



Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

A factor model for option returns

Matthias Büchner^a, Bryan Kelly^{b,c,d,*}^a Centre for Endowment Asset Management, Judge Business School, University of Cambridge, Trumpington St, Cambridge, CB2 1AG, United Kingdom^b Yale University, Yale School of Management, 165 Whitney Avenue, New Haven, CT 06511, USA^c AQR Capital Management, Two Greenwich Plaza, Greenwich, CT 06830, USA^d National Bureau of Economic Research, Cambridge MA 02138, USA

ARTICLE INFO

Article history:

Received 7 August 2021

Revised 22 December 2021

Accepted 24 December 2021

Available online 15 January 2022

JEL classification:

G11

G12

G13

Keywords:

Option return

Factor model

Return predictability

IPCA

ABSTRACT

Due to their short lifespans and migrating moneyness, options are notoriously difficult to study with the factor models commonly used to analyze the risk-return trade-off in other asset classes. Instrumented principal components analysis solves this problem by tracking contracts in terms of their pricing-relevant characteristics via time-varying latent factor loadings. We find that a model with three latent factors prices the cross-section of option returns and explains more than 85% of the variation in a panel of monthly S&P 500 option returns from 1996 to 2017. In particular, we show that the IPCA factors can be rationalized via an economically plausible three-factor model consisting of a level, slope and skew factor. Finally, out-of-sample trading strategies based on insights from the IPCA model have significant alpha over previously studied option strategies.

© 2022 Elsevier B.V. All rights reserved.

1. Introduction

Asset pricing aims to understand the risk-reward trade-off that investors face in financial markets. The most common empirical approach for evaluating this trade-off is to model returns with a low-dimensional common factor structure. The structure of option contracts makes it difficult to model their returns in this way. Instead, the literature primarily studies the risk-return trade-off in options using parametric no-arbitrage pricing models. These models require a full specification of underlying distributions and dynamics. While the parametric approach bene-

fits from arbitrage-free pricing and mathematical elegance, it is prone to model misspecification and likely too simplistic to describe empirically observed patterns of options returns. In practice, no-arbitrage models fail to account for a large part of the empirically observed variation in option returns (Israelov and Kelly, 2017).

Our objective is to develop an understanding for the risk-return trade-off in option markets using a factor pricing approach commonly applied in other asset classes. The motivation for factor modeling is independent of the asset class studied: from the asset pricing Euler equation and the assumption of no-arbitrage, a stochastic discount factor, m_{t+1} , exists that satisfies $E_t[m_{t+1}r_{i,t+1}] = 0$, and hence

$$E_t[r_{i,t+1}] = -\frac{\text{Cov}_t(m_{t+1}, r_{i,t+1})}{\text{Var}_t(m_{t+1})} \frac{\text{Var}_t(m_{t+1})}{E_t[m_{t+1}]} \quad (1)$$

For the first ratio in Eq. (1) we adopt the notation $\beta_{i,t}$ and note that it describes the conditional exposure of asset i

* Corresponding author at: Yale University, Yale School of Management, 165 Whitney Avenue, New Haven, CT 06511, USA.

E-mail address: bryan.kelly@yale.edu (B. Kelly).

URL: <http://mbuechner.com>

(M. Büchner),

<http://bryankellyacademic.org> (B. Kelly)

to systematic risk factors. For the second ratio we adopt the notation λ_t . It can be interpreted as the conditional price of risk associated with factors. When m_{t+1} is linear in factors f_{t+1} , as assumed in many asset pricing studies, the cross section of excess returns satisfies a linear factor model:

$$r_{i,t+1} = \alpha_{i,t} + \beta'_{i,t} f_{t+1} + \epsilon_{i,t+1}, \quad (2)$$

where for all i and t we have $E_t[\epsilon_{i,t+1}] = E[\epsilon_{i,t+1} f_{t+1}] = 0$, $E_t[f_{t+1}] = \lambda_t$, and $\alpha_{i,t} = 0$.

Typical factor models used in asset pricing are significantly less restrictive than parametric no-arbitrage option pricing models since they forgo the need to fully specify underlying distributions and dynamics. However, typical factor modeling approaches rely on pre-ordained factors and require a long time series of asset returns to estimate betas. This approach is difficult with options data for a few reasons. First, their short lives make it hard to estimate option betas with time series regression. Second, rapid migration of option attributes (such as moneyness and maturity) means that option risk exposures likewise migrate rapidly over time. Thus, it is important to incorporate conditional betas in option factor models, which further limits the viability of time series regression. Third, the factors are difficult to ascertain a priori. If instead one were to consider estimating a latent factor model, the same complications that make it difficult to estimate option betas via time series regression take the de facto latent factor method of principal components analysis off the table due to its reliance on static betas.

In this paper we take a different tack by treating risk factors as latent, but allow for time variation in the factor structure. We use the instrumented principal components analysis (IPCA) methodology of Kelly et al. (2019), which explicitly accounts for time variation in individual asset behavior by allowing risk factor loadings, $\beta_{i,t}$, to depend on observable asset characteristics. The characteristics serve as instrumental variables for conditional betas, which avoids the limitations of static betas in time series regression. In fact, a simple Taylor series decomposition of the option price demonstrates that option returns can be approximately explained via contract specific coefficients related to option Greeks and common factors.¹ However, any theoretically motivated decomposition is ultimately dependent on a specific model of option prices. An advantage of the IPCA approach is that it does not take a stand on the exact details of the underlying model and instead infers the relevant sources of risk from a large sample of realized option returns.

Previous literature studying the cross section of option returns has dealt with the challenges posed by options (short lifespans and time-varying risk attributes) by using empirical techniques such as portfolio sorts to construct ad hoc prespecified factors. This approach has had some success in describing option returns, but the resulting fraction of explained variation in option returns is underwhelm-

ing.² This is primarily due to standard equity-based asset pricing techniques being ill-equipped to deal with time-varying factor loadings. For example, rolling regressions are too slow to capture changes in contract risk. In contrast, IPCA offers an internally consistent approach to estimate conditional loadings and factors simultaneously. In addition, unlike their underlying assets which have an ambiguous set of asset characteristics, option contracts are special in the sense that some of their characteristics, like moneyness and maturity, are easy to measure and unambiguously relevant to the contract's factor risk. Hence, IPCA is particularly well-suited to the analysis of option returns.

1.1. Findings

As the laboratory to study the efficacy of the IPCA methodology in the context of option returns we focus on the panel of monthly S&P 500 option returns from 1996 to 2017. Most prior academic literature restricts analysis to particular subsets of option contracts, for example focusing on at-the-money contracts with one month to maturity, because it makes their analyses less sensitive to misspecification biases that are exacerbated when looking across contracts with different moneyness and maturities. In contrast, our goal is to achieve an accurate description of return variation and risk compensation for a wide range of contracts in a single model.

Taking the IPCA model to our sample of option returns, we assess the fit performance of option return factor models in two dimensions. The first is how well the estimated factors and betas capture the contemporaneous variation in realized returns. In particular, we evaluate candidate models in terms of their “total R^2 ” which measures the fraction of variance in individual option contract returns $r_{i,t+1}$ explained by $\hat{\beta}'_{i,t} \hat{f}_{t+1}$, where $\hat{\beta}_{i,t}$ are estimates of conditional loadings on estimated latent risk factors \hat{f}_{t+1} .

Second, we inspect a model's match of differences in expected returns across assets, which describes the accuracy of the model-implied risk-return trade-off. We measure this as a “predictive R^2 ” or the fraction of variance in realized returns $r_{i,t+1}$ explained by the model-implied conditional expected returns $\hat{\beta}'_{i,t} \hat{\lambda}$, where $\hat{\lambda}$ denotes the model's estimated risk prices.³ We compute both measures by aggregating over all option contracts and time periods.

We begin with a one-factor IPCA specification and with a set of characteristics consisting of option moneyness, time to maturity, implied volatility, embedded leverage, and Black-Merton-Scholes (BMS) “Greeks.” A single latent factor is sufficient to explain around 72% of the return variation in delta-hedged option returns. The prespecified factor models that we study, including variations of the Fama-French-Carhart model and two models that additionally include option-specific factors such as the Frazzini and Pedersen (2012) betting-against-beta (or embedded leverage)

² These papers include among others Coval and Shumway (2001), Goyal and Saretto (2009), Frazzini and Pedersen (2012), Zhan et al. (2021), Karakaya (2013).

³ Alternatively, one might entertain time-varying λ_t . We find that the gains from incorporating time-varying risk prices are small.

¹ We discuss this decomposition in Section 4.7.

factor and the Coval and Shumway (2001) option straddle factor, fail to achieve a total R^2 close to that of IPCA. Allowing for additional latent factors in IPCA further improves the model's ability to describe joint fluctuations in the panel of option returns. The total R^2 for the model with five factors reaches more than 90%.

Further, the IPCA model allows us to test whether average option returns are explained by factor risk or if contracts earn additional compensation (i.e., alpha) above and beyond that warranted by their risks. We find that at least three factors are necessary to price the cross-section of option returns on average. This finding is consistent with models such as Carr and Wu (2020) and Christoffersen et al. (2018) that advocate for low-dimensional multi-component models to capture option returns. However, we note that this finding does not rule out existence of significant alphas in subsets of contracts, similar to the findings of Jones (2006). Indeed, in an analysis of alphas in subsets of our sample (i.e. characteristic-managed portfolios) we find evidence for some significant managed portfolio alphas.

In terms of predictive R^2 , the IPCA model generally outperforms prespecified factor models by producing more accurate model-based option return forecasts (both in-sample and out-of-sample). The best performing prespecified factor models from earlier literature include option specific factors such as the embedded leverage factor of Frazzini and Pedersen (2012) and the Coval and Shumway (2001) straddle factor, though these underperform IPCA. When comparing IPCA with its static principal components analysis (PCA), we find that PCA is incapable of matching the behavior of option returns due to its counterfactual requirement of static betas for individual contracts. This highlights the important role played by time-varying factor loadings in IPCA. Since the individual option time series are short (only a handful of monthly observations) it is challenging to estimate static time series regression betas. IPCA overcomes this issue by restating the cross section in terms of managed portfolios which do not suffer from the extreme panel unbalancedness that plagues individual contracts. Yet IPCA still exploits vital cross-sectional variation related to individual contract characteristics. Compared to factor models with prespecified factors, we find that IPCA provides a more accurate description of the risk-return trade-off in index options. IPCA alphas are on average about one percentage point per annum smaller (i.e., less than half as large) compared to competing models.

IPCA is a statistical factor model, thus an important (and challenging) task is interpreting the factors that IPCA recovers. We find that a benchmark IPCA model can be roughly interpreted as capturing three flavors of risk. The first is the overall *level* of the volatility surface, the second is associated with *maturity* risk and summarized by the term structure slope of the volatility surface, and the third is associated with the *moneyness* skew of the volatility surface (and hence associated with index tail risk). In fact, when we construct observable factors by sorting options into level, maturity slope, and moneyness skew factors, we find that these explain between 70% and 90% of the time series variation in our set of three IPCA factors.

However, we stress that this factor interpretation is only a coarse approximation to IPCA. The fact that IPCA seeks statistically optimal factors, rather than relying on some ad hoc construction, is a key source of its superior and robust performance. And, as we learn from the IPCA model, time-varying loadings on the three observable factors are critical in order to come anywhere near the explanatory power of the IPCA model.

We examine which characteristics are the most important contributors to the accurate fits of IPCA. Option implied volatility and vega matter most for describing time variation in betas and produce the largest contributions to total R^2 . Option gamma, a measure of sensitivity to jump risk, is another important driver of options' factor betas. In addition to studying the importance of characteristics in our baseline model specification, we assess our choice of characteristics vis-à-vis an economically motivated list of characteristics from an option return decomposition. We find that all of our results are robust to the inclusion of additional option characteristics.

While most of our analysis uses monthly data, we also evaluate the IPCA model performance for daily returns. This analysis is motivated by Carr and Wu (2020), who propose a decomposition of daily option returns using a BMS approximation that attributes return performance to each of the BMS Greeks. Their attribution accurately describes returns of short-dated at-the-money contracts (where the approximation is most appropriate), but it struggles to describe returns of out-of-the money contracts and at longer maturities. We show that, in contrast to the Carr and Wu (2020) attribution, IPCA offers a uniformly accurate description of risk and return for options throughout the moneyness and maturity spectrum. A further robustness exercise demonstrates that IPCA still performs well when including contracts with up to two years in maturity.

1.2. Literature

Our paper relates to three main strands of literature. The first seeks to understand option returns by sorting options into portfolios based on a small number of characteristics, then measuring the full sample average returns of these portfolios. This literature includes among others Coval and Shumway (2001), Bakshi and Kapadia (2003), Goyal and Saretto (2009), Frazzini and Pedersen (2012), Cao and Han (2013), Karakaya (2013), Zhan et al. (2021), and Vasquez (2017).

A second strand of literature models option prices directly by enforcing no-arbitrage restrictions and specifying the full distributional properties and dynamics of the underlying asset. Starting from the seminal Black and Scholes (1973) model for vanilla options, this literature develops refinements that allow for various forms of stochastic volatility and jumps (Heston, 1993; Duffie et al., 2000; Carr and Wu, 2004).⁴ While this approach has the advantage of imposing economically meaningful no-arbitrage restrictions that guarantee consistent pricing across strikes and

⁴ Further papers studying models with price and volatility jump specifications are Eraker et al. (2003), Eraker (2004), Broadie et al. (2007).

maturities, they ultimately lead the researcher to sacrifice some realism for the sake of mathematical tractability. In fact, earlier empirical research demonstrates that arbitrage opportunities in option markets exist.⁵ As a result these models often have difficulty in matching the empirical behavior of option returns (Israelov and Kelly, 2017).

The third, and most closely related, literature acknowledges the limitations of no-arbitrage option pricing models, and directly models option returns with prespecified or latent factors. Jones (2006) estimates non-linear factor models for short-term deep out-of-the-money S&P 500 index options, allowing for potentially latent factors that eventually manifest as volatility and jump risks factors. Karakaya (2013) is focused on single-name equity options and proposes a three factor model consisting of a level, slope and value factor that captures much of the variation in delta-hedged option returns. Israelov and Kelly (2017) propose an approach for estimating conditional option return distributions using semi-parametric time series techniques. Carr and Wu (2020) develop a valuation framework that attributes daily option returns to variation in the first and second moments of the underlying price and underlying volatility. Brooks, Chance and Shafaati (2018) employ the LASSO estimator in the context of individual equity options to uncover the characteristics that provide independent information for the cross-section of option returns. Christoffersen, Fournier and Jacobs (2018) investigate a factor structure in single-name equity options using principal components of equity volatility, skews and term structures and document that these principal components explain a sizable fraction of the cross-sectional variation in option returns. In an analysis of single-name equity options Horenstein, Vasquez and Xiao (2020) use an asymptotic principal components approach and find strong evidence for a factor structure in equity option returns that is partly related to individual firm characteristics and partly to aggregate stock market risk.⁶

Our paper differs from previous literature by proposing an internally consistent factor-based approach to modeling option returns. It tackles the challenges of estimating factor betas by leveraging contract characteristics as conditioning instruments. Apart from delivering a richer factor model for returns through characteristics, our IPCA approach estimates latent factors without ex ante knowledge of the cross-section of returns, hence eliminating the need for the researcher to take a prior stance on the nature of the factors. And, in contrast to the prior literature, our approach models conditional factor loadings, which is critical due to the rapidly evolving risks at the individual option level. Notably, our approach does not enforce no-arbitrage across contracts, which allows us more flexibility in modeling the cross-section of index option returns and results in a better fit for realized returns. Importantly, the benefits of this flexibility are *not* an artifact of overfit. We show

that our model continues to excel in purely out-of-sample assessments.

While IPCA confers many advantages for modeling options returns, it also has some limitations. First, because factors are latent, it can be difficult to ascertain the interpretation of risks captured by our factor estimates. Second, IPCA does not impose no-arbitrage restrictions with the same stringency as more traditional option pricing models. Third, IPCA assumes that risk exposures are linear in option characteristics. Model improvements in each of these dimensions are interesting avenues for future work. For example, building in part on our paper, Fournier, Jacobs and Orłowski (2021) propose a non-parametric conditional factor model of option returns (assuming that factors are observable) to accommodate richer functional forms for risk exposures.

The paper proceeds as follows. Section 2 introduces the data and variables used in the analysis, Section 3 recaps the instrumented principal components methodology, Section 4 summarizes our empirical results, and Section 5 concludes.

2. Data

In this paper we focus on options on the S&P 500 index.⁷ Daily option data is obtained from OptionMetrics for the period of January 1996 to December 2017. The information provided by OptionMetrics includes contract specifications (exercise date, strike, etc.) as well as underlying index values, historical dividend yields, and option sensitivity measures such as the BMS delta, gamma, vega, and theta. Data for the VIX index is obtained through CBOE.

To introduce notation, each option contract i is defined by its strike price K_i and maturity date T_i (time-to-maturity is $T_i - t$). To measure the moneyness of a contract we use the forward BMS delta, $\Delta_{i,t}^{Fwd} = \Delta_{i,t} \cdot e^{d_{i,t}(T_i-t)/365}$, where $\Delta_{i,t}$ is the standard BMS delta and $d_{i,t}$ denotes the annualized dividend yield at time t . The forward delta has the advantage that an exactly at-the-money option (i.e., $K_i = S_t$) attains a forward delta of 0.5 in absolute value.

Our results focus on monthly holding period returns that are delta-hedged daily.⁸ Option returns are computed for periods defined by the expiration date in a given month which usually is the third Friday in a month, i.e. we compute monthly holding period returns from daily data over periods starting on the next trading day after the expiration date and ending with the expiration date.

Variation in the price of the underlying is the most important driver of option returns. By delta-hedging the option contracts daily, we obtain the portion of the option return that is not explained by local-linear exposure to the underlying index return. This choice is common in the academic context (e.g., see Cao and Han, 2013). The delta-hedged profit-and-loss (P&L) for a contract with value F

⁵ For example, see Ofek and Richardson (2003), Ofek et al. (2004), Constantinides et al. (2009), Chambers et al. (2014).

⁶ Our paper also relates to recent papers that apply IPCA in different asset classes, including Windmüller (2021) who studies international equities and Kelly, Palhares and Pruitt (2020) who study corporate bonds.

⁷ Note, the framework employed here also lends itself to the study of single-name options in a characteristics-rich environment. However, for the sake of simplicity and given their particular relevance to portfolio risk, in this paper, we focus on index options.

⁸ In Section 4.8, we perform a robustness exercise and assess the performance of our model at daily frequency.

over a period $t = 1, \dots, T$ is given by

$$\begin{aligned} \Pi_{[1,T]} = & \sum_{t=1}^{T-1} (F_{t+1} - F_t) - \sum_{t=1}^{T-1} \Delta_t (S_{t+1} - S_t) \\ & - \sum_{t=1}^{T-1} \frac{a_{t,t+1} r_t}{365} (F_t - \Delta_t S_t), \end{aligned} \quad (3)$$

where the first term is the raw P&L, the second term captures the adjustment from delta-hedging the position, and the last term adjusts for the cost of funding the delta-hedged portfolio at the risk-free rate where $a_{t,t+1}$ is the number of days between trading dates t and $t + 1$.

We compute returns not against the prevailing option mid-price, but against the mid-price of the underlying. This ensures that returns are well-behaved even in situations when the option mid-price is close to zero as is common for deep out-of-the-money (OTM) options. We denote the delta-hedged return against the spot price by r_{Spot}^Δ .

We apply a number of filters to the OptionMetrics data. We exclude observations in which i) the bid price is negative, ii) the bid exceeds the ask, iii) no-arbitrage conditions are violated, or iv) the OptionMetrics implied volatility is missing. To ensure the reliability of option price data, we only study observations for contracts with positive open interest. Following Karakaya (2013), we exclude observations with extreme embedded leverage by trimming data below (above) the 1st (99th) percentile of the embedded leverage distribution, where embedded leverage is defined as

$$\Omega = |\Delta \cdot S/F|. \quad (4)$$

Finally, we restrict our sample to call options with forward delta of 0.01 to 0.5 and -0.5 to -0.01 for put options, and require time-to-maturity (TTM) of one to 12 months. Focusing on this range of moneyness and maturity ensures the reliability of contract prices and follows the filters of (Israelov and Kelly, 2017). Outside of this range, contracts become severely illiquid which introduces noise in returns data. Figure 1 demonstrates how the trading activity in index options is concentrated in our sample range. Also note that, by put-call parity, the boundaries of ± 0.5 on our moneyness range are not restrictive. Although, we “limit” our sample in the way described above, it encompasses a much larger fraction of the outstanding contracts than many papers in the literature which often limit their sample to ATM contracts. In robustness tests, we also assess the validity of our findings for contracts with longer maturities up to two years (see Section 4.7).

To conclude the exposition of our data sample, Table 1 presents summary statistics. After applying all aforementioned filters, we arrive at a sample of approximately 77,000 option-month observations. Out of those observations a little more than two-thirds are put options. Since over the course of our sample period the underlying S&P 500 index appreciates considerably OTM put options are naturally more abundant than OTM calls. In addition, the put/call ratio for options on the S&P 500 is commonly higher than one as market participants demand more put options for insurance purposes.

Construction of the Implied Volatility Surface For analysis of the Carr and Wu (2020) P&L attribution in Section 4.8,

we construct an implied volatility surface to obtain moments of the underlying and volatility under the risk-neutral measure. We construct implied volatility surfaces using the daily OptionMetrics volatility file that contains volatilities over a prespecified moneyness and maturity grid. For the construction of the volatility surface we carry out a second order smooth bivariate spline interpolation in the TTM vs. moneyness space.⁹ Since it is common for implied volatilities of put and call options at the same moneyness and maturity to disagree as a consequence of differences in price pressure, we construct three different volatility surfaces: the first is constructed solely from observations of put options in the OptionMetrics volatility files, the second is constructed from observations of call options, and the third is constructed using both put and calls. In the latter case, the implied volatilities of put and call options are weighted by their market capitalization (defined as open interest times option mid-price).

3. Methodology

In this paper, we aim to uncover a factor structure in empirical option returns. To this end, we use the instrumented principal components model of Kelly et al. (2019, 2020). The model is specified for a general excess return $r_{i,t+1}$

$$\begin{aligned} r_{i,t+1} = & \alpha_{i,t} + \beta_{i,t} f_{t+1} + \epsilon_{i,t+1} \\ \alpha_{i,t} = & z'_{i,t} \Gamma_\alpha + v_{\alpha,i,t}, \quad \beta_{i,t} = z'_{i,t} \Gamma_\beta + v_{\beta,i,t} \end{aligned} \quad (5)$$


The system is estimated over a total of N assets and T periods. The loadings, $\beta_{i,t}$, are time-varying and partially depend on an $L \times 1$ vector of (option) characteristics $z_{i,t}$. We assume that $z_{i,t}$ includes a constant. The vector of factors, f_{t+1} , is dimension $K \times 1$ where the number of factors. Following Kelly et al. (2019), the IPCA model can be estimated by means of an alternating least squares procedure that iterates between the first order conditions of Γ^{10} and f_{t+1} .¹¹

There are a number of features that make IPCA ideal for analyzing the factor structure in options returns. First, observable characteristics feature centrally in the IPCA framework as they make the estimation of factors more efficient and improve model performance. Contrary to common equity, option contracts are defined through a set of precisely measured characteristics such as their strike, time-to-maturity, implied volatility, their location on the implied volatility surface, and dividend yields. Thanks to the dimensionality reduction embedded in IPCA, a large number of potential characteristics can be considered simultaneously. IPCA forms a linear mapping Γ_β between characteristics and factors that isolates the informative signal coming from characteristics and averages out the noise. This eliminates the need for the researcher to take an ad

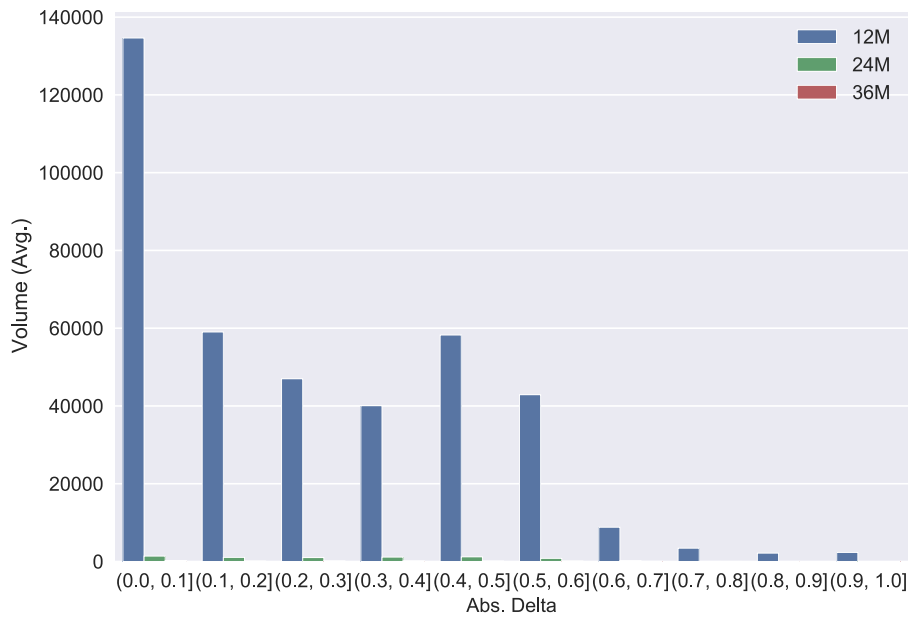
⁹ For details, see the SciPy function `SmoothBivariateSpline`, documented at <https://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.SmoothBivariateSpline.html>.

¹⁰ For brevity, we denote Γ_α and Γ_β jointly as just Γ . The IPCA model can equivalently be stated as $r_{i,t+1} = z'_{i,t} \tilde{\Gamma}_{t+1} + \epsilon_{i,t+1}$ where $\tilde{\Gamma} = [\Gamma_\beta, \Gamma_\alpha]$ is a horizontally stacked matrix and $\tilde{f}_{t+1} = [f_{t+1}, 1]$.

¹¹ For the Python implementation used to carry out the estimation we refer the reader to <https://github.com/bkelly-lab/ipca>.



(a) Total Open Interest



(b) Traded Volume

Fig. 1. Open Interest and Volume of S&P 500 Options. This figure shows the average total open interest and average traded volume for contracts sorted on their absolute forward BMS delta and time-to-maturity on the monthly rebalancing dates used in our study. The sample period is 1996 to 2017.

Table 1

Summary Statistics of Option Level Variables. This table provides summary statistics for the sample of option returns used in the analysis. Panel A contains results for call option contracts, while Panel B contains results for put option contracts. The sample period is 1996 to 2017. For option level variables the table shows the time-to-maturity *t*tm, moneyness *m*nness computed as $mness = \ln(K/S)/(IV \cdot \sqrt{ttm})$ with strike *K* and spot price *S*, Black-Scholes implied volatility *I*V, option delta, option hamma, option vega, option theta, and the annualized delta-hedged return versus spot price r_{Spot}^Δ .

Panel A: Call Option Contracts									
	t	m	e	i	d	g	v	t	r_{Spot}^Δ
Mean	129	1.05	32.02	0.16	0.20	0.002	196.85	-55.10	-3.14%
Median	91	0.97	25.99	0.15	0.18	0.002	155.54	-44.19	-2.31%
Std. Dev.	99	0.67	21.00	0.07	0.15	0.002	159.82	45.86	1.33%
No. Obs.	24,749	24,749	24,749	24,749	24,749	24,749	24,749	24,749	24,749
Panel B: Put Option Contracts									
	t	m	e	i	d	g	v	t	r_{Spot}^Δ
Mean	123	-1.16	19.10	0.26	-0.15	0.001	168.03	-70.60	-5.18%
Median	91	-1.20	16.73	0.24	-0.10	0.001	123.02	-59.06	-4.26%
Std. Dev.	99	0.66	11.08	0.10	0.14	0.002	148.34	48.94	1.48%
No. Obs.	52,341	52,341	52,341	52,341	52,341	52,341	52,341	52,341	52,341

hoc stance on the characteristics that matter. Second, time-varying loadings $\beta_{i,t}$ that depend on characteristics $z_{i,t}$ allow us to model the conditional behavior of option returns. The “identity” of an option contract is to a large extent pinned down by its location in TTM/moneyness space. IPCA tracks the migration of the asset in this space through the conditional betas. Third, by restricting the IPCA model (5) such that $\Gamma_\alpha = 0$, we can test whether risk compensation in option returns solely arises from exposure to systematic factors, f_t , or whether returns partially line up with characteristics directly (i.e. $\Gamma_\alpha \neq 0$), hence constituting compensation without risk.

Asset pricing performance To assess the performance of IPCA factors in pricing option returns, we use two measures following Kelly, Pruitt and Su (2019). Similar to a common R^2 goodness-of-fit measure we compute

$$R_{total}^2 = 1 - \frac{\sum_{i,t} (r_{i,t+1} - z'_{i,t}(\hat{\Gamma}_\alpha + \hat{\Gamma}_\beta \hat{f}_{t+1}))^2}{\sum_{i,t} r_{i,t+1}^2} \tag{6}$$

This total R^2 measures how well the set of factors and loadings captures realized returns. As a second measure, we compute

$$R_{pred}^2 = 1 - \frac{\sum_{i,t} (r_{i,t+1} - z'_{i,t}(\hat{\Gamma}_\alpha + \hat{\Gamma}_\beta \hat{\lambda}))^2}{\sum_{i,t} r_{i,t+1}^2} \tag{7}$$

where $\hat{\lambda}$ denotes the unconditional time-series mean of the factors. This predictive R^2 captures how well differences in average returns are explained through the model’s description of conditional expected returns, i.e. the model’s ability to describe risk.

In addition to studying the pricing performance for individual option contracts, it is insightful to examine the pricing performance for portfolios. The IPCA methodology incorporates a portfolio notion that circumvents a common problem in the asset pricing literature that is the choice of relevant test assets. IPCA is tightly linked to the idea of characteristic-managed portfolios. Let Z_t be an $N \times L$ matrix of characteristics at time t . Then managed portfolios

can be constructed via

$$x_{t+1} = \frac{Z'_t r_{t+1}}{N_{t+1}} \tag{8}$$

where N_{t+1} is the number of outstanding options at time $t + 1$. The managed portfolios, x_{t+1} , are a weighted average of option returns where the weights are determined by the characteristics in Z_t . Analogously to total and predictive R^2 at the individual option level, we can define performance measures for managed portfolios

$$R_{total,x}^2 = 1 - \frac{\sum_{l,t} (x_{l,t+1} - z'_{l,t} z_{l,t} (\hat{\Gamma}_\alpha + \hat{\Gamma}_\beta \hat{f}_{t+1}))^2}{\sum_{l,t} x_{l,t+1}^2} \tag{9}$$

$$R_{pred,x}^2 = 1 - \frac{\sum_{l,t} (x_{l,t+1} - z'_{l,t} z_{l,t} (\hat{\Gamma}_\alpha + \hat{\Gamma}_\beta \hat{\lambda}))^2}{\sum_{l,t} x_{l,t+1}^2} \tag{10}$$

where $l = 1, \dots, L$.

4. Empirical results

4.1. Data

From the data set described in Section 2, we collect a number of option characteristics and market variables. Option-level variables include the BMS forward delta (delta), time-to-maturity (ttm), embedded leverage (embed_lev), BMS implied volatility (impvol), BMS theta (theta), BMS gamma (gamma), BMS vega (vega), BMS volga (volga).¹² In all specifications we interact aforementioned variables with both a constant and an indicator variable that is equal one for put options and equal zero for calls. This interaction allows for the possibility that characteristics have different effects for calls versus puts.

¹² The sensitivity of vega to changes in volatility, here referred to as Volga = $\frac{\partial Vega}{\partial \sigma}$, is not part of the set of Greeks computed by OptionMetrics as standard and therefore computed from the Black-Scholes pricing formula.

In order to limit the impact of outliers and to aid the interpretation of the IPCA results, we re-scale all characteristics to the range $[-0.5, 0.5]$. By scaling the characteristics in this manner, we focus on the spacing of characteristics in the cross section which is the main determinant of differences in returns across contracts. In order to ensure similar leverage properties of the resulting characteristic-managed portfolios (see Eq. (8)) we enforce that resulting portfolio weights sum to zero as is common for long-short portfolios studied in the asset pricing literature.

4.2. IPCA Performance

To begin, we estimate IPCA using a baseline set of characteristics including `delta`, `ttm`, `embed_lev`, `impvol`, `theta`, `gamma` and `vega`, yielding a total of 15 characteristics after interaction with the put/call dummy and addition of a constant. The dependent variable is the monthly holding-period return that is delta-hedged daily, r_{Spot}^Δ , as defined in Section 2.

Table 8 details the performance of the restricted ($\Gamma_\alpha = 0$) and unrestricted ($\Gamma_\alpha \neq 0$) IPCA models with $K = 1, \dots, 5$ factors. With a single factor, IPCA explains more than 72% of the observed variation in option returns in the restricted model ($\Gamma_\alpha = 0$) and around 74% in the unrestricted model ($\Gamma_\alpha \neq 0$). When including an additional IPCA factor, the R_{total}^2 increases to around 80% for restricted and 81% for the unrestricted specification. We find that five factors are needed to explain more than 90% of the variation in monthly delta-hedged index option returns.

IPCA's least squares objective function directly targets the total R^2 . However, the IPCA factors do not explicitly target a high predictive R^2 . We find that the monthly predictive R^2 for a single factor is 5.47% in the restricted model and 7.59% in the unrestricted model. When including additional factors, the predictive R^2 increases in the restricted model, while it decreases slightly in the unrestricted model.¹³ A natural comparison for this result is a panel predictive regression of option returns on the same set of characteristics as used in IPCA. These regressions are identical to the estimation of IPCA with a single pre-specified constant factor. Table 2 details the results of the panel regression. We find that characteristics on their own predict only around 3.15% of the empirically observed variation in returns. Even adding time fixed effects, the R^2 is only 6.23%. In other words, IPCA simultaneously excels in explaining both realized option returns and future expected option returns.

It is common practice in the empirical asset pricing literature to examine the explanatory power of asset pricing models using portfolios such as, for example, the 5×5 Fama-French size and book-to-market sorted portfolios as test assets.¹⁴ The IPCA framework can be approximately stated in terms of managed portfolios $x_t = Z'_{t-1} r_t / N_{t-1}$ (see

Table 2

Panel Regression of Option Returns on Option Characteristics. The dependent variable is the monthly delta-hedged return versus the prevailing spot price r_{Spot}^Δ . Regression specification (1) uses a common intercept for all observations, while specification (2) uses time-fixed effects. As regressors we include the Black-Merton-Scholes (BMS) delta, gamma, vega, theta, time-to-maturity `ttm`, and the BMS implied volatility (`impvol`). We interact all characteristics with a put dummy variable that has value one for put options and zero for call options. For example, the interaction of delta and the put indicator is denoted `delta:put`. All predictors are standardized by their sample standard deviation. Standard errors are clustered by option contract and in parentheses we report t-statistics.

	r_{Spot}^Δ	
	(1)	(2)
delta	-0.0088 (-0.7169)	0.0560 (6.4211)
ttm	0.0244 (5.9405)	0.0327 (8.3471)
embed_lev	-2.011e-05 (-0.0069)	0.0063 (2.8575)
gamma	-0.0151 (-3.4482)	-0.0140 (-4.1134)
vega	-0.0091 (-1.8856)	-0.0254 (-6.3281)
theta	-0.0167 (-2.1788)	0.0118 (1.8536)
impvol	-0.0821 (-13.383)	-0.1467 (-14.711)
delta:put	0.0635 (3.7247)	0.0291 (2.6428)
ttm:put	-0.0141 (-3.4176)	-0.0295 (-7.2445)
embed_lev:put	0.0151 (4.4433)	0.0074 (2.1717)
gamma:put	-0.0056 (-0.8119)	0.0080 (1.7138)
vega:put	-0.0015 (-0.2155)	0.0311 (6.2716)
theta:put	0.0524 (5.7418)	0.0424 (6.0961)
impvol:put	0.0543 (11.800)	0.1109 (15.992)
Effects		Time
R^2	3.15%	6.23%
No. Obs.	77,090	77,090

Section 3), where N_t is the number of assets in the cross-section at time t . This construction yields an $L \times 1$ vector x_t , where L corresponds to the number of considered characteristics. We then compute performance measures for the managed portfolios as test assets. We notice that the R_{total}^2 is markedly higher for managed portfolios than for individual contracts. For example, using just a single factor the total R^2 is 95.48% in the unrestricted model and 94.41% in the restricted model. When moving from individual options to managed portfolios we find that the predictive R^2 increases to around 7% or 8% for both the restricted and unrestricted models.

In Panel C of Table 8 we display p-values for the hypothesis test $H_0 : \Gamma_\alpha = 0$ implemented with the bootstrap procedure outlined in Kelly et al. (2019) using 1000 bootstrap draws. For specifications with one and two factors, we reject the null hypothesis at the 90% confidence level, suggesting that the risk space spanned by one or two factor models leaves significant option alpha. However, for a model with three factors we find a pronounced jump in

¹³ In Section 4.2.1 we put these findings in context by comparing with models from previous literature.

¹⁴ The asset pricing literature has entertained a large range of potential test assets. This is a sign for the struggle that this literature has long faced in determining the set of assets that need to be priced in order for a model to be accepted.

Table 3

IPCA Performance by Bins of Option Delta, Maturity, and VIX. The table details the total R^2 for restricted IPCA models with $K = 1, \dots, 5$ by bins sorted on option characteristics. Note, the total R^2 is computed using only observations from the respective bin examined.

Panel A: Abs. Delta Bin						
Range	K=1	K=2	K=3	K=4	K=5	No. Obs
0.01 to 0.1	57.44%	75.41%	76.85%	80.96%	83.12%	34,957
0.1 to 0.2	71.11%	81.95%	85.66%	88.48%	89.58%	14,307
0.2 to 0.3	75.56%	81.05%	86.39%	91.64%	92.42%	10,393
0.3 to 0.4	75.43%	79.99%	87.11%	91.49%	92.56%	8920
0.4 to 0.5	74.24%	78.27%	85.35%	87.83%	89.69%	8513
Panel B: Time-to-Maturity Bin						
Range	K=1	K=2	K=3	K=4	K=5	No. Obs
1 Month	41.11%	52.16%	65.59%	77.27%	79.49%	16,598
2 Months	80.38%	88.61%	90.77%	92.03%	92.70%	19,087
3 to 6 Months	88.23%	92.67%	93.80%	94.40%	95.41%	22,723
6 to 12 Months	82.57%	88.31%	92.28%	93.12%	94.34%	18,682
Panel C: VIX						
Range	K=1	K=2	K=3	K=4	K=5	No. Obs
0% to 10%	25.31%	63.86%	65.96%	69.91%	78.96%	2783
10% to 20%	69.65%	77.56%	83.67%	87.63%	89.52%	47,061
20% to 30%	65.77%	73.93%	81.85%	86.66%	88.40%	20,736
30% to 90%	78.83%	84.98%	88.34%	91.47%	91.99%	6510

the p-value to 47.2%. This suggests that at least three factors are needed for the IPCA model to price the cross section of index option returns on aggregate. This is evidence that factor risk explains conditional expected option returns when $K = 3$. Jones (2006) notes that mispricing is reduced by increasing the number of factors, and we find evidence in a similar spirit: as we increase the number of factors the discrepancy between the restricted model with $\Gamma_\alpha = 0$ and the unrestricted model with $\Gamma_\alpha \neq 0$ vanishes. It is also important to point out that the bootstrap test captures the combined effects that characteristics have on returns. Hence, it is still possible that certain subsets of contracts have some alpha associated with characteristics. Jones (2006) finds that option alphas are on average close to zero, but also documents significant mispricing of specific sub-groups of option contracts such as deep OTM options. We examine mispricing of particular sub-groups of options next.

Performance in the Cross-Section of Contracts The IPCA framework is designed to target the aggregate R^2_{total} across all option contracts simultaneously. To understand the success of IPCA in capturing heterogeneity, we study behavior of contracts with different risk profiles and in different market regimes. In particular, we focus on sub-groups of contracts based on moneyness, TTM, and market volatility. For this analysis, we focus exposition on the restricted IPCA model with $K = 3$ factors (though we report results for other choices of K as well). We choose this specification based on a desire for a parsimonious model and because we fail to reject the null hypothesis of zero alpha for $K = 3$.

Table 3 breaks down IPCA performance by bins sorted on moneyness, TTM, and VIX index level. The results by bin are obtained by computing the total R^2 only from those option-month observations corresponding to a particular bin. Panel A of Table 3 shows that IPCA performance is

particularly strong in terms of both total and predictive R^2 for ATM and slightly OTM options, while the performance weakens to some extent for deep OTM options. By increasing the number of factors in IPCA, the model is able to better track contracts with low absolute deltas between 0.01 and 0.1. Panel B of Table 3 reports results sorted on the option's time-to-maturity. We find that IPCA performs best for options with maturities greater than one month in terms of the total R^2 . In order to assess the model's performance in different market environments (low vs. high volatility regimes), Panel C dissects the performance by level of the CBOE VIX index. We find that the IPCA model captures return variation across a range of market regimes, though low volatility regimes (VIX below 10%) are somewhat more challenging for the IPCA model.

Overall, these results suggest that IPCA successfully captures observed return variation not only in aggregate but also in specialized sub-groups of options. Due to their abundance in our sample, ATM and OTM contracts have a relatively higher weight in the IPCA estimation and in turn the model is especially well trained for describing these contracts. Nevertheless, option contracts with more extreme moneyness and implied volatility are reasonably well captured in an IPCA model using only standard option characteristics and a small number of factors. Furthermore, we find that allowing for more than $K = 3$ factors improves performance in these relatively small subsets of the sample, and indicating that additional factors serve to fit the differential behavior of options in the more extreme bins.

4.2.1. Comparison with extant factor models

We now turn to comparing the performance of the IPCA framework with a number of observable risk factor models entertained in earlier literature. We consider five different observable factor models. As a model with a single factor, the market factor, we use the capital asset pricing

ing model (“CAPM”). The Fama-French (1993) three factor model (“FF3”) adds the size (SMB) and value factors (HML). We add the momentum factor (UMD) to obtain the Carhart four factor model (“FFC4”). As demonstrated in earlier literature, the standard risk factors used to price equities are insufficient to price delta-hedged option returns.¹⁵ For this reason, we include two additional factors constructed specifically from S&P 500 index options.

The first is the betting-against-beta (BAB) factor of Frazzini and Pedersen (2012) constructed from embedded leverage, Ω , as defined in Eq. (4). Specifically, BAB is constructed by sorting options into a high- and low-embedded leverage portfolio using the cross-sectional median of embedded leverage as the breakpoint. We weight returns in the two portfolios using the option’s market capitalization (open interest times price). Then the self-financing, zero-beta factor return is given by

$$BAB_t = r_t^L / \Omega_{t-1}^L - r_t^H / \Omega_{t-1}^H,$$

where Ω_t^L and Ω_t^H are the weighted average embedded leverage of the low and high embedded leverage portfolios, respectively. We construct BAB separately for call and put options and then average them.

The second option risk factor we include is the straddle factor of Coval and Shumway (2001) as constructed in Frazzini and Pedersen (2012). For construction of the straddles we limit the set of considered contracts to those with 1 month to expiry. Then, we construct straddles from all pairs of put and call options with absolute BMS delta between 0.4 and 0.6. The straddle factor return is obtained by weighting the delta-hedged straddle returns with the market capitalization of the constituent options.

We add the BAB factor to the “FFC4” four factor model and refer to the resulting five factor model as “FFCB5”. Finally, we add the straddle factor to the “FFCB5” model and refer to the resulting model as “FFCBS6”.

To aid comparison of these observable risk factors with the IPCA latent factor model we allow for characteristics as instruments for time-varying factor loadings similar to IPCA. Note that the insight that dynamic betas are central to modeling option contract returns is an insight motivated by IPCA, and without doing so the observable factor models become nearly infeasible at the contract level and perform disastrously out-of-sample. We implement observable factor models by prespecifying the factors in the IPCA setting and then estimating the loadings matrix Γ_β by evaluating the associated first order condition.

Another natural model comparison for IPCA is the static PCA model without time-varying loadings. In order to estimate PCA for the individual option contracts where panel unbalancedness is a challenge, we utilize an alternating least squares scheme.¹⁶

Table 4 compares the different implementations. To provide a baseline, Panel A restates the performance

for the restricted ($\Gamma_\alpha = 0$) IPCA model with $K = 1, \dots, 5$ factors for both individual option returns as well as characteristic-managed portfolios. Panel B reports results for the implementation of IPCA with prespecified observable risk factors and characteristics as instruments. We find that the IPCA model with a single factor outperforms observable factor models in terms of the total R^2 for all combinations of observable factors that we examine. Adding the observable factors constructed from option returns (models “FFCB5” and “FFCBS6”) considerably improves the observable factor model fit. The R_{total}^2 increases by around 7% versus the Carhart model (“FFC4”) when the embedded leverage factor is included and by another 17% when the straddle factor is included, though still nearly half of the R_{total}^2 from IPCA. The best performing observable factor model, “FFCBS6,” generates a predictive R^2 similar to that of IPCA.

In Panel C, we summarize the fits from running latent factor models estimated with PCA. Similar to Kelly et al. (2019) we find that PCA is easily outperformed by IPCA when focusing on individual assets, with the predictive R^2 being negative for almost all specifications. When running PCA at the managed portfolio level the model fit matches IPCA in terms of total R^2 , while IPCA generally produces a higher predictive R^2 . This finding is plausible since at the managed portfolio level IPCA and PCA estimate the same number of coefficients. Nevertheless, the result demonstrates that IPCA is clearly superior in terms of tracking individual option contracts which are challenging to track in classical factor models without time-varying loadings given their short lifespans and swiftly changing risk attributes. Furthermore, IPCA provides the explicit modeling link between the individual underlying assets and the managed portfolios.

Taken together, this evidence suggests that the IPCA framework is superior to factor models from earlier literature. This includes both observable, prespecified factors (which benefit here from the inclusion of IPCA-like instrumented betas) and static latent factor models. IPCA better captures time variation in individual option returns as well as differences in average returns across assets. Comparing the IPCA and PCA models for individual contracts, we find clear evidence that time-varying factor loadings are essential for accurately capturing the cross section of option returns.

4.3. Unconditional and conditional alphas

So far, when testing for the significance of Γ_α we have found that for a model with $K = 3$ factors characteristics on their own do not line up with options returns. This suggests that our factors are conditionally mean-variance efficient. We further investigate alphas in the IPCA model and its competitors, both conditionally and unconditionally.

We test the model’s ability to price test portfolios unconditionally by running time-series regressions of test portfolio returns on the set of IPCA factors. As a benchmark, we pick the best performing observable factor model that includes the Fama-French three factors, momentum, the straddle factor as well as the betting-against-beta factor constructed from embedded leverage (“FFCBS6”). We

¹⁵ See, e.g., Coval and Shumway (2001), Cao and Han (2013), Frazzini and Pedersen (2012), Karakaya (2013) for factor models entertained in the earlier literature.

¹⁶ In particular, we use the MATLAB function PCA with the flag ‘Algorithm’ set to ‘als’. To aid convergence we limit the sample to contracts with at least 50% non-missing observations during their lifespan.

Table 4

IPCA versus Observable Factor Models. The table compares the total and predictive R^2 of the restricted IPCA model with $K = 1, \dots, 5$ factors (Panel A), a number of observable factor models (Panel B), and canonical principal components analysis (Panel C). In Panel B we start from the Fama-French three factor model and add the Carhart momentum factor (FFC4) the Frazzini and Pedersen (2012) 'betting against beta' factor (FFCB5), and the Covall and Shumway (2001) straddle factor (FFCBS6). Panel B contains the results when loadings are dynamic by instrumenting with characteristics. In Panel C we report the result from a standard principal components analysis (PCA) of the time-series of option returns. Given panel unbalancedness an alternating least squares algorithm is used to estimate PCA. To aid convergence we limit the sample to those contracts with at least six valid observations during the sample period. We report performance measures for individual options as test assets R^2_{total} / R^2_{pred} and for characteristics managed portfolios as test assets $R^2_{total,x} / R^2_{pred,x}$.

Panel A: IPCA					
	K				
	1	2	3	4	5
R^2_{tot}	72.28%	79.65%	85.07%	88.90%	90.22%
R^2_{pred}	5.47%	5.54%	6.39%	6.59%	6.77%
$R^2_{tot,x}$	94.41%	96.64%	98.86%	99.39%	99.61%
$R^2_{pred,x}$	7.20%	7.44%	7.90%	7.99%	8.06%

Panel B: Observable Factors - With Instruments					
	CAPM	FF3	FFC4	FFCB5	FFCBS6
R^2_{tot}	23.81%	25.64%	26.08%	33.15%	49.94%
R^2_{pred}	2.47%	2.38%	2.71%	4.45%	6.24%
$R^2_{tot,x}$	22.55%	25.83%	26.34%	32.68%	56.79%
$R^2_{pred,x}$	3.37%	3.32%	3.58%	5.60%	7.57%

Panel C: Principal Components Analysis					
	K				
	1	2	3	4	5
R^2_{tot}	18.08%	32.71%	41.70%	47.55%	52.11%
R^2_{pred}	-0.07%	-0.07%	-0.02%	-0.02%	0.15%
$R^2_{tot,x}$	94.09%	97.46%	98.74%	99.44%	99.72%
$R^2_{pred,x}$	7.29%	7.83%	7.87%	7.88%	7.90%

study two sets of anomaly portfolios. First, the managed portfolios of IPCA immediately provide us with a set of test assets that weight assets by their characteristics. Second, we use a set 20 of money and maturity double-sorted portfolios of call and put options. We re-leverage all portfolios to 10% annualized volatility for comparability. Note that by using portfolios as test assets, we are giving observable factor models a reasonable chance at competing with IPCA in terms of fit, while at the contract level observable factor models tend to perform dramatically worse than IPCA.

Panels (a) and (b) of Fig. 2 show unconditional portfolio alphas from time series regressions of managed portfolio returns on the set of factors. Significant alphas are denoted by filled in diamond markers. We find that the average absolute portfolio alpha in the IPCA model is more than a percentage point lower than that from the observable factor model FFCBS6. In panels (c) and (d) we make the same comparison in terms of average conditional alphas from the conditional beta version of each model (i.e., using IPCA in its intended form and allowing for instrumented betas on observable factors). Portfolio alphas are obtained as the time-series averages of period-by-period portfolio residuals in the conditional model. In this case we find that IPCA produces fewer significant portfolio alphas compared to FFCBS6, and in addition the average absolute alpha in IPCA is more than two percentage points smaller. Table 5 breaks out alphas for each characteristic-managed portfolio separately. This analysis shows that not

Table 5

Managed Portfolio Alphas - IPCA vs. Observable Factors. This table details the individual managed portfolios alphas in the conditional IPCA model with $K = 3$ and conditional observable factor model (FFCBS6) instrumented with the same set of characteristics. The portfolio alphas are obtained as time-series averages of the period-by-period model residuals. Absolute portfolio alphas with t-statistics greater than 2.0 are highlighted in bold print.

	IPCA		FFCBS6	
	alpha	t-stat	alpha	t-stat
delta	4.6%	5.01	6.9%	3.69
delta:put	2.4%	2.92	2.3%	1.23
ttm	0.4%	0.58	5.8%	3.11
ttm:put	1.0%	1.21	6.0%	3.25
embed_lev	0.0%	0.02	3.5%	2.04
embed_lev:put	0.2%	0.30	5.1%	2.74
theta	0.7%	1.21	2.7%	1.91
theta:put	3.3%	3.46	7.9%	5.36
impvol	-0.9%	-1.27	-3.0%	-1.31
impvol:put	0.5%	0.64	0.3%	0.16
gamma	-0.6%	-2.42	-0.4%	-0.27
gamma:put	2.2%	2.26	4.3%	3.08
vega	0.0%	0.10	-3.4%	-1.91
vega:put	1.0%	1.45	-1.4%	-0.70
const	-0.1%	-0.50	-1.8%	-1.15
Avg. Abs. Alpha	1.2%		3.7%	

only is the average IPCA alpha magnitude smaller than in the observable benchmark model, but also alphas are almost uniformly smaller across each of the characteristic-managed portfolios. Exceptions to this are delta:put and

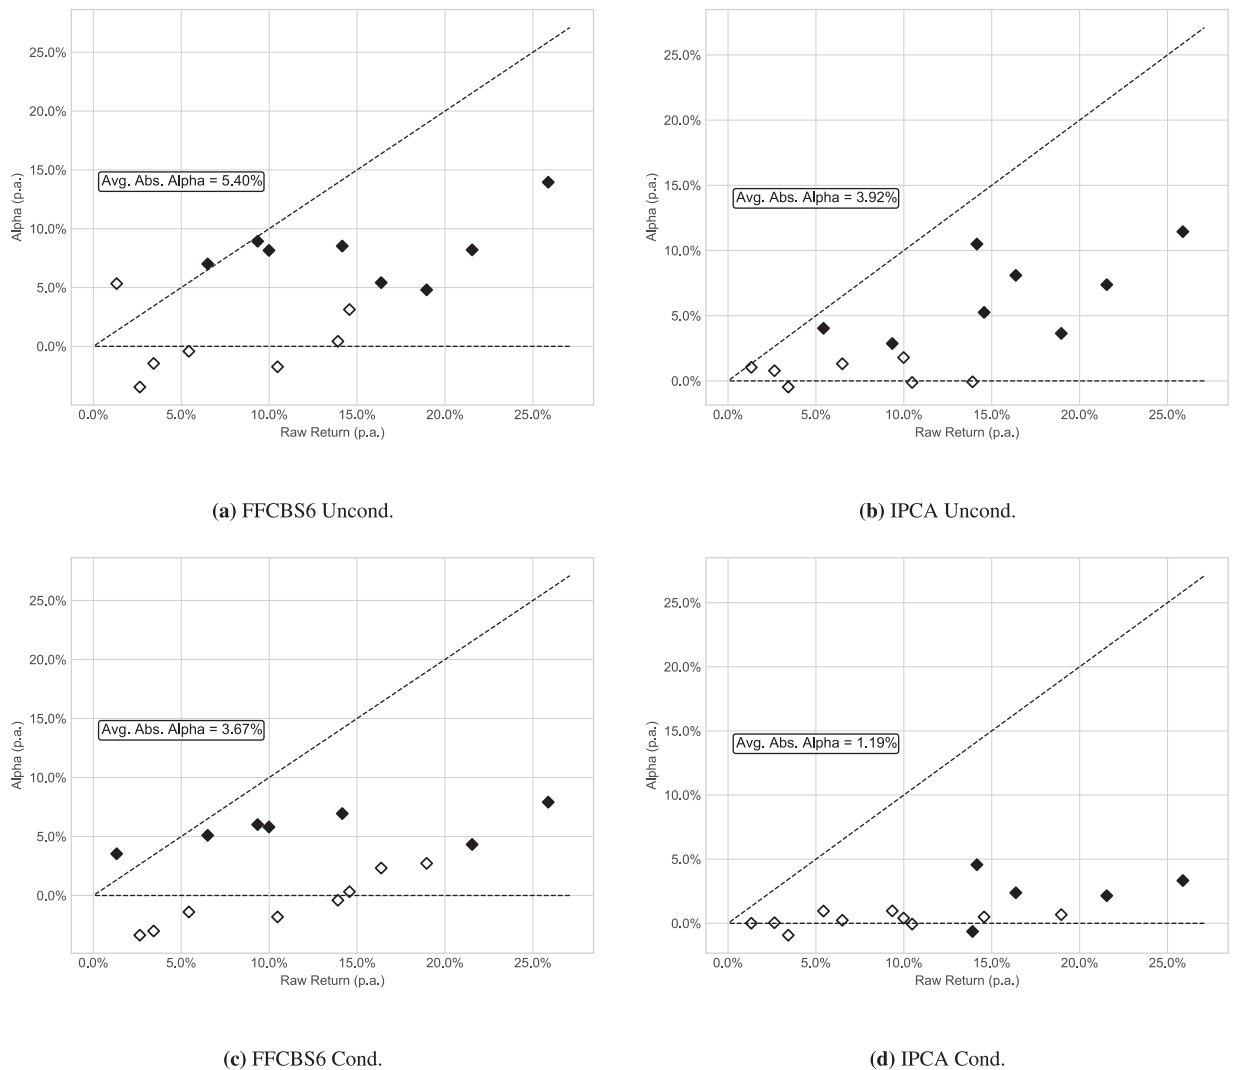


Fig. 2. Test of Mean-Variance Efficiency using Managed Portfolios. This figure plots portfolio alphas from a time series regression of portfolio returns on the set of factors from either an observable factor model (FFCBS6) and the IPCA model with $K = 3$ factors. The test portfolios are the managed portfolios constructed from the set of characteristics included in the estimated IPCA model: moneyness, maturity, implied volatility, embedded leverage, gamma and vega interacted with both a constant and the put/call dummy variable. This yields a total of 15 portfolios including the equally weighted portfolio. Portfolio alphas are plotted against the raw portfolio returns. All portfolios are re-leveraged to yield 10% annualized volatility. Filled-in diamond markers correspond to alphas with t-statistic greater than 2.0. Panels (a) and (b) show results from unconditional asset pricing tests, while panels (c) and (d) show results from the conditional FFCBS6 / IPCA model, i.e. alphas are computed as time-series averages of period-by-period portfolio residuals.

gamma portfolios, although even in these cases alphas are nearly the same as in the benchmark model.

We check the robustness of our previous results on the unconditional mean-variance efficiency of our factors by using a set of 20 BMS delta/TTM double-sorted portfolios that could provide a more challenging test case.

Figure 3 shows the results of this test. IPCA still outperforms the FFCBS6 model overall: average absolute alphas are around one percentage point lower using the IPCA factors than the observable factors. Furthermore, when visually assessing the distribution of alphas, we find that portfolio alphas from the FFCBS6 model exhibit a clear pattern: portfolio alphas increase with raw portfolio returns. This

suggests a systematic shortcoming of the observable factors in pricing the test portfolios.

In the previous analysis of mean-variance efficiency of our IPCA factors we have focused on the model with $K = 3$ factors. In order to inspect the behavior of average absolute portfolio alphas for varying number of latent factors, we now re-run the previous tests for the models with $K = 1, \dots, 5$ factors and compute the conditional and unconditional portfolio average absolute alphas as before. Table 6 summarizes the results. We find that unconditional and conditional alphas decrease as we increase the number of factors. This finding is *not* a guaranteed result of increasing the number of factors. Furthermore, we find a clear drop in portfolio alphas when moving from two to three latent fac-

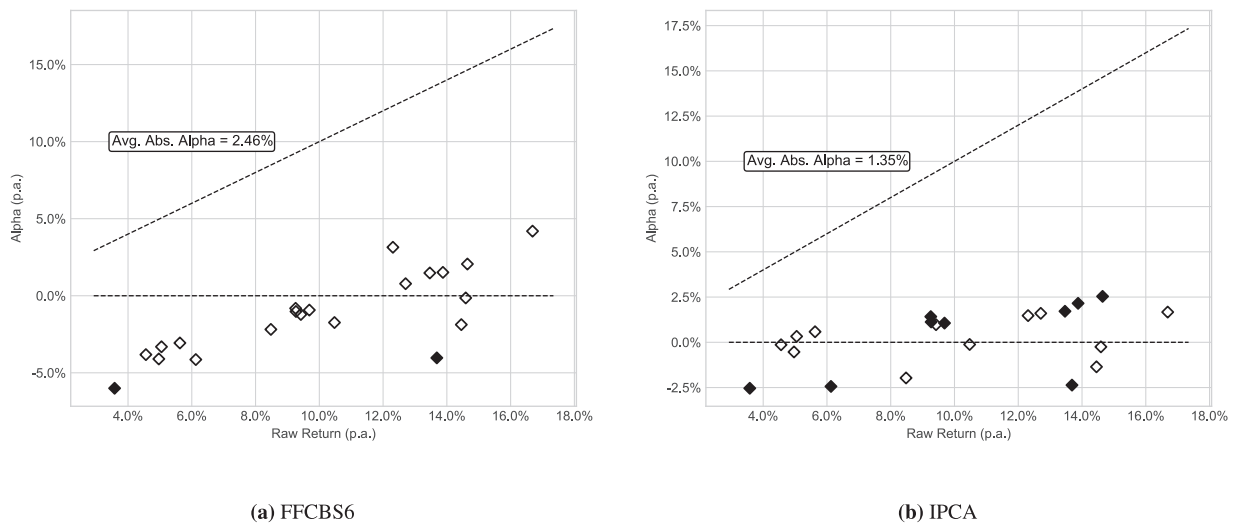


Fig. 3. Test of Unconditional Mean-Variance Efficiency using Double Sorted Portfolios. This figure plots portfolio alphas from a time series regression of portfolio returns on the set of factors from either an observable factor model (FFCBS6) and the IPCA model with $K = 3$ factors. The test portfolios are constructed by double sorting all option contracts on their time-to-maturity and option delta. Specifically, we construct portfolios of options with 1 month, 2 months, 3 to 6 months, 7 to 12 months, and 13 to 24 months to maturity and absolute option delta between zero and 0.1, 0.1 and 0.2, and so on up to an absolute delta of 0.5. This generates 20 double sorted portfolios. Portfolio alphas are plotted against the raw portfolio returns. All portfolios are re-leveraged to yield 10% annualized volatility. Filled-in diamond markers correspond to alphas with t-statistic greater than 2.0. Panels (a) and (b) show results from unconditional asset pricing tests.

Table 6

IPCA Portfolio Alphas. The table summarizes the conditional and unconditional portfolio average absolute alphas when the factors come from the restricted ($\Gamma_\alpha = 0$) IPCA model with $K = 1, \dots, 5$ factors. The test portfolios are the characteristics managed portfolios. Unconditional portfolio alphas are obtained from time series regressions of portfolio returns onto the set of factors. Conditional alphas are obtained as the time-series averages of period-by-period portfolio residuals in the main IPCA model. The reported values are the average absolute alphas across the set of test portfolios.

	K=1	K=2	K=3	K=4	K=5
Unconditional	9.32%	8.26%	3.92%	3.02%	2.26%
Conditional	7.85%	6.33%	1.19%	1.04%	0.57%

tors in agreement with the evidence from our earlier bootstrap exercise, and again favors a model with $K = 3$ factors.

4.4. Out-of-sample performance

So far, we have demonstrated that the IPCA model achieves a superior description of option return variation in-sample. We now turn to analyzing IPCA's out-of-sample fits.

We conduct a recursive estimation of IPCA in a backward-looking fashion. The model is estimated for an expanding window starting at half the available sample length, i.e. the first forecast is made in January 2007. The construction of out-of-sample realized factor returns, \hat{f}_{t+1} , follows Kelly et al. (2019, see Section 4.4).

We evaluate the out-of-sample fit of the restricted IPCA models with $K = 1, \dots, 5$ factors using the same performance metrics introduced above, i.e. the total and predictive R^2 using both individual options as well as managed portfolios as test assets. Table 7 details the results of our

Table 7

Out-of-Sample Performance The table shows the out-of-sample performance of the restricted IPCA model with $K = 1, \dots, 5$ factors. The out-of-sample exercise starts at 50% of our sample length, i.e. the first forecast is made for January 2007. Panel A contains results for individual option contracts, while Panel B examines the out-of-sample performance of managed portfolios.

	No. Factors				
	1	2	3	4	5
Panel A: Individual Options					
R^2_{total}	71.47%	76.18%	82.98%	86.53%	88.88%
R^2_{pred}	4.41%	3.31%	3.67%	4.34%	4.46%
Panel B: Managed Portfolios					
R^2_{total}	95.57%	97.01%	98.23%	98.89%	99.21%
R^2_{pred}	3.31%	3.00%	3.30%	3.43%	3.47%

analysis. We find that the strong in-sample performance of IPCA in terms of capturing observed return variation is also true on an out-of-sample basis. For the model with $K = 3$ factors that most of our analysis focuses on, the total R^2 only reduces from 85.07% in-sample to 82.98% out-of-sample. Out-of-sample IPCA fits continue to dramatically outperform even the in-sample fits of observable factor models such as FFCBS6. The predictive R^2 for IPCA at the individual contract level goes from 6.39% in-sample to 3.67% out-of-sample, while the predictive R^2 at the portfolio level drops from 7.90% in-sample to 3.30% out-of-sample.

4.5. Out-of-sample trading strategies

Most of our analysis to this point has focused on the statistical performance of IPCA. In this section we present

Table 8

IPCA Performance. This table reports the IPCA fit performance as measured by R^2_{total} and R^2_{pred} using both the restricted model ($\Gamma_\alpha = 0$) and the unrestricted model ($\Gamma_\alpha \neq 0$) with $K = 1, \dots, 5$. The employed option characteristics `delta`, `ttm`, `embed_lev`, `theta`, `impvvol`, `gamma`, `vega` are each interacted with a constant and an indicator variable that is equal one if the option is a put and is zero otherwise. Additionally, a constant is included in the set of characteristics. In Panel A performance measures are computed with respect to individual option contracts, while in Panel B performance measures are computed with respect to the characteristics-managed portfolios. Panel C reports p-values for the test with null hypothesis $H_0 : \Gamma_\alpha = 0$ from a bootstrap with 1000 draws for each time t .

		No. Factors				
		1	2	3	4	5
Panel A: Individual Options						
R^2_{total}	$\Gamma_\alpha = 0$	72.28%	79.65%	85.07%	88.90%	90.22%
	$\Gamma_\alpha \neq 0$	74.03%	81.08%	85.57%	89.32%	90.46%
R^2_{pred}	$\Gamma_\alpha = 0$	5.47%	5.54%	6.39%	6.59%	6.77%
	$\Gamma_\alpha \neq 0$	7.59%	7.58%	7.42%	7.21%	7.13%
Panel B: Managed Portfolios						
R^2_{total}	$\Gamma_\alpha = 0$	94.41%	96.64%	98.86%	99.39%	99.61%
	$\Gamma_\alpha \neq 0$	95.48%	97.00%	98.74%	99.41%	99.59%
R^2_{pred}	$\Gamma_\alpha = 0$	7.20%	7.44%	7.90%	7.99%	8.06%
	$\Gamma_\alpha \neq 0$	8.27%	8.27%	8.23%	8.18%	8.18%
Panel C: Bootstrap Test ($H_0 : \Gamma_\alpha = 0$)						
W_α p-value		7.4%	2.6%	47.2%	22.6%	3.6%

Table 9

Out-of-Sample Factor Portfolio Sharpe Ratios. The table summarizes the out-of-sample performance of the mean-variance optimal (tangency) portfolios formed from the IPCA factors. We assume a portfolio volatility target of 10% per year and rescale the portfolio weights accordingly using only backwards looking information. The performance measures are annualized expected return (ER), volatility, Sharpe ratio, skewness, and kurtosis. In addition, the table shows the time series alphas versus the embedded leverage factor (BAB) and the straddle factor with t -statistics in parentheses. We use an in-sample period of five years and then recursively expand the estimation window. The out-of-sample period runs from January 2001 to December 2017.

	ER	Vol	Sharpe	Skew	Kurtosis	$\alpha(BAB)$	$\alpha(Straddle)$	$\alpha(BAB + Straddle)$
IPCA K=1	0.094	0.096	0.986	-0.977	1.678	0.067 (2.218)	-0.01 (-0.125)	-0.010 (-0.768)
IPCA K=2	0.137	0.091	1.508	-0.598	0.899	0.109 (4.315)	0.041 (1.303)	0.037 (1.545)
IPCA K=3	0.166	0.099	1.673	-0.343	0.587	0.139 (4.920)	0.043 (1.438)	0.040 (1.669)
IPCA K=4	0.179	0.098	1.833	-0.326	0.520	0.151 (6.922)	0.066 (2.530)	0.063 (3.192)
IPCA K=5	0.197	0.109	1.802	-0.783	1.457	0.157 (5.689)	0.070 (1.827)	0.065 (2.310)

results from a trading strategy that aims to optimally combine the IPCA factors in a maximum Sharpe ratio sense. To this end, we study the investment performance of the tangency portfolio of IPCA factors. This provides a description of IPCA model performance in economic terms. The dual of factor pricing implies that a set of factors that can price all assets in the economy with zero alpha can also be combined to form the maximum Sharpe ratio portfolio in the economy.

To construct the out-of-sample tangency portfolio, we recursively estimate the IPCA model using only backward looking information and, based on the estimated factors, we then run a portfolio optimization with the [Ledoit and Wolf \(2004\)](#) covariance shrinkage estimator. This procedure produces tangency portfolio weights for each of the IPCA factors. Finally, we obtain the out-of-sample tangency portfolio return and analyze its performance. For ease of interpretation we re-leverage the tangency portfolio to target an annualized volatility of 10%. Further, we set the start of the out-of-sample period such that the first portfolio optimization makes use of five-years of in-sample data, hence leading to an out-of-sample period from January 2001 to December 2017.

We summarize the out-of-sample performance of the IPCA tangency portfolio based on $K = 1, \dots, 5$ factors in

Table 9. The table documents performance statistics including annualized expected returns and Sharpe ratios as well as portfolio alphas versus two benchmark models. The tangency portfolio for the IPCA model with three factors, which we focus on in many parts of our analysis, attains a Sharpe ratio of 1.67. Furthermore, the tangency portfolio yields a highly significant alpha of 15% annualized over and above a model consisting of the Fama-French three factors plus the momentum factor (“FFC4”). Vis-à-vis a model that in addition includes the [Frazzini and Pedersen \(2012\)](#) BAB factor and the [Coval and Shumway \(2001\)](#) straddle factor (“FFCBS6”) the IPCA tangency model still has an alpha of 4% per annum, albeit insignificant. However, IPCA models with four and five factors result in significant alphas versus the FFCBS6 model. These results show that not only is IPCA a successful model for pricing options returns, its trading strategy performance improves over successful strategies proposed in prior literature.

4.6. Interpreting the IPCA factors

Understanding the drivers of the IPCA model is as economically important as model performance itself. We attempt to provide an interpretation of the factors recovered by the IPCA model. As earlier, we focus on the IPCA model

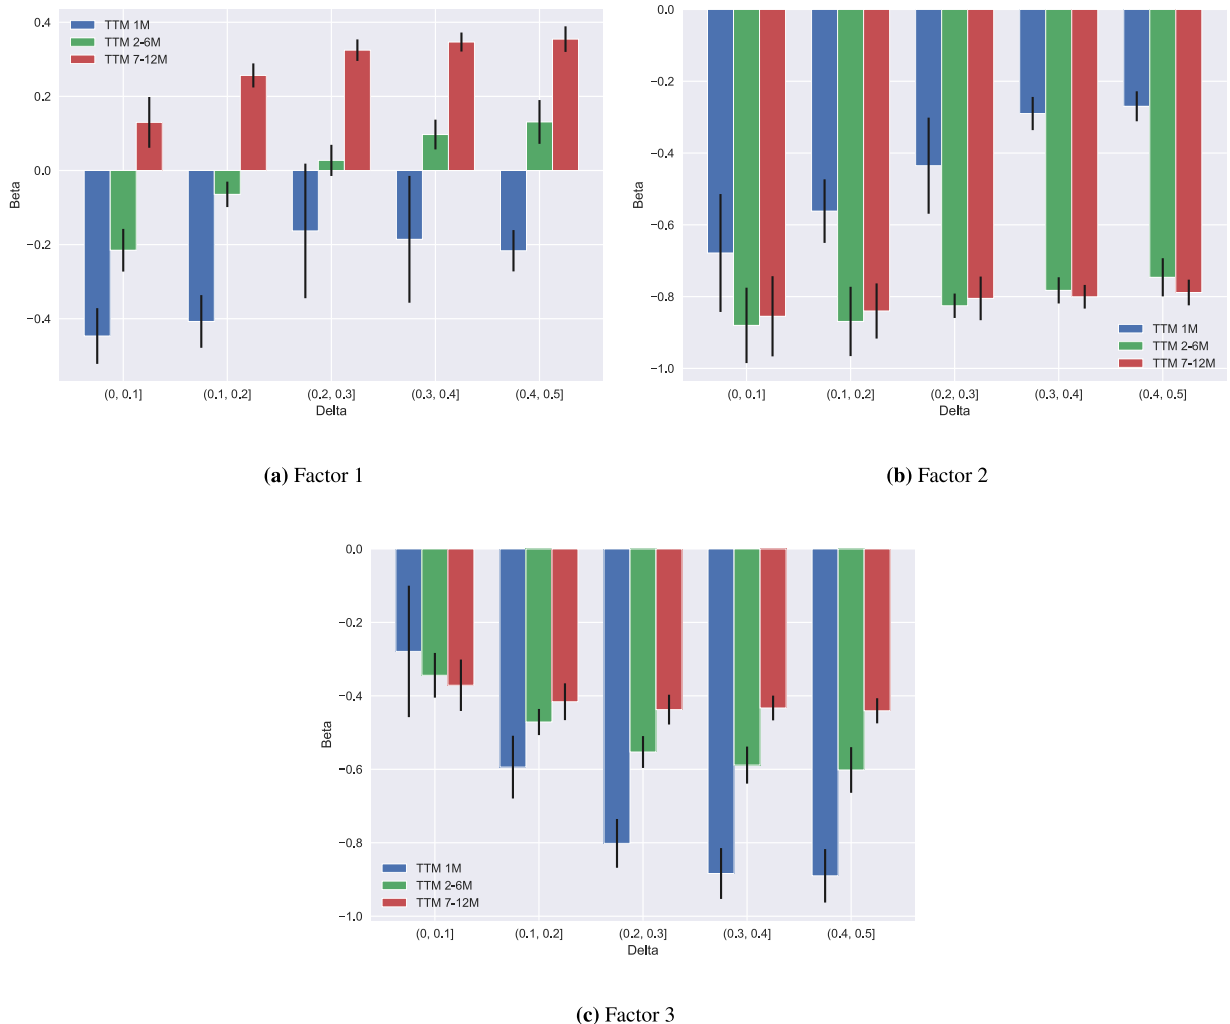


Fig. 4. IPCA Factor Regression Betas for Double Sorted Option Portfolios. This figure displays multiple regression betas of double sorted option portfolio returns onto the latent factors extracted by the restricted ($\Gamma_\alpha = 0$) IPCA model with $K = 3$ factors. The portfolios are constructed from equally weighting the returns on all options that fall in a given bin on the portfolio formation date. The option portfolios are constructed by double sorting all option contracts on their time-to-maturity and absolute forward option delta. The black overlaid bars indicate 95% confidence intervals.

with $K = 3$ factors. We adopt the convention to order IPCA factors by their variance such that “Factor 1” corresponds to the factor with highest time-series variance and “Factor 3” to the lowest. Our interpretation draws on four forms of evidence.

Relation to the implied volatility surface First, we study how the recovered IPCA factors relate to different regions of the implied volatility surface. To do so, we double-sort option contracts into portfolios for three maturity bins and five moneyness (absolute forward delta) bins each month, yielding 15 equally-weighted portfolios. We then run regressions of these portfolio returns on our three IPCA factors. Figure 4 plots the estimated regression betas. For Factor 1, we find a clear separation of regression coefficients by maturity. Short-dated contracts up to one month to maturity exhibit negative coefficients across different levels of moneyness, while longer dated contracts with maturity between seven and 12 months exhibit positive coefficients.

This spread suggests that Factor 1 captures variation in the term structure slope of the volatility surface. For Factor 2 we find two distinct patterns. The first is that longer dated contracts (maturities beyond one month maturity) exhibit stable coefficients across different levels of moneyness, in line with a volatility level effect. Second, for short-dated contracts (one month maturity), we find generally increasing coefficients, with deep OTM contracts being most negatively related to Factor 2. This second pattern is consistent with a short-dated moneyness skew factor and represents index tail risk (or, similarly, BMS gamma risk). Patterns are similar for Factor 3, where we find coefficients that are decreasing in absolute delta. The pattern is particularly pronounced for short-dated contracts, while for long-dated contracts coefficients are more stable. This evidence suggests that Factor 2 and Factor 3 together are a combination of a volatility level factor and moneyness skew factor.

Relation to observable option return factors In our next analysis, we attempt to build easily interpretable observable factors that align with the estimated IPCA factors. We construct three prespecified portfolios that are defined as a volatility level factor, and maturity slope factor, and moneyness skew factor. The level factor is an equally-weighted portfolio that shorts options with absolute delta between 0.4 and 0.5, and captures fluctuations in the overall level of the implied volatility surface. The maturity slope factor is constructed by buying options with maturities between six and 12 months and shorting options months with one month to maturity. Therefore, this factor captures fluctuations in option prices associated with shifts in the implied volatility term structure. Finally, the moneyness skew factor is long OTM call options with delta between 0.1 and 0.2 and short OTM put options with delta between -0.2 and -0.1 . This factor captures variation in the implied volatility skew, capturing shifts in the market's expectation of tail events.

Before comparing these observable factors with our IPCA factor, we first analyze their performance as a pricing model for individual contracts. Following the analysis of Table 4, we use the same set of option characteristics to instrument conditional loadings in this model as we use in the baseline IPCA specification. Among individual contracts, the observable factor model attains a total R^2 of 81.5% and predictive R^2 of 6.2%. For the managed portfolios, the observable option factor model attains a total R^2 of 94.7% and predictive R^2 of 7.8%. Comparing with the IPCA model with $K = 3$ in Table 4, we find that the observable option factor model underperforms IPCA by a small margin, but outperforms the benchmark observable factor models that we studied in Table 4.

We compare the level, maturity slope, and moneyness skew factors to IPCA factors via multiple regression, and report results in Table 10. These observable factors explain close to 90% of the variation in two of our three IPCA factors, and more than 70% of the variation in the third factor. Since the three observable factors are not orthogonal to each other, we follow Fournier et al. (2021) and augment the regression evidence by computing Shapley-Owen R^2 values to measure the marginal contribution of each regressor. Shapley-Owen R^2 compares the marginal increase in model fit from adding a given regressor to all potential subsets of regression models which do not include that regressor. The results indicate that the first IPCA factor is most closely related to the maturity slope factor, the second IPCA factor is most closely related to the moneyness skew factor, while the third IPCA factor is most closely related to the volatility level factor. Thus, an observable option factor model embedding the ideas of level, maturity slope, and moneyness skew fluctuations is useful for understanding the behavior of the IPCA model.

Factor behavior across market regimes Our third interpretation analysis examines the conditions under which each factor earns its risk compensation. We define market regimes following Karakaya (2013), and Table 11 reports annualized factor means, standard deviations and Sharpe ratios in each regime. Panel A collects the results for the three IPCA factors, while Panel B collects the results for the level, maturity slope, and moneyness skew factors.

Table 10

IPCA Factors versus Option Return Factors. The table relates the IPCA factors extracted by the restricted IPCA model with $K = 3$ factors to a set of observable option return factors. Specifically, in the spirit of Karakaya (2013), we reconstruct a level, maturity slope and moneyness skew factor. The level factor is constructed as the equally-weighted return to shorting at-the-money (absolute delta between 0.4 and 0.5) options with maturities of up to one year. The maturity slope factor is constructed by longing options with maturities between six and twelve months and shorting options with one month to maturity across all option deltas. Finally, the moneyness skew factor is constructed by longing options with deltas between 0.1 and 0.2 and shorting options with deltas between -0.2 and -0.1 . The table shows time series regressions results from regressing IPCA factors on the set of option return factors described above. Both IPCA factors and option return factors are standardised using their time series standard deviation. The bottom panel shows the Shapley-Owen value based decomposition of the adjusted R^2 . In parentheses we report t -statistics using Newey-West standard errors with four lags.

	F1	F2	F3
Level	-0.10 (-1.58)	0.32 (6.35)	1.13 (27.16)
Maturity Slope	0.71 (10.59)	-0.38 (-9.44)	0.59 (12.82)
Moneyness Skew	0.53 (8.80)	0.58 (7.63)	-0.53 (-8.61)
R^2_{adj}	72.3%	88.0%	89.3%
No. Obs.	261	261	261
	Shapley-Owen R^2		
Level	9.1%	28.4%	49.4%
Maturity Slope	38.6%	26.2%	20.2%
Moneyness Skew	24.7%	33.4%	19.8%

The first and third IPCA factors earn significant returns with Sharpe ratios of 1.04 and 1.42 respectively. Outside of volatility jump periods (period “Non Vol Jump”), we find that all three IPCA factors earn significant, elevated returns consistent with the evidence of a volatility risk premium in delta-hedged option returns documented in the literature. Comparing between Panels A and B we see a closer link between the IPCA factors and the three pre-specified factors. As suggested above, IPCA Factor 1 exhibits properties of a slope factor and indeed we find that its regime-dependent risk compensation broadly matches that of the maturity slope factor in Panel B. In particular, we find that the risk compensation of the slope factor is related to the business cycle (“Recession” vs “Non Recession”). For IPCA Factor 2 we find a close resemblance in risk properties to the moneyness skew factor. These share a significant negative price of risk during price jumps and negative compensation during volatility jumps. Likewise, for IPCA Factor 3, we see agreement with the volatility level factor in line with the regression evidence above.

IPCA factors, risk, and frictions In our fourth approach to interpretation, we regress IPCA factors on the realized return of the underlying, measures of aggregate risk, and a proxy for financial frictions. Regressing on the underlying market return helps us understand the extent to which option returns maintain residual exposure to the market after imperfect delta hedging (arising for example from heavy-tailed behavior of the underlying). We also regress on realized variance and implied variance (CBOE VIX index). As argued by Dew-Becker et al. (2021), contrasting

Table 11

IPCA Factor Summary Statistics. This table reports summary statistics of the restricted IPCA model for the full sample period as well as distinguished sub-periods. The sub-periods are: “Price Jump” indicated by S&P 500 return smaller than -4% in the past month, “Vol Jump” indicated by a VIX index increase of greater than 4%, “Recession” as indicated by NBER recession periods, “Severe Bear Mkt.” indicated by S&P 500 return smaller than -25% in the previous 12 months, “Rising Mkt.” constructed as the opposite of “Severe Bear Mkt.”, and the VIX regime (high, medium, low) constructed from terciles of the VIX level distribution over our sample period. The IPCA model is estimated using three factors and characteristics δ , t , m , embed_lev , impvol , γ , θ , vega . Reported measures are annualized. Significance of the Sharpe ratio is obtained from a regression of the factor return series on a constant where the regression is using Newey-West standard errors with up to 4 lags. Sharpe ratios printed in bold font indicate an absolute t-statistic exceeding 2.0. The sample period is from 1996 to 2017 covering 261 months.

Panel A: IPCA Factors										
Period	N. Obs	F1			F2			F3		
		Mean	StdDev	Sharpe	Mean	StdDev	Sharpe	Mean	StdDev	Sharpe
Full Sample	261	0.06	0.06	1.04	0.01	0.04	0.26	0.02	0.02	1.42
Price Jump	25	-0.14	0.15	0.93	-0.11	0.09	-1.19	0.01	0.02	0.48
Non Price Jump	236	0.05	0.04	1.32	0.02	0.02	0.98	0.02	0.01	1.60
Vol Jump	50	0.10	0.11	0.94	-0.06	0.07	-0.86	-0.01	0.02	-0.52
Non Vol Jump	211	0.05	0.04	1.33	0.03	0.02	1.28	0.03	0.01	2.10
Recession	26	0.06	0.13	0.50	-0.04	0.08	-0.46	0.03	0.02	1.44
Non Recession	235	0.06	0.05	1.32	0.02	0.03	0.52	0.02	0.01	1.42
High VIX	65	0.14	0.10	1.48	-0.02	0.06	-0.38	0.02	0.02	1.21
Medium VIX	130	0.04	0.04	0.94	0.02	0.03	0.85	0.02	0.01	1.49
Low VIX	66	0.02	0.02	1.22	0.02	0.01	1.39	0.02	0.01	1.96

Panel B: Option Level, Slope & Skew										
Period	N. Obs.	Level			Maturity Slope			Moneyness Skew		
		Mean	StdDev	Sharpe	Mean	StdDev	Sharpe	Mean	StdDev	Sharpe
Full Sample	261	0.02	0.02	1.03	0.01	0.01	0.91	0.01	0.01	0.58
Price Jump	25	-0.03	0.03	-0.82	0.03	0.02	1.63	-0.03	0.03	-0.95
Non Price Jump	236	0.02	0.01	1.66	0.01	0.01	0.81	0.01	0.01	1.29
Vol Jump	50	-0.03	0.03	-1.09	0.02	0.01	1.58	-0.01	0.02	-0.39
Non Vol Jump	211	0.03	0.01	2.46	0.01	0.01	0.70	0.01	0.01	1.39
Recession	26	0.01	0.03	0.32	0.01	0.01	0.98	-0.01	0.02	-0.52
Non Recession	235	0.02	0.01	1.25	0.01	0.01	0.90	0.01	0.01	0.98
High VIX	65	0.01	0.03	0.25	0.02	0.01	1.63	0.00	0.02	0.17
Medium VIX	130	0.02	0.01	1.75	0.01	0.01	0.62	0.01	0.01	0.81
Low VIX	66	0.02	0.01	2.23	0.00	0.00	0.53	0.01	0.00	1.70

realized versus implied variance allows us to differentiate which factors reflect jump/gamma risk (better captured through realized variance) versus the risk of shocks to future volatility (better captured through implied variance). Finally, we also regress on the intermediary capital risk factor of He et al. (2017), which measures financial frictions and thus constitutes a form of limits-to-arbitrage.

Table 12 displays the results of the regression. IPCA factors relate significantly not only to the implied volatility of option markets, but also to the dynamics of the underlying itself. The second factor is most closely aligned with shocks to VIX, and the third factor most related to realized variance. Most interestingly, we find that the intermediary capital risk factor is the dominant driver of IPCA factor 1 according to the Shapley-Owen R^2 ; it is responsible for half of the explanatory power of this factor. This suggests that the empirical success of IPCA derives in part from its ability to capture option price behavior associated with broad market frictions that are missing from more traditional no-arbitrage option pricing models.

Which Characteristics Matter? The IPCA model allows us to assess the statistical contribution of individual characteristics, which provides an economic interpretation for the source of model gains. Kelly et al. (2019) suggest assessing individual characteristic contributions by setting its associated Γ_β parameters to zero while keeping all other characteristic loadings fixed at their estimated values. They also

Table 12

IPCA Factors and the Dynamics & Liquidity of the Underlying. This table shows the results of multivariate regressions of the IPCA Factors from the model with $K = 3$ onto liquidity proxies and measures of underlying dynamics. The regressors are the changes in the CBOE VIX index, the changes in realized variance of the S&P 500 index, the realized return of the S&P 500 index, and the intermediary capital risk factor from He et al. (2017). The regressions are using Newey-West standard errors with four lags. Both the IPCA factors and regressors are standardized using their time series standard deviation. The bottom panel shows the Shapley-Owen value based decomposition of the adjusted R^2 . The sample period is 2000 to 2017. In parentheses we report t-statistics.

	F1	F2	F3
VIX	-0.17 (-0.88)	-0.46 (-3.17)	-0.02 (-0.22)
Realized Variance	-0.23 (-4.19)	-0.08 (-0.74)	-0.40 (-3.27)
Realized Return	-0.13 (-1.29)	0.14 (1.04)	-0.18 (-2.53)
Intermed. Cap. Risk	-0.48 (-2.45)	-0.08 (-1.07)	0.25 (3.02)
R^2_{adj}	22.1%	34.5%	19.4%
No. Obs.	215	215	215
Shapley-Owen R^2			
VIX	2.3%	16.2%	2.8%
Realized Variance	4.5%	6.3%	9.5%
Realized Return	4.7%	9.0%	1.4%
Intermed. Cap. Risk	10.6%	3.0%	5.8%

Table 13

IPCA Instrument Significance. This table contains results from the bootstrap test for individual characteristics contributions to overall fit in the restricted IPCA specification with $K = 3$ factors. The test sets all elements in Γ_β related to a given characteristics equal to zero and compares the overall model fit against the model with fully specified Γ_β . The first column summarizes the absolute reduction in total R^2 from setting the row in the matrix Γ_β pertaining to a given characteristic to zero. The second column contains p-values for the bootstrap test W_β with 1000 draws that tests $H_0 : \Gamma_\beta = [\gamma_{\beta,1}, \dots, \gamma_{\beta,l-1}, \mathbf{0}_{K \times 1}, \gamma_{\beta,l+1}, \dots, \gamma_{\beta,L}]$ against the alternative $H_1 : [\gamma_{\beta,1}, \dots, \gamma_{\beta,L}]$. The table is sorted by the reduction in total R^2 from largest to smallest.

	Reduction R^2_{total} (abs.)	W_β p-value
impvol:put	-26.80%	0.00
vega	-23.85%	0.00
impvol	-22.46%	0.00
gamma	-14.88%	0.00
theta	-14.78%	0.00
ttm	-7.03%	0.00
delta:put	-4.41%	0.00
delta	-2.90%	0.00
theta:put	-2.85%	0.00
vega:put	-2.43%	0.00
embed_lev	-2.24%	0.00
ttm:put	-2.17%	0.01
gamma:put	-1.64%	0.00
embed_lev:put	-1.51%	0.00

develop a bootstrap test of each characteristic's statistical significance in the model. We outline this procedure in Internet Appendix A.

Table 13 summarizes the contribution of each characteristic to the IPCA model measured as the reduction in total R^2 from excluding that characteristic, holding all other model estimates fixed. We find that the most important characteristics are implied volatility, BMS vega, and BMS gamma. The bootstrapped p-values confirm that these characteristics are highly significant contributors. For comparison, Fig. 5 plots the estimated factor loadings, Γ_β , and confirms that the characteristics identified as most important in the bootstrap analysis are also those that earn large coefficients in the beta specification.

Next, we investigate the functional shape of the factor exposures instrumented with our set of option characteristics. Although IPCA is a linear factor model at heart, non-linearities may be captured via choosing a relevant set of asset characteristics. This is particularly relevant in the context of options which have non-linear payoff profiles. To investigate this feature, we sort our sample into 10 bins, Δ_j , by absolute option deltas and compute the unconditional factor exposures $\beta^{\Delta_j} = \bar{z}'_j \Gamma_\beta$ where we fix characteristics at their per-bin means, \bar{z}_j . Fig. 6 plots the unconditional factor exposures for factors F1 to F3 against moneyness. Indeed, we find that unconditional factor exposures exhibit a non-linear association with option moneyness.

4.7. Robustness tests

This section reports a number of robustness checks for our main monthly IPCA analysis. Due to space constraints, the results are reported in the Internet Appendix B.

Alternative Sets of Characteristics

While IPCA allows for a large number of characteristics to be considered as instruments, the researcher must ultimately take a stance on the set of characteristics in the model. In our analysis so far the set of characteristics included moneyness, time-to-maturity, implied volatility, embedded leverage, gamma and vega, and all characteristics were interacted with both a constant and the put/call dummy variable. In order to put our choice of characteristics into perspective, we run two sets of analyses.

First, we compare the performance of the baseline model against a model in which the choice of characteristics is motivated economically from an option return decomposition. One way to approach the choice of characteristics problem is from the perspective of a no-arbitrage model such as BMS. In that context, a standard option return decomposition is obtained via a simple Taylor expansion of the option price F_t along its most relevant dimensions as follows

$$\Delta F_t = \frac{dF_t}{dS_t} \Delta S_t + \frac{dF_t}{dIV_t} \Delta IV_t + \frac{d^2 F_t}{dS_t^2} (\Delta S_t)^2 + \dots \quad (11)$$

This can be re-arranged to yield the relative change in the option price

$$\frac{\Delta F_t}{F_t} = \underbrace{\frac{dF_t}{dS_t} \frac{S_t}{F_t}}_{\Delta \text{Decomp}} \frac{\Delta S_t}{S_t} + \underbrace{\frac{dF_t}{dIV_t} \frac{IV_t}{F_t}}_{\nu \text{Decomp}} \frac{\Delta IV_t}{IV_t} + \underbrace{\frac{d^2 F_t}{dS_t^2} \frac{S_t^2}{F_t}}_{\Gamma \text{Decomp}} \frac{\Delta S_t^2}{S_t^2} + \dots \quad (12)$$

The terms ΔDecomp , νDecomp , and ΓDecomp can be understood as coefficients on a set of factors. Internet Appendix Table B.1 details the performance of an IPCA model that uses the three model-implied coefficients as characteristics. We find that the model with characteristics following the return decomposition in Eq. (12), achieves a slightly lower total R^2 than the baseline model (e.g. for $K = 3$ and $\Gamma_\alpha = 0$, 80.99% vs. 85.57%), while the predictive R^2 of the decomposition based model is comparable to the agnostic specification (e.g. for $K = 3$ and $\Gamma_\alpha = 0$, 6.36% vs. 6.39%). Next, we estimate a battery of models with varying sets of characteristics to further assess the dependence of our baseline model's performance on the choice of characteristics. There are a number of plausible characteristics that are available beyond our baseline set. Internet Appendix Table B.3 contains an overview of IPCA specifications using different sets of characteristics. The results demonstrate the robustness of our main findings because the model fits do not change dramatically with small changes in the characteristic set. One finding of this analysis is that higher order option Greeks such as BMS Volga (i.e., sensitivity to changes in BMS gamma) add little in terms of total and predictive R^2 . Overall, this evidence shows that the choice of characteristics in our baseline specification stacks up well against other plausible choices of characteristics.

Longer maturities Our baseline sample filters out contracts with maturities beyond one year due to their low liquidity. We conduct a robustness analysis that includes contracts with up to two years to maturity. We keep the

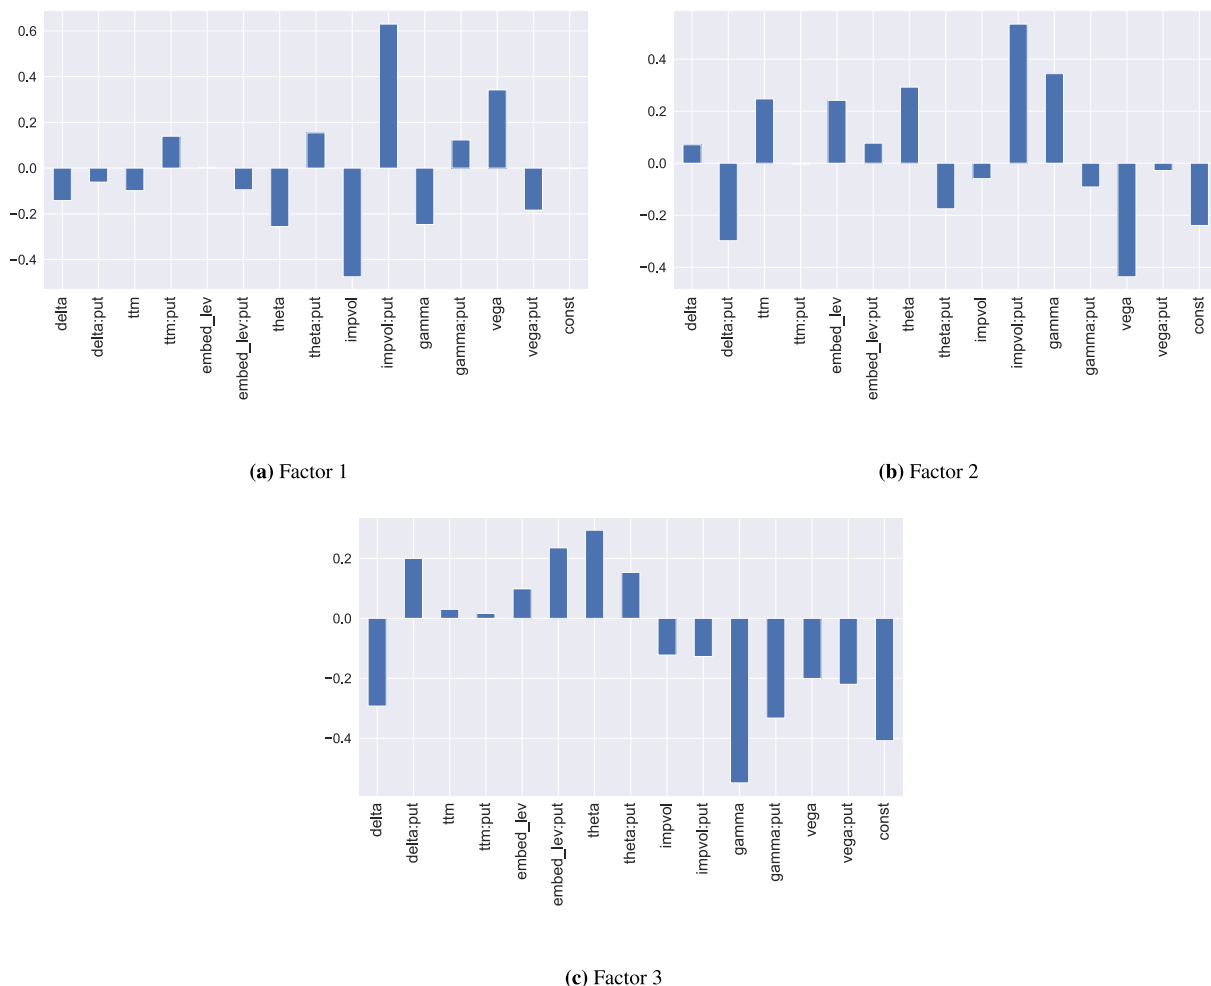


Fig. 5. Plots of loadings Γ_β . This figure shows the loadings, Γ_β , that create the mapping between characteristics and factors. Each subplot corresponds to one column of Γ_β , i.e. one IPCA factor. The IPCA model fitted is the restricted model with $K = 3$. The colon notation is used to indicate interaction with another variable, e.g. ttm:put refers to the interaction between an option's time-to-maturity with an indicator variable that is equal one for puts and zero for calls.

same set as instruments as in our baseline specification and re-estimate the IPCA model. Internet Appendix Table B.2 presents the results. Focusing on the restricted IPCA model with $K = 3$, we confirm that the total R^2 is relatively similar to baseline specification (85.57% vs. 86.11%). However, we note that the predictive R^2 is somewhat lower (5.41% vs. 6.39%) when adding longer maturity contracts. This finding may be related to the illiquidity associated with these kinds of contracts that makes precise measurement of expected returns more challenging.

4.8. Model performance at the daily frequency

Our analysis to this point focuses entirely on the monthly frequency. In this section, we estimate the IPCA model using daily data on delta-hedged returns. This horizon is of particular interest to market makers and traders that often only hold contracts for short periods of time and that place an emphasis on risk managing a book of derivatives.

As a benchmark, we compare IPCA against the P&L attribution framework proposed by Carr and Wu (2020) which is developed particularly with short horizon returns in mind. Their framework attributes option returns to variation in the first and second moments of the underlying price process and the implied volatility process, and is a formalization of the typical trader's "rule-of-thumb" model for managing option risk. It is obtained from a Taylor expansion of the BMS pricing formula, taking expectations under the risk-neutral measure and enforcing a no-arbitrage condition. A particularly neat version of their P&L attribution is obtained by scaling option gains by cash gamma:

$$\frac{dF}{F_{SS}S_t} = \frac{F_S}{F_{SS}} \frac{dS_t}{S_t} + IV_t^2 S_t \tau \left(\frac{dIV_t}{IV_t} - \mu_t dt \right) + \frac{1}{2} S_t \left(\left(\frac{dS_t}{S_t} \right)^2 - \sigma_t^2 dt \right)$$




Fig. 6. IPCA Factor Exposure by Level of Moneyness. This figure shows the unconditional factor exposures as a function of contract moneyness measured by the absolute value of the forward option delta. Specifically, we sort the sample of option returns into buckets on the absolute value of the option delta and for each bucket Δ_j compute the average factor exposure $\beta^{\Delta_j} = \bar{z}_j' \Gamma_{\beta}$ where \bar{z}_j is the mean of the characteristics in bucket Δ_j . The buckets of absolute forward delta are (0, 0.05], (0.05, 1], ..., (0.45, 5]. For comparability, we scale the factor exposure by multiplying with the factor standard deviations.

$$\begin{aligned}
 &+ \frac{1}{2} S_t z_+ z_- \left(\left(\frac{dIV_t}{IV_t} \right)^2 - \omega_t^2 dt \right) \\
 &+ z_+ S_t \left(\left(\frac{dS_t}{S_t} \frac{dIV_t}{IV_t} \right) - \gamma_t dt \right), \tag{13}
 \end{aligned}$$

where F_S denotes the BMS option delta, F_{SS} denotes the BMS option gamma, $\tau = (T_i - t)/365$ is the standardized time-to-maturity, and $z_{\pm} = (\log(K_t/S_t) \pm \frac{1}{2} IV_t^2 \tau)$. Further, the moments under the risk-neutral measure are

$$\begin{aligned}
 \mu_t &= \mathbb{E}_t \left[\frac{dIV_t}{IV_t} \right] / dt, & \sigma_t^2 &= \mathbb{E}_t \left[\left(\frac{dS_t}{S_t} \right)^2 \right] / dt, \\
 \omega_t^2 &= \mathbb{E}_t \left[\left(\frac{dIV_t}{IV_t} \right)^2 \right] / dt, & \gamma_t^2 &= \mathbb{E}_t \left[\left(\frac{dS_t}{S_t}, \frac{dIV_t}{IV_t} \right) \right] / dt,
 \end{aligned}$$

i.e. μ_t is the expected rate of change in the BMS implied volatility, σ_t^2 and ω_t^2 are the conditional variance rate of underlying and implied volatility, respectively, and γ_t is the conditional covariance between the two processes. Note that in contrast to the BMS framework itself, the Carr and Wu (2020) framework does not require specifying the full dynamics of the underlying.

The risk-neutral moments above can be obtained from the implied volatility surface. However, since we are not limiting the analysis to ATM options on a grid of fixed maturities as in Carr and Wu (2020), a few generalizations are necessary. We estimate μ_t from the implied volatility term-structure, but instead of using fixed maturity points to obtain the term structure slope, we approximate the slope by taking a symmetric 60-day window around the

location of a given contract on the implied volatility surface. For contracts close to the edges of the surface, we shift the window from which we compute the slope to ensure a window size of 60 days. The variance rate ω_t is estimated from a 63-day rolling window variance of daily implied volatility changes where we ensure the availability of at least 21 observations. The covariance rate γ_t is estimated from the rolling 21-day covariance between spot changes and implied volatility changes and we ensure availability of at least 10 observations. Finally, we estimate σ_t^2 from the ATM implied volatility at a maturity of one month.

We measure the performance of the Carr and Wu (2020) approach individually for each of the risk attributions. To this end we define the series

$$\begin{aligned}
 R_0 &= \frac{dF}{S_t} - F_S \frac{dS_t}{S_t} \\
 R_1 &= R_0 - IV_t^2 S_t F_{SS} \tau \left(\frac{dIV_t}{IV_t} - \mu_t dt \right) \\
 R_2 &= R_1 - \frac{1}{2} S_t F_{SS} \left(\left(\frac{dS_t}{S_t} \right)^2 - \sigma_t^2 dt \right) \\
 R_3 &= R_2 - \frac{1}{2} S_t F_{SS} z_+ z_- \left(\left(\frac{dIV_t}{IV_t} \right)^2 - \omega_t^2 dt \right) \\
 R_4 &= R_3 - z_+ S_t F_{SS} \left(\left(\frac{dS_t}{S_t} \frac{dIV_t}{IV_t} \right) - \gamma_t dt \right).
 \end{aligned}$$

We then compute the fraction of variation in delta-hedged option returns r_{spot}^{Δ} explained by sequentially accounting

Table 14

Comparison of IPCA against a No-Arbitrage Model at Daily Frequency. This table details the performance of both the Carr and Wu (2020) model and the restricted IPCA model with $K = 1, \dots, 5$ factors at daily frequency. The daily returns are delta-hedged. For the no-arbitrage model the total R^2 is computed as follows: for a series $R_i, 1, \dots, 4$ the R-Squared is computed as $R^2_{total,i} = 1 - \text{Var}(R_i)/\text{Var}(R_0)$ where R_0 is the series of delta-hedged daily returns. The computation of the total R^2 in the IPCA model follows the usual form.

Panel A: Average Performance									
	Carr & Wu - R^2_{total}				IPCA - R^2_{total}				
	R_1	R_2	R_3	R_4	K=1	K=2	K=3	K=4	K=5
All Options	72.22%	85.29%	82.94%	85.20%	69.98%	85.89%	91.53%	92.92%	93.76%
Panel B: Average Performance by Moneyness Bin									
abs. delta	Carr & Wu - R^2_{total}				IPCA - R^2_{total}				
	R_1	R_2	R_3	R_4	K=1	K=2	K=3	K=4	K=5
0 to 0.1	65.78%	70.96%	58.38%	69.32%	64.5%	67.7%	73.1%	81.5%	85.3%
0.1 to 0.2	72.65%	83.54%	81.54%	85.26%	77.3%	84.8%	90.6%	92.0%	93.0%
0.2 to 0.3	73.46%	87.30%	86.64%	88.08%	75.6%	88.0%	93.8%	94.2%	94.7%
0.3 to 0.4	73.05%	89.08%	88.89%	89.35%	70.4%	88.5%	94.3%	94.8%	95.2%
0.4 to 0.5	73.50%	90.12%	90.09%	90.18%	62.9%	88.0%	93.4%	94.2%	94.7%
Panel C: Average Performance by Time-to-Maturity Bin									
ttm	Carr & Wu - R^2_{total}				IPCA - R^2_{total}				
	R_1	R_2	R_3	R_4	K=1	K=2	K=3	K=4	K=5
1 Month	51.9%	72.5%	67.7%	72.0%	69.0%	80.2%	91.7%	94.1%	95.1%
2 Months	76.8%	92.1%	91.2%	92.6%	77.9%	91.9%	93.1%	94.7%	95.6%
3 to 6 Months	87.4%	94.6%	94.0%	94.7%	72.8%	90.8%	92.8%	93.7%	94.2%
6 to 12 Months	93.2%	96.4%	96.3%	96.5%	58.0%	78.8%	87.8%	88.4%	89.4%
Panel D: Average Performance by VIX bin									
VIX	Carr & Wu - R^2_{total}				IPCA - R^2_{total}				
	R_1	R_2	R_3	R_4	K=1	K=2	K=3	K=4	K=5
0% to 10%	65.90%	74.41%	69.07%	69.52%	56.2%	76.1%	85.7%	90.1%	92.3%
10% to 20%	70.36%	86.33%	84.64%	86.91%	71.5%	85.4%	91.4%	93.4%	94.2%
20% to 30%	76.79%	82.85%	79.13%	81.36%	63.9%	85.6%	90.8%	91.9%	92.6%
30% to 90%	71.95%	95.33%	93.85%	97.32%	74.5%	86.5%	92.3%	93.6%	94.5%

for vega risk, gamma risk and so on as

$$R^2_i = 1 - \text{Var}[R_i]/\text{Var}[r_{Spot}^\Delta], \quad i = 1, \dots, 4. \tag{14}$$

Since there is no direct analogue of R^2_{pred} in this framework, we focus our analysis on the total R^2 .

We summarize the results from a comparison of the Carr and Wu (2020) model and the restricted IPCA framework in Table 14. Before diving into the comparison of the two approaches, we note that IPCA also performs well at daily frequency and delivers high total R^2 across a large range of option contracts. In analogy to our earlier exercise, we present results for $K = 1, \dots, 5$ factors to highlight how the IPCA performance increases as we allow for more factors to mirror progressively hedging additional sources of risk in the Carr and Wu (2020) model. Internet Appendix Table B.4 reports a more detailed set of results for the IPCA model at daily frequency.

With two or more factors, IPCA performs at least as well as Carr and Wu (2020) in describing the behavior of daily option returns (Table 14, Panel A). This is further illustrated in sub-samples of our data that isolate observations by moneyness, maturity, and the level of the VIX (Panels B, C, and D). While the Carr and Wu (2020) attribution framework performs reasonably well for options near the money, its comparative performance versus IPCA dete-

riorates for OTM options (Panel B) and for short-dated options with one month to maturity (Panel C). Finally, Panel D demonstrates that both models are able to capture option return variation over a wide range of market conditions as proxied by the VIX index.

5. Conclusion

In this paper we study a latent factor model (IPCA) for option returns. We demonstrate that a coherent factor-based description of option returns is possible in a model with time-varying factor loadings. The problems typically posed by short option lifespans and rapidly changing risk attributes for factor beta estimation is resolved by IPCA's use of time-varying betas that are instrumented by option contract characteristics.

We find that a low dimensional latent factor model is successful in capturing variation in option returns and describing differences in risk across a wide range of options. It substantially outperforms observable factor models in terms of its ability to capture the variation in realized option returns.

The model also provides an accurate description of the risk-return trade-off in options markets. When using three IPCA factors, we do not reject the hypothesis of zero condi-

tional alpha in our data. We also find that a trading strategy designed to efficiently capture the risk-return trade-off (as estimated from IPCA) earns an annualized Sharpe ratio as high as 1.8 and has positive alpha versus previously proposed investment strategies using index options.

We test the robustness of our findings in an out-of-sample exercise and find that IPCA's performance remains strong, still outperforming the even in-sample results from observable factor models. The risk factors recovered by IPCA can be interpreted as capturing fluctuations in the level of the volatility surface, in the maturity slope, and in the short-dated moneyness skew. While most of our analysis focuses on monthly data, we also find that the IPCA model matches the behavior of options returns at the daily frequency across a wide range of option contracts.

Our model and findings suggest a number of promising directions for future research. One is to extend this modeling scheme to the single name options markets. In particular, the IPCA framework may present the solution to identifying a unified factor model that can accurately price all single name options simultaneously, taking both stock characteristics and option contract attributes as inputs to the beta specification. Recently, Kelly et al. (2020) use IPCA to study the integration of stock and bond markets. Likewise, our framework may be used to pursue a joint asset pricing model of options and the underlying stocks.

Acknowledgements

The editors on this manuscript were William Schwert and Nikolai Roussanov.

For helpful comments and suggestions we thank Torben Andersen, Bjorn Eraker, Scott Murray, Simon Schmickler, Peter van Tassel, Yinan Su, Kaushik Vasudevan, Stefan Windmüller, and conference participants at the FMA Derivatives & Volatility 2019 Workshop, the WFA Annual Meeting 2020, the Econometric Society World Congress 2020, and seminar participants at Warwick University. Matthias Büchner gratefully acknowledges the support of the Centre for Endowment Asset Management (CEAM) at Judge Business School, University of Cambridge. AQR Capital Management is a global investment management firm, which may or may not apply similar investment techniques or methods of analysis as described herein. The views expressed here are those of the authors and not necessarily those of AQR.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.jfineco.2021.12.007](https://doi.org/10.1016/j.jfineco.2021.12.007).

References

Bakshi, G., Kapadia, N., 2003. Delta-hedged gains and the negative market volatility risk premium. *Rev. Financ. Stud.* 16 (2), 527–566.

- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *J. Polit. Economy* 81 (3), 637–654.
- Broadie, M., Chernov, M., Johannes, M., 2007. Model specification and risk premia: evidence from futures options. *J. Finance* 62 (3), 1453–1490.
- Brooks, R., Chance, D., Shafaati, M., 2018. The Cross-Section of Individual Equity Option Returns. Unpublished working Paper. Louisiana State University.
- Cao, J., Han, B., 2013. Cross-section of option returns and idiosyncratic stock volatility. *J. Financ. Econ.* 108 (1), 231–249.
- Carr, P., Wu, L., 2004. Time-changed Lévy processes and option pricing. *J. Financ. Econ.* 71, 113–141.
- Carr, P., Wu, L., 2020. Option profit and loss attribution and pricing: a new framework. *J. Finance* 75 (4), 2271–2316.
- Chambers, D.R., Foy, M., Liebner, J., Lu, Q., 2014. Index option returns: still puzzling. *Rev. Financ. Stud.* 27 (6), 1915–1928.
- Christoffersen, P., Fournier, M., Jacobs, K., 2018. The factor structure in equity options. *Rev. Financ. Stud.* 31 (2), 595–637.
- Constantinides, G.M., Jackwerth, J.C., Perrakis, S., 2009. Mispricing of S&P 500 index options. *Rev. Financ. Stud.* 22 (3), 1247–1277.
- Coval, J.D., Shumway, T., 2001. Expected option returns. *J. Finance* 56 (3), 983–1009.
- Dew-Becker, I., Giglio, S., Kelly, B., 2021. Hedging macroeconomic and financial uncertainty and volatility. *J. Financ. Econ.* 142 (1), 23–45.
- Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68 (6), 1343–1376.
- Eraker, B., 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *J. Finance* 59 (3), 1367–1404.
- Eraker, B., Johannes, M., Polson, N., 2003. The impact of jumps in volatility and returns. *J. Finance* 58 (3), 1269–1300.
- Fournier, M., Jacobs, K., Orłowski, P., 2021. Modeling Conditional Factor Risk Premia Implied by Index Option Returns. Unpublished Working Paper. HEC Montreal.
- Frazzini, A., Pedersen, L.H., 2012. Embedded Leverage. NBER Working Paper 18558. NBER Working Paper Series.
- Goyal, A., Saretto, A., 2009. Cross-section of option returns and volatility. *J. Financ. Econ.* 94 (2), 310–326.
- He, Z., Kelly, B., Manela, A., 2017. Intermediary asset pricing: new evidence from many asset classes. *J. Financ. Econ.* 126 (1), 1–35.
- Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.* 6 (2), 327–343.
- Horenstein, A.R., Vasquez, A., Xiao, X., 2020. Common Factors in Equity Option Returns. Unpublished Working Paper. University of Miami.
- Israelov, R., Kelly, B., 2017. Forecasting the Distribution of Option Returns. Unpublished Working Paper. University of Chicago.
- Jones, C.S., 2006. A nonlinear factor analysis of S&P 500 index option returns. *J. Finance* 61 (5), 2325–2363.
- Karakaya, M., 2013. Characteristics and Expected Returns in Individual Equity Options. Unpublished Working Paper. University of Chicago.
- Kelly, B.T., Palhares, D., Pruitt, S., 2020. Modeling Corporate Bond Returns. Unpublished Working Paper. Yale School of Management.
- Kelly, B.T., Pruitt, S., Su, Y., 2019. Characteristics are covariances: a unified model of risk and return. *J. Financ. Econ.* 134, 501–524.
- Kelly, B.T., Pruitt, S., Su, Y., 2020. Instrumented Principal Component Analysis. Unpublished Working Paper. Yale University.
- Ledoit, O., Wolf, M., 2004. Honey, I shrunk the sample covariance matrix. *J. Portfolio Manage.* 30 (4), 1–22.
- Ofek, E., Richardson, M., 2003. DotComMania: the rise and fall of internet stock prices. *J. Finance* 58 (3), 1113–1137.
- Ofek, E., Richardson, M., Whitelaw, R.F., 2004. Limited arbitrage and short sales restrictions: evidence from the options markets. *J. Financ. Econ.* 74 (2), 305–342.
- Vasquez, A., 2017. Equity volatility term structures and the cross-section of option returns. *J. Financ. Quant. Anal.* 52 (6), 2727–2754.
- Windmüller, S., 2021. Firm characteristics and global stock returns: a conditional asset pricing model. *Rev. Asset Pricing Stud.*
- Zhan, X.E., Han, B., Cao, J., Tong, Q., 2021. Option return predictability. *Rev. Financ. Stud.*