Institut canadien des dérivés
Canadian Derivatives Institute

L'Institut bénéficie du soutien financier de l'Autorité des marchés financiers ainsi que du ministère des Finances du Québec

Document de recherche

DR 20-03

Making Better Use of Option Prices to Predict Stock Returns

Publié Avril 2020

Ce document de recherche a été rédigée par :

Dmitriy Muravyev, University of Michigan
Aurelio Vasquez, ITAM
Wenzhi Wang, Boston College

L'Institut canadien des dérivés n'assume aucune responsabilité liée aux propos tenus et aux opinions exprimées dans ses publications, qui n'engagent que leurs auteurs. De plus, l'Institut ne peut, en aucun cas être tenu responsable des conséquences dommageables ou financières de toute exploitation de l'information diffusée dans ses publications.

Making Better Use of Option Prices to Predict Stock Returns*

Dmitriy Muravyev[a], Aurelio Vasquez[b], and Wenzhi Wang[c]

April 29, 2020

Abstract

A typical stock has hundreds of listed options. To study their information content, researchers usually aggregate option prices in ad hoc ways such as volatility skew. We introduce a more general approach to reduce option prices' dimensionality while preserving their rich information. We apply principal component analysis (PCA) to implied volatility surfaces. The first five principal components explain most of IV surface deviations from its average across stocks. The first three components predict future stock returns in-sample and out-of-sample. Moreover, they drive out popular option-based return predictors in joint regressions. The first principal component predicts returns due to its strong correlation with shorting costs. The second and third components' predictability is due to market inefficiencies that disappeared post-2002. Our results shed light on why option prices predict stock returns.

JEL Classification: G12, G13, G14

Keywords: Option prices, stock returns, PCA

---

* We would like to thank David H. Solomon, Ricardo Lopez Aliouchkin and seminar participants at ITAM and Boston College for helpful comments. We thank the Canadian Derivatives Insitute and Asociaci6n Mexicana de Cultura A.C. for financial support. We are responsible for any remaining errors. We thank Barry Schaudt and Boston College Research Services for computing support.

[a] Department of Finance, Michigan State University, Eppley Center, 667 North Shaw Lane, East Lansing, MI 48824. Tel.: +1 (517) 353-5793. E-mail: muravyev@msu.edu. Corresponding author.

[b] Department of Finance, ITAM, Río Hondo No. 1, Col. Progreso Tizapán, México, D.F., México, 01080. Tel.: + 52 (55) 5628 4000x6518 E-mail: aurelio.vasquez@itam.mx

[c] Carroll School of Management, Boston College

Making Better Use of Option Prices to Predict Stock Returns

April 29, 2020

Abstract

A typical stock has hundreds of listed options. To study their information content, researchers usually aggregate option prices in ad hoc ways such as volatility skew. We introduce a more general approach to reduce option prices' dimensionality while preserving their rich information. We apply principal component analysis (PCA) to implied volatility surfaces. The first five principal components explain most of IV surface deviations from its average across stocks. The first three components predict future stock returns in-sample and out-of-sample. Moreover, they drive out popular option-based return predictors in joint regressions. The first principal component predicts returns due to its strong correlation with shorting costs. The second and third components' predictability is due to market inefficiencies that disappeared post-2002. Our results shed light on why option prices predict stock returns.

JEL Classification: G12, G13, G14

1 Introduction

A typical stock has hundreds of listed options of different type (calls and puts), moneyness, and time-to-expiration.[1] More than 800,000 option contracts are currently listed on more than 4,000 U.S. stocks making it the largest market by the number of listed securities.[2] Thus, option prices offer rich multi-dimensional information; however, the curse of dimensionality makes it difficult to study. For example, it is not feasible to use the entire implied volatility surface to predict returns or other variables of interest. To overcome this problem, researchers typically aggregate option prices into a single ratio such as the volatility skew or the call-put spread.[3] These ratios are intuitive and proved useful for forecasting stock returns, volatility, and higher moments. However, such ad hoc aggregation can miss important variation in option prices. Is there a better way to reduce option price dimensionality while preserving their rich information? Well-known statistical methods have been designed specifically for this purpose. Indeed, we show that such parsimonious ways to aggregate option prices preserve more information than the conventional ad-hoc counterparts both in terms of explaining surface variation and predicting stock returns. Most importantly, this feature allows our approach to shed light on why option prices predict future stock returns. Option prices predict stock returns because they reflect shorting costs. We also find some evidence of option prices reflecting risk factors that require premium in the cross-section of stock returns. In the first half of the sample, option prices also provided pure alpha signals about stock mispricing, but those signals stopped working after 2002, which coincides with the time

[1] E.g. Apple had 3,902 listed options while CBOE Group had 268 options listed on its stock in January 2014.
[2] According to Trade Alert, LLC estimates.
[3] Christoffersen, Jacobs, and Chang (2013) provide a comprehensive survey of option-implied information for forecasting. E.g., the implied volatility spread of Cremers and Weinbaum (2010) is an open-interest weighted average between implied volatilities for put and calls on a given day.

when many popular equity anomalies lost their profitability. Overall, our results provide important insights about economics of the options market.

In this paper, we introduce a parsimonious approach that summarizes in few numbers implied volatility and thus option prices on a given stock-day using conventional dimensionality reduction techniques. Let us briefly describe it. Following the convention, option prices are first converted into an interpolated implied volatility surface. Implied volatilities carry the same information but are more convenient than option prices as implied volatilities can be easily compared across stocks and are less sensitive to changes in the underlying price. We then perform a principal component analysis (PCA) on the cross-section of interpolated implied volatility surfaces to extract their factor structure. We de-mean a volatility surface for each stock and day before applying PCA to make it comparable across stocks and to focus on the IV surface shape rather than volatility level.[4] That is, the PCA input matrix has stock-by-day rows and interpolated volatility surface columns. The principal components quantify how much a given volatility surface deviates from the average. We pick PCA because it is simple and popular. We are among the first to apply PCA to the options market.

How much does the shape of volatility surface vary across stocks? The surface cross-section reveals a strong factor structure. The first five principal components explain 78% of the variation across de-meaned volatility surfaces, with the first and second components explaining 32% and 18% respectively. Thus, PCA describes the input surface of 112 interpolated volatilities with just handful of numbers, which greatly reduces the dimensionality. The first component has

[4] If we do not de-mean the implied volatility surfaces, the first component would be the level of the surfaces. In that case we could drop the first component and focus on the remaining principal components to study the shape of volatility surfaces. This is analogous to performing PCA to stock returns and finding that the first component is the market return.

a 66% correlation with the put-call spread of Cremers and Weinbaum (2010) and is weakly negatively correlated with the skew. The second component is negatively correlated with the time-slope, the difference between long-term and short-term implied volatility as in Vasquez (2017). Overall, the main principal components aggregate option information differently from level, slope, and curvature or from conventional ratios such as the call-put spread, skew, and time slope. Thus, these principal components do not simply linearly combine conventional measures, which can explain their superior information content that we document next.

After analyzing the surface factor structure, we explore its information content. Specifically, we study whether this condensed option information, reflected by the main principal components, predicts future stock and option returns. The first three components strongly predict stock returns next week and month. To facilitate a fair comparison of our approach against other option-based predictors, we aggregate the first three components into a single predictor, which we call the surface deviation, by simply adding their loadings ($PC1 + PC2 + PC3$). As the name suggests, this variable reflects an average deviation of a current implied volatility (IV) surface from the overall sample average. This surface deviation robustly predicts future stock returns in-sample and out-of-sample. Its Fama-French-Carhart monthly alphas are 1.32% for equal-weighted and 1.25% for value-weighted portfolios with *t*-statistics of 11.7 and 8.7, respectively. The predictability remains significant up to four months after portfolio formation. The results change little out-of-sample, when principal components are computed in a rolling window using only volatility surface data available prior to the returns prediction date.

To further test the information content of the IV surface deviation, we benchmark it against other option-based predictors including the call-put implied volatility spread (Bali and Hovakimian (2009), Cremers and Weinbaum (2010), and Yan (2011)) and measures or risk-neutral skew,

skewness, and kurtosis (Xing, Zhang, and Zhao (2010), Rehman and Vilkov (2012), Conrad, Dittmar, and Ghysels (2013), Stilger, Kostakis, and Poon (2016), and Bali, Hu, and Murray (2016)). Univariate regressions confirm that these variables are by far the main return predictors that rely only on current option prices. However, the surface deviation makes each of these predictors insignificant in bivariate regressions. For example, the *t*-statistics for the call-put volatility spread decreases and changes sign from 6 to -1 after including the surface deviation, which remains highly significant. Thus, the surface deviation subsumes the forward-looking information from conventional ad-hoc ratios and thus confirms that PCA approach preserves information content of the IV surface.

Why do option prices predict stock returns? The above results help us answer this question by balancing two extremes: on the one hand, using the entire IV surface to predict returns is not feasible; on the other hand, conventional ad-hoc ratios miss most of surface information. Thus, explaining why a particular ad-hoc ratio predicts returns cannot be generalized to the entire surface.[5] We address this challenge by first summarizing IV surface information with PCA and then explaining why the PCA-based surface deviation predicts returns. The surface deviation effectively represents entire surface's information: it (i) explains most of surface variation across stocks, (ii) reflects main principal components that predict returns while non-main components do not, and (iii) subsumes predictive power of popular option-based predictors. Thus, we can benchmark the surface deviation against proxies for alternative explanations to uncover the economic driver of the return predictability. We test three main explanations: (i) informed trading (Cremers and Weinbaum (2010), Roll, Schwartz, and Subrahmanyam (2010), and Xing, Zhang,

[5] For example, Muravyev, Pearson, and Pollet (2020) argue that the volatility spread and skew predict returns only through their ability to reflect shorting costs.

and Zhao (2010)), (ii) jump risk (Bali and Hovakimian (2009), and Yan (2011)), and (iii) short sale constraints (Johnson and So (2012), and Ofek, Richardson, and Whitelaw (2004)). We find that while proxies for jump risk and informed trading do not affect surface deviation's predictive ability, the short-sale constraints do. Individually, surface deviation and stock lending fee strongly predict returns. But, when we put them together, one remains significant and the other does not depending on specification. However, the first principal component remains a significant return predictor even after excluding hard-to-short stocks suggesting that short sale is important but not the only reason this PC1 predicts returns. We find some evidence that the remaining predictability can be due to compensation for taking risk. Specifically, surface deviation remains significant in the sample of S&P 500 index constituents, which are characterized by negligible trading and shorting costs and are rarely mispriced. Thus, risk premium is a prevailing explanation in this subsample. Finally, the second and third principal components (PC2 and PC3) only predict returns until early 2002. Remarkably, PC3 made most of its profits during the internet bubble burst. Thus, these two components' predictability is due to market inefficiencies that disappeared post-2002.

Lastly, we also show that the surface deviation predicts cross-section of option returns. It negatively and monotonically predicts future weekly straddle returns in portfolio sorts and is by far the strongest predictor compared to popular option return predictors including the volatility time slope of Vasquez (2017), historical minus implied volatility of Goyal and Saretto (2009), and idiosyncratic volatility as in Cao and Han (2013). This exercise confirms that the PCA approach preserves most of relevant IV surface information for predicting option returns.

We contribute to the literature in several ways. First, principal component analysis and similar methods have been widely used in finance to understand the term structure of interest rates (Litterman and Scheinkman (1991)), the term structure of credit and CDS spreads (Collin-

Dufresne, Goldstein and Martin (2001) and Pan and Singleton (2008)), and more recently, the equity volatility levels, skews, and term structures for S&P500 index options (Christoffersen, Fournier, and Jacobs (2017) and Israelov and Kelly (2017)). We are the first to apply PCA to options on a large stock panel and provide a new application of PCA that summarizes information contained by the implied volatility surface by reducing it dimensionality.[6] More importantly, our results shed light on why option prices predict stock returns. During early years, option prices could be used to identify mispricing in the equity market. However, the inefficiencies disappeared post 2002. On the other hand, the first principal component consistently predicts returns during the entire sample period. Its predictive ability can be attributed to option prices reflecting shorting costs and to some extend to option prices capturing the true risk factors. While mispricing, short-sale constraints, and risk-premium were suggested as possible explanations for why certain option ratios predict stock returns. The difference is that our PCA approach let us analyze why the entire surface rather than particular ratios predict returns, and this analysis reveals the dynamic and multifaceted nature of option-based return predictability.

The paper is organized as follows. Section 2 describes the data and the methodology. Section 3 explores how the principal components predict stock returns, compares them to other option-based predictors, and identifies the channel for return predictability. Section 4 provides several robustness tests and studies option return predictability. Section 5 concludes.

[6]Christoffersen, Fournier, and Jacobs (2017) is the only other paper we are aware of that applies PCA to equity options. They use equity options for 30 firms in the Dow Jones while we use all optionable stocks, and their approach differs from ours in its methodology as well as the purpose.

2 Data and Methodology

In this section, we first describe the data and then explain how we perform the principal component analysis on volatility surfaces and aggregate principal components into the surface deviation. In short, we apply a classic PCA to standard option and stock data.

2.1 Data

We use the cross-section of volatility surfaces from the OptionMetrics database which provides end-of-day option prices and Greeks on all exchange-listed option on U.S. equities from 1996 to 2014. OptionMetrics provides the interpolated volatility surfaces obtained with kernel smoothing for each stock and day, which is widely used in the options literature. OptionMetrics computes implied volatility using "the industry-standard Cox-Ross-Rubinstein (CRR) binomial tree model". The surface contains implied volatility for calls and puts across standardized maturities and deltas. The standardized expirations are 30, 60, 91, 122, 152, 182, 273, and 365 calendar days at absolute deltas 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, and 0.8. Thus, 112 standardized implied volatilities across different option types, maturities, and moneyness levels are reported. Recently, longer maturities of 547 and 730 calendar days became available, we exclude them to make our results comparable across the entire sample period. The data are also available for absolute deltas of 0.25, 0.35, 0.45, 0.55, 0.65, and 0.75, which we drop for computational feasibility as PCA on large matrices is computationally intensive. These intermediate deltas unlikely to provide incremental information to the main deltas that we use.

Options differ in their parameters (time-to-maturity and moneyness) across stocks, thus the interpolated volatility surface is crucial to facilitate comparison of option prices across stocks. The OptionMetrics volatility surface makes our results easy to replicate and compare to the existing literature that overwhelmingly uses it. We eliminate close-end funds, real estate investment trusts,

American depository receipts, and stocks with price below one dollar. To avoid stocks with highly illiquid options, we follow Johnson and So (2012) and only use stocks that traded at least 50 calls and 50 put contracts on a given day.

We compute popular option-based stock return predictors in a standard way. The call-put volatility spread is the open-interest weighted difference between call and put implied volatilities with the same strike and expiration as in Cremers and Weinbaum (2010). For robustness, we also use the call-put spread in Bali and Hovakimian (2009) and Yan (2011), who compute it as the difference between one-month implied volatilities of a put and a call with absolute delta of 0.5, i.e. at-the-money. Risk-neutral volatility, skewness, and kurtosis are calculated as in Bakshi, Kapadia, and Madan (2003) for the last trading day before the testing period as used in Rehman and Vilkov (2012) and Stilger, Kostakis, and Poon (2016). Option skewness is the difference between OTM put and ATM call implied volatilities as computed by Xing, Zhang, and Zhao (2010). The time slope of the implied volatility term structure is defined as the difference between long and short-term ATM implied volatilities as in Vasquez (2017). We also control for firm characteristics that we extract from Center for Research and Security Prices (CRSP) and Compustat. The final sample requires data from OptionMetrics, CRSP, and Compustat. From CRSP, we extract the stock market capitalization (size) and use daily returns to calculate Tuesday-to-Tuesday weekly returns. From Compustat, we take book values to calculate book-to-market ratios of individual firms. We skip one day between the return predictors and portfolio formation (stock returns) to avoid micro-structure effects.

2.2 Methodology

In theory, we can predict stock returns using the entire volatility surface, which consists of 112 implied volatilities. However, such approach fails because of the overfitting in-sample and

poor performance out-of-sample. Also, most of these implied volatilities are highly correlated causing the multicollinearity problem. Thus, the surface dimensionality must be reduced first to condense 112 implied volatilities to just few numbers. Essentially, we want to describe how volatility surface for one stock, say Apple or Google, differs from the average IV surface across all stocks. Luckily, there are well-known statistical methods that solve this problem.

We apply the principal component analysis (PCA), by far the most popular method for dimensionality reduction, to a large matrix with rows corresponding to stocks by week and columns corresponding to 112 interpolated implied volatilities. In our sample, this input matrix includes de-meaned implied volatility for 1,897,536 firm-week rows (936 weeks and 2,027 unique firms in a typical week) and 112 columns corresponding to standardized option parameters. Cross-stock comparison is possible because OptionMerics interpolates implied volatility surface for standardized values by option type, days to expiration, and option delta (moneyness) that are the same across stocks. The goal is to reduce IV surface to a few linear combinations of these 112 implied volatilities that describe most of the surface variation across stocks.

We first average Tuesday to Monday daily volatility surfaces to a weekly level to make the PCA analysis computationally feasible and to reduce the effect of outliers and microstructure noise. Next, we de-mean the surfaces for each stock and week by subtracting the average implied volatility of the whole surface (calls/puts, moneyness, and maturity). This step let us focus on how the shape of the surface rather than volatility level varies across stocks. Extensive literature already studies variation in volatility across stocks and how it relates to stock returns. We decompose a correlation matrix of the de-meaned volatility surfaces of all stocks across time using standard principal component analysis. PCA describes 112 implied volatilities representing the surface with few linear combinations capturing most of the surface variation and ranks them in decreasing

order. By construction, these principal components (PCs) are one-by-112 vectors that are orthogonal to each other. From a linear algebra perspective, PCs are a set of unit-length vectors (eigenvectors) that point in direction of most variation in the data (thus, their signs are not well-defined).

2.3 Describing IV Surface Principal Components

The first five principal components (PCs) explain 78% of the total variance of the de-meaned implied volatility surfaces (calculated by dividing the eigenvalue corresponding to the eigenvector by the sum of all eigenvalues). Figure 1 shows that the first five principal components explain 32%, 21%, 13%, 8%, and 4% variation, respectively. If we didn't de-mean IV surfaces, the volatility level would be responsible for most of the variation.

Each principal component is a 1-by-112 weight vector corresponding to standardized option parameters: calls/puts, maturity, and moneyness. Figure 2 plots PCs as a function of option parameters to gain visual intuition. The first principal component (PC1), which explains 32% of the variance, looks like the call-put IV spread combined with the option skew. Indeed, PC weights are positive for calls and negative for puts and increase as puts (calls) become more in-the-money ITM (out-of-the-money, OTM). This pattern is homogeneous across option time to expiration. Correlations reported in Table 2 confirm this interpretation. PC1 exposure has a correlation of 66% and -39% with the call-put volatility spread defined in Cremers and Weinbaum (2010) and the option skew defined in Xing, Zhang, and Zhao (2010). According to PCA, PC1 defines the most common direction in which volatility surfaces differ across stocks and over time. The second principal component (PC2) explains 18% of total surface variance. While PC2 loadings are positive for short-term maturities, they change sign at about 150 days-to-expiration and become negative for long-term maturities. PC3 explains 11% of variation and is most related to the option

skew. For puts (calls), the factor loadings are positive (negative) for ITM options and negative (positive) for OTM options. PC3 has a correlation of 28% with the option skew and -22% with risk-neutral skewness. Finally, PC4 and PC5 account for 8% and 4% of the variation and their interpretation is not straightforward. The highest absolute correlations are between PC4 and risk-neutral skewness at -15%, and between PC5 and the implied volatility spread at -26%. Correlations in Table 2 shows that existing factors such as size, book-to-market, momentum, reversal, and illiquidity are largely unrelated with the PC surface components. Thus, the surface PCs' ability to predict stock returns is unlikely to due to these firm characteristics. Overall, the main principal components aggregate option information differently from level, slope, and curvature or conventional ratios such as the put-call IV spread, skew, and time slope. Thus, these principal components are not simply a linear combination of conventional measures, which can explain their superior information content that we document next.

Principal components are orthogonal to each other and point in the direction of maximum variation in the data. We use them to compute PC exposures that reflect how much a given surface deviates in the direction of a given PC eigenvector. Every week, we compute PC exposure of each stock's IV surface to each of the first five PCs by multiplying a vector of 112 de-meaned implied volatilities for this stock-day by 112 weights of a given principal component. That is, PCA reduces the original implied volatility panel with stock-by-week rows and 112 implied volatility columns representing volatility surface to the same number of rows but only five columns reflecting exposures to the first five PCs. We focus on this new panel through the rest of the paper.

We next explore how these PC exposures depend on the standard measures computed on the same volatility surface and estimate the following Fama-MacBeth (1973) regression:

$$PC_{j,t} = \gamma_{0,t} + \gamma_1 IVTSlope_{i,t} + \gamma_2 IVSkew_{i,t} + \gamma_3 IVSpread_{i,t} + \gamma_{4-6} RNMoments + \epsilon \quad (1)$$

where $PC_{j,t}$ is the j^{th} principal component exposure on week *t*. Independent variables include the volatility time slope, the volatility skew, the put-call spread, risk-neutral volatility, skewness, and kurtosis. Table 3 shows that the adjusted R^2 are 60% for the PC1 regression, 89.5% for PC2, 39% for PC3, and well below 50% for PC4 and PC5. As expected, most of the coefficients are significant in all regressions because the explanatory variables and PCs are computed from the same volatility surface.

To facilitate a fair comparison of our approach against other option-based predictors, we aggregate the first three components into a single variable by simply adding their loadings ($PC1 + PC2 + PC3$). It reflects average deviation of a current IV surface from the overall sample average. Figure 3, Panel A shows that the surface deviation weights depend on option parameters in a non-linear way. The weights are the most negative for low levels of moneyness, defined as strike over stock price, and short maturities for both calls and puts. The loadings increase with the maturity and the moneyness.

Panel B of Figure 3 shows the time series average of the surface deviation (aggregate PC). On each day, we compute the average of surface deviation exposures across all stocks. By construction, the average surface deviation across the entire sample period is zero. It is positive in calm periods. However, in crash periods such as internet burst of September 2001 or the financial crisis, the aggregate surface deviation displays big negative spikes similar to VIX. The aggregate surface deviation has very similar patterns as aggregate volatility and liquidity. However, most of our results focus on the cross-section of stocks rather than aggregate time series.

Overall, the first five principal components explain most of the variation of the data and are largely independent of the existing option-based return predictors such as the implied volatility

spread, the option skew, and the risk neutral skewness. We next explore the ability of these principal components to predict stock returns.

3 Information Content of Volatility Surface for Future Stock Returns

After showing that a volatility surface on a given stock-day can be represented by its exposures to the main principal components, we explore its information content. This PCA-based approximation has many potential applications including predicting IV surface changes, a crucial element of managing risk of option portfolios. However, this paper focuses on how approximated surface deviation predicts stock returns. If principal components indeed summarize value-relevant information in the surface, then they will make redundant other standard option-based return predictors. Importantly, PCA does not use stock return information in any way and thus avoids data mining or overfitting return predictors.

We first show that the first PCs, which quantify how a given surface deviates from the average, predict stock returns in cross-sectional regressions. Table 4 presents the results of the Fama-MacBeth regressions of weekly and monthly stock returns on the volatility surface PC exposures for the first four PCs, option-based variables, and firm characteristics. We include the PCs simultaneously because their exposures are by construction orthogonal. We then add stock characteristics that are known to predict returns including lagged one-month and six-month return (reversal and momentum), size, book-to-market ratio, and the Amihud illiquidity measure. Finally, the last specification also includes option-based predictors including the implied volatility put-call spread, skew, time slope, and open-to-stock volume ratio.

It is well known that price in the option and the stock markets are connected. The four main variables extracted from the option market that are known to predict stock returns are the call-put implied volatility spread, risk-neutral skewness, the option to stock volume ratio, and the volatility

of implied volatility. Various versions of the implied volatility spread are used to predict returns in Cremers and Weinbaum (2010), Bali and Hovakimian (2009), and Yan (2011). Risk-neutral skewness predictability is documented in Bali, Hu, and Murray (2016), Conrad, Dittmar, and Ghysels (2013), Rehman and Vilkov (2012), Stilger, Kostakis, and Poon (2016), and Xing, Zhang, and Zhao (2010). The predictability of the option to stock volume ratio has been documented by Johnson and So (2012) and Roll, Schwartz, and Subrahmanyam (2010).

The first three PCs significantly predict stock returns, and the predictability becomes stronger after including controls. PC1 is a stronger predictor than PC2 and PC3, and its coefficients of 0.0045 and 0.0126 for weekly and monthly return regressions, respectively, correspond to about a 0.13% and 0.35% higher return per standard deviation increase in PC1 (which is 0.28). PC1 and PC2 have coefficients that are not much lower than for PC1. PC4 does not predict returns, which is natural as higher principal components capture smaller variations in the surface some of which are just noise.

Interestingly, conventional option-based predictors are not significant if principal components are included while the coefficients for PCs change little. PCs subsume all the predictive information of other option-based variables. To make this point even stronger, we define the surface deviation as PC1 + PC2 + PC3 so that PC information is aggregated in just one variable. Given that the first three principal components predict stock returns and that by construction the PC factors are uncorrelated, we create an aggregate principal component (PC) factor. The goal of the surface deviation factor is to combine the information from the three PCs that predict stock returns into a single variable. The surface deviation factor is a linear combination of the first three PCs. Since PC1, PC2, and PC3 have a positive relation with future stock returns, we define the surface deviation factor as PC1 + PC2 + PC

Table 5 reports how this volatility deviation predicts weekly and monthly stock returns, and the results are consistent with the individual PC results in Table 4. A one standard deviation increase in the surface deviation (which is 0.42) correspond to 0.17% and 0.55% higher weekly and monthly stock returns, respectively. Importantly, surface deviation makes other option-based predictors insignificant. In fact, none of other predictors except for lagged weekly stock return are significant.

Portfolio sorts in Table 6 confirm that the surface deviation is a strong return predictor. Results are similar for raw returns and four-factor Fama-French alphas, so we only describe the latter. The 10-minus-1 weekly alpha is 0.40% and 0.38% for equally-weighted and value-weighted sorts with t-statistics of 6.8 and 6.0 respectively. Monthly return portfolio sorts make the results stronger. The 10-minus-1 alphas are 1.32% and 1.25% for equally and value-weighted sorts with t-statistics of 11.7 and 10.6. Weekly alpha mostly originates in the first and tenth portfolios, while 8^{th} and 9^{th} deciles contribute meaningfully to monthly alpha (first and last deciles still dominate). Panel B of Table 5 reports portfolio sorts for PC1. Weekly and monthly alphas are 0.29% and 0.89%, respectively, and are highly statistically significant.

After documenting strong predictability for next week and month, we explore return predictability horizon for the surface deviation. It significantly predicts returns up to four months in the future. Figure 4 shows weekly four-factor alphas for a 10-minus-1 portfolio sorted on the surface deviation for up to thirty weeks after portfolio formation. The alphas remain statistically significant up to week 15 and becomes zero around week 21. Cumulative alpha reaches 4% in the first 12 weeks, then 6% around week 22 and then stabilizes around 6.5%. Thus, this weekly signal is remarkably persistent. We later show that the surface deviation is persistent especially for the 10^{th} portfolio that produces most of the alpha.

Portfolios sorted on the surface deviation earn significant alpha. How much turnover the strategy that buys the top and sells the bottom deciles would generate? To assess the persistence of the surface deviation, we analyze the transition probabilities across decile portfolios. Table 7 reports the probability that a stock moves up or down one portfolio, remains in the same portfolio, or moves to any other portfolio next month. We report these transition probabilities separately for all portfolios and for portfolio 10, because this portfolio is responsible for most of the return predictability. For all portfolios, we find that the probability of remaining in the same portfolio is much higher than the probability of moving up or down by one portfolio. Stocks stay in the same portfolio next month in 36% cases, which is obviously higher than the 10% unconditional probability implied by pure chance. However, the probability that a stock remains in portfolio 10 next month is much higher: 59%. Thus, a trading strategy based on 10-minus-1 portfolio does not have high turnover over time and the same stocks remain in portfolio 10 for multiple weeks and even months.

To perform the PCA analysis, we use the entire sample period. We address the concerns that PCA has a look-ahead bias and that PCs are unstable by performing an out-of-sample analysis. We perform a PCA analysis using a 3-year rolling window of volatility surfaces. We use implied volatility information up to time t to forecast returns for week $t+1$ and beyond. As before, the surface deviation is computed as a sum of the first three PCs. Table 8 shows how this out-of-sample surface deviation forecasts weekly and monthly returns in the Fama-MacBeth regressions. The regressions confirm positive and significant relation between the surface deviation factor and future stock returns. As expected, the coefficients and t-statistics decrease slightly compared to the in-sample results in Table 5: coefficient of 0.0030 (t-statistic of 4.5) decrease to 0.0029 (3.8) for out-of-sample weekly results with all controls.

We conclude that the surface deviation factor successfully combines the information of the first three PCs to predict future stock returns. To further validate this conclusion, we benchmark the surface deviation against main option-based predictors. We explore five popular option predictors: 1) the implied volatility spread, the average difference between call and put options, as defined by Cremers and Weinbaum (2010), 2) the volatility smirk or option skew, the difference between OTM put and ATM call implied volatilities, as defined by Xing, Zhang, and Zhao (2010), 3) the risk-neutral skewness and 4) risk-neutral kurtosis as defined in Conrad, Dittmar, and Ghysels (2013), and 5) the O/S option to stock volume ratio as defined in Johnson and So (2012).

Table 9 shows that all these variables except for risk-neutral kurtosis strongly predict weekly and monthly stock returns in univariate regression controlling for stock-based predictors. For example, the IV spread has t-statistics of 4.6 and 6.1 in weekly and monthly regressions. Strikingly, these option predictors lose their predictive power after surface deviation is added to the regressions. E.g., t-statistics for the IV spread and IV skew drop from 6.1 to -1.1 and from -4.4 to -0.8 in monthly regressions. At the same time, the surface deviation is significant in all specifications with t-statistic ranging from 7.2 to 3.7 for monthly returns and its coefficient does not change across specifications. The coefficient for option-to-stock volume is least affected by adding the surface deviation because option volume information is not spanned by implied volatility (options prices) that PCs and the surface deviation are based on. The PCA approach can be extended include option volume and changes in implied volatility surface to better reflect other sources of option market information.

Overall, these remarkable results imply that the surface deviation encompasses return predictability contained in other option predictors. The surface deviation factor outperforms existing stock return predictors that rely on ad-hoc implied volatility ratios such as the volatility

spread, the option skew, risk-neutral skewness, and volatility of implied volatility. In contrast, the surface deviation is data-driven, as PCA preserves maximum information from the entire IV surface.

After establishing that the surface deviation and the principal components that define it encompasses all well-known option predictors of stock returns, the surface deviation can help us understand why option prices predict stock returns. While it is useful to understand why a particular ad-hoc ratio of implied volatilities predict returns, the surface deviation allows for a more general insights because it potentially reflects IV ratios that are not yet studied in the literature.

Several hypotheses have been proposed to explain why conventional option ratios predict stock returns. First, demand pressure in the options market by informed investors affects option prices, and this signal is slowly reflected in the underlying price (Cremers and Weinbaum (2010), Roll, Schwartz, and Subrahmanyam (2010), and Xing, Zhang, and Zhao (2010)). Second, option prices reflect jump risk that should incur a risk premium which is reflected in the underlying stock returns (Bali and Hovakimian (2009), and Yan (2011)). Option prices reflect short-sale constraints, and as well-known, stocks that are hard or expensive to sell short underperform the market (Johnson and So (2012), and Stilger, Kostakis, and Poon (2016)). For example, Ofek, Richardson, and Whitelaw (2004) and Muravyev, Pearson, and Pollet (2018) show that the put-call parity violations are proportional to the stock lending fees and thus affect the put-call volatility spread and skew. The stock lending fee is like a shadow continuous dividend and OptionMetrics implied volatilities ignore it leading to apparent put-call parity violations. Finally, Conrad, Dittmar, and Ghysels (2013) argue that risk-neutral skewness reflects underlying risk factors, and thus stock returns could be simply compensation for taking risk.

We first study informed trading and jump risk hypotheses by creating proxy variables for them. PIN, the probability of informed trading, measures informed trading and adverse selection. The down and up risk-neutral jump risk using the methodology proposed by Bollerslev and Todorov (2011) captures the jump risk. Table 10 shows how these explanatory variables impact the relation between the surface deviation and future weekly and monthly stock returns. We first put each proxy in the Fama-MacBeth regression individually and then add the surface deviation in the second specification. Measures of informed trading and jump risk do not affect the predictive ability of the surface deviation. Of course, it could be the case that the available proxies are just not sufficiently strong.

We next study short sale constraints because they directly affect the IV surface and an extensive literature connects options with shorting costs. Short sellers pay stock lending fee to short stock that reduces after-cost profitability of their trades and reflects the difficulty of locating the shares. Thus, we study how lending fee interacts the surface deviation. Comprehensive daily data on lending fees are only available starting from July 2006, which reduces our sample period. However, fee is available almost for all optionable stocks after that day, thus stock cross-section is not affected. Lending fee is itself a strong negative return predictor as hard-to-short stocks underperform the market. Table 11 confirms this known fact: the coefficient for the fee has a t-statistics of -3.3 and -4.0 for weekly and monthly returns. Surface deviation remains a significant predictor in the subsample with available lending fees. The surface deviation and the lending fee are strongly related. If both are included in the regression, the coefficients and statistical significance drop, but the surface deviation is more affected. For weekly returns, the coefficient and its t-statistic for surface deviation drop from 0.0017 (2.8) to 0.0010 (1.7), while lending fee experiences a small drop from -0.02 (-3.3) to -0.016 (-3.05). The surface deviation remains

significant (t-statistics of 2.2) for monthly regressions after including lending fee. To better understand this relation, we replace surface deviation with its components in last column in the table. PC1 remains a strong predictor after including lending fees.

You may notice that PC2 and PC3 fail to predict returns in Table 11 that covers post-2006 period with available lending fee. We examine this observation closer by studying how the return predictability by surface deviation and its components changed over the sample period. Each week, we sort stocks into deciles, compute equally-weighted returns I the following week, and then take the ten-minus-one return difference. Figure 5 shows a four-week moving average for the ten-minus-one weekly return for portfolios sorted on the surface deviation and the first three PCs. Returns are obviously noisy, but the trends are easy to spot. Average weekly returns drops from 0.65% up through early 2002 to 0.29% afterwards. As surface deviation is simply a sum of the first three PCs, we can identify which PCs are responsible for decline in profitability. Weekly returns for PC1 fluctuate between -0.5% and 1% but remain relatively stable with a 0.26% average. PC2 profits a little bit during good times, but loses a lot during market crashes such as the burst of the internet bubble and the 2008 financial crisis. PC3 is remarkable, it made all its money during the burst of the internet bubble.

To examine whether the return predictability is due to compensation for taking risk, we study sample of S&P 500 constituents. Financial frictions are especially small for S&P 500 stocks; trading costs are negligible; investors can easily accumulate large positions and short; lending fees are almost zero as index funds provide an unlimited supply of shares. Due to extensive research and constant investor attention, these stocks are rarely mispriced. Because of low frictions and mispricing, returns for S&P 500 stocks are largely unpredictable, but any predictability is likely due to compensation for taking risk. Table 12 shows that the surface deviation continue to predict

weekly stock returns for S&P 500 stocks; moreover, the coefficient of 0.0029 (t-statistics of 2.7), is only slightly lower than 0.036 for the full sample in Table 5. Like for the full sample, the only other significant predictor is weekly stock reversal with t-statistic of -3.9. Individual PCs continue to predict stock return but, like for surface deviation, statistical significance is lower due to a smaller sample size. Overall, these results tentatively suggest that a significant chunk of surface return predictability is due to compensation for taking risk.

In sum, these results do not fully resolve the puzzle but point that a lot of option-based return predictability is due to compensation for taking risk and also related to shoring costs.

4 Option Returns

Another natural application for our approach is to study the relation between the surface deviation and the cross-section of future option returns. Specifically, we study how the surface deviation predicts weekly straddle returns and show that it compares favorably to strongest option return predictors. A straddle is an option strategy that simultaneously buys an at-the-money call option and an at-the-money put option with the same strike price and expiration date. This position is almost delta neutral at the origin and thus is immune to small directional changes in the underlying stock price. We perform portfolio sorts and Fama-MacBeth regressions of one-week straddle returns for options that expire between three and six weeks. Average straddle returns are proportional the difference between realized return variance and the variance reflected in option prices. The surface deviation reflects the most important features of the de-meaned implied volatility surface, and thus we expect it to predict straddle returns well as both are derived from option prices.

Every week, we form decile portfolios based on the surface deviation and report average next-week straddle returns in Table 13. The surface deviation varies from -0.77 for the bottom

decile portfolio to 0.54 for the top one. Option returns monotonically decrease from 1.7% per week for portfolio one to -2.6% for portfolio ten. The portfolio that buys decile ten and sells decile one earns a weekly return of -4.2% with a t-statistic of -14.8. This long-short portfolio has positive skew and fat tails. Four-factor alpha for the 10-minus-1 portfolio is similar to raw returns. Value-weighted portfolios display very similar predictability results to the equally-weighted case. In untabulated results, we supplement portfolio sorts with Fama-MacBeth regressions. In addition to the surface deviation, we include three strongest known option return predictors: the slope of the volatility term structure as in Vasquez (2017), historical minus implied volatility as in Goyal and Saretto (2009), and idiosyncratic volatility as in Cao and Han (2013). The multivariate analysis confirms that the surface deviation strongly predicts option returns, and the coefficient for the surface deviation is larger and more statistically significant in the multivariate regression. Moreover, the surface deviation is more significant than all other option return predictors.

A growing literature studies the determinants of option returns. We contribute to it by identifying another powerful return predictor: the surface deviation, which optimally compresses information in implied volatility surface.

5 Conclusion

In this paper, we introduce two ideas: a methodological and an empirical. Methodologically, we suggest PCA a simple way to describe the entire volatility surface with just few numbers reflecting exposure to main principal components. We find that five principal components (PC) explain more than 70% of the variability of the (de-meaned) volatility surfaces across stocks and time. The main principal components cannot be explained in terms of simple volatility ratios such as IV spread between puts and call, volatility skew and time slope. After reducing surface dimensionality, we study its information content. Why do option prices predict

stock returns? First, principal components subsume other popular option-based predictors of stock returns such as the difference between put and call implied volatility, the option skew, and the term structure of implied volatilities. Perhaps, PCs indeed keep all the important information from volatility surface. Our PCA approach let us analyze why the entire surface rather than particular ratios predict returns, and this analysis reveals the dynamic and multifaceted nature of option-based return predictability. During early years, option prices could be used to identify mispricing in the equity market. However, the inefficiencies disappeared post 2002. On the other hand, the first principal component consistently predicts returns during the entire sample period. Its predictive ability can be attributed to option prices reflecting shorting costs and to some extend to option prices capturing the true risk factors.

6 References

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.

Bakshi, G., N. Kapadia, and D. Madan, 2003, Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options, Review of Financial Studies 16, 101-143.

Bali, T. G., and A. Hovakimian, 2009, Volatility Spreads and Expected Stock Returns, Management Science 55, 1797-1812.

Bali, T.G., J. Hu, and S. Murray, 2016, Option Implied Volatility, Skewness, and Kurtosis and the Cross-Section of Expected Stock Returns, Working paper.

Baltussen, G., S. Van Bekkum, and B. Van der Grient, 2012, Unknown Unknowns: Uncertainty About Risk and Stock Returns. Journal of Financial and Quantitative Analysis 54, 1615-1651.

Bollerslev, T., and V. Todorov, 2011, Tails, Fears, and Risk Premia, Journal of Finance 66, 2165-2211.

Cao, J., and B. Han, 2013, Cross section of Stock Option Returns and Stock Volatility Risk Premium, Journal of Financial Economics 108, 231-249.

Carhart, M., 1997, On Persistence in Mutual Fund Performance, Journal of Finance 52, 57-82.

Christoffersen, P., M. Fournier, and K. Jacobs, 2018, The Factor Structure in Equity Options, Review of Financial Studies 31, 595-637.

Christoffersen, P., Jacobs, K. and Chang, B.Y., 2013. Forecasting with option-implied information. In Handbook of economic forecasting (Vol. 2, pp. 581-656). Elsevier.

Collin-Dufresne, P., R. S. Goldstein, and J. S. Martin, 2001, The Determinants of Credit Spread Changes, Journal of Finance 56, 2177-2207.

Conrad, J., R.F. Dittmar, and E. Ghysels, 2013, Ex Ante Skewness and Expected Stock Returns, Journal of Finance 68, 85-124.

Cremers, M., and D. Weinbaum, 2010, Deviations from Put-Call Parity and Stock Return Predictability, Journal of Financial and Quantitative Analysis 45, 335-367.

Fama, E., and K. French, 1993, Common Risk Factors in the Returns on Stocks and Bonds, Journal of Financial Economics 33, 3-56.

Fama, E., and K. French, 2008, Dissecting Anomalies, Journal of Finance 63, 1653-1678.

Fama, E., and M. J. MacBeth, 1973, Risk, Return, and Equilibrium: Empirical Tests, Journal of Political Economy 81, 607-636.

Goyal, A., and A. Saretto, 2009, Cross Section of Option Returns and Volatility, Journal of Financial Economics 94, 310-326.

Gutierrez Jr, R., and E. Kelley, 2008, The Long-Lasting Momentum in Weekly Returns, Journal of Finance 63, 415-447.

Johnson, T., and E. So, 2012, The Option to Stock Volume Ratio and Future Returns, Journal of Financial Economics 106, 262-286.

Litterman, R., and J. Scheinkman, 1991, Common Factors Affecting Bond Returns, Journal of Fixed Income 1, 54-61.

Neuberger, A., 2012, Realized Skewness, Review of Financial Studies 25, 3423-3455.

Pan, J., and K.J. Singleton, 2008, Default and Recovery Implicit in the Term Structure of Sovereign CDS Spreads, Journal of Finance 63, 2345-2384.

Rehman, Z., and G. Vilkov, 2012, Risk-Neutral Skewness: Return Predictability and Its Sources, Working Paper, BlackRock and Goethe University.

Roll, R., E. Schwartz, and A. Subrahmanyam, 2010, O/S: The Relative Trading Activity in Options and Stock, Journal of Financial Economics 96, 1-17.

Stilger, P.S., A. Kostakis, and S.H. Poon, 2016, What Does Risk-Neutral Skewness Tell Us About Future Stock Returns?, Management Science 63, 1814-1834.

Vasquez, A., 2017, Equity Volatility Term Structures and the Cross Section of Option Returns, Journal of Financial and Quantitative Analysis, 52, 2727-2754.

Xing, Y., X. Zhang, and R. Zhao, 2010, What Does the Individual Option Volatility Smirk Tell Us about Future Equity Returns?, Journal of Financial and Quantitative Analysis 45, 641-662.

Yan, S., 2011, Jump Risk, Stock Returns, and Slope of Implied Volatility Smile, Journal of Financial Economics 99, 216-233.

Figure 1: Surface principal component contributions

This figure shows the relative contribution of the first ten principal components, i.e., how much variation of the cross-section of implied volatility surfaces each component explains. Each implied volatility surface contains 112 volatilities for calls and puts, absolute option deltas of 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, and 0.8 and option maturities of 30, 60, 91, 122, 152, 182, 273, and 365 calendar days. After demeaning each volatility surface, we apply the standard principal component analysis to the matrix with stock-by-week rows and implied volatility columns.

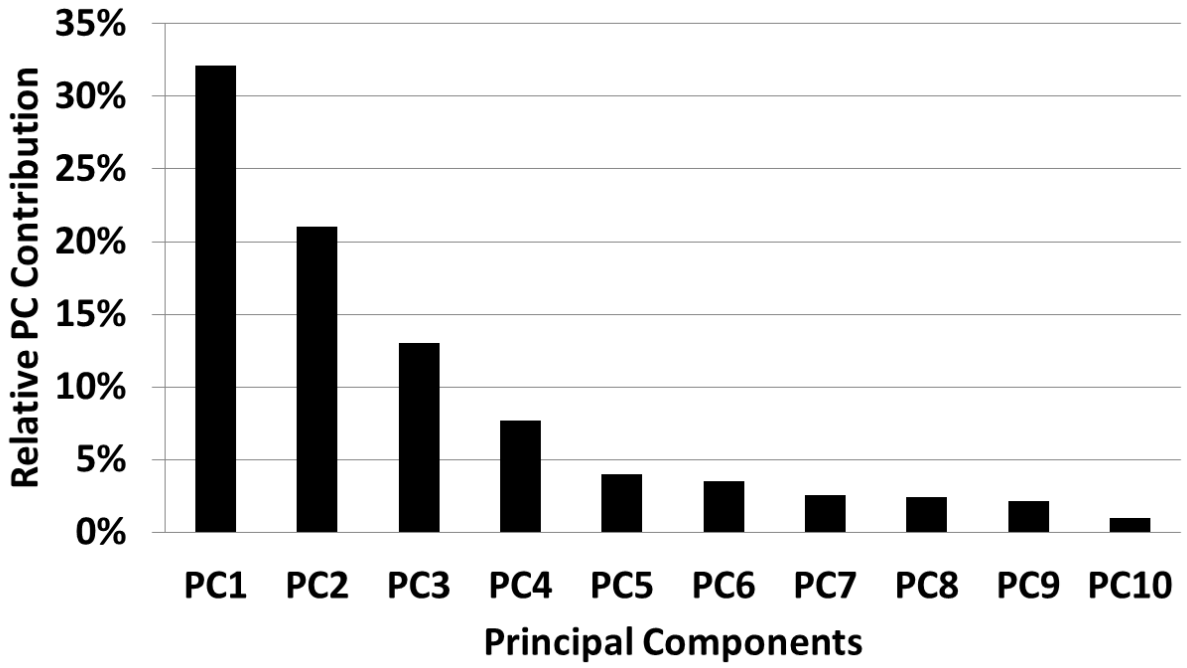

Figure 2: Surface principal component weights

This figure shows the first five principal components as a function of option parameters: calls (upper panel) vs. puts (lower panel), moneyness defined as the strike price divided by stock price, and option maturity in calendar days. Principal components are eigenvectors that point in the direction of maximum variation in the implied volatility surfaces across stock-weeks. Each implied volatility surface contains 112 volatilities for calls and puts, absolute option deltas of 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, and 0.8 and option maturities of 30, 60, 91, 122, 152, 182, 273, and 365 calendar days. After demeaning each volatility surface, we apply the standard principal component analysis to the matrix with stock-by-week rows and implied volatility columns.

Panel A: 1st Principal Component (PC1)

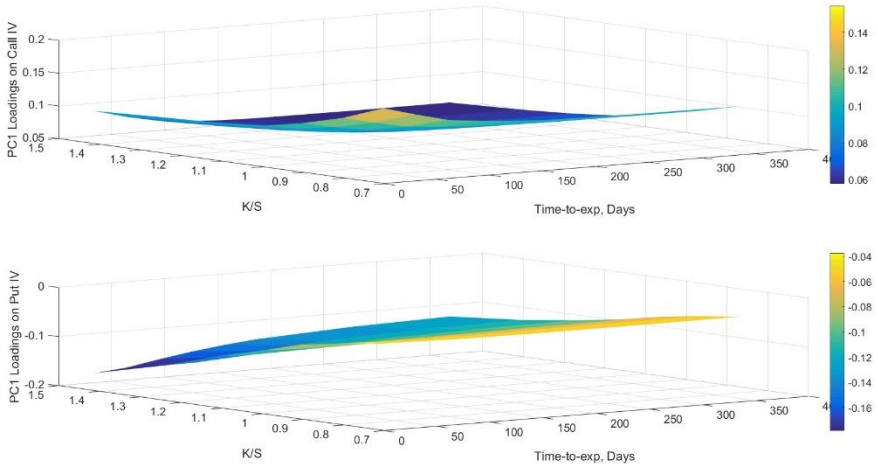

Panel B: 2d Principal Component (PC2)

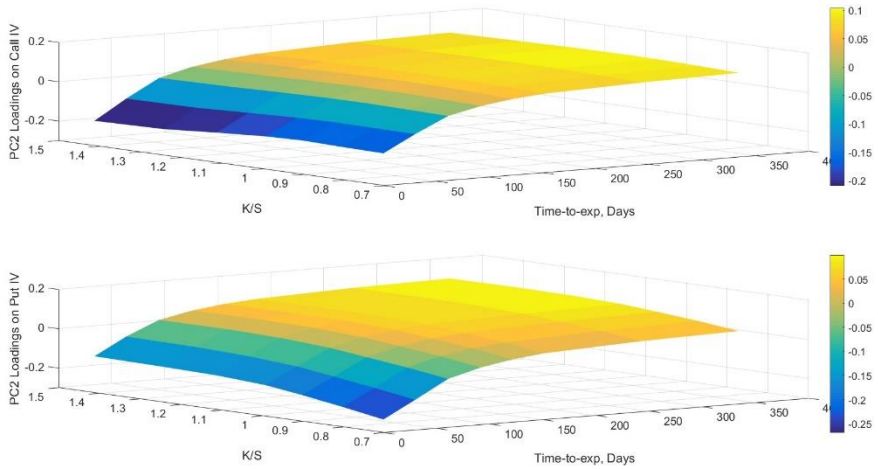

Panel C: 3d Principal Component (PC3)

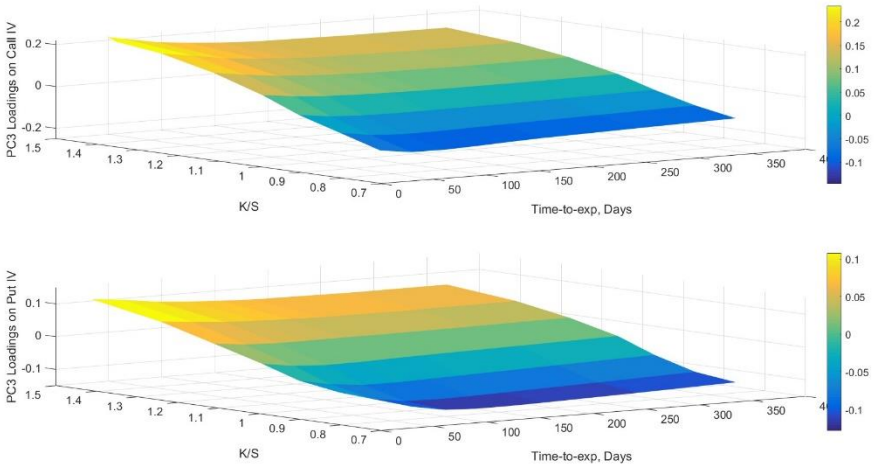

Panel D: 4th Principal Component (PC4)

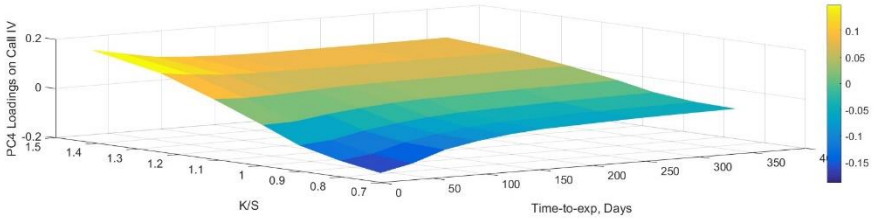

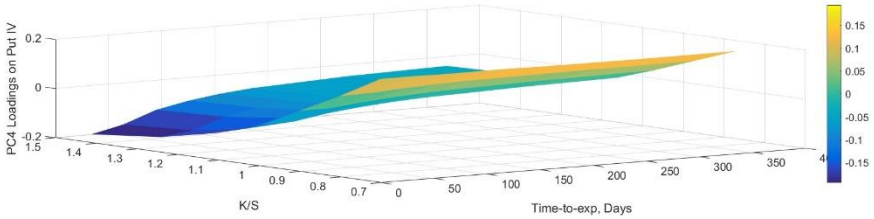

Panel E: 5th Principal Component (PC5)

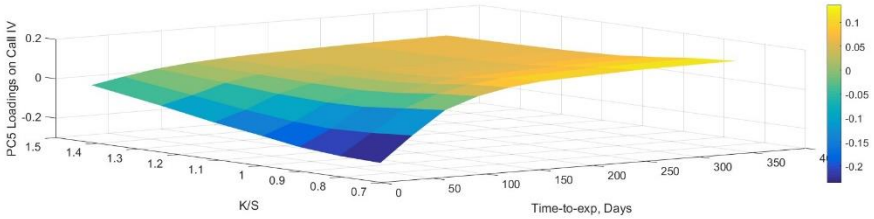

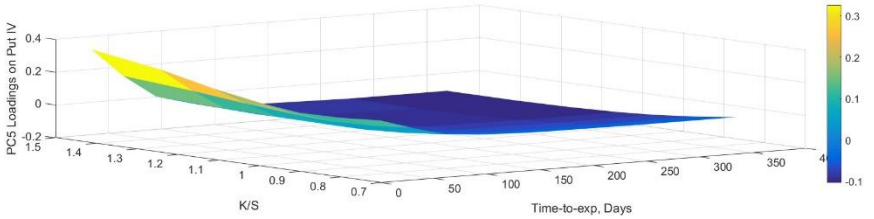

Figure 3: Surface deviation weights

Panel A of this figure shows the loadings/weights for the surface deviation as a function of option parameters: call (upper panel) vs. puts (lower panel), moneyness defined as the strike price divided by stock price, and option maturity in calendar days. The surface deviation is a sum the first three PCs: PC1 + PC2 + PC3. Panel B the time-series average across stocks of the surface deviation. Principal components are eigenvectors that point in the direction of maximum variation in the implied volatility surfaces across stock-weeks. Each implied volatility surface contains 112 volatilities for calls and puts, absolute option deltas of 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, and 0.8 and option maturities of 30, 60, 91, 122, 152, 182, 273, and 365 calendar days. After demeaning each volatility surface, we apply the standard principal component analysis to the matrix with stock-by-week rows and implied volatility columns.

Panel A: Surface Deviation Loadings

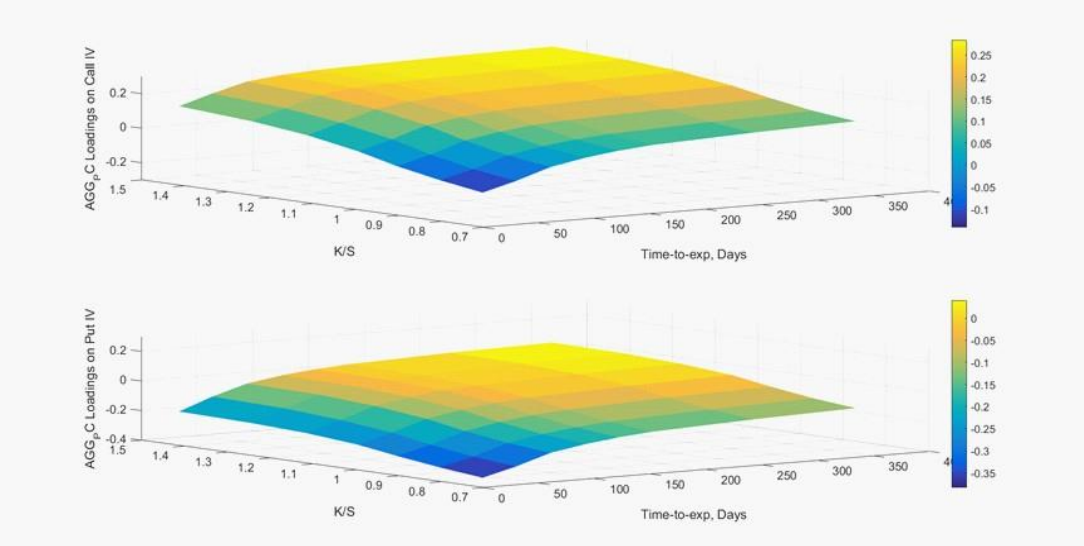

Panel B: Time-Series Average

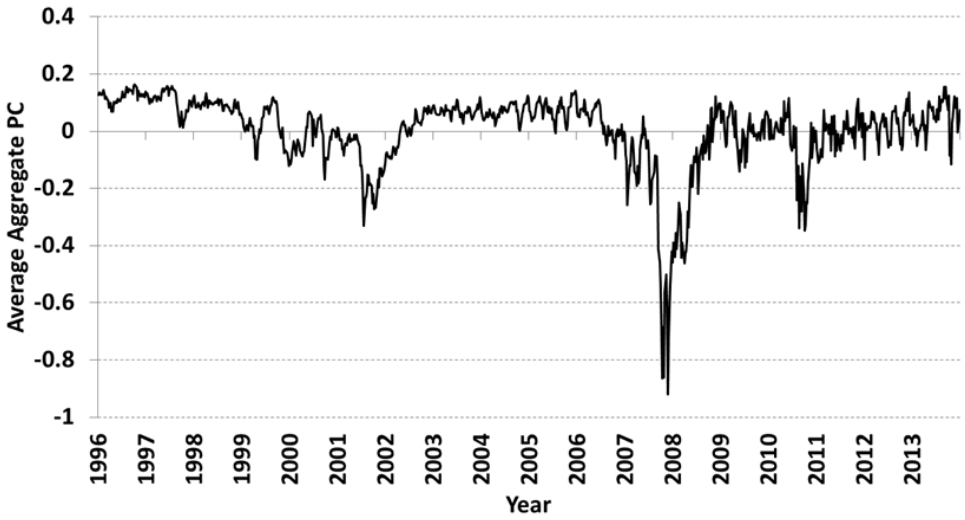

Figure 4: Return predictability horizon

This figure shows the return predictability horizon for the surface deviation. Panel A reports the four-factor alpha of the 10-minus-1 equally-weighted portfolio in week $(t+\tau)$ on the surface deviation in week t. Panel B reports cumulative alphas from week $(t+1)$ to week $(t+\tau)$. We report the 95% confidence interval.

Panel A: Weekly Returns

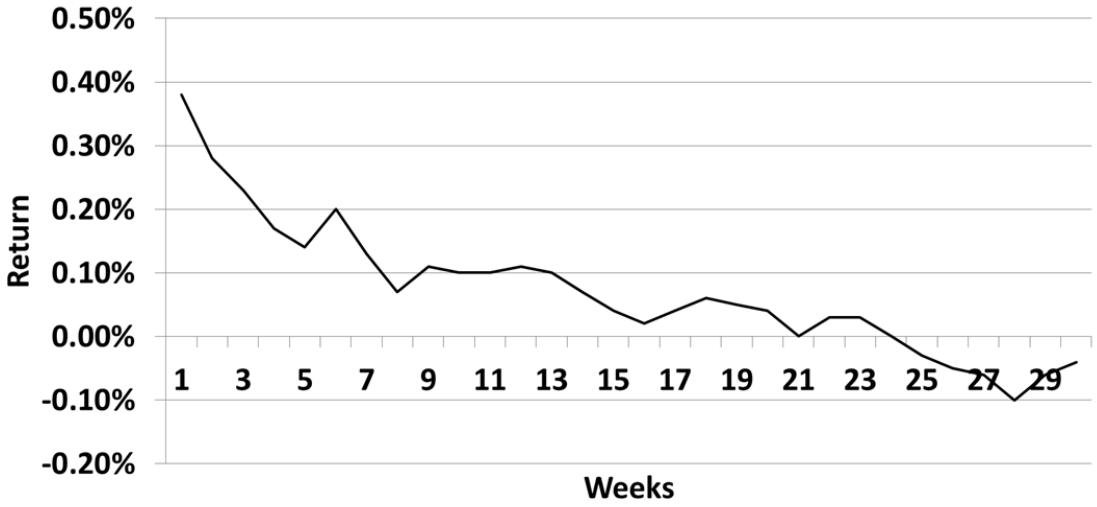

Panel B: Cumulative Weekly Returns

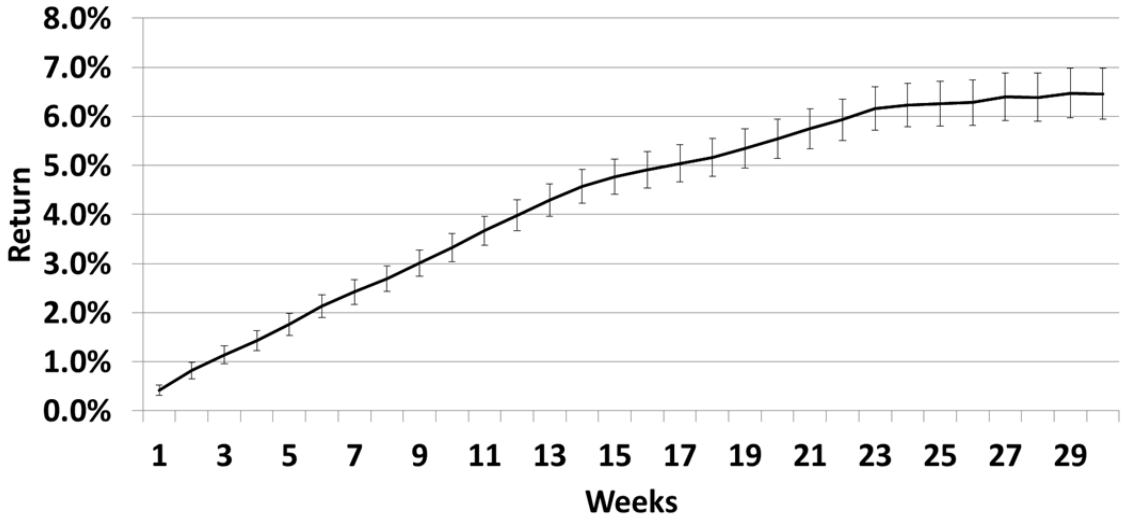

Figure 5: Return predictability over time

This figure shows how return predictability changes over the sample period. Each week, we sort stocks into deciles by a given principal component and compute subsequent ten-minus-one portfolio difference in weekly stock returns. Figures show 25-week moving average of these ten-minus-one weekly stock returns.

Panel A. Surface deviation

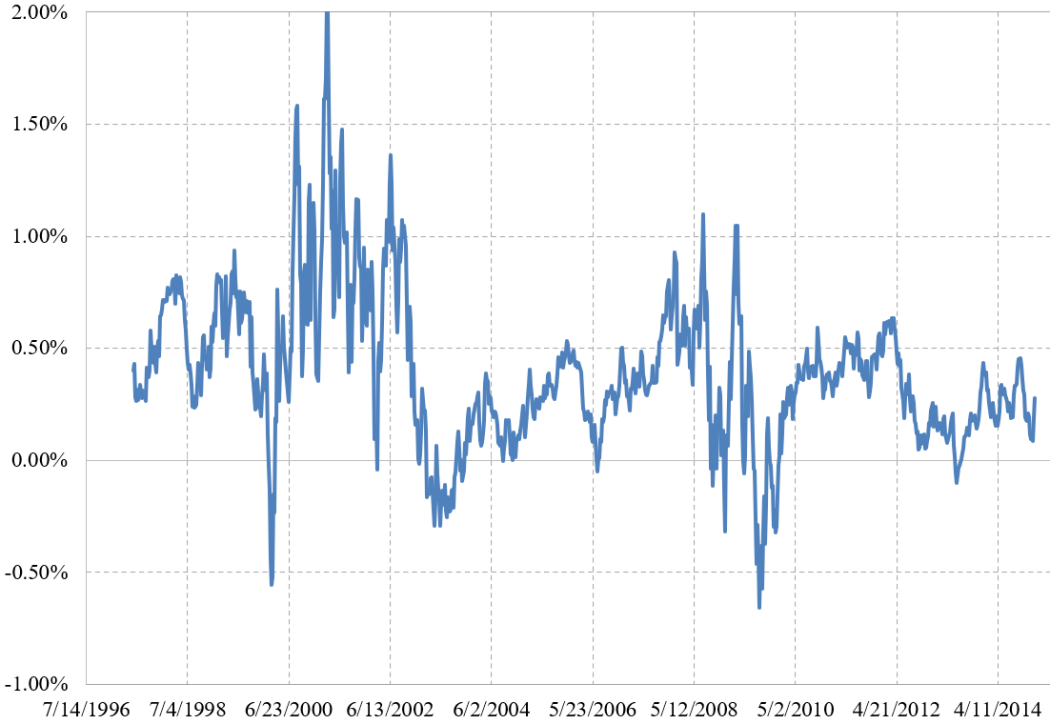

Panel B. PC 1

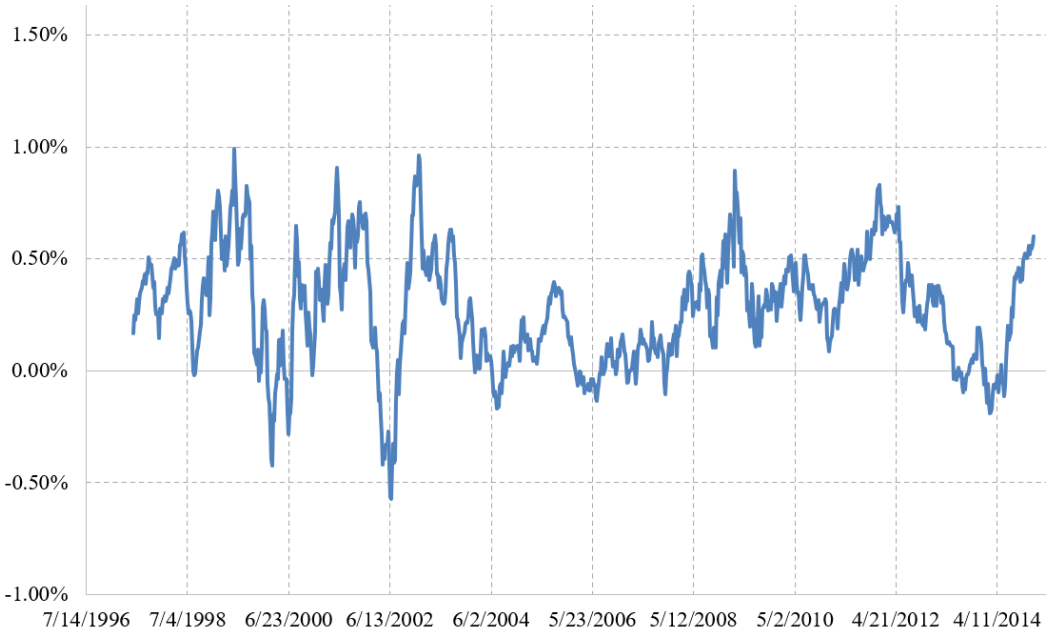

Panel C. PC 2

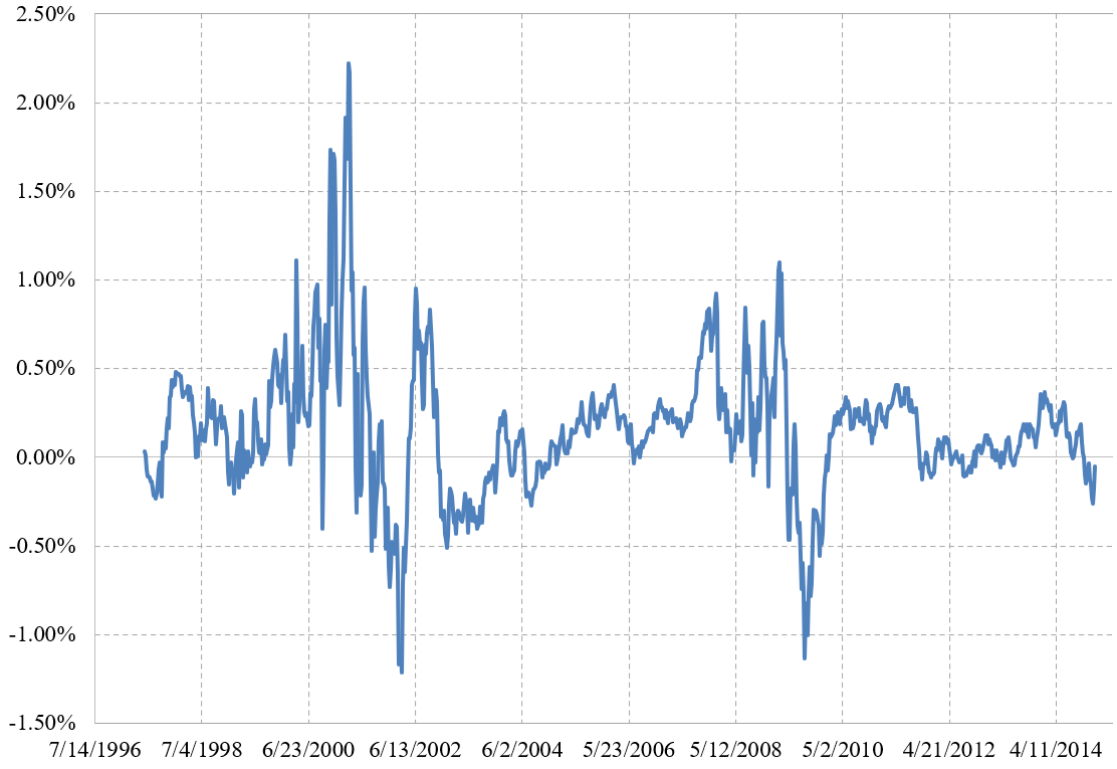

Panel D. PC 3

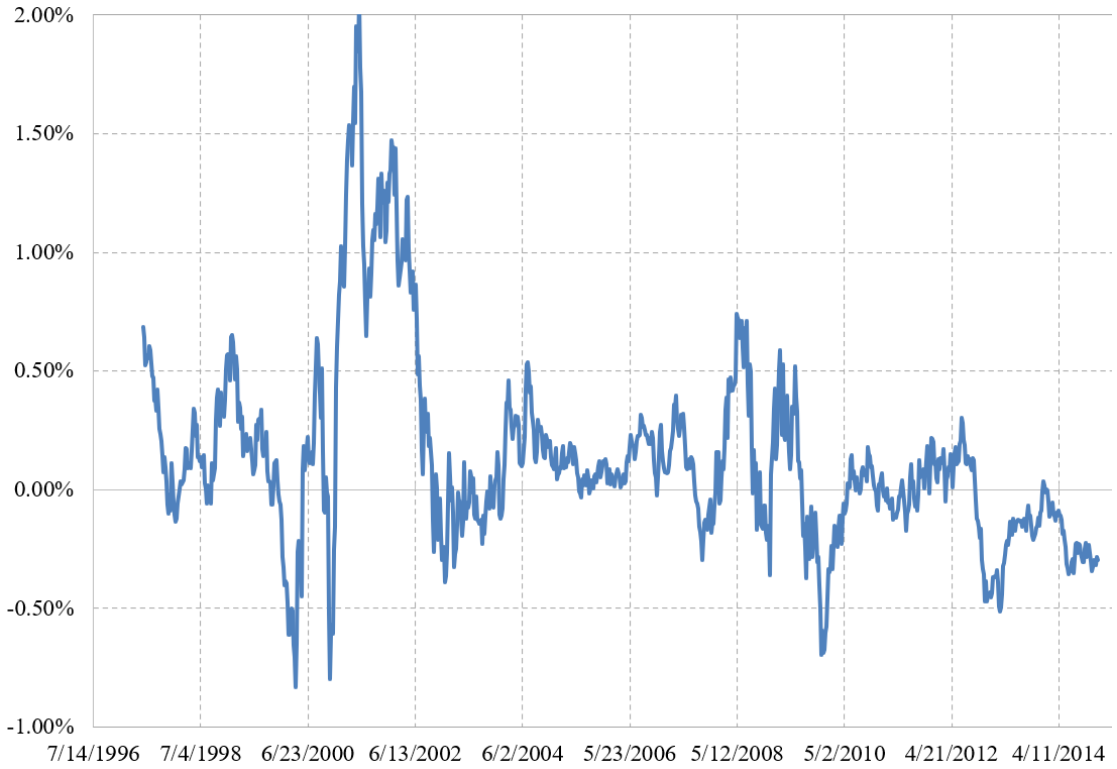

Table 1: Summary statistics

This table reports summary statistics for the first five IV surface principal components, the surface deviation, which is the sum of the first three PCs. Other variables include the implied volatility skew as in Xing, Zhang, and Zhao (2010), implied volatility spread as in Cremers and Weinbaum (2010), the IV time slope defined as the difference between long and short term ATM implied volatilities, risk-neutral moments as in Conrad, Dittmar, and Ghysels (2013), O/S option to stock volume ratio as in Johnson and So (2012), last-week return (Ret_{1w}) as a proxy for stock reversal, lagged six-month return (Ret_{6M}) as a proxy for momentum, logarithm of market-capitalization (Size), book-to-market (BE/ME), and the Amihud illiquidity measure.

Variable	Mean	Std.Dev.	5% Percentile	50% Percentile	95% Percentile
PC1	0.000	0.287	-0.292	0.019	0.256
PC2	0.000	0.246	-0.401	0.044	0.243
PC3	0.000	0.179	-0.218	-0.006	0.248
PC4	0.000	0.136	-0.149	-0.009	0.185
PC5	0.000	0.096	-0.092	-0.016	0.153
Surface Dev	0.000	0.419	-0.547	0.048	0.395
$\overline{IV}$ Skew	0.055	0.074	-0.009	0.040	0.182
IV Spread	-0.008	0.067	-0.078	-0.005	0.055
IV Time Slope	-0.029	0.079	-0.149	-0.017	0.046
RN Vol	0.469	0.153	0.253	0.452	0.745
RN Skew	-0.361	0.306	-0.893	-0.336	0.087
RN Kurt	3.064	0.649	2.193	3.002	4.150
O/S	3.374	7.707	0.000	1.000	13.907
Ret_{1w}	0.002	0.065	-0.089	0.000	0.099
Ret_{6M}	0.151	0.801	-0.571	-0.002	1.325
Size	7.261	1.584	4.927	7.105	10.134
BE/ME	0.549	0.634	0.067	0.447	1.370
Amihud	0.032	0.345	0.000	0.003	0.075

Table 2: Correlations

This table reports the correlation matrix of the principal component exposures, option variables, and firm characteristics. We report correlations for the first five IV surface principal components, PC1 to PC5, the surface deviation, which is a sum of the first three component exposures, implied volatility spread, implied volatility skew, risk-neutral skewness and risk-neutral kurtosis, O/S option-to-stock volume ratio, last-week return (Ret_{1w}), lagged six-month return (Ret_{6M}), market-capitalization (Size), book-to-market (BE/ME), and the Amihud measure of illiquidity.

	PC1														
PC2	0.9	PC2													
PC3	-1.4	-4.2	PC3												
PC4	1.0	1.7	-1.7	PC4											
PC5	-0.9	-2.7	0.7	-6.8	PC5										
SurfaceDev	71.1	-55.6	-40.8	0.4	0.6	SurfaceDev									
IVSkew	18.2	30.0	14.9	47.5	23.2	-10.8	IVSkew								
IVSpread	66.1	-8.5	-2.1	-9.1	-3.7	51.7	4.3	IVSpread							
RNSkew	4.7	9.8	-23.2	-16.0	-0.4	6.1	-21.0	9.1	RNSkew						
RNKurt	0.7	-9.6	1.9	10.3	5.5	5.9	6.6	0.7	-69.7	RNKurt					
O/S	-8.1	3.3	7.3	-7.2	-3.9	-10.8	-4.6	-6.5	-9.2	8.2	O/S				
Ret_{1w}	-5.7	-2.0	1.0	2.0	0.0	-3.3	1.0	-10.7	-2.5	1.7	1.4	Ret_{1w}			
Ret_{6M}	-0.8	0.0	3.6	-1.3	-3.1	-2.1	-1.2	-0.6	5.3	-5.4	5.3	0.0	Ret_{6M}		
Size	3.6	-15.8	16.9	-6.2	-10.2	4.6	-1.7	4.1	-38.2	34.1	18.6	3.4	3.0	Size	
BE/ME	3.2	5.1	-4.8	0.9	6.9	1.4	3.5	1.2	1.3	-2.6	-6.6	-4.6	-10.9	-14.3	BE/ME
Amihud	-6.5	1.6	-2.8	1.3	-2.6	-4.4	-2.6	-7.2	3.9	-4.6	-0.5	0.0	5.9	-7.3	-0.9

Table 3: Understanding principal components

This table presents the Fama-MacBeth results from regressing each of the first five principal components on existing predictors of the cross-section of stock returns. The remaining variables are the slope of the implied volatility term structure defined as the difference between long and short term ATM implied volatilities, implied volatility spread, implied volatility skew, and the risk-neutral volatility, skewness, and kurtosis.

	PC1	PC2	PC3	PC4	PC5
Intercept	-0.0003	0.173***	-0.074***	-0.009	-0.107***
	(-0.04)	(23.78)	(-4.57)	(-1.29)	(-29.65)
IV Time Slope	0.247***	-0.895***	-0.566***	0.436***	0.430***
	(18.21)	(-149.56)	(-13.69)	(-23.09)	(48.95)
IV Skew	4.291***	1.140***	-0.615***	-0.076***	-0.270***
	(40.21)	(45.87)	(-15.94)	(-4.07)	(-15.10)
IV Spread	0.131***	2.889***	-0.327***	0.675***	0.187***
	(7.23)	(339.29)	(-26.31)	(-69.51)	(38.32)
RN Vol	0.001	-0.017***	-0.168***	0.016***	0.084***
	(0.18)	(-6.15)	(-11.02)	(-4.27)	(32.28)
RN Skew	0.016***	-0.027***	0.217***	-0.034***	0.022***
	(4.07)	(-22.64)	(45.75)	(-19.37)	(15.58)
RN Kurt	0.008***	-0.008***	0.072***	-0.011***	0.016***
	(3.42)	(-9.61)	(28.12)	(-12.59)	(14.28)
Adj R^2	0.601	0.895	0.389	0.471	0.204

Table 4: IV surface principal components and stock returns

This table presents the Fama-Macbeth regressions of weekly and monthly stock returns on the first four principal components, PC1 to PC4, of the cross-section of volatility surfaces. Independent variables include PC1 to PC4, the IV time slope (the difference between long and short term ATM implied volatilities), the implied volatility spread and skew, option-to-stock volume ratio, last-week return (Ret_{lw}), lagged six-month return (Ret_{6M}), market-capitalization (Size), book-to-market (BE/ME), and Amihud illiquidity measure. We report the t-statistics in parenthesis that are based on robust standard errors and the significance levels are indicated with * (10%), ** (5%), *** (1%). Sample size is 1,914,768 for weekly returns and 1,896,942 for monthly returns.

	Weekly Returns			Monthly Returns		
	1	2	3	4	5	6
Intercept	0.002	0.0014	0.0021	0.0078*	0.0083	0.0099
	(1.65)	(0.55)	(0.82)	(1.83)	(0.92)	(1.13)
PC1	0.0045***	0.0042***	0.0035***	0.0126***	0.0122***	0.0130***
	(8.12)	(8.63)	(4.92)	(7.40)	(8.17)	(5.93)
PC2	0.0028***	0.0036***	0.0026**	0.0099***	0.0121***	0.0099***
	(3.12)	(5.30)	(2.29)	(3.11)	(5.58)	(2.96)
PC3	0.0038***	0.0031***	0.0023**	0.0120***	0.0091***	0.0074**
	(3.58)	(3.10)	(2.36)	(3.25)	(2.69)	(2.20)
PC4	-0.0008	-0.0010	-0.0011	0.0024	0.0003	-0.0004
	(-0.46)	(-0.70)	(-0.84)	(0.39)	(0.06)	(-0.09)
Ret_{1w}		-0.0202***	-0.0200***		-0.0272***	-0.0278***
		(-6.20)	(-6.21)		(-3.72)	(-3.89)
Ret_{6M}		-0.0002	-0.0001		-0.0017	-0.0014
		(-0.43)	(-0.30)		(-0.92)	(-0.79)
Size		-0.0000	0.0000		-0.0004	-0.0002
		(-0.13)	(0.06)		(-0.52)	(-0.24)
B/M		0.0008	0.0007		0.0021	0.0016
		(1.63)	(1.50)		(1.15)	(0.98)
Amihud		-0.0017	-0.0023		-0.0077	-0.0096
		(-0.28)	(-0.38)		(-0.37)	(-0.46)
IVTSlope			0.0024			0.0059
			(0.72)			(0.58)
IVSkew			-0.0024			-0.0077
			(-1.13)			(-1.15)
IVSpread			0.0035			-0.0076
			(1.00)			(-0.74)
O/S			-0.0001			-0.0004
			(-1.48)			(-1.46)
Adj. R^2	0.013	0.0483	0.0535	0.0135	0.0516	0.0573

Table 5: Surface deviation and stock returns

This table presents the Fama-Macbeth regressions of weekly and monthly stock returns on the surface deviation. Independent variables also include the IV time slope (the difference between long and short term ATM implied volatilities), the implied volatility spread and skew, option-to-stock volume ratio, last-week return (Ret_{1w}), lagged six-month return (Ret_{6M}), market-capitalization (Size), book-to-market (BE/ME), and Amihud illiquidity measure. We report the t-statistics in parenthesis that are based on robust standard errors and the significance levels are indicated with * (10%), ** (5%), *** (1%).

	Weekly Returns			Monthly Returns		
	1	2	3	4	5	6
Intercept	0.0020	0.0016	0.0019	0.0078*	0.0082	0.0093
	(1.65)	(0.60)	(0.77)	(1.83)	(0.91)	(1.07)
Surface Dev	0.0035***	0.0036***	0.0030***	0.0115***	0.0118***	0.0110***
	(5.57)	(7.10)	(4.52)	(5.27)	(6.94)	(4.90)
Ret_{1w}		-0.0209***	-0.0205***		-0.0282***	-0.0294***
		(-6.22)	(-6.31)		(-3.67)	(-4.06)
Ret_{6M}		-0.0002	-0.0001		-0.0017	-0.0014
		(-0.44)	(-0.27)		(-0.91)	(-0.78)
Size		-0.0001	0.0000		-0.0004	-0.0002
		(-0.24)	(0.09)		(-0.50)	(-0.19)
B/M		0.0008*	0.0007		0.0023	0.0017
		(1.70)	(1.53)		(1.22)	(1.01)
Amihud		-0.0028	-0.0030		-0.0117	-0.0140
		(-0.43)	(-0.49)		(-0.54)	(-0.65)
IVTimeSlope			0.0008			0.0031
			(0.33)			(0.41)
IVSkew			-0.0023			-0.0038
			(-1.24)			(-0.63)
IVSpread			0.0051			-0.0047
			(1.56)			(-0.42)
O/S			-0.0001			-0.0004
			(-1.49)			(-1.42)
Adj. R^2	0.0054	0.0432	0.0508	0.0057	0.0467	0.0547

Table 6: Portfolio returns

This table reports weekly and monthly equal-weighted and value-weighted stock returns for decile portfolios formed on the surface deviation (Panel A) and for the first principal component (Panel B). The portfolio alpha is the intercept of the regression of portfolio excess return on the Fama-French-Carhart factors. We report the t-statistics in parenthesis that are based on robust standard errors and the significance levels are indicated with * (10%), ** (5%), *** (1%).

Panel A Surface deviation

	Weekly Returns				Monthly Returns			
	EW		VW		EW		VW	
Decile	Return	Alpha	Return	Alpha	Return	Alpha	Return	Alpha
1	0.01%	-0.15%	0.01%	-0.15%	0.11%	-0.54%	0.11%	-0.54%
	(0.03)	(-2.13)	(0.04)	(-2.07)	(0.32)	(-3.18)	(0.32)	(-3.20)
2	0.17%	0.00%	0.16%	0.00%	0.62%	-0.05%	0.60%	-0.06%
	(1.11)	(-0.04)	(1.11)	(-0.07)	(2.11)	(-0.36)	(2.08)	(-0.45)
3	0.22%	0.05%	0.21%	0.05%	0.80%	0.12%	0.77%	0.10%
	(1.63)	(1.15)	(1.60)	(1.02)	(3.06)	(0.97)	(3.03)	(0.84)
4	0.19%	0.03%	0.18%	0.02%	0.79%	0.12%	0.77%	0.11%
	(1.53)	(0.59)	(1.53)	(0.50)	(3.36)	(1.10)	(3.37)	(1.00)
5	0.18%	0.02%	0.18%	0.02%	0.84%	0.20%	0.82%	0.18%
	(1.56)	(0.49)	(1.59)	(0.50)	(3.82)	(1.81)	(3.83)	(1.66)
6	0.20%	0.04%	0.19%	0.03%	0.78%	0.13%	0.74%	0.10%
	(1.82)	(1.01)	(1.79)	(0.87)	(3.77)	(1.26)	(3.72)	(0.98)
7	0.18%	0.02%	0.17%	0.02%	0.76%	0.13%	0.72%	0.10%
	(1.72)	(0.62)	(1.73)	(0.55)	(3.87)	(1.28)	(3.84)	(1.06)
8	0.22%	0.07%	0.21%	0.06%	0.87%	0.24%	0.83%	0.22%
	(2.23)	(1.96)	(2.21)	(1.77)	(4.49)	(2.41)	(4.51)	(2.24)
9	0.22%	0.06%	0.20%	0.06%	0.88%	0.28%	0.85%	0.25%
	(2.09)	(1.60)	(2.07)	(1.44)	(4.41)	(2.65)	(4.46)	(2.50)
10	0.40%	0.25%	0.38%	0.23%	1.37%	0.79%	1.31%	0.72%
	(3.03)	(4.59)	(2.97)	(4.36)	(5.12)	(5.51)	(5.12)	(5.28)
10-1	0.39%	0.40%	0.37%	0.38%	1.27%	1.32%	1.20%	1.25%
	(5.65)	(6.81)	(4.91)	(6.01)	(9.86)	(11.71)	(8.71)	(10.58)

Panel B PC1

	Weekly Returns				Monthly Returns			
	EW		VW		EW		VW	
Decile	Return	Alpha	Return	Alpha	Return	Alpha	Return	Alpha
1	0.07%	-0.10%	0.07%	-0.10%	0.39%	-0.32%	0.39%	-0.32%
	(0.52)	(-1.93)	(0.51)	(-1.86)	(1.36)	(-2.29)	(1.38)	(-2.39)
2	0.14%	-0.02%	0.14%	-0.02%	0.70%	0.03%	0.67%	0.00%
	(1.19)	(-0.47)	(1.17)	(-0.44)	(2.93)	(0.24)	(2.90)	(0.02)
3	0.15%	0.00%	0.15%	0.00%	0.67%	0.01%	0.65%	0.00%
	(1.40)	(-0.03)	(1.38)	(-0.00)	(3.10)	(0.12)	(3.12)	(0.03)
4	0.15%	0.00%	0.16%	0.00%	0.67%	0.02%	0.65%	0.02%
	(1.43)	(-0.12)	(1.45)	(0.00)	(3.22)	(0.20)	(3.29)	(0.18)
5	0.19%	0.03%	0.19%	0.03%	0.78%	0.14%	0.75%	0.12%
	(1.78)	(0.76)	(1.72)	(0.70)	(3.82)	(1.38)	(3.84)	(1.22)
6	0.20%	0.04%	0.21%	0.05%	0.83%	0.18%	0.80%	0.16%
	(1.81)	(0.93)	(1.84)	(1.13)	(3.92)	(1.68)	(3.95)	(1.58)
7	0.21%	0.05%	0.22%	0.06%	0.87%	0.25%	0.85%	0.23%
	(1.86)	(1.23)	(1.89)	(1.37)	(3.97)	(2.29)	(3.99)	(2.19)
8	0.23%	0.08%	0.23%	0.08%	0.85%	0.24%	0.83%	0.23%
	(1.93)	(1.88)	(1.90)	(1.81)	(3.56)	(1.97)	(3.60)	(1.99)
9	0.23%	0.09%	0.24%	0.09%	0.90%	0.29%	0.86%	0.27%
	(1.74)	(1.74)	(1.77)	(1.78)	(3.37)	(2.22)	(3.30)	(2.10)
10	0.33%	0.19%	0.36%	0.21%	1.16%	0.57%	1.09%	0.50%
	(2.22)	(3.09)	(2.37)	(3.48)	(3.76)	(3.63)	(3.60)	(3.23)
10-1	0.26%	0.29%	0.28%	0.31%	0.77%	0.89%	0.71%	0.82%
	(5.96)	(6.96)	(6.65)	(7.60)	(8.40)	(9.96)	(7.53)	(9.13)

Table 7: Portfolio transition probabilities

This table reports the monthly transition probabilities for the portfolios sorted on the surface deviation. A stock can remain in the same portfolio, move one portfolio up or down, or move by more than one portfolio the next month. We report probabilities for all portfolios and separately for portfolio ten (that is responsible for most for stock return predictability). E.g., a stock is 59% likely to remain in the 10^{th} portfolio the next month.

	Portfolios	
	All	10th
Move up one portfolio, ↑	14%	0%
Remain in same portfolio, =	24%	59%
Move down one portfolio, ↓	14%	17%
Move to other portfolios, ↔	48%	24%

Table 8: Out-of-sample predictability

This table presents the Fama-MacBeth regressions of weekly and monthly stock returns on the surface deviation. To construct the surface deviation, we use a 3-year rolling window of implied volatility surfaces to compute surface principal components. The surface deviation is the sum of the first three PC exposures. The remaining variables are the IV time slope (the difference between long and short term ATM implied volatilities), the implied volatility spread and skew, option-to-stock volume ratio, last-week return (Ret_{lw}), lagged six-month return (Ret_{6M}), market-capitalization (Size), book-to-market (BE/ME), and Amihud illiquidity measure. We skip a day between return predictors and stock return formation. We report the t-statistics in parenthesis that are based on robust standard errors and the significance levels are indicated with * (10%), ** (5%), *** (1%).

	Weekly Returns			Monthly Returns		
	1	2	3	4	5	6
Intercept	0.0018	0.0015	0.0013	0.0073*	0.0080	0.0072
	(1.53)	(0.56)	(0.47)	(1.73)	(0.88)	(0.77)
Surface Dev	0.0027***	0.0029***	0.0029***	0.0092***	0.0095***	0.0104***
	(3.96)	(5.19)	(3.84)	(3.84)	(5.00)	(3.93)
IVTimeSlope			-0.0020			-0.0082
			(-0.72)			(-0.89)
IVSkew			-0.0015			-0.0018
			(-0.70)			(-0.24)
IVSpread			0.0068**			0.0007
			(2.18)			(0.06)
O/S			-0.0001			-0.0002
			(-1.58)			(-1.53)
Ret_{1w}		-0.0211***	-0.0206***		-0.0292***	-0.0295***
		(-6.34)	(-6.35)		(-3.85)	(-4.06)
Ret_{6M}		-0.0002	-0.0002		-0.0017	-0.0014
		(-0.44)	(-0.32)		(-0.90)	(-0.80)
Size		-0.0001	0.0000		-0.0004	-0.0003
		(-0.25)	(0.02)		(-0.55)	(-0.32)
B/M		0.0009*	0.0007		0.0024	0.0020
		(1.77)	(1.59)		(1.30)	(1.12)
Amihud		-0.0026	-0.0023		-0.0115	-0.0120
		(-0.41)	(-0.37)		(-0.53)	(-0.55)
Adj.R2	0.0066	0.0439	0.0492	0.0067	0.0473	0.0531

Table 9: Surface deviation vs. option-based predictors

This table presents the Fama-Macbeth results from future stock returns on existing option implied predictors and the aggregate implied volatility (IV) principal component (PC). Panel A reports next-week returns and Panel B next-month returns. Independent variables include the implied volatility spread and skew, option-to-stock volume ratio. Each regression also controls for last-week return (Ret_{lw}), lagged six-month return (Ret_{6M}), market-capitalization (Size), book-to-market (BE/ME), and Amihud illiquidity measure. We report the t-statistics in parenthesis that are based on robust standard errors.

Panel A: Weekly returns

	1		2		3		4		5	
X:	IV Spread		IV Skew		RN Skew		RN Kurt		O/S	
Intercept	0.001	0.001	0.0015	0.0020	0.000	0.001	0.001	0.002	0.0011	0.0014
	(0.33)	(0.56)	(0.66)	(0.85)	(0.11)	(0.42)	(0.26)	(0.49)	(0.41)	(0.55)
X	0.009	-0.003	-0.0116	-0.0031	0.003	0.001	-0.000	-0.000	-0.0001	-0.0001
	(4.69)	(-0.83)	(-3.85)	(-0.93)	(3.27)	(1.53)	(-0.60)	(-0.57)	(-2.41)	(-1.87)
Surface Dev		0.003		0.0026		0.003		0.003		0.0024
		(4.47)		(2.89)		(3.60)		(4.29)		(6.52)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj.R^2	0.049	0.055	0.0584	0.0654	0.057	0.063	0.054	0.06	0.0488	0.0486

Panel B: Monthly returns

	1		2		3		4		5	
X:	IV Spread		IV Skew		RN Skew		RN Kurt		O/S	
Intercept	0.006	0.008	0.0069	0.0085	0.004	0.007	0.007	0.009	0.0062	0.0076
	(0.68)	(0.93)	(0.89)	(1.11)	(0.45)	(0.69)	(0.56)	(0.79)	(0.75)	(0.92)
X	0.031	-0.010	-0.0374	-0.0071	0.009	0.005	-0.001	-0.001	-0.0003	-0.0002
	(6.18)	(-1.14)	(-4.43)	(-0.78)	(3.35)	(1.77)	(-1.17)	(-1.14)	(-2.52)	(-1.93)
Surface Dev		0.011		0.0091		0.010		0.011		0.0081
		(4.99)		(3.49)		(3.71)		(4.49)		(7.21)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj.R^2	0.049	0.055	0.0549	0.0549	0.06	0.066	0.058	0.063	0.0488	0.0486

Table 10: Why option prices predict stock returns?

This table explores alternative explanations for why option prices predict stock returns. PIN (probability of informed trading), risk-neutral up and down risk are proxies for these explanations. The table reports Fama-MacBeth regressions for weekly and monthly returns. Each regression also controls for last-week return (Ret_{lw}), lagged six-month return (Ret_{6M}), market-capitalization (Size), book-to-market (BE/ME), and Amihud illiquidity measure. We report the t-statistics in parenthesis that are based on robust standard errors and the significance levels are indicated with * (10%), ** (5%), *** (1%).

Panel A: Weekly Return Regressions

X:		PIN		Risk-Neutral Jump Up		Risk-Neutral Jump Down	
	1	2	3	4	5	6	7
Intercept	0.0017	0.0035	0.0030	0.0036**	0.0026	0.0040**	0.0029*
	(1.48)	(0.89)	(0.75)	(2.23)	(1.58)	(2.50)	(1.81)
X		-0.0012	-0.0020	-0.0015	-0.0011	-0.0024	-0.0017
		(-0.23)	(-0.38)	(-0.72)	(-0.53)	(-0.99)	(-0.69)
Surface Dev	0.0035***		0.0034***		0.0037***		0.0035***
	(5.98)		(4.30)		(5.72)		(5.90)
Adj.R2	0.0151	0.0556	0.0568	0.0697	0.0711	0.0701	0.0712
# Obs	1,914,768	1,510,167	1,510,167	1,510,167	1,510,167	1,510,167	1,510,167

Panel B: Monthly Return Regressions

X:		PIN		Risk-Neutral Jump Up		Risk-Neutral Jump Down	
	1	2	3	4	5	6	7
Intercept	0.0071*	0.0135	0.0115	0.0165***	0.0129**	0.0166***	0.0128**
	(1.76)	(0.99)	(0.82)	(2.90)	(2.26)	(2.93)	(2.28)
X		-0.0012	-0.0042	-0.0082	-0.0069	-0.0108	-0.0086
		(-0.07)	(-0.23)	(-1.08)	(-0.92)	(-1.24)	(-0.99)
Surface Dev	0.0115***		0.0137***		0.0127***		0.0119***
	(5.84)		(5.01)		(5.92)		(6.39)
Adj.R2	0.0151	0.0603	0.0619	0.0713	0.0727	0.0718	0.0728
# Obs	1,914,768	1,510,167	1,510,167	1,510,167	1,510,167	1,510,167	1,510,167

Table 11: Return predictability and stock lending fee

This table presents the Fama-Macbeth regressions of weekly and monthly stock returns on the surface deviation and principal components and study how they interact with stock lending fee paid by short sellers. We report the t-statistics in parenthesis that are based on robust standard errors and the significance levels are indicated with * (10%), ** (5%), *** (1%). The lending fee data are available starting in July 2006, which reduces the sample size to 887,150 for weekly and 876,745 for monthly regressions.

	Weekly Returns				Monthly Returns			
	1	2	3	4	5	6	7	8
Intercept	0.0029*	0.0031*	0.0031*	0.0030*	0.0104*	0.0113*	0.0111*	0.0109*
	(1.65)	(1.80)	(1.75)	(1.72)	(1.71)	(1.89)	(1.83)	(1.82)
Surface Dev	0.0017***		0.0010*		0.0063***		0.0041**	
	(2.80)		(1.71)		(3.42)		(2.22)	
Lending Fee		-0.0200***	-0.0158***	-0.0111*		-0.0825***	-0.0618***	-0.0546***
		(-3.30)	(-3.05)	(-1.95)		(-3.99)	(-3.44)	(-2.72)
PC1				0.0018***				0.0043***
				(3.68)				(2.85)
PC2				0.0007				0.0047*
				(0.72)				(1.73)
PC3				0.0004				0.0022
				(0.42)				(0.67)
PC4				-0.0008				-0.0035
				(-0.59)				(-0.76)
Adj. R2	0.0046	0.004	0.0077	0.0157	0.0041	0.0043	0.0073	0.0148

Table 12: Predicting stock returns for S&P 500 constituents

This table presents the Fama-Macbeth regressions of weekly stock returns on the surface deviation and individual principal components for the sample of S&P 500 stocks. Independent variables also include last-week return (Ret_{1w}), lagged six-month return (Ret_{6M}), market-capitalization (Size), book-to-market (BE/ME), and Amihud illiquidity measure. We report the t-statistics in parenthesis that are based on robust standard errors and the significance levels are indicated with * (10%), ** (5%), *** (1%). The sample includes 446,332 observations.

Weekly Return	1	2
Intercept	0.0050*	0.0044
	(1.72)	(1.57)
Surface Dev.	0.0029***	
	(2.70)	
PC1		0.0027*
		(1.82)
PC2		0.0027**
		(2.12)
PC3		0.0039*
		(1.79)
PC4		-0.0073**
		(-2.15)
Ret_{1w}	-0.0168***	-0.0157***
	(-3.90)	(-3.81)
Ret_{6M}	-0.0000	0.0000
	(-0.03)	(0.06)
Size	-0.0004*	-0.0003
	(-1.68)	(-1.40)
B/M	0.0008	0.0007
	(1.64)	(1.46)
Amihud	-0.3356	-0.2150
	(-0.42)	(-0.28)
Adj. R2	0.0859	0.1049

Table 13: Predicting option returns

This table shows how the surface deviation predicts future options returns. We form equal-weighted and value-weighted decile portfolios based on the surface deviation and reports next-week straddle returns. Straddle returns are computed from ATM call and put options with about 30 days to expiration. We report the t-statistics in parenthesis that are based on robust standard errors and the significance levels are indicated with * (10%), ** (5%), *** (1%).

	P1	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9**	**P10**	**P10 - P1**
SurfaceDev	-0.772	-0.22	-0.1	-0.029	0.024	0.069	0.111	0.156	0.219	0.541	0.541
Equal Weighted Straddle Returns											
Average	0.017	0.006	-0.002	-0.004	-0.009	-0.01	-0.016	-0.018	-0.02	-0.026	-0.042
T-Stat	(6.20)	(2.43)	(-0.92)	(-1.66)	(-3.72)	(-4.32)	(-6.76)	(-7.22)	(-7.83)	(-7.89)	(-14.83)
StdDev	0.076	0.073	0.072	0.072	0.068	0.069	0.065	0.07	0.074	0.093	0.081
Skew	0.873	0.68	0.747	0.99	0.634	0.969	0.393	0.352	1.251	1.341	0.386
Kurt	4.10	5.18	7.49	6.99	7.02	8.45	5.10	5.30	11.21	10.48	2.26
Value Weighted Straddle Returns											
Average	0.018	0.007	-0.002	-0.003	-0.009	-0.01	-0.016	-0.018	-0.021	-0.026	-0.045
T-Stat	(6.73)	(2.73)	(-0.60)	(-1.32)	(-3.65)	(-4.20)	(-6.80)	(-7.36)	(-8.25)	(-8.14)	(-15.64)
StdDev	0.078	0.074	0.072	0.073	0.068	0.069	0.065	0.07	0.073	0.092	0.081
Skew	0.944	0.797	0.809	1.067	0.636	1.071	0.409	0.455	1.266	1.496	0.268
Kurt	4.26	5.74	7.50	6.78	6.27	8.28	4.80	5.51	11.08	11.74	1.83