)	(9.42)	(8.25)	(4.33)	(0.01)	(-3.67)	(-3.75)	(-5.38)	(-6.03)	(-4.64)	(-4.62)
99% conf. int.	[1.3, 3.2]	[2.5, 4.9]	[4.1, 6.7]	[4.0, 7.1]	[1.7, 5.7]	[-2.8, 2.7]	[-5.9, -1.3]	[-4.8, -1.1]	[-6.1, -2.5]	[-5.0, -2.3]	[-3.0, -1.0]	[-3.9, -1.3]
Avg. no. of calls	453	484	473	466	452	1220	990	761	806	471	290	172
Panel B: 3.5 weeks												
Avg. ret. diff. to prev. bin (%)	1.5***	1.9***	2.3***	2.7***	3.4***	3.5***	-2.1**	-1.7***	-2.9***	-1.6***	-1.4**	-1.4***
(t-stat)	(4.66)	(4.99)	(5.48)	(5.24)	(5.10)	(4.24)	(-2.21)	(-2.53)	(-4.58)	(-2.78)	(-1.93)	(-2.22)
99% conf. int.	[0.8, 2.3]	[1.0, 2.8]	[1.4, 3.4]	[1.5, 3.9]	[1.9, 5.0]	[1.6, 5.4]	[-4.4, 0.1]	[-3.2, -0.2]	[-4.4, -1.5]	[-2.9, -0.2]	[-3.1, 0.2]	[-2.9, 0.0]
Avg. no. of calls	434	461	449	441	429	1159	943	722	770	459	294	191
Panel C: 4.5 weeks												
Avg. ret. diff. to prev. bin (%)	0.5*	-0.3	0.3	1.9***	1.7***	4.1***	2.3***	0.7	-1.2**	-2.2***	-0.3	-1.5***
(t-stat)	(1.93)	(-1.05)	(0.97)	(4.35)	(3.14)	(6.09)	(3.43)	(1.23)	(-2.05)	(-4.16)	(-0.58)	(-2.76)
99% conf. int.	[-0.1, 1.2]	[-1.0, 0.4]	[-0.4, 1.1]	[0.9, 2.9]	[0.4, 3.0]	[2.5, 5.6]	[0.7, 3.9]	[-0.6, 2.0]	[-2.5, 0.1]	[-3.4, -1.0]	[-1.5, 0.9]	[-2.7, -0.2]
Avg. no. of calls	417	442	427	415	409	1123	907	700	745	441	287	188
Panel D: 5.5 weeks												
Avg. ret. diff. to prev. bin (%)	0.9***	-0.6**	0.7**	1.6***	1.8***	3.7***	2.2***	0.3	-1.4**	-1.8***	-0.5	-1.6***
(t-stat)	(3.17)	(-2.26)	(2.25)	(3.56)	(4.03)	(6.25)	(2.68)	(0.46)	(-2.28)	(-3.61)	(-0.70)	(-2.57)
99% conf. int.	[0.3, 1.6]	[-1.3, 0.0]	[0.0, 1.5]	[0.6, 2.7]	[0.8, 2.9]	[2.3, 5.1]	[0.2, 4.1]	[-1.2, 1.8]	[-2.8, 0.0]	[-3.1, -0.7]	[-2.0, 1.2]	[-3.2, -0.3]
Avg. no. of calls	389	410	395	387	380	1057	856	659	707	420	276	180
Panel E: 6.5 weeks												
Avg. ret. diff. to prev. bin (%)	0.6**	0.5*	0.9***	1.1***	1.8***	3.8***	2.8***	1.5***	-0.8	-2.2***	-0.7	-0.7
(t-stat)	(2.24)	(1.93)	(2.85)	(3.16)	(4.35)	(7.95)	(4.70)	(2.95)	(-1.51)	(-4.46)	(-1.30)	(-1.12)
99% conf. int.	[0.0, 1.1]	[-0.1, 1.2]	[0.2, 1.6]	[0.3, 2.0]	[0.9, 2.8]	[2.7, 4.9]	[1.4, 4.3]	[0.3, 2.6]	[-2.1, 0.4]	[-3.4, -1.1]	[-2.0, 0.6]	[-2.0, 0.7]
Avg. no. of calls	367	387	375	368	355	1000	811	624	678	398	268	171

The table presents the average intermediate call return differences between moneyness bins for options with different initial time to maturity. The holding period for the intermediate returns in one week. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 pairwise draws of differences. The first bin (Bin 1) is restricted to include only calls with $X \in [0.95, 0.975]$.

Table 16: Differences of Intermediate Call Options Returns Across Bins

Bin	2	3	4	5	6
Range for X	(1.00, 1.05]	(1.05, 1.2]	(1.2, 1.4]	(1.4, 1.6]	(1.6, 2]
Panel A: 2.5 weeks					
Avg. ret. diff. to prev. bin (%)	2.8***	8.5***	-6.1***	-8.1***	-4.5***
(t-stat)	(3.15)	(6.59)	(-3.54)	(-5.49)	(-4.47)
99% conf. int.	[0.8, 4.9]	[5.5, 11.4]	[-10.2, -2.2]	[-11.7, -4.9]	[-7.0, -2.3]
Avg. no. of calls	303	638	1,050	651	422
Panel B: 3.5 weeks					
Avg. ret. diff. to prev. bin (%)	0.9	8.5***	-0.7	-3.1***	-1.9***
(t-stat)	(1.40)	(8.08)	(-0.49)	(-3.19)	(-2.93)
99% conf. int.	[-0.5, 2.5]	[6.2, 11.1]	[-4.4, 2.7]	[-5.5, -0.9]	[-3.5, -0.5]
Avg. no. of calls	293	606	995	616	413
Panel C: 4.5 weeks					
Avg. ret. diff. to prev. bin (%)	-1.4***	6.0***	5.3***	0.7	-1.3
(t-stat)	(-2.94)	(8.49)	(5.26)	(0.87)	(-1.58)
99% conf. int.	[-2.5, -0.3]	[4.3, 7.6]	[3.0, 7.6]	[-1.3, 2.6]	[-3.2, 0.5]
Avg. no. of calls	279	582	957	593	398
Panel D: 5.5 weeks					
Avg. ret. diff. to prev. bin (%)	0.3	5.5***	6.1***	0.2	-1.4
(t-stat)	(0.71)	(6.10)	(5.34)	(0.26)	(-1.52)
99% conf. int.	[-0.7, 1.4]	[3.7, 7.9]	[3.4, 8.7]	[-2.0, 2.4]	[-3.7, 0.7]
Avg. no. of calls	258	540	897	560	378
Panel E: 6.5 weeks					
Avg. ret. diff. to prev. bin (%)	1.1***	6.4***	6.0***	-0.4	-1.6***
(t-stat)	(3.00)	(13.72)	(6.93)	(-0.53)	(-2.60)
99% conf. int.	[0.3, 2.0]	[5.3, 7.5]	[4.0, 8.0]	[-2.3, 1.4]	[-3.2, -0.2]
Avg. no. of calls	240	507	845	532	359

The table presents the average intermediate call return differences between moneyness bins for options with different initial time to maturity, with the additional requirement that the same stocks are included in neighboring bins at the same date. The holding period for the intermediate returns in one week. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 pairwise draws of differences.

Table 17: Differences of Call Options Returns for Subsamples

Moneyness bin	(0.95, 0.975]	(0.975, 1]	(1, 1.025]	(1.025, 1.05]	(1.05, 1.1]	(1.1, 1.2]	(1.2, 1.3]	(1.3, 1.4]	(1.4, 1.6]	(1.6, 2]
ex ante VIX<18 (in-sample median)										
Ret. 1^{st} differences	1.0	2.3**	2.2*	0.7	-4.9**	-7.4***	-23.6***	-3.9	-22.7	-26.5*
(t-stat)	(1.15)	(2.11)	(1.70)	(0.45)	(-2.33)	(-2.48)	(-4.80)	(-0.27)	(-1.21)	(-1.88)
99% conf. int.	[-0.9, 2.9]	[-0.2, 4.9]	[-0.7, 5.4]	[-2.8, 4.3]	[-9.7, 0.1]	[-14.2, -0.4]	[-34.8, -12.0]	[-32.7, 32.5]	[-68.9, 18.6]	[-60.5, 5.5]
ex ante VIX>18 (in-sample median)										
Ret. 1^{st} differences	2.3***	2.2**	2.3	-1.4	-2.5*	-11.4***	-16.1***	-12.3**	-12.2	-25.4***
(t-stat)	(2.55)	(2.04)	(1.61)	(-0.87)	(-1.69)	(-4.26)	(-4.44)	(-2.10)	(-1.32)	(-2.20)
99% conf. int.	[0.3, 4.4]	[-0.3, 4.8]	[-0.9, 5.6]	[-5.0, 2.4]	[-5.9, 1.0]	[-17.6, -5.2]	[-24.5, -7.7]	[-25.8, 1.5]	[-31.5, 11.3]	[-57.2, -3.8]
01/1996-06/2001										
Ret. 1^{st} differences	2.0	1.0	0.3	0.8	-1.3	-12.2***	-15.7***	-7.1	-22.5***	-33.6***
(t-stat)	(1.42)	(0.64)	(0.12)	(0.29)	(-0.53)	(-3.00)	(-3.49)	(-1.13)	(-2.96)	(-3.84)
99% conf. int.	[-1.1, 5.3]	[-2.6, 4.7]	[-4.7, 5.8]	[-5.4, 7.0]	[-7.1, 4.6]	[-21.5, -2.7]	[-26.0, -4.9]	[-20.4, 8.5]	[-41.0, -5.7]	[-56.4, -16.1]
07/2001-12/2006										
Ret. 1^{st} differences	1.7*	2.5	2.1	-2.5	-4.5*	-9.6***	-35.3***	-9.4	-22.0	-14.7
(t-stat)	(1.82)	(1.52)	(1.18)	(-1.17)	(-1.88)	(-2.72)	(-5.49)	(-0.67)	(-1.13)	(-1.33)
99% conf. int.	[-0.4, 4.0]	[-1.1, 6.5]	[-1.9, 6.4]	[-7.4, 2.5]	[-9.9, 1.2]	[-17.5, -1.1]	[-49.9, -20.2]	[-38.6, 27.7]	[-71.1, 20.1]	[-43.5, 7.2]
01/2007-06/2012										
Ret. 1^{st} differences	0.4	2.5*	0.2	-2.1	-2.7	-6.2*	-10.7	-20.8*	0.9	-25.8
(t-stat)	(0.30)	(1.90)	(0.15)	(-1.06)	(-1.20)	(-1.97)	(-1.62)	(-1.79)	(0.05)	(-1.14)
99% conf. int.	[-2.5, 3.4]	[-0.5, 5.6]	[-3.1, 3.8]	[-6.5, 2.6]	[-7.8, 2.7]	[-13.2, 1.3]	[-25.5, 5.3]	[-48.1, 6.0]	[-31.7, 45.0]	[-89.1, 12.1]
07/2012-12/2017										
Ret. 1^{st} differences	2.8**	3.0*	6.6***	2.2	-5.9**	-10.6**	-15.7***	4.0	-24.5	-28.5
(t-stat)	(2.13)	(1.82)	(3.09)	(1.08)	(-2.06)	(-2.06)	(-2.84)	(0.19)	(-0.83)	(-1.28)
99% conf. int.	[-0.2, 5.8]	[-0.6, 7.0]	[1.9, 11.9]	[-2.3, 7.0]	[-12.2, 1.0]	[-22.8, 0.9]	[-28.6, -3.1]	[-34.9, 61.5]	[-98.0, 39.7]	[-81.4, 22.9]
Only S&P 500 constituents										
Ret. 1^{st} differences	1.7***	2.3***	2.2***	-0.5	-3.6***	-9.6***	-19.4***	-8.6	-16.8*	-25.9***
(t-stat)	(2.73)	(2.91)	(2.31)	(-0.42)	(-2.86)	(-4.83)	(-6.54)	(-1.21)	(-1.73)	(-2.90)
99% conf. int.	[0.3, 3.1]	[0.5, 4.1]	[0.0, 4.6]	[-3.0, 2.1]	[-6.5, -0.6]	[-14.2, -5.0]	[-26.3, -12.5]	[-23.6, 9.2]	[-40.4, 5.3]	[-48.4, -6.7]

The table presents differences in average call returns between adjacent moneyness bins for various subsamples, and otherwise use our benchmark approach. The first bin is restricted to include only calls with moneyness$X \in [0.9, 0.95]$. In the first two panels, we condition of the level of the VIX at formation date and differentiate whether the ex ante VIX is above or below its median of 18. In the following four panels the analysis is conducted separately for subsequent periods of 5.5 years. For the last panels, only stocks that where constituents of the S&P 500 index at their time are included. The stars indicate that the values are statistically significant at the 10%, 5% and 1% level. Significance is based on confidence intervals computed from bootstraps with 100,000 pairwise draws of differences. The t-stats are additionally reported in parenthesis.

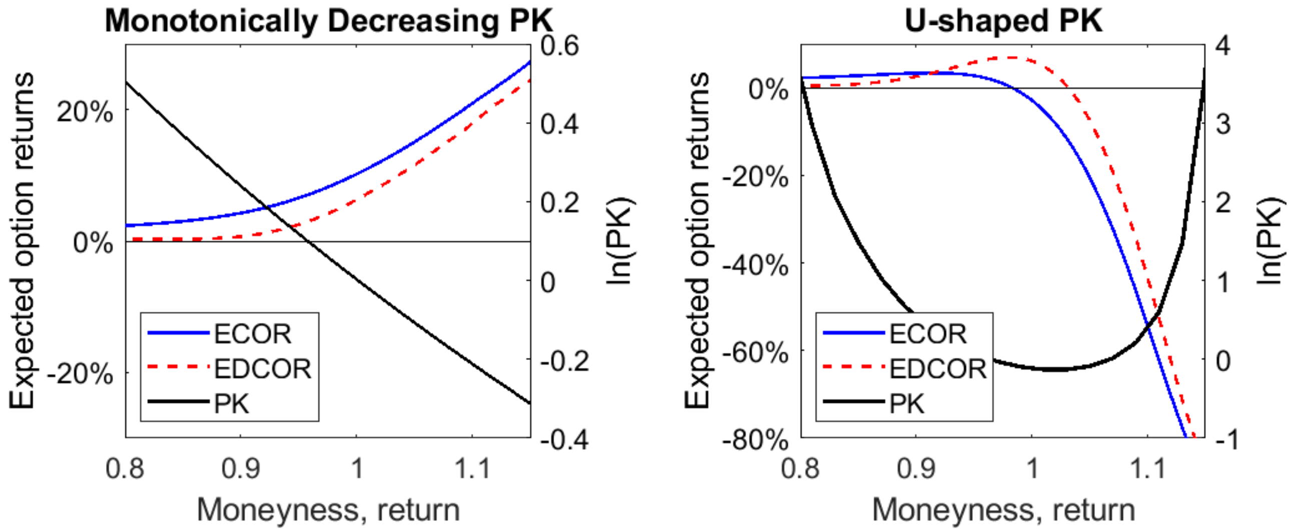

Figure 1: Expected Option Returns for Alternative Shapes of the Pricing Kernel

The graphs show the expected returns of different types of options under a monotonically decreasing and a U-shaped PK, respectively. ECOR and EDCOR stand for expected call option returns and expected digital call option returns, respectively. The left graph is based on the Black-Scholes-Merton model, with $\mu = 8.0\%$ p.a., $\sigma = 18.86\%$ p.a., and $r = 1.51\%$ p.a. The right graph shows the results from a simulation of the Christoffersen et al. (2013) model, which is a special case of the model from Section 3.2, with $b = 1$ and $c = 0$. Moneyness is defined as K/S_0.

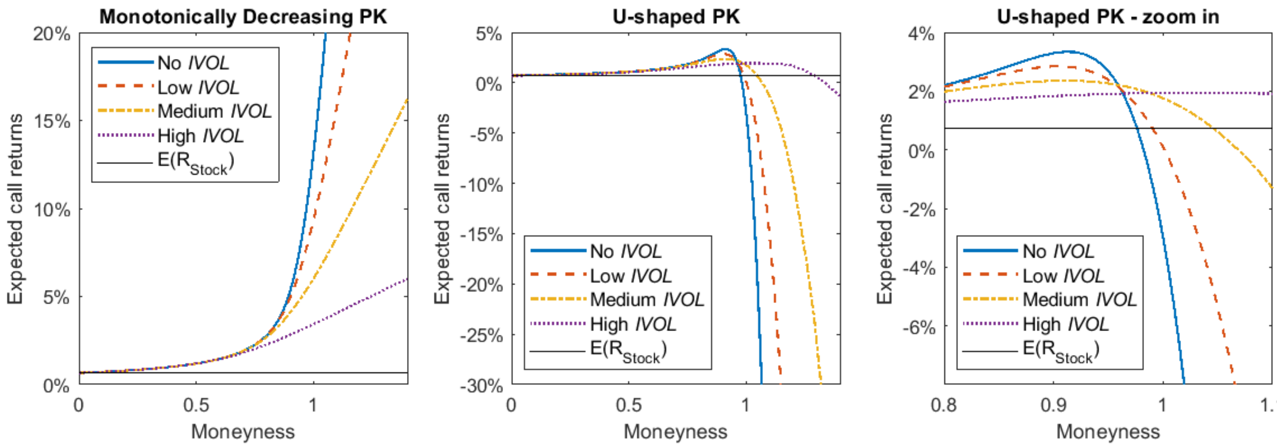

Figure 2: Call Option Returns in the Black-Scholes-Merton Model and the Model with the U-shaped PK for Different Levels of $IVOL$

The left plot illustrates average one month hold-until-maturity call returns in the Black-Scholes-Merton model for different levels of $IVOL$. The middle and right plots show average one month hold-until-maturity call returns in the model with the U-shaped PK by moneyness for different levels of idiosyncratic volatility ($IVOL$). The model specification is outlined in Section 3.2, and in both cases, we keep $b_i = 1$, and vary c_i in $\{0, 1, 2, 4\}$. Details on the parameters and simulation can be found in Section 3.2 and Appendix C.2.

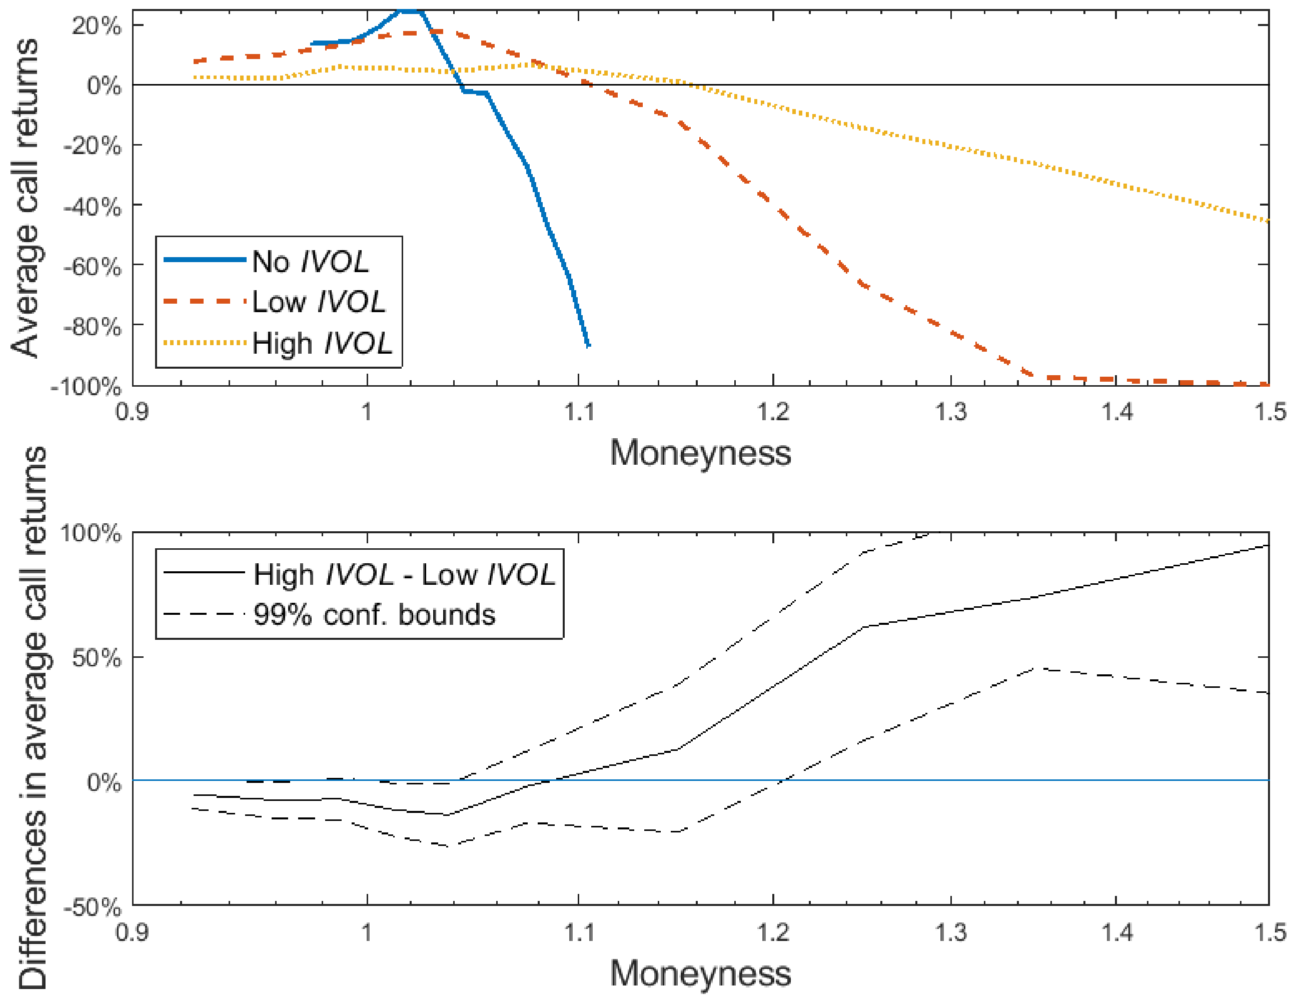

Figure 3: Call Option Returns for the S&P 500 index and $IVOL$ Sorted Portfolio

We perform an independent double sort on moneyness and $IVOL$. The upper graph shows average call returns for the high and low $IVOL$ portfolios, as well as for an asset with zero $IVOL$ represented by the S&P 500 index. The lower graph shows the average return difference between the high and the low $IVOL$ portfolios, i.e., the difference between the dotted and the dashed lines in the upper graph. Moneyness is defined as K/S_0. In the lower graph dashed lines represent the bounds of the 99% confidence interval generated via a bootstrap with 100,000 pairwise draws. The x-axis are log-scaled.

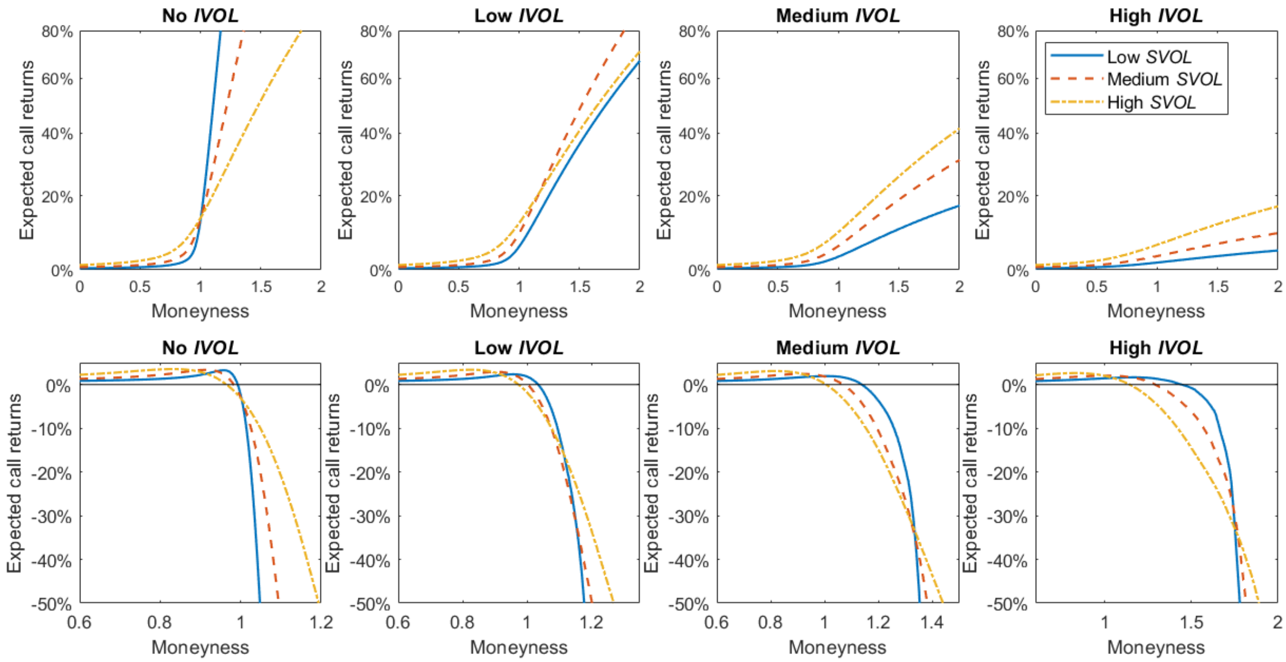

Figure 4: Call Option Returns in the Black-Scholes-Merton Model and the Model with the U-shaped PK for Different Levels of $SVOL$ and $IVOL$

The plot shows average one month hold-until-maturity call returns for different levels of $SVOL$ and $IVOL$ in the Black-Scholes-Merton model (top) and in the model with the U-shaped PK (bottom) from Equ. (1)-(5) in Section 3.2. In each plot, we vary the level of $SVOL$, while keeping $IVOL$ fixed (i.e., fixing one c_i from $\{0, 1, 2, 4\}$ in Equ. (1)). $SVOL$ is modeled in a "CAPM-style", where b_i takes on a value from $\{0.5, 1, 2\}$. Further details on the parameters and the simulation can be found in Section 3.2 and Appendix C.2.

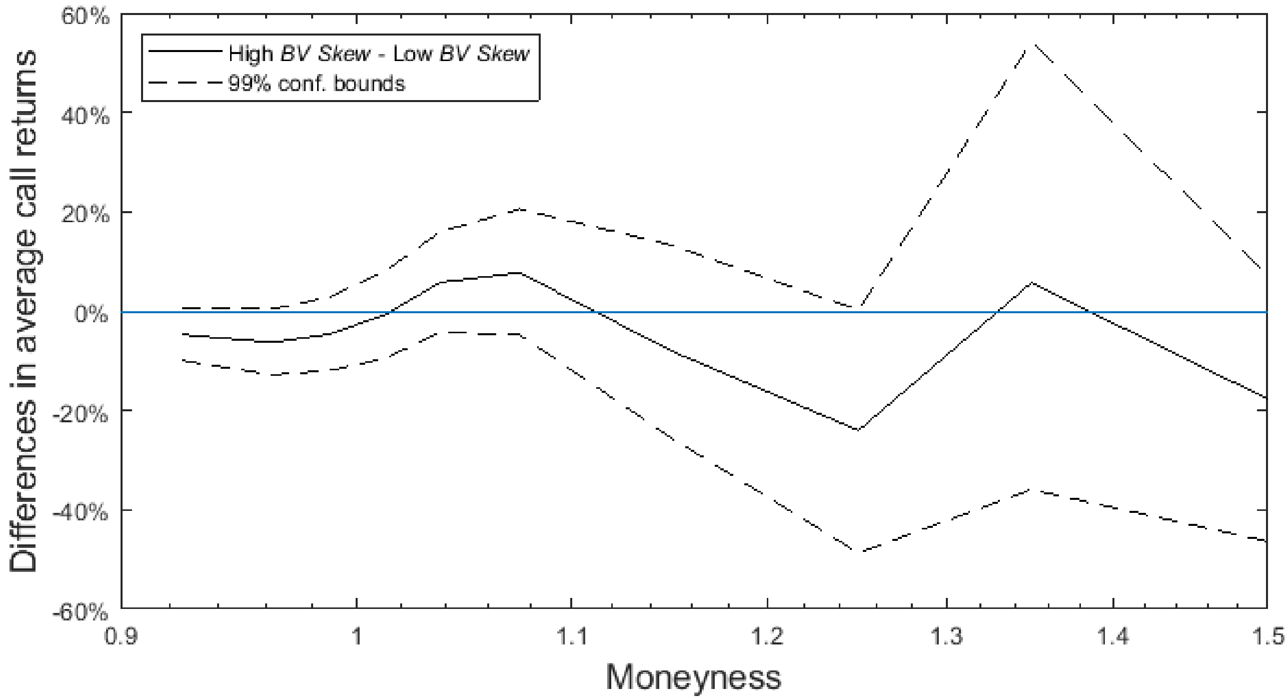

Figure 5: Call Option Returns for High Minus Low $BVSkew$ Sorted Portfolios

We perform a dependent double sort on moneyness and $BVSkew$, i.e., within a given moneyness bin, we sort call options in five portfolios based on their $BVSkew$ measure. The figure shows the difference in average returns between low and high $BVSkew$ portfolios (see Section 4) across moneyness bins. Moneyness is defined as K/S_0. The dashed lines represent the bounds of the 99% confidence interval generated via a bootstrap with 100,000 pairwise draws. The x-axis are log-scaled.

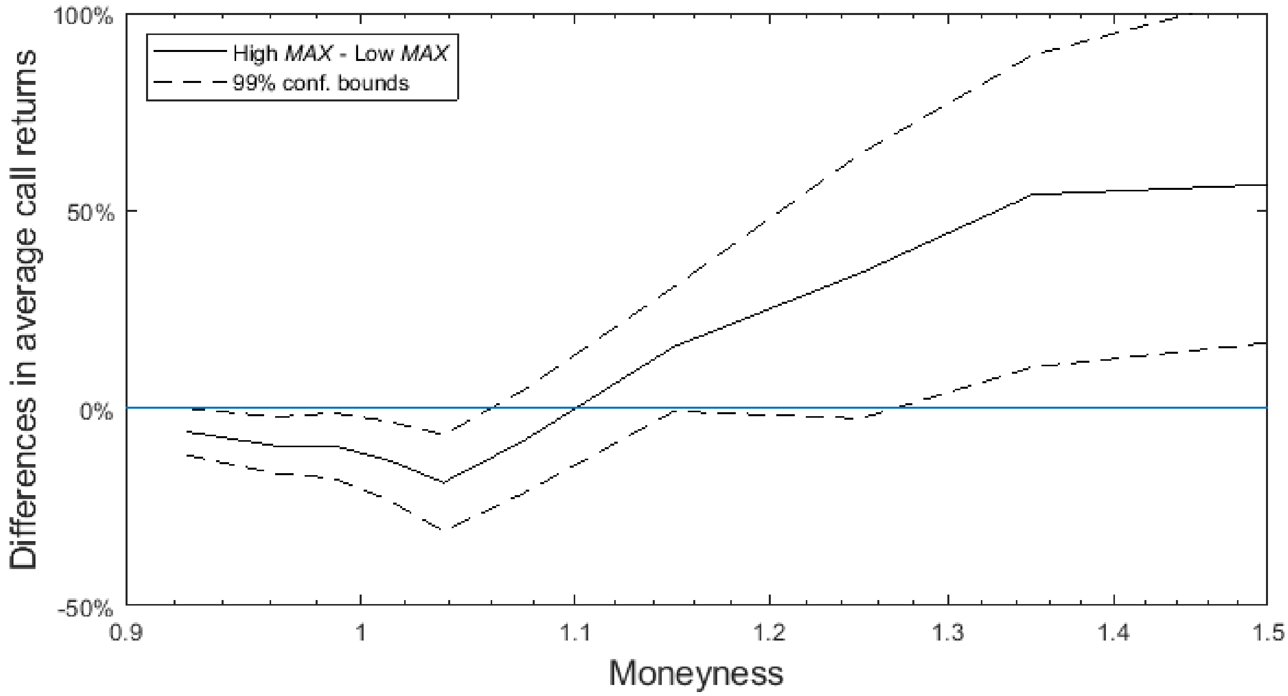

Figure 6: Call Option Returns for High Minus Low MAX Sorted Portfolios

We perform an independent double sort on moneyness and MAX. The figure shows the difference in average returns between call portfolios on high and low MAX stocks (see Section 4) across moneyness bins. Moneyness is defined as K/S_0. The dashed lines represent the bounds of the 99% confidence interval generated via a bootstrap with 100,000 pairwise draws. The x-axis are log-scaled.

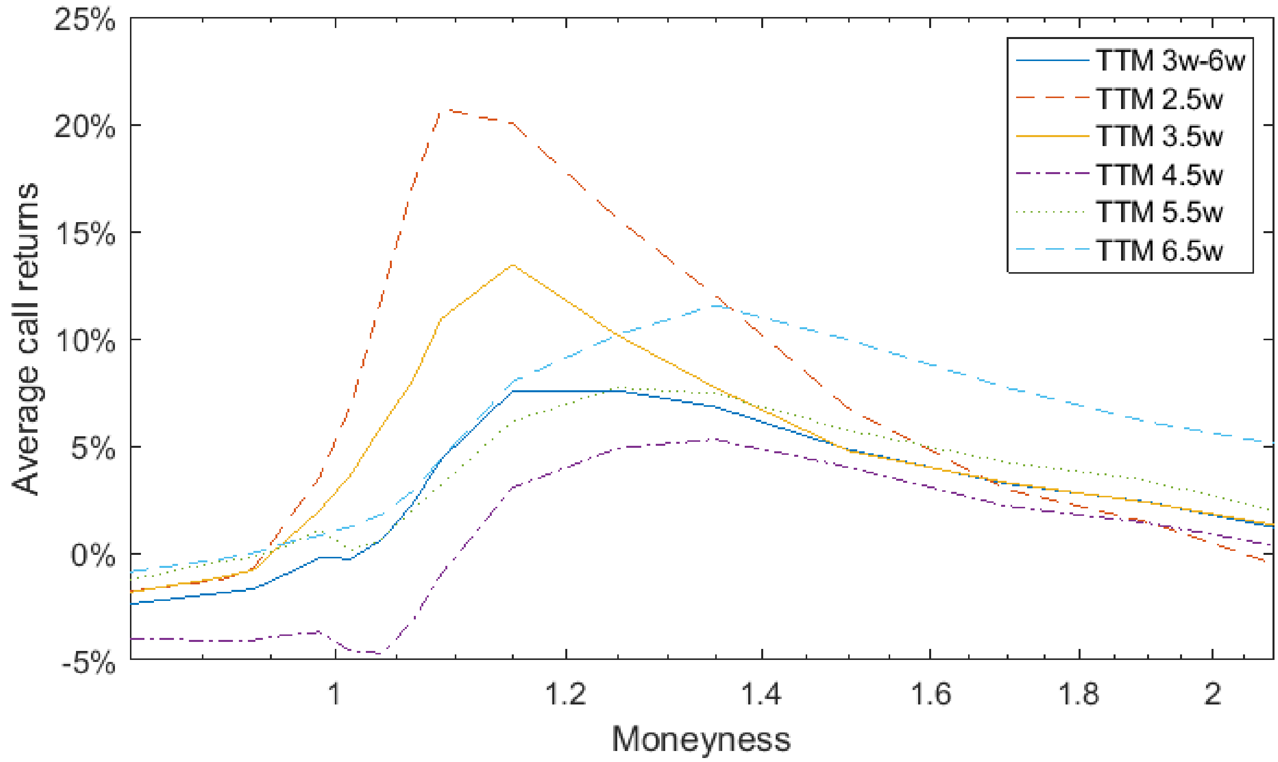

Figure 7: Intermediate Call Returns by Moneyness and Time to Maturity

The figure plots the average intermediate call returns as a function of moneyness for different option maturities. The holding period for the intermediate returns in one week. The '3w-6w' line is the average of returns for options with 3.5, 4.5, and 5.5 weeks to maturity, as in Chaudhuri and Schroder (2015). Moneyness is defined as K/S_0. Points on the moneyness axis represent the bin midpoint, and the axis are log-scaled.

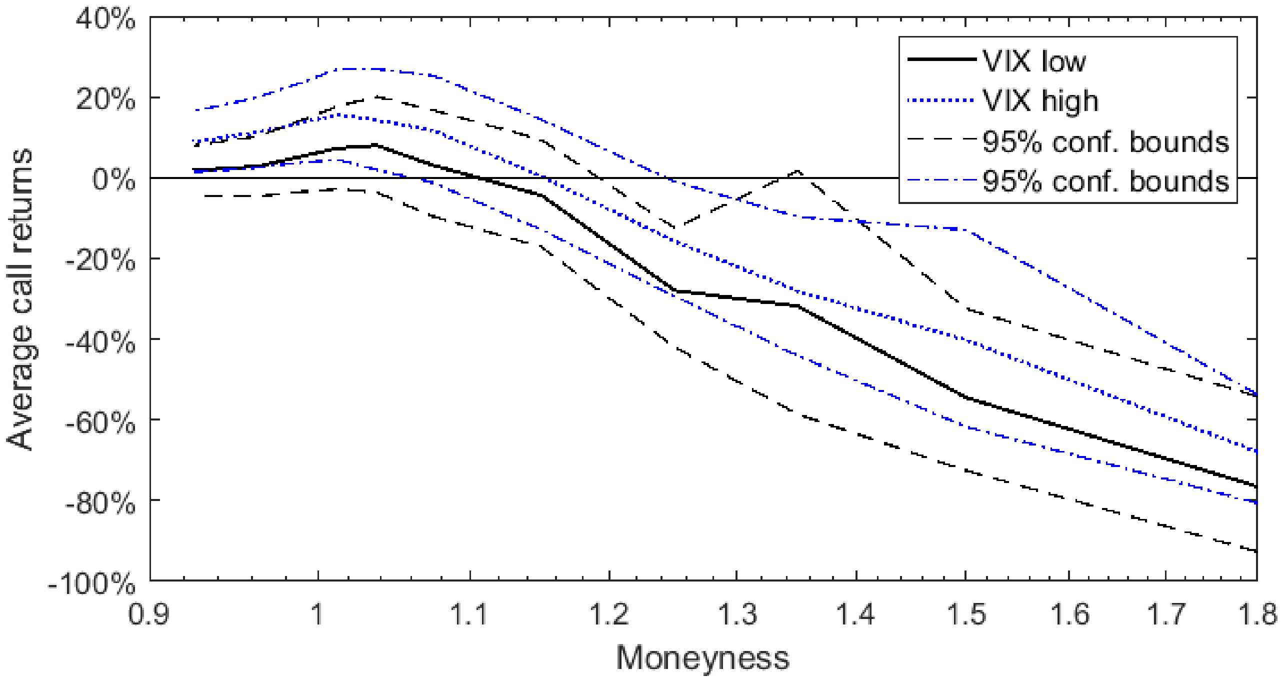

Figure 8: Call Option Returns for High and Low Volatility Periods (Measured by VIX)

The figure shows average one month hold-until-maturity call returns for two subsamples where the volatility index VIX is above or below a level of 18 (its full-sample median) at the portfolio formation date. Moneyness is defined as K/S_0. The dashed lines represent 95% confidence intervals generated via a bootstrap with 100,000 pairwise draws. Points on the moneyness axis represent the bin midpoint, and the axis are log-scaled.

References

AMIHUD, Y. (2002): "Illiquidity and stock returns: cross-section and time-series effects," *Journal of Financial Markets*, 5, 31–56.

ANG, A., R. J. HODRICK, Y. XING, AND X. ZHANG (2006): "The cross-section of volatility and expected returns," *Journal of Finance*, 61, 259–299.

ARETZ, K., M.-T. LIN, AND S.-H. POON (2018): "Moneyness, Volatility, and the Cross-Section of Option Returns," Worlking Paper.

BABAOĞLU, K., P. CHRISTOFFERSEN, S. HESTON, AND K. JACOBS (2017): "Option valuation with volatility components, fat tails, and nonmonotonic pricing kernels," *Review of Asset Pricing Studies*, 8, 183–231.

BAELE, L., J. DRIESSEN, S. EBERT, J. M. LONDONO, AND O. G. SPALT (2019): "Cumulative prospect theory, option returns, and the variance premium," *Review of Financial Studies*, 32, 3667–3723.

BAKSHI, G., J. CROSBY, AND X. GAO BAKSHI (2019): "The Dark Matter in Equity Index Volatility Dynamics: Assessing the Economic Rationales for Unspanned Risks," Working Paper.

BAKSHI, G., D. MADAN, AND G. PANAYOTOV (2010): "Returns of claims on the upside and the viability of U-shaped pricing kernels," *Journal of Financial Economics*, 97, 130–154.

BALI, T. G., N. CAKICI, AND R. F. WHITELAW (2011): "Maxing out: Stocks as lotteries and the cross-section of expected returns," *Journal of Financial Economics*, 99, 427–446.

BANSAL, R. AND A. YARON (2004): "Risks for the long run: A potential resolution of asset pricing puzzles," *Journal of Finance*, 59, 1481–1509.

BARONE-ADESI, G., N. FUSARI, A. MIRA, AND C. SALA (2020): "Option market trading activity and the estimation of the pricing kernel: A Bayesian approach," *Journal of Econometrics*, 216, 430–449.

BLACK, F. AND M. SCHOLES (1973): "The pricing of options and corporate liabilities," *Journal of political economy*, 81, 637–654.

BLAU, B. M., T. B. BOWLES, AND R. J. WHITBY (2016): "Gambling preferences, options markets, and volatility," *Journal of Financial and Quantitative Analysis*, 51, 515–540.

BOYER, B. H. AND K. VORKINK (2014): "Stock options as lotteries," *Journal of Finance*, 69, 1485–1527.

BROADIE, M., M. CHERNOV, AND M. JOHANNES (2009): "Understanding Index Option Returns," *Review of Financial Studies*, 22, 4493–4529.

Byun, S.-J. and D.-H. Kim (2016): "Gambling preference and individual equity option returns," *Journal of Financial Economics*, 122, 155–174.

Chabi-Yo, F. (2012): "Pricing kernels with stochastic skewness and volatility risk," *Management Science*, 58, 624–640.

Chan, K. and S. X. Ni (2018): "Idiosyncratic Skewness Preference or Measurement Error? Evidence from the Options Market," Working Paper.

Chaudhuri, R. and M. Schroder (2015): "Monotonicity of the stochastic discount factor and expected option returns," *Review of Financial Studies*, 28, 1462–1505.

Chaudhury, M. (2017): "Volatility and expected option returns: A note," *Economics Letters*, 152, 1–4.

Christoffersen, P., S. Heston, and K. Jacobs (2013): "Capturing option anomalies with a variance-dependent pricing kernel," *Review of Financial Studies*, 26, 1963–2006.

Coval, J. and T. Shumway (2001): "Expected option returns," *Journal of Finance*, 56, 983–1009.

Driessen, J., P. J. Maenhout, and G. Vilkov (2009): "The price of correlation risk: Evidence from equity options," *The Journal of Finance*, 64, 1377–1406.

Duarte, J., C. S. Jones, and J. L. Wang (2019): "Very noisy option prices and inferences regarding option returns," Working Paper.

Fama, E. F. and K. R. French (1993): "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 33, 3–56.

Filippou, I., P. Garcia-Ares, and F. Zapatero (2018): "Demand for lotteries: The choice between stocks and options," Working Paper.

——— (2019): "Short Squeeze Uncertainty and Skewness," Working Paper.

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

Heston, S. L. (1993): "A closed-form solution for options with stochastic volatility with applications to bond and currency options," *Review of Financial Studies*, 6, 327–343.

Hu, G. and K. Jacobs (2020): "Volatility and Expected Option Returns," *Journal of Financial and Quantitative Analysis*, 55, 1025–1060.

Lemmon, M. and S. X. Ni (2014): "Differences in trading and pricing between stock and index options," *Management Science*, 60, 1985–2001.

Linn, M., S. Shive, and T. Shumway (2018): "Pricing kernel monotonicity and conditional information," *Review of Financial Studies*, 31, 493–531.

Merton, R. C. (1973): "Theory of rational option pricing," *The Bell Journal of economics and management science*, 141–183.

Muravyev, D. (2016): "Order flow and expected option returns," *Journal of Finance*, 71, 673–708.

Newey, W. K. and K. D. West (1987): "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix," *Econometrica*, 55, 703–708.

Ni, S. X. (2008): "Stock option returns: A puzzle," Working Paper.

Polkovnichenko, V. and F. Zhao (2013): "Probability weighting functions implied in options prices," *Journal of Financial Economics*, 107, 580–609.

Rubinstein, M. (1976): "The valuation of uncertain income streams and the pricing of options," *Bell Journal of Economics*, 7, 407–425.

Sichert, T. (2020): "The Pricing Kernel is U-shaped," Working Paper.

Vasquez, A. (2017): "Equity volatility term structures and the cross section of option returns," *Journal of Financial and Quantitative Analysis*, 52, 2727–2754.

A Data filters

The options data are obtained from the OptionMetrics database. Following the standard in the literature, we apply the following filters, i.e., we exclude option price observations with one of the following characteristics:

- best bid equal to zero
- best bid greater than or equal to best ask
- average between bid and ask price less than 1/8$
- bid ask spread is smaller than the minimum tick size
- no p.m. settlement
- not American style
- non-standard expiration day
- non-standard settlement
- zero open interest
- missing implied volatility
- delta greater than one or less than minus one
- violation of standard no arbitrage bounds
- dividends paid during the remaining lifetime of the option.

After merging OptionMetrics with CRSP using the mapping table provided by WRDS, we only keep those stocks with a matching score of 1.

Deviating from above, when using data of options written on the S&P 500 index, we (in addition) use these respective standard requirements:

- volume> 0
- option is European style
- option is a.m. settled on the standard expiration cycle.

B Calculation of Variables

- We use one year ($N = 252$ trading days) as the benchmark horizon for all variables where applicable, but our results also hold up for longer and shorter horizons (21, 42, 63, 126, 504 and 1260 days).
- For $N \neq 252$, we annualize all variance measures by $(252/N)$.
- All accounting variables are lagged by at least five months.

Variable	Definition
X	Moneyness is defined as option strike price (K) divided by current stock price (S_0).
$IVOL$	We measure idiosyncratic volatility as $IVOL_{i,t} = \sqrt{\sum_{j=0}^{N-1} e^2 i, t-j}$, where $e_{j,t}$ is the daily residual from a regression on N past daily returns on the three Fama-French factors (MKT, SMB, HML) (see Ang et al. (2006)).
$TVOL$	We measure total volatility as $TVOL_{i,t} = \sqrt{\sum_{j=0}^{N-1} r_{i,t-j}^2 - \left(\sum_{j=0}^{N-1} r_{i,t-j}\right)^2}$.
$SVOL$	We measure systematic volatility of as $SVOL_{i,t} = \sqrt{TVOL_{i,t}^2 - IVOL_{i,t}^2}$.
$BVSkew$	We calculate this measure following Boyer and Vorkink (2014).
MAX	MAX denotes the average of the 10 highest returns over the past 3 months (as in Bali et al., 2011) .
β	CAPM beta, measured using the past N daily returns.
VRP	The variance risk premium denotes the difference between historic volatility and implied option volatility (HV-IV). We compute this measure as described in Goyal and Saretto (2009).
MOM	Return over the past 12 months excluding the past month.
REV	Return over the past month.
$TURN$	Average number of shares traded the past N trading days.
D/A	Debt to asset ratio.
B/M	Book to market ratio.
PRC	Price of the stock.
$SIZE$	Price of the stock times number of shares outstanding.
$ILLIQ$	Amihud (2002) illiquidity measure.

C Details on the Model

C.1 Model solution

The model introduced in Equ. (1) - (5) reads

$$\ln\left(\frac{S_{t,i}}{S_{t-1,i}}\right) = r + \left(b_i\,\mu - \frac{b_i^2}{2}\right) h_{z,t} - \frac{c_i^2}{2} h_{y,t} + b_i\sqrt{h_{z,t}} z_t + c_i\sqrt{h_{y,t}} y_t, \tag{C.1}$$

$$h_{z,t} = \omega_z + \beta_z h_{z,t-1} + \alpha_z \left(z_{t-1} - \gamma_z\sqrt{h_{z,t-1}}\right)^2, \tag{C.2}$$

$$h_{y,t} = \omega_y + \beta_y h_{y,t-1} + \alpha_y \left(y_{t-1} - \gamma_y\sqrt{h_{y,t-1}}\right)^2, \tag{C.3}$$

$$y_t, z_t \sim N(0,1).$$

The one period SDF reads

$$\frac{M_t}{M_{t-1}} = \left(\frac{S_{t,I}}{S_{t-1,I}}\right)^{\phi_i} \exp\left[\delta_i + \eta_i h_{z,t} + \xi\left(h_{z,t+1} - h_{z,t}\right)\right]. \tag{C.4}$$

The the parameters of the SDF are given as:

$$\begin{aligned}
\delta_i &= -(\phi_i+1)r - \xi\omega_z + \frac{1}{2}\ln(1-2\xi\alpha_z)\\
\eta_i &= -(\mu - \frac{1}{2})\phi - \xi\alpha_z\gamma^2 + (1-\beta_z)\xi - \frac{(\phi_i - 2\xi\alpha_z\gamma_z)^2}{2(1-2\xi\alpha_z)}\\
\phi_i &= -(\mu - \frac{b_i}{2} + \gamma_z)(1-2\alpha_z\xi) + \gamma_z - \frac{b_i}{2}.
\end{aligned} \tag{C.5}$$

To arrive there, we use the two pricing conditions

$$E_{t-1}\left[\frac{M_t}{M_{t-1}}\right] = \exp(-r)$$

$$E_{t-1}\Big[\frac{M_t}{M_{t-1}}\frac{S_{t,i}}{S_{t-1,i}}\Big] = 1.$$

Note that the first condition does not depend on any specific stock i. As a consequence, the structure of the coefficients of δ and η is the same as in Christoffersen et al. (2013). Note however, that in our setting δ and η depend on a the stock-specific ϕ_i and thus become stock-specific themselves.

Plugging the general structure of the SDF in (C.4) and the return dynamics from (C.1)-(C.3) into the second condition, we obtain

$$\begin{aligned}
E\left[\frac{M_t}{M_{t-1}}\frac{S_{t,i}}{S_{t-1,i}}\right] = E\Big[\exp\Big(&(\phi+1)r + \phi(\mu - \frac{1}{2})h_z + (b\mu - \frac{b^2}{2})h_z + (\phi+b)\sqrt{h_z}z + \delta + \eta h_z\\
&+ \xi\omega + \xi(\beta-1)h_z + \xi\alpha(z - \gamma\sqrt{h_z})^2 + c\sqrt{h_y}y - \frac{c^2}{2}h_y\Big)\Big],
\end{aligned} \tag{C.6}$$

where all subscripts are dropped to simplify notation.

Note that we can calculate the expected value of the z and y terms in (C.6) separately, as the the two variables are independent. Expanding the square and collecting terms gives

$$
\begin{aligned}
E\left[\frac{M_t}{M_{t-1}}\frac{S_{t,i}}{S_{t-1,i}}\right] =& E\Big[\exp\Big((\phi+1)r+\delta+\xi\omega \\
&+[(\phi+b)(\mu-\frac{1}{2})+\frac{b}{2}-\frac{b^2}{2}+\eta+\xi(\beta-1)+\xi\alpha\gamma^2]h_z \\
&+[(\phi+b)-2\xi\alpha\gamma]\sqrt{h_z}z+[\xi\alpha]z^2\Big)\Big]+E\Big[\exp\Big(c\sqrt{h_y}y-\frac{c^2}{2}h_y\Big)\Big].
\end{aligned}
\tag{C.7}
$$

To calculate the expected value of the right-hand side of (C.7), we use the following result, which we take from Christoffersen et al. (2013):

$$
E\left[\exp(dz^2+2dez)\right]=\exp\Big(-\frac{1}{2}\ln(1-2d)+\frac{2d^2e^2}{1-2d}\Big). \tag{C.8}
$$

Here we have

$$
d=\xi\alpha
$$

$$
e=\left(\frac{(\phi+b)-2\xi\alpha\gamma}{2\xi\alpha}\right)\sqrt{h_z},
$$

and thus

$$
2d^2e^2=2\xi^2\alpha^2\left(\frac{(\phi+b)-2\xi\alpha\gamma}{2\xi\alpha}\right)^2 h_z=\frac{1}{2}\left((\phi+b)-2\xi\alpha\gamma\right)^2 h_z.
$$

Using (C.6) with above specification for d and e gives

$$
\begin{aligned}
E\left[\exp\Big([(\phi+b)-2\xi\alpha\gamma]\sqrt{h_z}z+\xi\alpha z^2\Big)\right] &= \exp\left(-\frac{1}{2}\ln(1-2\xi\alpha)+\frac{\left((\phi+b)-2\xi\alpha\gamma\right)^2}{2(1-2\xi\alpha)}h_z\right) \\
&= \exp\left(-\frac{1}{2}\ln(1-2\xi\alpha)+\frac{(\phi-2\xi\alpha\gamma)^2+b^2+2b\left(\phi-2\xi\alpha\gamma\right)}{2(1-2\xi\alpha)}h_z\right).
\end{aligned}
$$

Furthermore we have

$$
E\left[\exp\left(cy\sqrt{h_y}-\frac{c^2}{2}h_y\right)\right]=1.
$$

Taken together, we arrive at:

$$
\begin{aligned}
E\left[\frac{M_t}{M_{t-1}}\frac{S_{t,i}}{S_{t-1,i}}\right] = &\exp\Big((\phi+1)r+\delta+\xi\omega \\
&+\left[(\phi+b)(\mu-\frac{1}{2})+\frac{b}{2}-\frac{b^2}{2}+\eta+\xi(\beta-1)+\xi\alpha\gamma^2\right]h_z \\
&-\frac{1}{2}\ln(1-2\xi\alpha)+\frac{(\phi-2\xi\alpha\gamma)^2+b^2+2b\left(\phi-2\xi\alpha\gamma\right)}{2(1-2\xi\alpha)}h_z\Big).
\end{aligned} \tag{C.9}
$$

Setting the right-hand side of (C.9) to one, taking logs and using the expressions for δ and η from (C.5) gives

$$\left[b(\mu-\frac{1}{2})+\frac{b}{2}-\frac{b^2}{2}+\frac{b^2+2b\left(\phi-2\xi\alpha\gamma\right)}{2(1-2\xi\alpha)}\right]h_z=0.$$

Solving for ϕ yields

$$\phi=-(\mu-\frac{b}{2}+\gamma)(1-2\alpha\xi)+\gamma-\frac{b}{2}.$$

C.2 Parameters and Simulation

As model parameters, we use the numbers from Table 4, column "Equity and volatility premia" in Christoffersen et al. (2013). To simulate a cross-section of underlying stocks with different levels of $IVOL$ and $SVOL$, we need to choose values for b and c. For this, we use the distribution of these values in our sample, as displayed in Table F.5. To match the model to the data, we proceed as follows. For each month t and each stock i, we measure the $IVOL$ exposure of a stock as the ratio $c_i = IVOL_{i,t}/TVOL_{market,t}$, where we use the CRSP value-weighted index to represent the market. For our simulation, we choose b and c such that they are close to the 10%, 50% and 90% quantiles of the distribution. We use either no $IVOL$ exposure (corresponding to the market, $c = 0$), or low, medium and high $IVOL$ exposure, with $c = 1, 2, 4$, respectively. Furthermore, we choose $b = 0.5$, $b = 1$, and $b = 2$ to represent low, medium, and high levels of $SVOL$. The risk-free rate is set to $r = 3\%$ p.a., the return horizon is one month (21 out of 252 trading days).

We simulate 10^7 paths of the model and calculate expected option returns for all parameter combinations and for a variety of moneyness values. We use the SDF to price options, since there is no closed-form solution for the option price. This requires to use a high number of paths, as the SDF is very volatile. We start the simulations with $h_{y,0}$ and $h_{z,0}$ equal to their respective long-run means.

We also test a wide set of alternative parameters, which are displayed in Table F.6, and find that the model predictions remain qualitatively the same as for the benchmark parametrization. In particular, we test several alternative parameter estimates from Sichert (2020), who documents that a GARCH model estimated over a period with both high and low volatility clusters can have fundamentally different properties than one estimated over high and low volatility periods separately. Relevant in our context is that an estimation of one fixed model over heterogeneous volatility periods generates biased volatility forecasts, under both the physical and the risk-neutral probability measure. In particular such a model overpredicts volatility in calm periods, and vice versa, thus overpricing and underpricing options, respectively, in the different regimes. The solution is to augment the Heston-Nandi (2001) GARCH model by structural breaks. The estimation of this so-called change-point GARCH model results in alternating high and low volatility regimes, with a typical duration of such a regime of around five years.

D Effects of Volatility on Expected Call Option Returns

D.1 General relationship in the BSM model

Here, we give a more detailed analysis of the impact of different types of volatility on expected option returns. We start from a well-know decomposition of the expected call return $E[R_i^C]$ in the BSM model:

$$E[R_i^C] = r + \Omega_i \left[E[R_i] - r \right], \tag{D.1}$$

where $\Omega_i > 1$ is the elasticity of the call price w.r.t. the stock price. Next, we take the partial derivative with respect to b_i and obtain:

$$\begin{aligned} \frac{\partial E[R_i^C]}{\partial b_i} &= \frac{\partial \Omega_i}{\partial \sigma_i} \frac{\partial \sigma_i}{\partial b_i} \left[E[R_i] - r \right] + \Omega_i \frac{\partial E[R_i]}{\partial b_i} \\ &= \frac{\partial \Omega_i}{\partial \sigma_i} \frac{b_i \sigma_z^2}{\sqrt{b_i^2 \sigma_z^2 + c_i^2 \sigma_y^2}} [b_i(\mu - r) - r] + \Omega_i(\mu - r). \end{aligned} \tag{D.2}$$

Similarly, for the $IVOL$ exposure c_i we obtain

$$\frac{\partial E[R_i^C]}{\partial c_i} = \frac{\partial \Omega_i}{\partial \sigma_i} \frac{\partial \sigma_i}{\partial c_i} \left[E[R_i] - r \right]. \tag{D.3}$$

D.2 Idiosyncratic volatility

In the case of $IVOL$, equation (D.3) gives us a nice insight: $\frac{\partial \Omega_i}{\partial \sigma_i}$ is always negative for all calls with positive strikes, while all other terms are positive, and hence the effect of an increase in $IVOL$ on expected call returns in the BSM model is negative. Put differently, increasing σ_i without changing μ_i, always leads to an decrease in expected call returns. With b_i fixed, this is in our setting equivalent to $c_i \to \infty$. In fact, it is well known that, the limit is the expected return of the underlying, i.e.:

$$\lim_{\sigma_i \to \infty} E[R_i^C] = E[R_i]. \tag{D.4}$$

In case of the model with the U-shaped PK, there are no more analytic solutions. However, the results in Figure 2, as well an extensive simulation analysis, suggest that the intuition concerning the effect of $IVOL$ on expected call returns shown in D.4 also carries over to our model. The fundamental difference is that expected call returns are positive for low, and negative for high strikes. This in turn implies, that with increasing $IVOL$ expected call returns approach the expected call return from above for low strikes, and from below for high strikes.

D.3 Systematic volatility

When increasing b_i in Equations (1) and (7)-(8), respectively, one increases both the expected return of the underlying and its volatility. For the BSM case, consider again Equation (D.2), which highlights that an increase in b_i leads to two opposing effects. Analogous to our discussion of $IVOL$ above, the first summand in Equation (D.2) is always negative, while the second one is always positive. Hence it becomes a quantitative question which effect dominates. In the following, we will discuss this question and we will refer to the first summand as the volatility effect, and the second summand as the drift effect.

The graphs in the upper row of Figure 4 show a rich pattern for the BSM model. The top left graph shows that when the call's moneyness is low, the call behaves more like the underlying asset and the drift effect dominates. When moneyness is high, however, the volatility effect dominates. The remaining graphs in the top row show that this pattern in addition depends on the level of $IVOL$ of the underlying stock. When $IVOL$ is high, the marginal effect of increasing $SVOL$ is low, and therefore the drift effect dominates. As a result, expected call returns are increasing in $SVOL$ across the full moneyness spectrum.

Next, consider the U-shaped PK model. Here the volatility and drift effect described above are again relevant, and in addition we have what we call the variance risk premium (VRP) effect. Ignoring the drift and VRP effect for the moment, increasing b_i is equivalent to increasing $IVOL$, which means that expected call returns are pushed towards the expected return of the underlying (either from above or from below). Through the drift effect, an increase in b_i always leads to an increase in expects call returns. With respect to the VRP effect, note that a higher b_i, all else equal, means a higher exposure of stock returns to the priced variance risk h_z. Due to the negative price of market variance risk, and the positive correlation between stock returns and systematic variance for positive stock returns, this leads to increasingly negative out-of-the-money call returns. Hence for out-of-the-money calls, when returns are negative, the volatility and drift effect increase returns, while the VRP effect decreases returns. Which of of the three effects dominates can be seen in the lower row of Figure 4.

E Additional results

Tables F.1 and F.2 here.

F Replication of Papers

F.1 General remark

When we replicate results from other papers, we generally follow the original approach as closely as possible. Sometimes, minor deviations are necessary due to data restrictions, unclear details of the original methodology, or for the purpose of remaining consistent with our own approach.

F.2 Chaudhuri and Schroder (2015)

F.2.1 Methodological issues

We generally follow the data cleaning and return calculation procedure outlined in the original paper whenever possible. From 2010 on weekly expiration cycles ("weekly options") are available for options on single stocks. To be consistent with the original methodology and the first part of the sample, we also only consider options with a time to maturity of three to six weeks expiring on the standard expiration cycle (3rd Friday of each month). Chaudhuri and Schroder (2015) use data until 2012, but do not explicitly state how they treat returns on those weekly options in the period for the period from 2010 to 2012.

Chaudhuri and Schroder (2015) calculate returns from Tuesday to Tuesday. Since they do not specify their procedure in case Tuesday is not a trading day, we assume the following. If Tuesday is not a trading day, we use Wednesday data instead, either for the start date, or the end date, or both. If for either start or end date neither Tuesday nor Wednesday data is available, we discard the week.

Chaudhuri and Schroder (2015) state that they use options with a maturity between 3 and 6 weeks. As they use Tuesday data, and expiration is on Fridays (or formally Saturdays in the first year of the sample), we understand this as taking options with about 3.5, 4.5 and 5.5 weeks to maturity.

Chaudhuri and Schroder (2015) state that they report and study "average-return differences average differences (across both stocks and time)" (p.1472). There is no explicit statement such as each date portfolios are formed, and then the time-series average is calculated. Therefore it seems that the author's report and study the simple average of all pooled returns. This also roughly matches the values we obtain if we follow that approach. Furthermore, this calculations leads to much larger t-statistics, which are in a very similar magnitude as those reported in their paper. In contrast, we follow the standard asset pricing approach to form portfolios at each date and report time-series averages of portfolio returns. This appears to be important also in the context of option returns, since they do not represent a balanced sample, and the number of options available appears to be correlated with economic conditions (e.g., it is lower in times of financial turmoil) and hence returns. Therefore it is not surprising that the average of all pooled returns sometimes deviates significantly from the time-series average.

Instead of using a strict at-the-money requirement, the original papers discards options (and hence underlying stocks), if the call with a strike closest to the current stock price is more than one strike increment away from the current stock price. In particular, they "remove those stock-week combinations for which the absolute difference between the current stock price and nearest strike is \$2.5 or higher when current stock prices are \$25 or less, \$5 or higher when the stock price is \$25 through \$200, and \$10 or higher when the stock price is over \$200. This does not consistent a tight at-the-money bin as the typical [0.95 1.05] or [0.90 1.10] moneyness requirement, but the moneyness can actually be very far away from one, especially for stocks with small prices. This explains why in the tables the moneyness range of the first bin is rather large.

We could not find a clear statement in Chaudhuri and Schroder (2015) whether an underlying stock at one point in time must be present in all bins or not. To be strict and in the save side, we always require that an underlying stock must have a call in all used bins in the ordinal sorts below. Furthermore, we deviate from Chaudhuri and Schroder (2015) in such that we do not study the returns of in-the-money call options. Since this moneyness region is not of our central interest, this weaker requirement allows us to include more options into the analysis and reduce the presented numbers. The results are fully robust to also include the in-the-money options (as also the cardinal sorts strongly suggest).

Finally, the original paper does not specify how options are treated that have a start price, but not an end price. While one could treat them in several possible ways, we decide to exclude them. As Chaudhuri and Schroder (2015) do not specify any procedure, which suggests the same approach. We would like to point out that this can lead to a potential bias, as missing prices usually are, among other reasons, associated with a large move of the underlying's stock price.

F.2.2 Results for the 1996−2012 sample

Figure F.4 shows weekly call returns by moneyness for the 1996-2012 sample, i.e., the analogue to Figure 7 for the sample used in Chaudhuri and Schroder (2015). The patterns strongly mimic those for the full sample.

F.3 Boyer and Vorkink (2014)

For the results presented in the main text we follow our benchmark filter approach, whit the exception that we now do not restrict the options' moneyness. The results are virtually the same if one applies the slightly different filters used by Boyer and Vorkink (2014). In addition, Boyer and Vorkink (2014) form their option portfolios either a week before the options' expiration, at the beginning of the calendar month or at the beginning of the calendar month preceding the expiration, and then hold the options until maturity. Hence our monthly returns are somewhere between their middle and largest maturity bucket. Last, Boyer and Vorkink (2014) do allow that multiple options written on the same stock are included in the same portfolio. We also perform the analysis after including only one option per stock in each portfolio and find no substantial difference in the results.

F.4 Byun and Kim (2016)

F.4.1 Methodological issues

As in the original paper, at-the-money options are defined as moneyness (K/S_0) between 0.9 and 1.1, and for each stock we use only the call that has a moneyness closest to 1. We add the out-of-the-money[26] test, and define out-of-the-money call options here as those with moneyness between 1.1 and 1.3, and for each stock and date use the call that has a moneyness closest to 1.2.

F.4.2 Fama-MacBeth Regressions

As in Byun and Kim (2016), in each month, we estimate the following cross-sectional regression:

$$r_t^i = \gamma_{0,t} + \gamma_{1,t}\ V_{t-1}^i + \gamma_{2,t}' Z_{t-1}^i + \epsilon_t^i,$$

where r_t^i is the observed call return for stock i at date t, V is MAX, and Z is a vector of controls. The controls include those from the original paper (except for IOR, due to data limitations) plus option moneyness. The definition of variables is as in Byun and Kim (2016). Table F.3 reports the time-series averages of the γ coefficients, together with Newey and West (1987) t-statistics for at-the-money calls, and Table F.4 for out-of-the-money calls.

Finally, all the results are very similar when we restrict the sample to the period 1996-2013 used in Byun and Kim (2016).

[26]The robustness section of the original paper contains also an test with out-of-the-money options. However, the respective moneyness range, described as "calls with moneyness higher than one but closest to 1.1" (p. 168) is not clearly defined, as it does not specify an upper limit. The results in Figure 6 suggest that the turning point in the MAX-average call returns relationship is at a moneyness level around 1.1. We therefore use our different out-of-the money bin.

Table F.1: Fama-MacBeth Regressions of In-The-Money Call Returns with $IVOL$ and $SVOL$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$IVOL$	-0.107^{***}	-0.099^{**}	-0.100^{***}			-0.113^{***}		-0.129^{***}	
	(-2.82)	(-2.57)	(-2.60)			(-3.32)		(-3.80)	
$\ln(IVOL)$				-0.036^{*}			-0.043^{**}		-0.046^{***}
				(-1.67)			(-2.48)		(-2.82)
$SVOL$					-0.079	0.071		0.047	
					(-0.85)	(0.79)		(0.52)	
$\ln(SVOL)$							0.010		0.005
							(1.13)		(0.56)
X		0.276^{***}	-2.610	-2.418	-1.686	-3.170	-3.162	-5.620^{**}	-5.366^{**}
		(2.78)	(-0.85)	(-0.78)	(-0.65)	(-1.18)	(-1.15)	(-2.34)	(-2.18)
X^2			1.637	1.523	1.150	1.960	1.952	3.312^{**}	3.168^{**}
			(0.93)	(0.86)	(0.77)	(1.27)	(1.23)	(2.39)	(2.23)
MOM								0.006	0.005
								(0.31)	(0.30)
REV								-0.072^{**}	-0.071^{**}
								(-2.39)	(-2.37)
VRP								0.066^{***}	0.061^{***}
								(4.62)	(4.24)
$ILLIQ$								-3.095	-3.752
								(-0.65)	(-0.80)
$R^2(\%)$	1.9	2.5	2.6	2.9	3.5	4.5	4.6	6.6	6.6

This table shows the results for Fama-MacBeth regressions of one month hold-until-maturity call returns on stock and option characteristics plus additional controls. For each month, we estimate the cross-sectional regression

$$R^C_{t,i} = \gamma_{0,t} + \gamma_{1,t}\ IVOL_{t-1,i} + \gamma_{2,t}\ SVOL_{t-1,i} + \gamma'_{3,t} Z_{t-1,i} + \epsilon_{t,i},$$

where $R^C_{t,i}$ is the return for call option i at time t, and Z is a vector of controls. The definitions for the different variables are provided in Appendix B. For all in-the-money calls with a moneyness $X = K/S_0$ between 0.80 and 0.95, for each underlying stock, we keep only the call which has an X closest to 0.875. The table reports the time-series averages of the γ-coefficients and the associated Newey and West (1987) adjusted t-statistics (with five lags) in parentheses. The last row gives the averages of the adjusted-R^2 from the regressions. ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively. The total number of observations is $267,983$ in all regression specifications.

Table F.2: Fama-MacBeth Regressions of At-The-Money Call Returns with $IVOL$ and $SVOL$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$IVOL$	-0.205^{***}	-0.203^{***}	-0.198^{***}			-0.189^{***}		-0.211^{***}	
	(-2.84)	(-2.82)	(-2.75)			(-2.72)		(-3.09)	
$\ln(IVOL)$				-0.069^{*}			-0.066^{*}		-0.071^{**}
				(-1.90)			(-1.96)		(-2.19)
$SVOL$					-0.266	-0.003		-0.049	
					(-1.63)	(-0.02)		(-0.32)	
$\ln(SVOL)$							0.000		-0.007
							(0.02)		(-0.49)
X		0.470^{*}	28.889^{**}	25.681^{**}	31.396^{***}	25.529^{**}	22.992^{**}	19.803^{*}	17.989
		(1.81)	(2.40)	(2.13)	(2.70)	(2.20)	(1.97)	(1.75)	(1.59)
X^2			-14.208^{**}	-12.608^{**}	-15.450^{***}	-12.523^{**}	-11.264^{*}	-9.681^{*}	-8.783
			(-2.35)	(-2.09)	(-2.66)	(-2.15)	(-1.93)	(-1.71)	(-1.55)
MOM								0.011	0.010
								(0.37)	(0.34)
REV								-0.182^{***}	-0.182^{***}
								(-2.95)	(-2.91)
VRP								0.128^{***}	0.118^{***}
								(4.30)	(4.09)
$ILLIQ$								2.161	1.635
								(0.32)	(0.25)
$R^2(\%)$	1.3	1.6	1.8	2.1	2.1	2.9	3.2	4.3	4.6

This table shows the results for Fama-MacBeth regressions of one month hold-until-maturity call returns on stock and option characteristics plus additional controls. For each month, we estimate the cross-sectional regression

$$R^C_{t,i} = \gamma_{0,t} + \gamma_{1,t}\ IVOL_{t-1,i} + \gamma_{2,t}\ SVOL_{t-1,i} + \gamma'_{3,t} Z_{t-1,i} + \epsilon_{t,i},$$

where $R^C_{t,i}$ is the return for call option i at time t, and Z is a vector of controls. The definitions for the different variables are provided in Appendix B. For all at-the-money calls with a moneyness $X = K/S_0$ between 0.95 and 1.05, for each underlying stock, we keep only the call which has an X closest to 1.00. The table reports the time-series averages of the γ-coefficients and the associated Newey and West (1987) adjusted t-statistics (with five lags) in parentheses. The last row gives the averages of the adjusted-R^2 from the regressions. ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively. The total number of observations is $281,210$ in all regression specifications.

Table F.3: Fama-MacBeth Regressions of At-The-Money Call Returns with MAX

	(1)	(2)	(3)
MAX	-2.257^{***}	-0.803	-0.914
	(-3.55)	(-0.89)	(-1.10)
$IVOL$		-0.151^{*}	-0.160^{**}
		(-1.95)	(-2.18)
$BVSkew$			0.014
			0.014
$ln(SIZE)$			0.001
			(0.14)
$\ln(PRC)$			-0.000
			(-0.01)
$\ln(B/M)$			0.011
			(1.38)
D/A			0.044
			(1.27)
$TURN$			0.001
			(1.47)
REV			-0.162^{***}
			(-3.16)
MOM			0.022
			(1.11)
X			0.103
			(0.48)
$R^2(\%)$	1.2	1.6	5.0

This table replicates the regression results from Table 6 in Byun and Kim (2016).

Table F.4: Fama-MacBeth Regressions of Out-Of-The-Money Call Returns with MAX

	(1)	(2)	(3)
MAX	-1.006	2.785^{*}	2.445
	(-1.27)	(1.72)	(1.63)
$IVOL$		-0.415^{***}	-0.370^{***}
		(-2.74)	(-3.15)
$BVSkew$			0.001
			(0.19)
$\ln(SIZE)$			-0.054^{***}
			(-3.60)
$\ln(PRC)$			0.079^{**}
			(2.24)
$\ln(B/M)$			0.022
			(1.10)
D/A			-0.060
			(-0.73)
$TURN$			0.005^{***}
			(3.46)
REV			-0.265^{**}
			(-2.44)
MOM			0.009
			(0.23)
X			-2.244^{***}
			(-7.00)
$R^2(\%)$	0.4	0.7	3.7

This table is an analogue to the regression results shown in Table 6 of Byun and Kim (2016), but using out-of-the-money calls instead of at-the-money calls.

Table F.5: Distribution of CAPM-β and $IVOL$ Exposure in the Stock Return Sample

	Mean	Min	1%	5%	10%	25%	50%	75%	90%	95%	99%	Max
CAPM-β	1.2	-14.6	0.2	0.4	0.6	0.8	1.1	1.5	1.9	2.2	2.9	9.4
$IVOL/TVOL_{Market}$	2.7	0.0	0.6	0.8	1.0	1.5	2.2	3.3	4.8	5.9	9.0	142.0

Table F.6: Alternative model parameters

Panel A: Our model						
ω	α	β	γ	μ	ξ	Source
0	8.89E-07	0.756	515.6	1.59	117,439	Christoffersen et al. (2013), period 1996 - Oct 2009; benchmark case
0	8.89E-07	0.756	515.6	2.71	117,439	Christoffersen et al. (2013), μ increased to match the index return of the longer sample
0	1.41E-06	0.755	409.6	1.59	0	Christoffersen et al. (2013), restricted estimation with $\xi = 0$
1.90E-14	1.56E-06	0.713	417.1	4.33	46,248	Sichert (2020), period 1996 - Aug 2015
8.97E-08	2.89E-07	0.699	999.0	17.02	499,023	Sichert (2020), low vol period Jan 1996 - Oct 1996
2.20E-10	2.47E-06	0.826	238.6	5.90	75,168	Sichert (2020), high vol period Oct 1996 - Aug 2003
6.87E-07	1.78E-06	0.733	341.7	0.11	80,716	Sichert (2020), low vol period Aug 2003 - Jun 2007
1.82E-11	1.42E-06	0.696	441.0	10.47	128,518	Sichert (2020), high vol period Jun 2007 - Nov 2011
2.08E-12	4.57E-06	0.820	174.4	7.83	18,541	Sichert (2020), low vol period Nov 2011 - Aug 2015

Panel B: Black-Scholes-Merton model			
σ_z	μ	r	
18.86% p.a.	8.0% p.a.	1.51% p.a.	Estimated from daily S&P500 returns, and daily t-bill rate from 1996-2018, respectively
16% p.a.	10% p.a.	3% p.a.	Hu and Jacobs (2020)

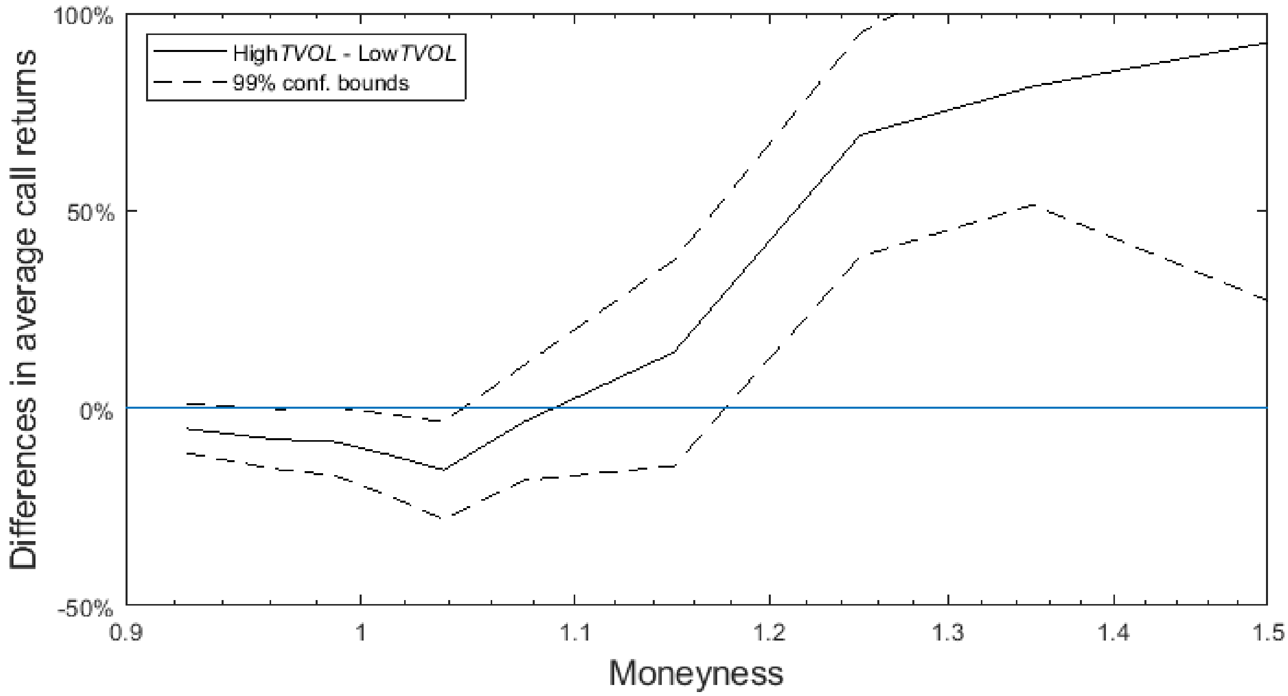

Figure F.1: Call Option Returns for High Minus Low $TVOL$ Sorted Portfolios

We perform an independent double sort on moneyness and $TVOL$. The figure shows the average return difference between calls on stocks in the high and the low $TVOL$ portfolio across moneyness bins. Moneyness is defined as K/S_0. The dashed lines represent the bounds of the 99% confidence interval generated via a bootstrap with 100,000 pairwise draws. The x-axis are log-scaled.

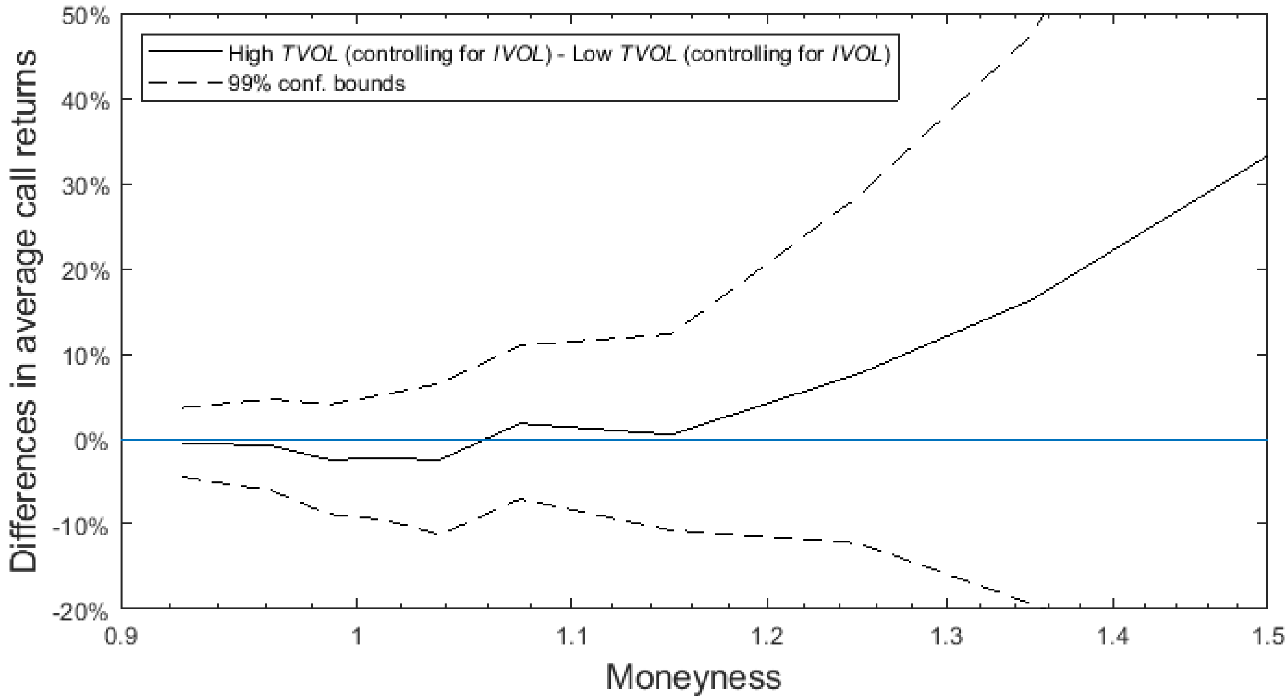

Figure F.2: Call Option Returns for High Minus Low $TVOL$ Sorted Portfolios, Controlling For $IVOL$

The figure shows the difference between the high and the low $TVOL$ call portfolio, after controlling for $IVOL$, across moneyness bins (dependent double sort). Moneyness is defined as K/S_0. The dashed lines represent the bounds of the 99% confidence interval generated via a bootstrap with 100,000 pairwise draws. The x-axis are log-scaled.

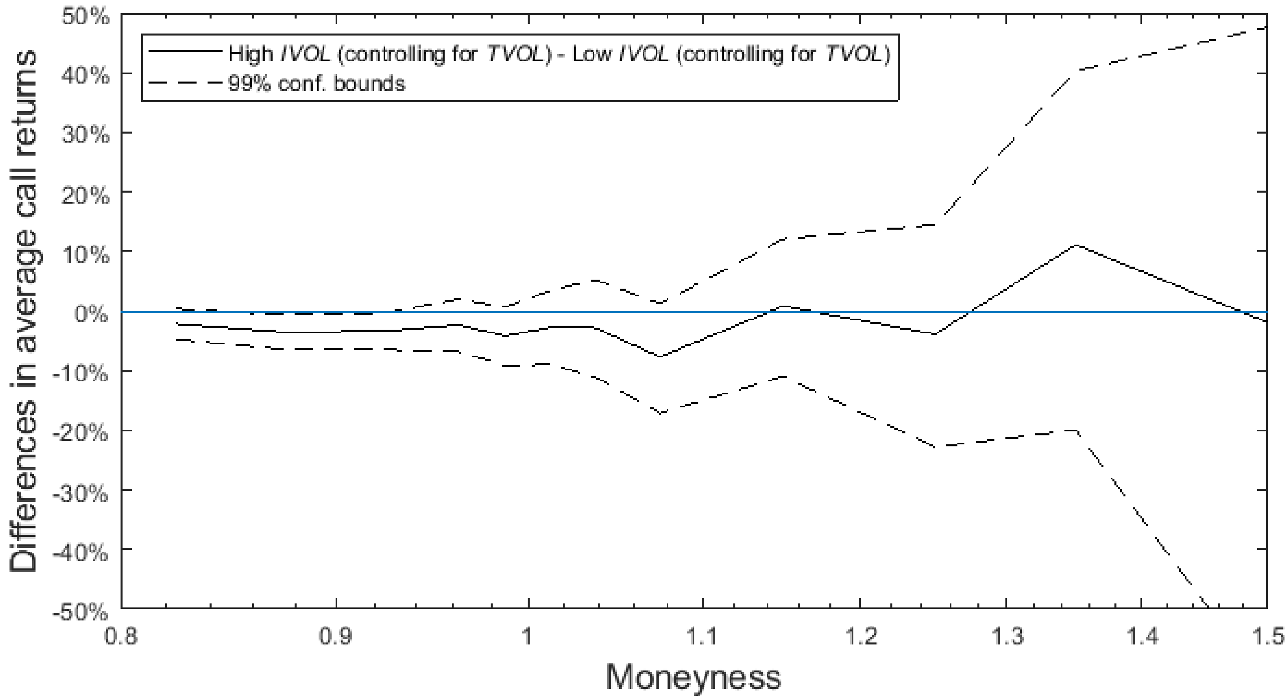

Figure F.3: Call Option Returns for High Minus Low $IVOL$ Sorted Portfolios, Controlling For $TVOL$

The figure shows the difference of the high minus low $IVOL$ call portfolio, after controlling for $TVOL$ for each moneyness bin (dependent double sort). Moneyness is defined as K/S_0. The dashed lines represent the bounds of the 99% confidence interval generated via a bootstrap with 100,000 pairwise draws. The x-axis are log-scaled.

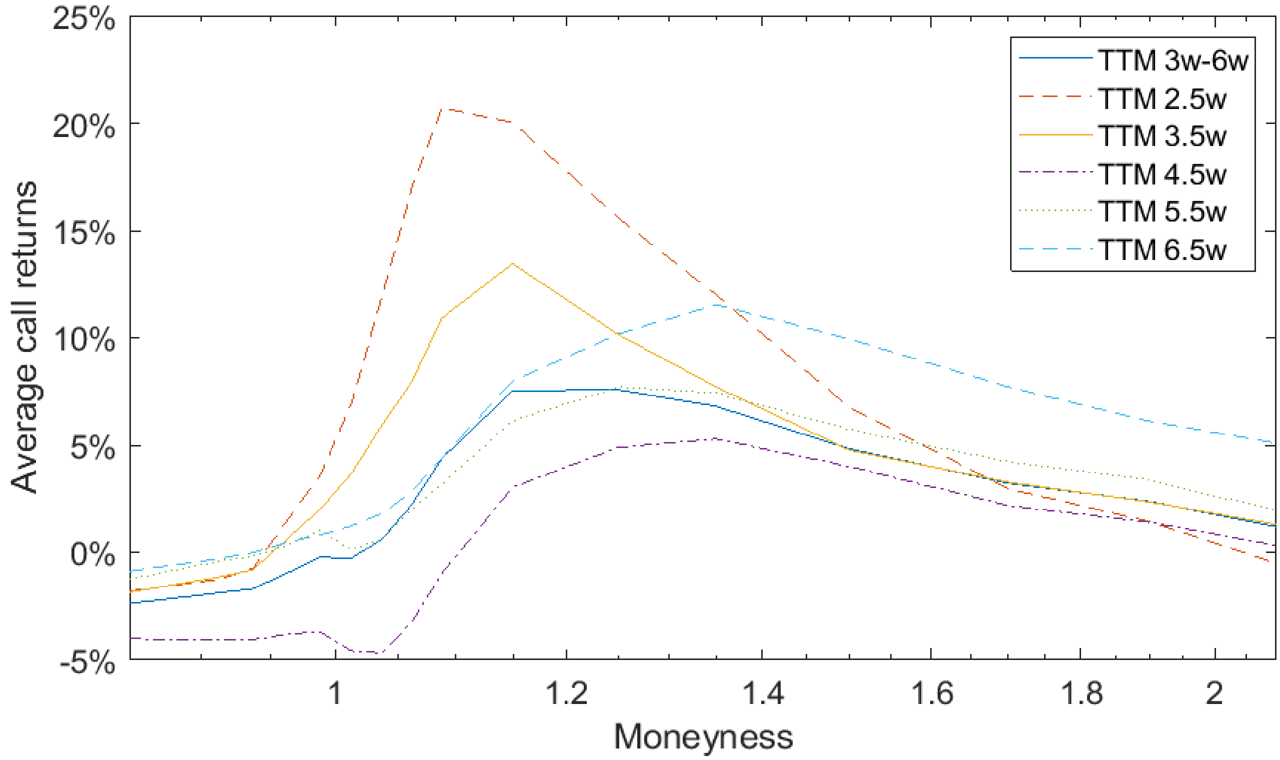

Figure F.4: Intermediate Call Returns by Moneyness and Time to Maturity (1996-2012)

The figure is the analogue to Figure 7, where the data is from 1996-2012, as in Chaudhuri and Schroder (2015). The figure plots the average intermediate call returns as a function of moneyness for different option maturities. The holding period for the intermediate returns in one week. The '3w-6w' line is the average of returns for options with 3.5, 4.5, and 5.5 weeks to maturity, as in Chaudhuri and Schroder (2015). Moneyness is defined as K/S_0. Points on the moneyness axis represent the bin midpoint, and the axis are log-scaled.