A Non-Linear Market Model*

Tobias Sichert

Department of Finance

Stockholm School of Economics

tobias.sichert@hhs.se

This draft: October 19, 2023

First draft: May 31, 2022

Abstract

This paper shows that most market-data based stock portfolio sorts have a strong tail risk exposure in both tails, while accounting data-based sorts do not. An option-inferred model with non-linear pricing of market risk captures the tail risk patterns and explains many prominent cross-sectional stock return anomalies such as momentum, betting-against-beta, idiosyncratic volatility, and liquidity. A key feature of the model is a sizable upside risk premium of approximately 4% per annum. Finally, the pricing results can be explained with compensation for exposure to systematic variance risk.

JEL classification: G10, G12, G13

Keywords: Pricing kernel, anomalies, upside risk, variance risk, risk premia, skewness, tail risk

*I thank Daniel Andrei, Magnus Dahlquist, Bob Dittmar, Mathieu Fournier (discussant), Stefano Giglio, Amit Goyal, Cam Harvey, Michael Hasler, Burton Hollifield, Jan Pieter Krahnen, Lars-Alexander Kühn, Lasse Pedersen, Simon Rottke, Riccardo Sabbatucci, Christian Schlag, Maik Schmeling, David Schreindorfer, Julian Thimme, Christian Wagner, Rüdiger Weber, Irina Zviadadze (discussant) conference participants of the AFA, DGF, JEF, NFA, SED, and seminar participants at the universities of Arizona State, Carnegie Mellon, Collegio Carlo Alberto, Göteborg, Goethe, Nova SBE, Princeton, Stockholm School of Economics, Toronto, Zürich and Warwick for valuable comments and suggestions.

I. Introduction

Although the linear Capital Asset Pricing Model (CAPM) is clearly one of the most popular models of modern financial economics, abundant empirical evidence suggests that the model is misspecified. This has spurred the search for alternative sources of systematic risk. One obvious approach is to consider a non-linear model, but it requires finding ways to discipline the potentially infinite number of non-linear pricing functions. A common approach is to assume a representative agent and derive the non-linearity implied by the agent's marginal utility function. Prominent examples include coskewness risk (Harvey & Siddique 2000), cokurtosis risk (Dittmar 2002) and downside risk (Ang et al. 2006a, among others). However, these parametric approaches often cannot explain many of the existing return anomalies. Moreover, standard utility functions rule out non-monotonic pricing functions that have empirical support from the options data (Jackwerth 2000, Aït-Sahalia & Lo 2000, Rosenberg & Engle 2002).

This paper suggests an alternative approach to obtain a non-linear market model directly from the data without imposing any structure on the pricing relationship. For this, I rely on the S&P 500 options market, where non-linear payoffs are traded and priced. To infer the pricing of market risk from the data, I use standard tools from the literature on the estimation of the option-implied pricing kernel by combining index options data with a return density forecast. My paper is the first to connect this pricing kernel to the cross-sectional pricing of stock returns.

The empirical results show that the non-linear pricing of market risk inferred from option prices can correctly price a large set of prominent stock return anomalies, both in US and European data. In particular, this includes stock portfolios sorted on CAPM-beta, volatility, idiosyncratic volatility, momentum, turnover, zero-trade, sensitivity to aggregate volatility, past maximum return and industry, which all have large anomalous returns relative to the CAPM. The pricing kernel that sets the anomaly alphas to zero is only a function of the market return, and its functional form is estimated fully out-of-sample using only index options and index return data, without using any cross-sectional return information.

A well-known property of the estimated option-implied pricing kernel is its non-monotonic U-shape (Bakshi et al. 2010, Christoffersen et al. 2013, Song & Xiu 2016), which is illustrated in Figure 1. The interesting feature is that the pricing kernel is increasing in the area of positive returns, which implies an upside risk premium. I show that this upside risk premium, which is

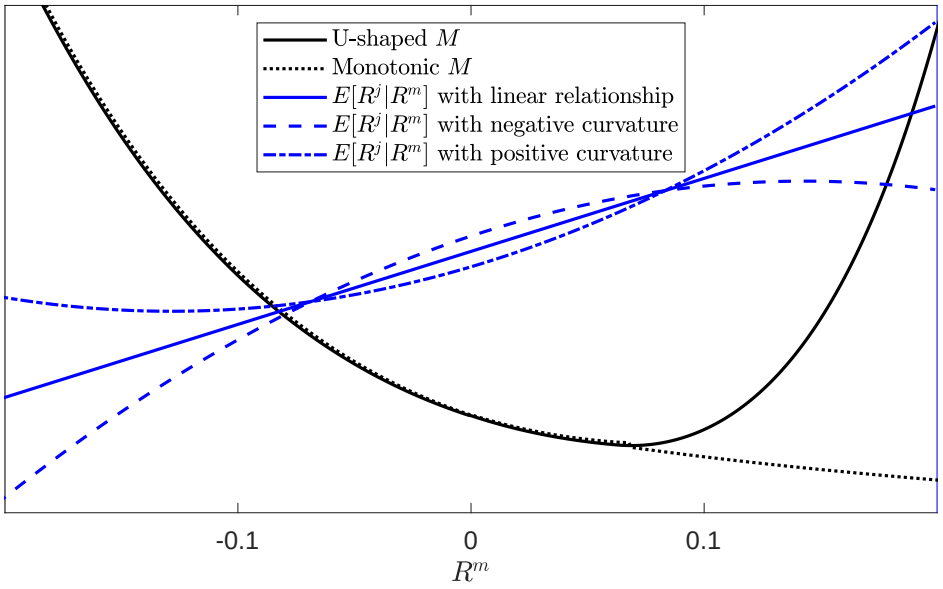

Figure 1: **Stylized pricing kernels and return relationships with curvature.** The figure illustrates a stylized U-shaped pricing kernel M, monotonic M, together with the returns of asset j that have either a linear relationship, or positive and negative return curvature relative to the market return R^m.

absent in existing factor models, explains on average half of the alphas of both the CAPM and the Fama & French (1993) three factor model, and up to 89% in case of the low beta anomaly. Furthermore, upside risk is on average associated with a statistically significant risk premium of 4.0% p.a.

The prevalence and relative magnitude of upside risk premium is strikingly consistent across other asset classes. In particular, when using the same pricing kernel to price several options and commodity portfolios, I again find a significant upside risk premium of about 50% of the CAPM alpha. Finally, applying the analogous approach to European equities portfolios delivers very similar results than for the US and confirms the results out-of-sample.

To interpret my findings, I provide a structural model that generates a non-monotonic pricing kernel due to exposure to variance risk. An empirical test of this channel that uses variance risk as a second (non-linear) factor can reconcile the observed anomalies.

The theoretical analysis starts with the derivation of two properties of the data-generating process that are necessary for non-linear pricing of risk to outperform linear pricing. These two conditions are, first, a systematic non-linear dependence between the market return and anomaly returns, and second, that dependence must be inversely related to the pricing kernel.

Turning to the data, I document a striking feature of many prominent anomalies which I refer to as "curvature patterns". This pattern is graphically illustrated in Figure 1. Those portfolios that on average generate positive CAPM alphas exhibit a negative curvature of their returns relative to the market, and vice versa. As can be seen from Figure 1, portfolios with negative curvature underperform the linear relationship both in times of large positive and large negative market returns.

Schneider et al. (2020) have documented this pattern for the low beta and the idiosyncratic volatility anomaly. I show that this pattern is much more widespread, by studying all 95 portfolio sorts considered in Gu et al. (2020) in my sample ranging from 1996-2019. I find that most portfolio sorts where the sorting variable is based on market data (stock prices or trading volume) have strong curvature patterns. These are 16 sorts in total, and include past maximum return, various measures of momentum and stock liquidity, and industry sorts. In contrast, portfolio sorts where the sorting variable is based on accounting data, which are the majority of commonly used variables, do not display any such curvature patterns.

To show the importance of this finding, recall that the risk premium (i.e., the expected excess return) R^j on any asset j can be written as:

$$E\left[R^j\right] = -\text{Cov}\left[R^j, M\right] \times R_f. \tag{1}$$

Therefore, all else equal, a more negative curvature implies that the portfolio returns are more negatively correlated with a U-shaped pricing kernel M, and therefore require a higher risk premium. Focusing on the upside, the negative curvature together with the U-shaped M induce an additional upside risk premium. Those portfolios that have a positive curvature, on the contrary, are hedge assets in both tails and therefore earn low returns on average.

While numerous studies show that downside risk is priced in the cross-section of stock returns (e.g., Ang et al. 2006a, Lettau et al. 2014, Kelly & Jiang 2014, Farago & Tédongap 2018), upside risk has received much less attention in the literature. Although a few papers have extended their stock-level downside risk measure to an upside risk measure, they generally find weak and insignificant upside risk premia on the individual stock level (Ang et al. 2006a, Chabi-Yo et al. 2018). In contrast, I find that the upside risk is significantly priced at the portfolio level. This divergence is related to the finding that many well-know portfolios have

strong curvature patterns and hence tail risk exposure, while measures of tail exposure have low persistence on the stock level.

To interpret my pricing results economically, I link them to an explanation based on variance risk, which is arguably the most important risk factor for option pricing. The finding of a non-monotonic option-implied pricing kernel is often referred to as the "pricing kernel puzzle", since it is at odds with the monotonicity predicted by many classical theories. A promising explanation for this puzzle is variance risk, as suggested by Chabi-Yo (2012), Christoffersen et al. (2013), and Bakshi et al. (2022). In particular, Christoffersen et al. (2013) propose a pricing kernel which is decreasing in the market return and increasing in return variance. The existence of a variance risk premium in combination with a U-shaped relationship between returns and variances then induces a U-shaped projection of the pricing kernel onto index returns. This explanation is appealing, since the variance risk premium is empirically well-documented, and the enormous trading volumes in options and other volatility related derivatives show that investors are clearly concerned with variance risk.

However, from the perspective of existing equilibrium models the U-shaped pricing kernel is still a puzzle, and it is unclear whether these relationships from reduced form models can also be obtained in a structural model. To close the gap and illustrate the link, I show that a modified Bansal & Yaron (2004) long-run risks model can generate a U-shaped pricing kernel, by introducing GARCH type dynamics for the variance of consumption growth. The key feature of these dynamics is that they have a feedback mechanism where consumption shocks impact future variance, leading to an increase in variance following both large positive and negative consumption shocks.[1] I provide supporting evidence for this property in consumption data. As dividends are modeled as levered consumption, the new feature makes large positive market returns associated with increases in variance on average. Due to her Epstein & Zin (1989) preferences, the representative agent dislikes high variance, and hence these states of the world are associated with high marginal utility. In other words, a high positive return is indicative of an increase in the variance of fundamentals, which the agent dislikes more than the large positive consumption growth associated with it.

Building on these arguments, I empirically test whether my pricing results are consistent

[1] Tédongap (2015) also studies a long-run risks model with GARCH dynamics, but does not study the projected pricing kernel or upside risk.

with a compensation for investors' exposure to systematic variance. Specifically, I use the returns to a variance contract, i.e., the asset that costs VIX^2 and pays the monthly realized variance, as an additional pricing factor. This extended and U-shaped market model again delivers low and insignificant pricing errors, which shows that exposure to systematic variance can indeed explain the otherwise anomalous portfolio returns. In other words, referring back to the curvature patterns from Figure 1, many anomalies exhibit a payoff profile that resembles a short straddle, and - just like the options position - they earn positive excess returns relative to the CAPM in compensation for losses in both tails. As extreme returns in both tails tend to coincide with spikes in realized variance, a variance-averse investor dislikes these states of the world and hence requires an additional risk premium for stocks with negative curvature.[2]

This explanation differs from existing studies on the pricing of market variance risk in equities in two ways. First, based on an ICAPM argument, several studies show that shocks to future expected variance are priced in the cross-section (Ang et al. 2006b, Bansal et al. 2014, Campbell et al. 2018). In contrast, I show that using realized variance as an additional factor can explain numerous return anomalies. Second, I document that realized variance has a U-shaped relationship relative to the contemporaneous index return, and hence the model generates a U-shaped pricing kernel. In contrast, shocks to future expected variance have a monotonically decreasing relationship relative to the index return and therefore do not generate upside risk.

The variance risk explanation brought forward in this paper is related to the coskewness model of Kraus & Litzenberger (1976) and Harvey & Siddique (2000). In a more recent empirical test, Schneider et al. (2020) show that the model can explain the low beta and the idiosyncratic volatility anomaly. I differ from these studies in several ways. First, my model can explain a substantially larger number of anomalies as well as the returns of other asset classes. In contrast, I show that the coskewness factor cannot explain most of these anomalies. Second, my approach allows me to study the effect of upside and downside risk separately, while most existing approaches cannot differentiate between the two. In fact, while the coskewness literature usually emphasizes the downside risk, I highlight that upside risk is substantial and relevant

[2]The findings further suggest that the pricing of risks in the options and stock markets are consistent, and that the two markets are integrated. This is in contrast to several papers that argue that the two markets are partially segmented (e.g., Bollen & Whaley 2004, Garleanu et al. 2008, Frazzini & Pedersen 2022, Dew-Becker & Giglio 2022).

for understanding anomaly returns. Third and related, I provide an economic explanation for upside risk via variance risk. In the coskewness model the upside risk emerges from a second order Taylor-series expansion of a utility function over returns. However, as long as utility is monotonically increasing (nonsatiation) - which is the case for standard utility functions - the true pricing kernel cannot be U-shaped.[3] Fourth, my approach does not rely on a factor mimicking portfolio to measure the market price of risk, but obtains it directly from the options data.

The remainder of the paper proceeds as follows. Section II studies non-linear pricing of market risk theoretically and provides empirical evidence for curvature patterns in returns. Section III introduces the estimation methodology for the pricing kernel along with the estimation results. The pricing results and their robustness are presented in Section IV. Section V shows that a variance risk-based explanation can empirically rationalize the results. Section VI outlines a equilibrium framework that generates an upside risk premium, and Section VII concludes. The appendix collects technical details, and the internet appendix contains derivations, further results and robustness checks.

[3]Several theories that depart from rational expectation also imply qualitatively that stocks with positive skewness of returns on average have low returns. Examples include the endogenous optimal beliefs model of Brunnermeier et al. (2007), the heterogeneous skewness preference model of Mitton & Vorkink (2007), and the cumulative prospect theory model of Barberis & Huang (2008). In contrast to these theories, I provide a quantitative explanation based on variance risk. In addition, I focus on systematic risk on a portfolio level, while the theories focus on idiosyncratic skewness on the stock level, which quickly diversifies away when aggregating stocks into portfolios.

II. Non-linear Pricing and Curvature Patterns in Asset Returns

This section discusses theoretically under which conditions non-linear pricing of risk matters. I then verify the first condition in the data, namely, that returns of trading strategies have a systematic, non-linear relationship relative to the index returns. For this, I first present a detailed empirical example using the idiosyncratic volatility anomaly, and then generalize to many equity portfolios and other asset classes. The second condition is verified in Sections III and IV.

A. Conditions for non-linear pricing to matter

Let $M_{t+1} = M_t(R^m_{t+1})$ denote a pricing kernel that is a potentially non-linear function of the market excess return R^m such that the Euler equation

$$E_t[M_{t+1}R^m_{t+1}] = 0 \tag{2}$$

holds. In principle, the pricing kernel can be a function of many sources of risk, and hence $M(R^m)$ denotes only the pricing kernel projected onto R^m. As discussed in Cochrane (2005), the projected pricing kernel has the same pricing implications as the original kernel for all assets which payoffs depend only on R^m. Under which conditions it prices any other asset's excess return R^j depends on the structure of the economy.

To address this question, one can formally ask if also

$$E_t[M_{t+1}R^j_{t+1}] = 0 \tag{3}$$

holds. To proceed, I have to make further assumptions on the data-generating process. To get an intuition on the relevance of non-linear pricing, consider the following setup:

$$R^j_{t+1} = \beta^j R^m_{t+1} + \epsilon_{t+1}, \tag{4}$$

and three potential assumptions on the CAPM pricing errors ϵ:

$$E_t[\epsilon_{t+1}|R^m_{t+1}] = 0, \tag{5}$$

$$E_t[\epsilon_{t+1}|R^m_{t+1}] = \alpha \neq 0, \tag{6}$$

$$E_t[\epsilon_{t+1}|R^m_{t+1}] = g(R^m), \quad E_t[\epsilon_{t+1}] = \alpha \neq 0, \quad \mathrm{Cov}_t(\epsilon_{t+1}, R^m_{t+1}) = 0, \tag{7}$$

where $g(\cdot)$ denotes a non-linear function. In the first case, any $M(R^m)$ that satisfies (2) will also satisfy (3), regardless of its functional form (all proofs are provided in the internet appendix IA.1.1). In the second case, ϵ contains one or several sources of systematic risk that are included in the unprojected pricing kernel and command additional risk premia beyond the market. But these risks are not related to R^m, and hence (3) will not hold. In the empirical test below, book-to-market or profitability-sorted stock portfolios appear to be such an asset.

The third case is the relevant one for this paper. Here, the conditions specify that ϵ is non-zero in expectations and has a systematic but non-linear relationship with R^m. The conditions rule out that a linear $M(R^m)$ can price R^j, and $g(R^m)$ will show up in the pricing error alpha. A simple example for this case would be $\epsilon_{t+1} = (R^m_{t+1})^2$, with R^m following a symmetric distribution. However, a non-linear $M(R^m)$ could potentially capture the dependence $g(R^m)$ and hence price R^j. Formally, plugging the structure of R^j from (4) and (7) into (3) yields:

$$E_t[M_{t+1}R^j_{t+1}] = E_t[M_{t+1}\beta^j R^m_{t+1}] + E_t[M_{t+1}g(R^m)] \tag{8}$$

$$= 0 + E_t[M_{t+1}g(R^m)]. \tag{9}$$

The last term in (9) will be non-zero for a linear pricing model under assumptions (7), but can be zero if $M(R^m)$ is inversely related to $g(R^m)$(IA.1.1 provides an example).

In other words, in order for non-linear pricing to be relevant, one has to find systematic, non-linear relationship $g(R^m)$ between the CAPM pricing errors ϵ and R^m, that are in addition inversely related to $M(R^m)$, i.e., covary with plausible sources of risk. The following sections will show the first property in the data, while the second property is studied in the Sections IV and V.

Finally, the third case (7) can have two different economic mechanisms and interpretations. Either purely $M(R^m)$ captures how market risk is priced in a non-linear way, or there is a

second priced risk factor that has a non-linear dependence with R^m such as, e.g., variance risk. The empirical results in Sections V and VI point towards the second interpretation.

B. Empirical curvature patterns: example

I will illustrate the case of (7) together with the negative dependence between $g(R^m)$ and $M(R^m)$ using the example of the idiosyncratic volatility (Ivol) anomaly of Ang et al. (2006b) which refers to the empirical fact that portfolios comprised of stocks with high Ivol have negative CAPM alphas, and vice versa. To motivate my measure for the non-linearity $g(R^m)$ I draw on the key stylized property from the option-implied pricing kernel estimation in Section B below, namely, that the empirical estimates of $M(R^m)$ are predominately U-shaped. The simplest way to represent this properties is:

$$M_{t+1}^{approx} = a + b \times R^m + c \times (R_{t+1}^m - d)^2, \tag{10}$$

where the parameter d can capture the finding that the U-shape is not symmetric. In (10), the third term specifies the deviation from the linear CAPM, and hence the implied non-linearity is $g(R^m) \equiv (R_{t+1}^m)^2$, which measures curvature in returns.[4]

To quantify the curvature patterns in the data, I first run the standard CAPM regression for each portfolio j of the ten Ivol-sorted portfolios:

$$R_{t+1}^j = \alpha^j + \beta^j \times R_{t+1}^m + \epsilon_{t+1}, \tag{11}$$

and second, the CAPM regression extended by a square term on the market return:

$$R_{t+1}^j = constant + \tilde{\beta}^j \times R_{t+1}^m + \gamma^j \times \left(R_{t+1}^m\right)^2 + \epsilon_{t+1}. \tag{12}$$

[4]It is evident that my measure of curvature is related to the coskewness measure of Harvey & Siddique (2000). I emphasize that it is merely a means to illustrate both the pervasive deviations from a linear relationship and the systematic relationship between tail exposure and CAPM alphas. Using the measure does not imply or assume that coskewness it the true pricing model. In fact, one obtains similar results when using alternative measures for $g(R^m)$, such as a piece-wise linear function, a quartic $(R^m)^4$-term, or sensitivity to realized variance. I prefer to use the quadratic $g(R^m)$ due to its simplicity to capture co-movement in both tails. In contrast to prior studies, I do not use any of these measures for pricing or for quantifying risk premia. Instead, the risk premia are inferred from the options market. Hence I use the term curvature instead of coskewness to highlight the different interpretation of the relationship.

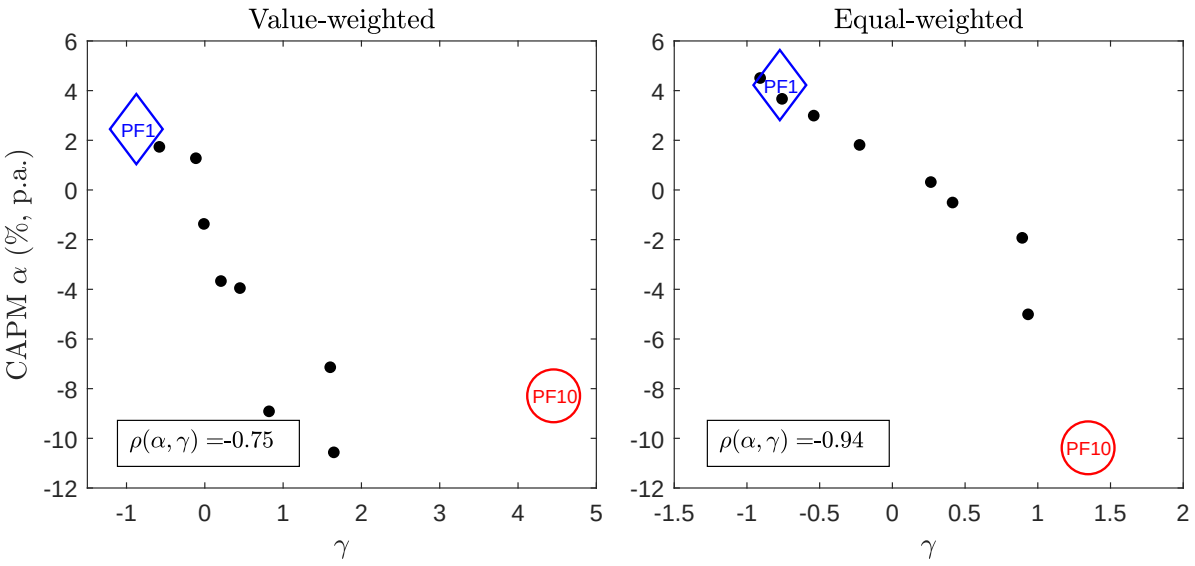

Figure 2: **CAPM alphas and curvature for Ivol-sorted portfolios.** The figure plots the estimated CAPM alphas from (11) against the estimated γ coefficients from (12). The left plot shows results for value-weighted returns of the Ivol-sorted stock portfolios, and the right shows the equal-weighted analogue. The blue diamonds mark the first portfolio comprised of the low Ivol stocks, and the red circles the tenth portfolio that contains the high Ivol stocks at each date. The legend reports the correlation between the alphas and γ's of each plot. Returns are monthly and the time period is 1996-2019.

The additional γ parameter measures the curvature of the relationship between R^j and R^m.

Figure 2 illustrates the results, and one observes a strong negative relationship between a portfolio's CAPM alpha and its curvature γ. The high Ivol portfolio has a highly positive curvature and hence tends to outperform relative to the CAPM both when the market has large negative and large positive returns, while the opposite is true for the low Ivol portfolio. Hence, the CAPM errors indeed have a non-linear relationship $g(R^m)$ that is inversely related to the pricing kernel in (10).

This curvature pattern, which is illustrated graphically in Figure 1, has an analogue in the options market: The low Ivol portfolio qualitatively resembles the payoff of a short straddle, while the high Ivol portfolio resembles a long straddle. Empirically, it is well-documented that the Ivol anomaly as well as short straddles earn excess returns relative to the CAPM on average, while the opposite is true for high Ivol and long straddle, respectively. In the presence of a U-shaped $M(R^m)$ the low Ivol portfolio is particularly risky, while the high Ivol portfolio offers a hedge. To foreshadow the results, this property is exactly what the option-implied pricing kernel below will capture and price correctly.

C. Evidence for equities and other asset classes

C.1. Approach and Data

Starting with US equities, my goal is twofold. On the one hand, I want to identify and study prominent anomalies that have strong curvature patterns. For this, I start with the 95 sorting variables considered in Gu et al. (2020). While I will discuss the results in more detail later, for now, to motivate my sample selection, I note that most portfolio sorts where the sorting variable is based on market data (stock prices and trading volume) have strong negative curvature patterns, while accounting data based sorts do not. Out of those sorts with negative curvature patterns, I select the most prominent anomalies, and include the remainder in the internet appendix. On the other hand, I want to study other prominent anomalies without any curvature pattern, i.e., well-known accounting based anomalies.

As a result, my benchmark data includes portfolios of stocks sorted on their CAPM beta (sometimes referred to as "betting-against-beta"), volatility (Vol) and idiosyncratic volatility (Ivol) as in Ang et al. (2006b), past maximum daily returns (Max) as in Bali et al. (2011), 12m-1m momentum as in Jegadeesh & Titman (1993), turnover (Turn) as in Datar et al. (1998), and days with zero trading volume (Ztrade) as in Liu (2006). Finally, I add beta w.r.t. changes in the VIX index ($\beta_{\Delta \text{VIX}}$) as in Ang et al. (2006a), since the idea behind their variable is to measure exposure to systematic variance risk, which is related to the variance-risk based explanation in this paper.

As further common test assets that largely do not have any curvature patterns, I include portfolios sorted on size, book-to-market (B/M), investment (Inv), profitability (Prof), accruals (Acc) and industry portfolios, each calculated following the definition on Kenneth French's website. Return data are from CRSP, accounting data from COMPUSTAT, and factor returns are taken from Kenneth French's website. Returns are monthly and calculated such that they match the timing of the option expiration. For each variable I sort stocks in decile portfolios. Further details regarding the calculation of characteristics and the portfolio formation are provided in Appendix D.

To get more out-of-sample evidence I use international stock market data. The sample is from the Compustat Global Securities database and comprises 16 developed European markets.[5]

[5]The countries are Austria, Belgium, Denmark, Finland, France, Germany, Great Britain (the United

I focus on European countries, since my approach requires a representative stock market index as well as an active options market to calculate the option-implied pricing kernel. All values are in USD, and the European factor data is again obtained from the Kenneth French's website. The data is from 2002-2019, which matches the available EuroStoxx 50 options data.

For options, I construct four straddle portfolios with initial moneyness $K/S_0 \in \{0.90, 0.95, 1.00, 1.05\}$. At each date, I use the S&P 500 call and put option with moneyness closest to the target moneyness. For the portfolios where one of the two options is in-the-money, I verify that the recorded option prices are very close to a synthetic option price calculated using the implied volatility of the second, out-of-the-money option. Since both excess returns of long straddles and their CAPM alphas are negative on average, I consider a short position. Straddles are particularly suited to study the tail exposure, since they naturally have a strong convex return relationship with the market.[6]

For commodities, I use returns to commodity spot indices from the Commodities Research Bureau, obtained from Bloomberg. This includes the Spot Index, and six sub-indices on Metals, Textiles, Industrials, Foodstuffs, Fats and Oils, and Livestock.

C.2. Results

To quantify the curvature patterns, I first run regressions (11) and (12) for each portfolio. Then, for equities, I calculate the correlation between α^j from (11) and γ^j from (12) across the ten portfolios of each strategy, both for value-weighted (VW) and equal-weighted (EW) portfolios. For options and commodities I calculate the correlation across the respective four and seven portfolios.

The results in Panel A of Table 1 show that for portfolios sorted on Beta, Vol, Ivol, Max, Mom, Turn, Ztrade, and $\beta_{\Delta \text{VIX}}$ there is a strong negative correlation between the portfolios' α and γ. Furthermore, the correlation is stronger for equal-weighting and often close to -1.

To appreciate the economic magnitude of these patterns, Panel F in Table 1 presents the γ^{HML} of the high-minus-low (HML) portfolio. Consider for example the case of Ivol. The

Kingdom), Italy, Ireland, Luxembourg, the Netherlands, Norway, Portugal, Spain, Sweden, and Switzerland.

[6]Note that the fact that M is derived from the same options data does not necessarily imply that M prices these returns correctly. Only if the return density $f_t(R_{t+1})$ used for the estimation in (13) below is the true one, or at least close to it, M will also price option returns.

Table 1: **Curvature Patterns for Various Asset Classes**

Panel A-E show the correlation coefficient $\rho(\alpha, \gamma)$ between the CAPM alpha from (11) and the curvature parameter γ from (12) across portfolios in the first row of each panel. For equities, the correlation is calculated across the ten portfolios of each sort. VW denotes value-weighted portfolios, and EW denotes equally-weighted portfolios. For options and commodities, the curvature is calculated across the respective four and seven portfolios. For a more detailed description of the portfolio calculation, see Appendix D. Panel F shows the magnitude of curvature parameters γ from (12) for the value-weighted portfolios from Panel A. The first two rows show the lowest and highest γ for each sort, while the last row shows the γ for the high-minus-low (HML) portfolio. t-statistics (in parenthesis) are adjusted for heteroscedasticity and autocorrelation (Newey & West 1987).

Panel A: US Equities with curvature patterns								
	Beta	Vol	Ivol	Max	Mom	Turn	Ztrade	$\beta_{\Delta \text{VIX}}$
VW	-0.75	-0.85	-0.73	-0.94	-0.76	-0.68	-0.70	-0.75
EW	-0.99	-0.96	-0.93	-0.95	-0.89	-0.97	-0.96	-0.94
Panel B: Other Equity								
	Size	B/M	Inv	Prof	Acc	Industry		
VW	0.61	-0.10	-0.43	0.29	-0.28	-0.54		
EW	0.06	-0.08	-0.58	-0.25	-0.12	0.03		
Panel C: European Equity								
	Beta	Vol	Ivol	Max	Mom	Turn	Ztrade	$\beta_{\Delta \text{VIX}}$
VW	-0.78	-0.89	-0.71	-0.83	-0.80	-0.66	-0.70	-0.77
EW	-0.97	-0.97	-0.96	-0.95	-0.94	-0.89	-0.89	-0.76
Panel D: Options				Panel E: Commodities				
	-0.63			-0.88				
Panel F: Magnitude of Curvature Patterns for VW from Panel A								
	Beta	Vol	Ivol	Max	Mom	Turn	Ztrade	$\beta_{\Delta \text{VIX}}$
Min(γ)	-1.50	-1.57	-0.87	-1.57	-0.79	-1.32	-1.00	-0.78
	(-4.01)	(-4.25)	(-3.15)	(-5.61)	(-4.63)	(-4.56)	(-3.58)	(-2.32)
Max(γ)	2.05	2.75	4.45	2.51	3.01	1.48	1.41	1.23
	(2.36)	(3.21)	(6.11)	(3.93)	(2.28)	(2.51)	(2.45)	(2.21)
γ^{HML}	-3.55	-4.32	-5.33	-4.02	-2.59	-2.16	-1.89	-1.30
	(-3.75)	(-4.63)	(-6.83)	(-5.46)	(-1.66)	(-3.22)	(-2.81)	(-1.75)

$\gamma = -5.3$ implies that for a for a $R^m = \pm 10\%$ the Ivol HML portfolio is expected to underperform the linear relationship by $-5.3 \times 0.1^2 = -5.3\%$, and for a relatively large $R^m = \pm 20\%$ the non-linearity amounts to $-5.3 \times 0.2^2 = -21.1\%$ per month.

The documented correlation between alpha and curvature across portfolios confirms that there is a systematic pattern, and that these sorts are systematically related to an additional

source of risk besides the linear market factor. If a negative γ^{HML} would result just from chance, it would likely not appear jointly with the correlation pattern across all ten portfolios. On the flip side, there are a few sorts that exhibit a correlation pattern, but the γ^{HML} is insignificant. This could stem from a non-monotonic exposure across portfolios, and the online appendix presents an example.

Another interesting insight from Panel F of Table 1 is that the sorts based on $\beta_{\Delta\text{VIX}}$ exhibit the lowest in γ^{HML}. This finding is surprising, since $\beta_{\Delta\text{VIX}}$ should measure an exposure to systematic variance, and one might hence expect strong dependence in the tails of the return distribution. Section V revisits this point and shows to what extent $\beta_{\Delta\text{VIX}}$ indeed measures exposure to (monthly) systematic variance risk.

Panel B of Table 1 shows that for size, book-to-market, profitability, and accruals portfolios, there is no, or even a positive, relationship between α and γ. For investment, and industry-sorted portfolios there is a negative relationship for the value-weighted portfolios, but it is not as strong as for the sorts in Panel A. Turning to European stocks, the results in Panel C confirm the $\alpha-\gamma$ correlation pattern found in US data. Lastly, Panel D and E document the same results for straddles and commodities. While several studies have documented downside-dependence for some anomalies[7], the novel result here is the upside-dependence. Moreover, Table IA.4 documents that the curvature patterns are also obtained when estimating a piece-wise linear function that allows for different coefficients for upside and downside dependence, emphasizing, that the relationship is really convex/concave in both tails and not mechanically induced by the $(R^m)^2$-term.

Finally, turning to the curvature patterns in the full set of all 95 sorting variables considered in Gu et al. (2020), the results paint an interesting picture. All accounting based sorts, which are the vast majority of sorts, display no strong negative curvature patterns. In contrast, most sorts based on market data (prices and volume) do have strong negative curvature patterns. In addition to the variables studied in Panel A of Table 1, this includes several other measures of momentum and stock liquidity. In total, 16 out of 22 sorts that use only market data have strong curvature patterns. Moreover, all of the remaining six sorts of the group have low and insignificant alphas in the sample, and hence my model would also not predict them to have significant curvature patterns. For brevity, all details are relegated to Table IA.1.

[7]See, e.g., Boguth et al. (2011), Lettau et al. (2014).

III. Pricing Kernel Estimation

This section briefly explains the approach for estimating the pricing kernel as a function of stock market returns, estimation of conditional risk-neutral and physical return densities, discusses data sources, and presents the estimation results.

A. Approach

It is well known that the absence of arbitrage implies the existence of a pricing kernel, or stochastic discount factor, that prices all assets. This paper studies the (non-linear) projection of the pricing kernel onto stock market returns, defined as:

$$E_t[M_{t+1}|R^m_{t+1}] = \frac{1}{R^f_t}\frac{f^*_t(R^m_{t+1})}{f_t(R^m_{t+1})}, \tag{13}$$

where $f^*_t(R^m_{t+1})$ and $f_t(R^m_{t+1})$ denote the conditional risk-neutral and physical density of the ex-dividend market R^m_{t+1}, respectively, and R^f_t is the gross risk-free rate. This object is the expected value of the pricing kernel in $t+1$, conditional on observing a given return R^m_{t+1}, and hence $M(R^m)$ is only a function of R^m. For brevity, I use the shorthand notation M to denote the ex ante expectation, and M_{t+1} to denote realizations. Furthermore, M without superscript refers to the option-implied estimate from (13), and I use specific superscripts to denote alternative specifications.

To estimate f^* and f I use standard tools from the literature which I next briefly summarize. The interested reader is referred to Appendix A for more details and descriptive statistics. First, for each month, I extract f^* from S&P 500 option prices, using the classical results of Banz & Miller (1978) and Breeden & Litzenberger (1978). This method is standard, and the obtained densities are truly conditional as they reflect only option information from a given point in time. The data is from OptionMetrics and ranges from 1996-2019. Second, to construct f, I use a method often referred to as "filtered historical innovations", which is semi-parametric and the most common methodology in the literature on pricing kernel estimation.[8] Data is daily S&P 500 returns obtained from CRSP.

[8]See, e.g., Rosenberg & Engle (2002), Barone-Adesi et al. (2008), Christoffersen et al. (2013), Faias & Santa-Clara (2017) or Christoffersen et al. (2022).

Third, to make f_t a conditional density, I require a conditional volatility forecasts. Since this is a key ingredient for the estimation of M (Sichert 2023), this paper spends some effort to obtain a good model to forecast volatility.

There is an extensive econometric literature on volatility modeling and forecasting. By now it is generally understood that models based on high frequency realized variances (RV) outperform other models, as e.g., the standard GARCH models. I follow Bekaert & Hoerova (2014) and compare out-of-sample performance of numerous state-of-the art RV models for the S&P 500 at the one month horizon. The best RV model contains realized jumps and the VIX as predictors. The details of the model specification and estimation are provided in the Appendix B, and alternative RV models are included in the robustness Section F.

Throughout the paper, I will refer to the variance forecast from t to $t+1$ (21 trading days) as expected variance over the next month, denoted as σ_t^2. Lastly, it is important to note that in the following, all quantities in (13) are calculated in a rolling window fashion, i.e., using only data available up to date t, and hence M is estimated in real time.

The benchmark approach estimates f^* only over the range of strikes where options data exists. In the sample, the realized return is never lower than the lowest observed strike, and hence the extrapolation of the left tail is not relevant for the pricing results below. However, high positive returns do exceed the highest observed strike nine times in the sample, on average by 1.2 percentage points. Therefore a robust way of extrapolating either f^* or M in the right tail of the distribution is required. My benchmark approach is to simply extrapolate the observed M linearly. Alternatively, I use the ratio of the cumulative return density in the tails that is not covered by the options data. Formally:

$$M(R^{right\ tail}) = \Big(1 - F_t^*\big(R_{t,t+1}^{max}\big)\Big) / \Big(1 - F_t\big(R_{t,t+1}^{max}\big)\Big), \tag{14}$$

where R_{t+1}^{max} correspond to the highest available strike at date t, and F denotes the cumulative density function (CDF). This ratio provides an indication of the behavior of M in the tail, and it can be interpreted as the average M in that region.

These extrapolation approaches of M are more robust then extrapolating f^* by "tail fitting", which makes the estimated M in the tail strongly dependent on guessing the right parametric distribution for the tails. I test alternative ways to extrapolate M - but also f^* - in the

robustness section.

B. Estimation results

Figure 3 presents the estimated pricing kernels M for each month, grouped by the 24 years in the sample. The pricing kernel estimates are relatively steep on the left, and predominately U-shaped. The U-shape is most pronounced in times of high market volatility, e.g., around the 2000's and during the financial crisis and its aftermath 2008-2011. In times of low volatility, the estimates sometimes have a hump around zero, which is likely caused by a too high volatility forecast (see Sichert (2023) for a more detailed discussion of this issue). The argument is consistent with the observation that this hump is most pronounced in the years 2017-2018, where both market volatility and the VIX were at their all time low. Such an all time low of volatility is hard to capture with an out-of-sample prediction, where the model never saw these low levels in the estimation. Nevertheless, most estimated M tend upwards towards their right tail.

The steepness of M on the left is consistent with the highly negative returns to S&P 500 put options (Broadie et al. 2009, among others). The increasing part on the right is consistent with the negative average returns of out-of-the-money S&P 500 call options (Bakshi et al. 2010). In the literature on the cross-section of stock returns, the steepness of M on the left is consistent with evidence on the importance of downside risk. On the contrary, the implications for stock returns of the increasing part on the right has received little attention, but will be the focus of the following analysis.

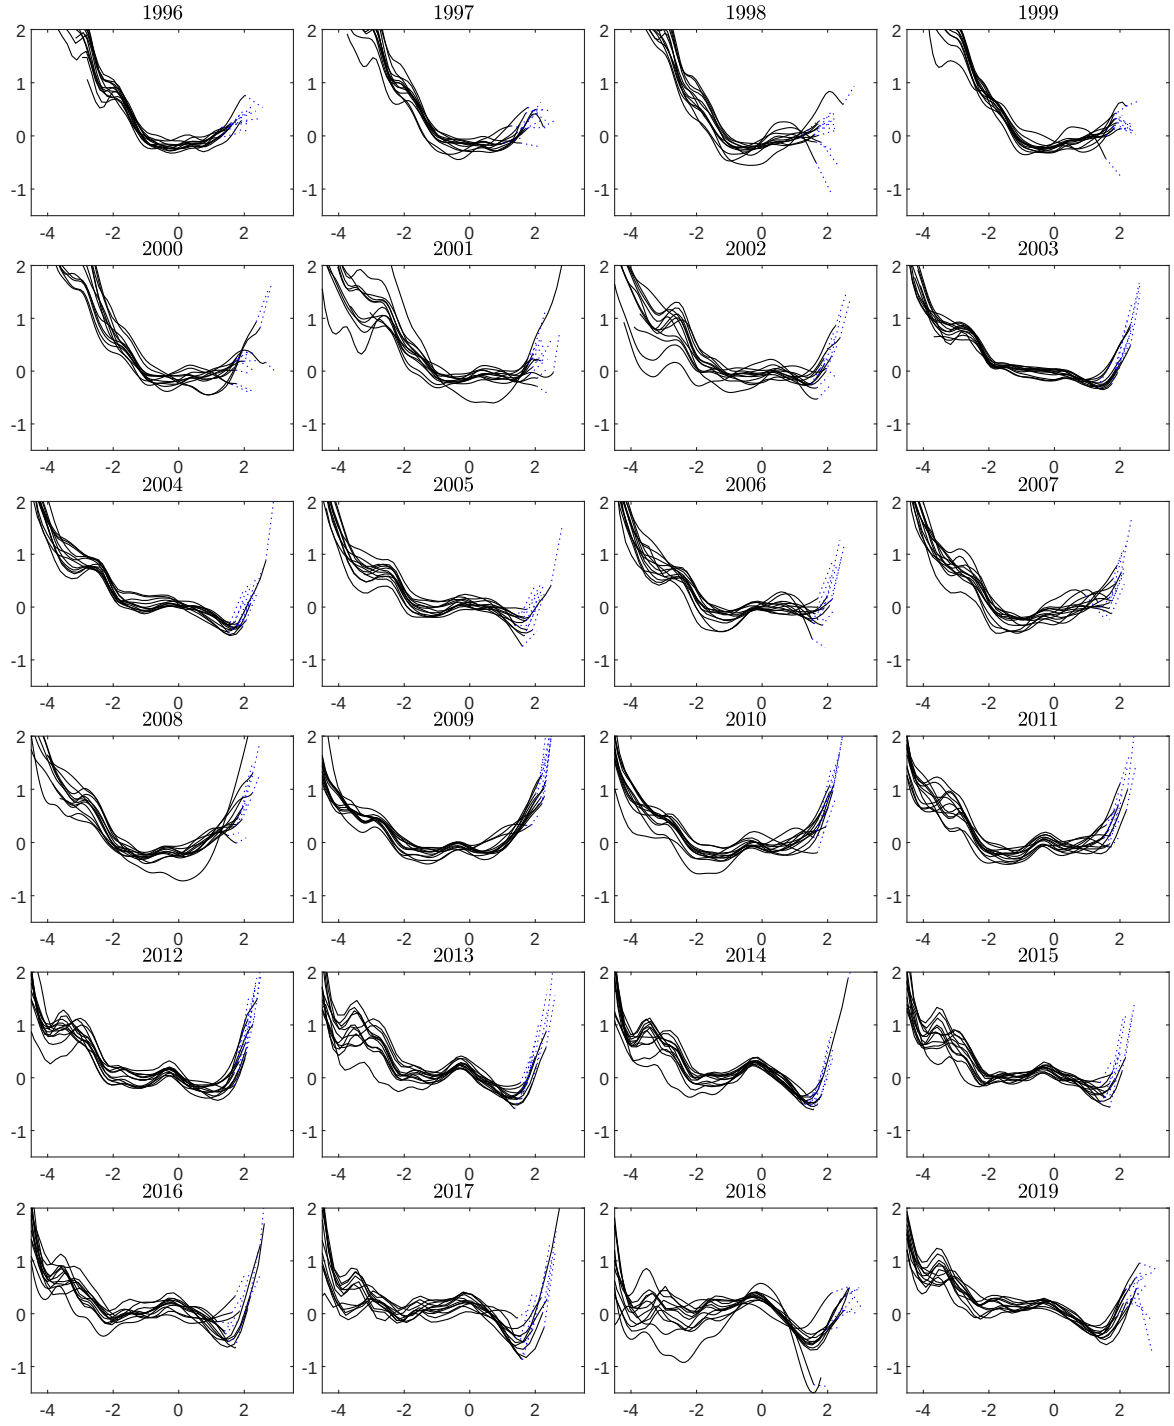

Figure 3: **Pricing kernel estimates.** The figure shows the natural logarithm of estimated pricing kernels. Standardized returns $= R^m_{t+1}/\sigma_t$ are on the horizontal axis. The horizon is one month. The dotted blue line connects the points, which depict the ratio of the CDFs of the tail from (14), with the corresponding pricing kernels.

IV. Pricing Results

A. Approach

To test the pricing performance of M, I study Euler equation errors in the following, which is a non-parametric and standard way to test a candidate pricing kernel (see, e.g., Cochrane 2005). Euler equation errors can economically be interpreted as pricing errors, and are commonly referred to as "alphas". The Euler equation reads

$$E_t[M_{t+1}R^j_{t+1}] = 0 \tag{15}$$

for the excess return R^j of each asset j. Empirically, I calculate alphas as:

$$\alpha^j = \frac{1}{T}\sum_{t=1}^{T} M_{t+1}R^j_{t+1}. \tag{16}$$

The series of realized M_{t+1} is calculated as described in Section C below. Note that the alphas for factor models from (16) are virtually the same as from the more common regression-based model tests.[9]

To compare the pricing results to standard benchmark models, I also calculate the alphas of the CAPM and the Fama & French (1993) 3-factor model (FF3). The pricing kernel in the CAPM is a linearly decreasing function of the market return. Also the FF3 model is virtually monotonically decreasing in the index, as the other two factors are virtually orthogonal to the market (depicted in Figure 5 below). Finally, I consider the coskewness (CoSkew) factor of Harvey & Siddique (2000) in addition to the market. For the interested reader, the details on the calculation of the pricing kernel in the linear factor models can be found in Appendix C.2.

As test assets, I use the portfolios described in Section C.1. For each strategy, I form the high-minus-low (HML) portfolio such that the strategy produces a positive CAPM alpha. In most cases this also maximizes the curvature exposure. Since industry portfolios have no natural HML sorting, I go long (short) the portfolio with the highest (lowest) in-sample CAPM alpha, to keep the analysis simple. All portfolios are value-weighted in the benchmark analysis.

[9] To be precise: $\alpha^{\text{Euler}} = \alpha^{\text{regression}}/R_f$. IA.1.4 shows this relationship for the CAPM. A general proof can be found in Dahlquist & Söderlind (1999).

For commodities I form a long-short portfolio that maximizes CAPM alpha (as well as curvature) which is long the portfolio with the highest CAPM alpha, and short the one with the lowest CAPM alpha. Finally, since all straddles have negative curvature, I take a short position in the at-the-money straddle financed at the risk-free rate, since it has both the highest CAPM alpha and the most negative curvature.

B. Monotonically decreasing pricing kernel and upside risk

To isolate the increasing part of M, and therefore the upside-risk premium, I calculate a strictly monotonic decreasing counterpart M^{mon} of each estimated M. A graphical illustration can be found in Figure 1, where the solid black line is extended by the dotted line to the right of the minimum. To be precise, for a given estimate M, I first find the global minimum in the area of positive returns. Next, graphically speaking, I discard the entire estimate to the right of this minimum, and then linearly extrapolate the remainder using the slope coefficient from the CAPM. The difference in pricing errors between M^{mon} and M for an asset measures the its upside risk premium. Formally:

$$\text{Upside risk premium}^j = \frac{1}{T}\sum_{t=1}^{T}\Big(M_{t+1}^{mon} - M_{t+1}\Big)R_{t+1}^{j}. \tag{17}$$

A formal derivation of the relationship is provided in the IA.1.2.

C. Realized pricing kernel

In order to assess the pricing performance of M using (16), one needs to calculate a realized M. Each line in Figure 3 is an ex ante functional form, which describes how the random future realized R_{t+1}^{m} maps into a realized M_{t+1}. Hence, for each observed return R_{t+1}^{m} from date t to $t+1$, there is a corresponding realized M_{t+1}. Figure 4 illustrates this mapping for a selected date. The ex ante M was calculated on June 18, 2008. The observed log return on July 19, 2008, was -5.50%, which corresponds to a realized $M_{t+1} = \exp(-0.22) = 0.80$.

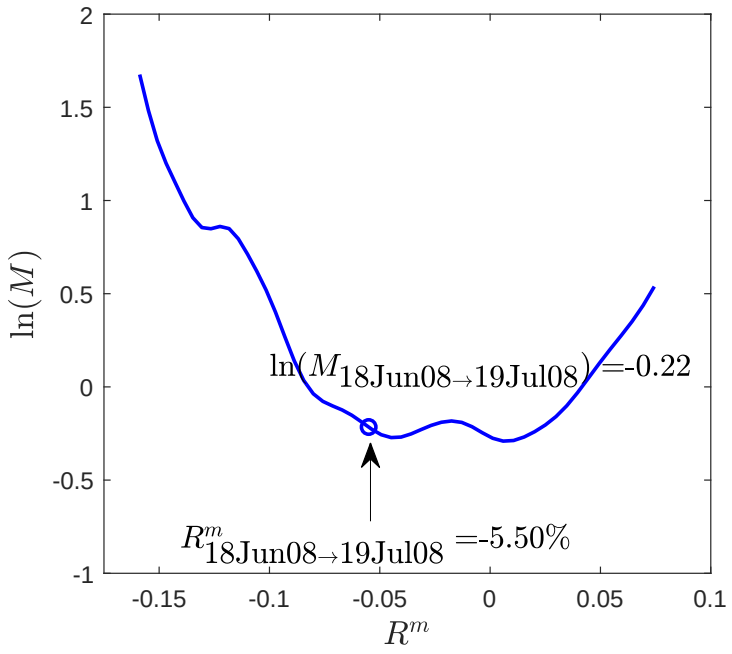

Figure 4: **Calculation of realized pricing kernel.** The figure illustrates the calculation of the realized pricing kernel on a selected date.

D. Results for US equities

For the pricing of equities, I split the analysis into those strategies that have a strong $\alpha - \gamma$ correlation, and those that do not. Table 2 present the main pricing results for both groups. I follow Burnside (2011) and bootstrap statistical significance.[10] The first column in Panel A shows that the option-implied M prices all tabulated anomalies. The alphas are economically low and statistically insignificant. The second column column shows the estimated upside risk premium. For all anomalies, this upside risk premium is large, and between 29% of the CAPM alpha in the case of $\beta_{\Delta\text{VIX}}$ and 89% for the low beta anomaly. On average, the upside risk premium is 4.0% p.a., and almost half of the CAPM alpha. Furthermore, the upside risk premium is always statistically significant. Hence, it is an important risk factor for understanding the returns to those portfolio sorts.

The third and fourth column show the CAPM and FF3 alpha, respectively. Not all of them are statistically significant for two reasons. First, the level of alphas is lower than in previous

[10]In particular, N=50,000 i.i.d. bootstrap draws are used. The draws are i.i.d., since the Ljung-Box test at lags up to 30 rejects auto-correlation in the time series. I verified for the CAPM and the FF3 that the bootstrapped p-value is similar to a regression-based p-value obtained using Newey & West (1987) adjusted standard errors, and often more conservative.

Table 2: **Pricing Results for Equity Portfolios**

The table shows the pricing errors (alphas) for M, the CAPM, the Fama & French (1993) 3 factor model FF3, and the CAPM plus the coskewness factor of Harvey & Siddique (2000) all calculated via (16), and the upside risk premium (RP) calculated via (17). Returns are from the high-minus-low value-weighted portfolios. The symbols ***, ** and * indicate that values are significantly different from zero at the 1%, 5%, and 10% significance levels, respectively. The significance levels are obtained from 50,000 (pairwise) bootstrap draws from the sample of alphas (differences in alphas). All numbers are annualized in percent.

	M	Upside RP	CAPM	FF3	CoSkew
Panel A: Equity Portfolios with Curvature Patterns					
Beta	-0.65	6.26***	7.02	8.03*	4.83
Vol	3.01	5.15***	11.29**	11.18**	8.83*
Ivol	2.92	4.44***	10.74**	10.30*	8.57*
Max	1.14	3.80***	8.05*	7.43*	6.24
Mom	1.65	3.73***	10.39*	13.09**	10.15*
Turn	3.88	3.41***	8.54**	7.33*	5.86
Ztrade	4.00	3.14***	8.36**	7.33**	6.24*
$\beta_{\Delta \text{VIX}}$	3.17	2.04***	6.91**	7.00**	5.40*
Avg	2.39	4.00***	8.91**	8.96**	7.02**
Panel B: Other Equity Portfolios					
Size	1.20	0.10	-0.01	-0.39	0.24
B/M	1.53	-1.19**	1.19	-2.75	0.47
Inv	3.60	-0.68	3.50*	2.20	3.09
Prof	5.90***	0.25	6.44***	6.74***	6.11***
Acc	3.55*	0.04	3.65**	3.59**	3.77**
Industry	5.90	2.81***	7.90*	8.60**	5.69

studies, since it is well-known that alphas of many anomalies decrease in the later sample. In addition, the bootstrap tends to be conservative relative to using Newey & West (1987) adjusted standard errors from regressions, and this is particularly true for the FF3 model.[11] Nevertheless, the average alpha is significant at the 5% level for both the CAPM and FF3. Furthermore, the pricing error of M is always much lower than those of either the CAPM or FF3 (and the difference is statistically significant, see Table IA.5).

The last column shows that the CAPM plus coskewness factor cannot price most of the

[11]This can for example be seen when considering that a regression based FF3 test leads to a t-statistic of 2.77 for Vol, 2.88 for Ivol, and 3.10 for the average alpha, which would all be significant at the 1% level. Furthermore, this study uses open prices from the third Friday to match the option expiry timing. Using closing prices at the settlement dates would increases annual alphas on average by 2 percentage points and t-statistics by 1. These results are consistent with Polk et al. (2019), who document that anomalies accrue mostly during trading hours, and have the opposite sign overnight.

anomalies. The decrease in alpha for Beta, Vol and Ivol is consistent with the results in Schneider et al. (2020). Nevertheless, many other alphas do not decrease relative to the CAPM, and the average alpha is significant at the 5% level.

Switching the attention to Panel B of Table 2, one can make several observations. First, some prominent sorts such as size, B/M or investment have relatively low alphas in the recent sample, which has been already been documented in the literature. Second, M prices all portfolios similarly well as the factor models. For Size, B/M, Inv, Prof, and Acc the pricing errors for M are very close to those for CAMP and FF3. For the industry portfolios, the pricing error of M is insignificant, but relatively large. Furthermore, there is a significant upside risk premium of 2.8% p.a., which is about 32% of the CAPM alpha. This finding is consistent with both the negative correlation between α and γ in Table 1, and the findings of Dittmar (2002).[12]

In sum, given that M is estimated from options data only and without any cross-sectional stock return information, the results are very encouraging. They further suggest that many anomalies are only anomalies for linear models, and do not exist if the pricing kernel is allowed to be non-linear, and even non-monotonic to capture upside risk.

E. Results for other asset classes

Next, I extend the analysis to the other asset classes, and the results are presented in Table 3. For European equities I calculate M analogously to the US, using options and return data for the EuroStoxx 50, and the resulting M is again largely U-shaped (see Figure IA.1).

Panel A of Table 3 shows that M again prices returns of portfolio sorted on Beta, Vol, Ivol, Max, Turn, and Ztrade. All alphas are close to zero and insignificant. In contrast, these portfolios have large and significant alphas relative to the CAPM and FF3. Momentum is an exception, but the pricing error of M is still half of that of the factor models. The last column again shows a large and significant upside risk premium. On average, it is 5.6% p.a., or almost half of the CAPM and FF3 alphas.

The results in Panel B show that M also prices returns to short at-the-money S&P 500

[12]Dittmar (2002) shows that a cubic $M(R^m)$ can price industry portfolios, but notes that the resulting shape of M is inconsistent with standard utility theory. He further shows that imposing a monotonically decreasing $M(R^m)$ that is consistent with standard utility theory significantly worsens the pricing performance. Both results together imply that the upside risk is crucial for explaining the returns, consistent with my results.

Table 3: **Pricing Results for European Equity, Option Straddles and Commodity Portfolios**

The table shows the pricing errors (alphas) for M, the CAPM, the Fama & French (1993) 3 factor model FF3, and the CAPM plus the coskewness factor of Harvey & Siddique (2000) all calculated via (16), and the upside risk premium (RP) calculated via (17). Returns are from the high-minus-low value-weighted portfolios. The symbols ***, ** and * indicate that values are significantly different from zero at the 1%, 5%, and 10% significance levels, respectively. The significance levels are obtained from 50,000 (pairwise) bootstrap draws from the sample of alphas (differences in alphas). All numbers are annualized in percent.

	M	Upside RP	CAPM	FF3	CoSkew
Panel A: European Equity					
Beta	-2.51	6.72***	9.34**	8.83**	9.09**
Vol	1.61	5.91***	13.08***	14.42***	11.57***
IVOL	1.54	5.47***	12.68***	14.13***	11.15***
Max	-0.61	5.66***	10.50***	11.69***	9.18***
Mom	12.74**	6.04***	24.35***	25.40***	24.16***
Turn	3.52	6.19***	15.11***	15.39***	14.08***
Ztrade	3.35	6.19***	14.96***	15.25***	13.95***
$\beta_{\Delta \text{VIX}}$	0.98	3.65***	9.28***	10.79***	9.40***
Avg	2.58	5.73***	13.66***	14.49***	12.82***
Panel B: Options					
Straddle	-0.74	43.70***	81.42*	77.49*	78.37*
Panel C: Commodities					
HML	-1.16	3.50**	3.67	2.84	4.04*

straddles well. The pricing errors are virtually zero and insignificant. For the CAPM and FF3, however, they are large and significant. Moreover, the last column reveals a large upside risk premium, which is more than half of the CAPM alpha. While it is well known that the CAPM does not price option returns well, it is a new result that a substantial part of the pricing error stems from upside risk. Finally, for commodities in Panel C one finds insignificant pricing errors for M, and a sizable and mildly significant upside risk premium.

F. Robustness

The main parametric input in the estimation of M is the return density and the variance forecast. I therefore present robustness with respect to various modeling choices made for the estimation of M and present the results in Table 4. To keep the exposition concise I focus on the

equity portfolios with a pronounced curvature pattern and summarize the pricing performance by presenting the average alpha, i.e., the analogue to the line "Avg" in Panel A of Table 2. Overall, the results are fully robust to any of the choices made above. In all specifications, the pricing error for M is insignificant, and upside risk premium is estimated to be between 2.4% and 4.9% p.a. Furthermore, in the cases where the upside risk premium is estimated lower than the 4.0% of the benchmark approach, the difference stems from the fact that the M of the alternative approaches are less U-shaped (not tabulated).

First, I test alternative methods to extrapolate M in the tail. The first alternative uses (14) to extrapolate M in the region of high positive returns. The second alternative uses the method suggested by Figlewski (2010) to "complete the tails" of $f^*(R)$ using a generalized extreme value (GEV) distribution.

Second, I test alternative approaches to model the physical density $f(R)$. One alternative is to use an extending window of shocks for the shocks Z in (33). A second alternative is to use the skew t distribution of Azzalini & Capitanio (2003). The parameters of the distributions are estimated with maximum-likelihood, and in addition, I make the parameters that control skewness and kurtosis linear functions of σ_t, to account for the finding that higher moments vary with market volatility (Gormsen & Jensen 2022).

Third, I present the results for alternative variance models, which rank behind the benchmark model in terms of variance forecasting accuracy (for details, see Appendix B). They are variations of the benchmark model, where either only the VIX or jumps are included in the model, or the estimation methodology is changed from weighted-least squares to OLS. Table 4 shows that the pricing errors for the alternative RV models slightly increase and the upside risk premium (as well as the U-shape of the estimated M, not tabulated) slightly decreases with the variance forecasting performance of the RV model.

Fourth, I calculate a weakly monotone M^{mon}, i.e., extrapolate to the right of the global minimum with a slope of 0. This decreases the average estimated upside risk premium (M^{mon}-M in the last column) by about 1.2 percentage point p.a. In other words, of the average 4.0% risk premium, 2.8% stem from the increasing part of M, and 1.2% stems from the difference of a flat versus a decreasing M in the right tail.

Finally, I present the results for equal-weighted portfolios, instead of value-weighted as in the benchmark approach. This increases all alphas, but otherwise does not change the results

Table 4: **Robustness Pricing Results for Equity Portfolios**

The table shows the pricing errors (alphas) for M, the CAPM, the Fama & French (1993) 3 factor model FF3, and the CAPM plus the coskewness factor of Harvey & Siddique (2000) all calculated via (16), and the upside risk premium (RP) calculated via (17). Returns are from the high-minus-low value-weighted (last row: equal-weighted) of all portfolios in Panel A of Table 2. For a description of the model details for each line see the main text. The symbols ***, ** and * indicate that values are significantly different from zero at the 1%, 5%, and 10% significance levels, respectively. The significance levels are obtained from 50,000 (pairwise) bootstrap draws from the sample of alphas (differences in alphas). All numbers are annualized in percent.

	M	Upside RP	CAPM	FF3	CoSkew
CDF extrapolation	3.44	3.81^{***}	8.91^{**}	8.96^{**}	7.02^{**}
GEV	3.50	2.54^{***}	8.91^{**}	8.96^{**}	7.02^{**}
$f(R)$ - extending window	2.39	3.05^{***}	8.91^{**}	8.96^{**}	7.02^{**}
$f(R)$ - Skew t	0.37	4.92^{***}	8.91^{**}	8.96^{**}	7.02^{**}
RV model w. VIX	3.71	3.18^{***}	8.91^{**}	8.96^{**}	7.02^{**}
RV model w. Jumps	3.39	3.51^{***}	8.91^{**}	8.96^{**}	7.02^{**}
RV model w. VIX, Jumps, OLS	3.54	3.71^{***}	8.91^{**}	8.96^{**}	7.02^{**}
$M^{weakly\ mon}$	2.27	2.83^{***}	8.91^{**}	8.96^{**}	7.02^{**}
EW portfolios	5.69	3.19^{***}	10.77^{***}	10.89^{***}	9.22^{**}

qualitatively.

V. A Variance Risk-Based Explanation

Variance risk is arguably the most important risk factor for options. This section shows empirically that variance risk is also important for explaining in the cross-section of stock returns and that it induces a non-monotonic projected pricing kernel. The following Section VI provides an equilibrium foundation for this reduce-from approach.

A. Variance risk and U-shaped M

On a general level, the U-shaped estimates for M in Figure 3 show that investors are averse against R^m realizations in both tails. In the option pricing literature stochastic variance is arguably the most important risk factor for explaining option prices. Christoffersen et al. (2013) show that priced variance risk is a viable explanation for the U-shaped pricing kernel. Their model assumes a pricing kernel which is decreasing in R^m, and increasing in variance. The latter feature generates the empirically well-documented variance risk premium. Furthermore, their model assumes a U-shaped relationship between variance and returns, which I confirm empirically below. Both assumptions together gives rise to the U-shaped projection. The structural model in Section VI provides an equilibrium foundation for this reduced-form modeling approach.

To test the mechanism empirically, I augment the CAPM with a variance factor. Fortunately, variance risk is a traded asset. In particular, it is well known that the VIX^2 is the price of a variance swap with 30 days to maturity.[13] Since the VIX is calculated as the forward price, I convert it into the spot price VIX^2/R_f, which is equivalent to the price of the replicating options portfolio. The pricing kernel then reads:

$$M_{t+1}^{VS} = a_t + b_t \times R_{t+1}^m + c_t \times \left(\frac{RV_{t+1}^{(21)}}{VIX_t^2/R_{f,t}} - R_{f,t} \right), \tag{18}$$

where $RV_{t+1}^{(21)}$ denotes the realized monthly variance of the market index from time t to $t+1$,

[13] Strictly speaking it is only a close approximation. It has become standard in the literature to use the VIX as a proxy for the price of the swap. See e.g. Cheng (2019) for a detailed discussion of these issues. Furthermore, using a proprietary data set of quoted prices for S&P 500 variance swaps with one month maturity ranging from January 1996 to October 2013, I verify that the VIX^2 is a close approximation of actual variance swap rate. The two series have a correlation of 99.1%, and the average level of swap rates is 21.2% compared to an average VIX of 21.6% (both in VIX units).

and hence the last term mimics the excess return on a variance swap (VS).[14]

After plugging (18) into (1) and conditioning down one obtains an expression for expected returns, which one can then test by running the following regression:

$$R^j_{t+1} = \alpha^j + \beta^j \times R^m_{t+1} + \gamma^j \times \left(\frac{RV^{(21)}_{t+1}}{VIX^2_t/R^f_t} - R^f_t \right) + \epsilon_{t+1}. \tag{19}$$

Panel A of Table 5 presents the results for each HML equity portfolio portfolio, where t-statistics are based on Newey & West (1987) adjusted standard errors. The results show that adding the variance factor to the CAPM decreases the alphas significantly: all presented alphas are statistically insignificant and the average decreases to 3.42% p.a. from 8.91% p.a. in Table 2. The last column shows that the estimated γ coefficients are negative for each long-short portfolio. To illustrate the economic magnitude, first note that the average excess return of the variance swap factor is -26.44% per month. Hence the γ for the average anomaly return of -0.0157 implies a premium for systematic variance exposure of $-0.0156 \times -26.44\% = 0.42\%$ per month, or 4.98% p.a. The highest variance risk premium is found for momentum with $-0.0406 \times -26.44\% = 1.07\%$ per month, or 12.88% p.a.

When turning to European equities in Panel B, one again finds a sizable decrease in alphas which renders all of them insignificant. The estimated γ coefficients are similar in magnitude as in the US return data. Furthermore, the pricing error for at-the-money S&P 500 straddles in Panel D is an insignificant -45.75% p.a., relative to a significant 81.42% p.a. for the CAPM in Table 2. The estimated exposure towards variance risk, γ, is large and significant, which is probably not surprising given the abundant evidence that options, and in particular straddles, are heavily exposed to variance risk.

Overall, a simple extension of the CAPM with a priced variance factor leads to a pricing performance which is similar to the one of the option-implied M. To dissect the pricing performance, Figure 5 sheds light on the relationship between R^m, and different measures for realized variance. The top left plot shows that R^m has a U-shaped relationship with the contemporane-

[14]While there are different approaches to calculate the payoff of a variance swap, the most common one in the literature is to use squared daily log returns. This paper uses five-minute returns throughout, since this measurement is widely recognized to deliver a more precise estimate of RV. Since the average RV from high-frequency returns is lower than the one obtained from squared daily returns, I scale my RV to be at the same level on average, and hence the results are comparable to other studies. Nevertheless, using squared daily returns to calculate the payoff delivers very similar results (see Table IA.6).

Table 5: **CAPM + Variance Swap Pricing Results for Equity, European Equity, Commodity and Option Portfolios**

The table shows alphas from regression (19), annualized in percent. t-statistics (in parenthesis) are adjusted for heteroscedasticity and autocorrelation (Newey & West 1987). The estimated γ coefficients are multiplied by 100 for readability.

	α		β		$\gamma \times 100$	
Panel A: Equity						
Beta	-0.04	(-0.01)	-1.28	(-10.35)	-2.15	(-1.58)
Vol	5.64	(1.09)	-1.32	(-8.21)	-1.57	(-1.59)
Ivol	4.74	(0.72)	-1.16	(-6.11)	-1.79	(-1.41)
Max	2.94	(0.52)	-1.00	(-7.44)	-1.52	(-1.40)
Mom	-0.92	(-0.10)	-0.88	(-4.22)	-4.06	(-2.50)
Turn	6.70	(1.55)	-0.86	(-8.84)	-0.28	(-0.37)
Ztrade	5.01	(1.32)	-0.85	(-9.04)	-0.89	(-1.39)
$\beta_{\Delta \text{VIX}}$	4.60	(1.43)	-0.52	(-6.30)	-0.65	(-0.96)
Avg	3.58	(0.97)	-0.98	(-8.98)	-1.61	(-2.11)
Panel B: European Equity						
Beta	-6.83	(-0.57)	-0.85	(-9.73)	-2.48	(-1.48)
Vol	2.24	(0.16)	-0.53	(-3.41)	-1.48	(-0.69)
Ivol	2.67	(0.19)	-0.45	(-2.78)	-1.35	(-0.61)
Max	-0.81	(-0.07)	-0.48	(-3.86)	-1.59	(-0.92)
Mom	-3.20	(-0.17)	-0.50	(-2.31)	-4.15	(-1.54)
Turn	5.59	(0.63)	-0.71	(-8.96)	-1.32	(-0.98)
Ztrade	5.58	(0.63)	-0.71	(-8.98)	-1.31	(-0.97)
$\beta_{\Delta \text{VIX}}$	2.01	(0.17)	-0.17	(-1.32)	-1.11	(-0.66)
Avg	0.91	(0.09)	-0.55	(-4.72)	-1.85	(-1.22)
Panel C: Commodities						
HML	2.84	(0.48)	-0.00	(-0.01)	-0.33	(-0.43)
Panel D: Options						
HML	-45.75	(-0.75)	-2.24	(-1.16)	-49.67	(-3.67)

ously realized monthly variance. The two black lines are regression lines fitted separately for the domain of positive and negative returns. It is evident that the relationship is not monotonically decreasing, but has a stylized U-shape. The bottom left plot shows that the same pattern is found for the realized pricing kernel implied by (18).[15]

[15]The correlation between the realized M and the M^{VS} is 49% and hence the replications is far from perfect. However, to put this into perspective, the correlation between the full pricing kernel and its projection in the equilibrium model below is 30%. In addition, other factors might contribute to the replication error, such as noise in measurement, the pricing of variance risk could be non-linear, and

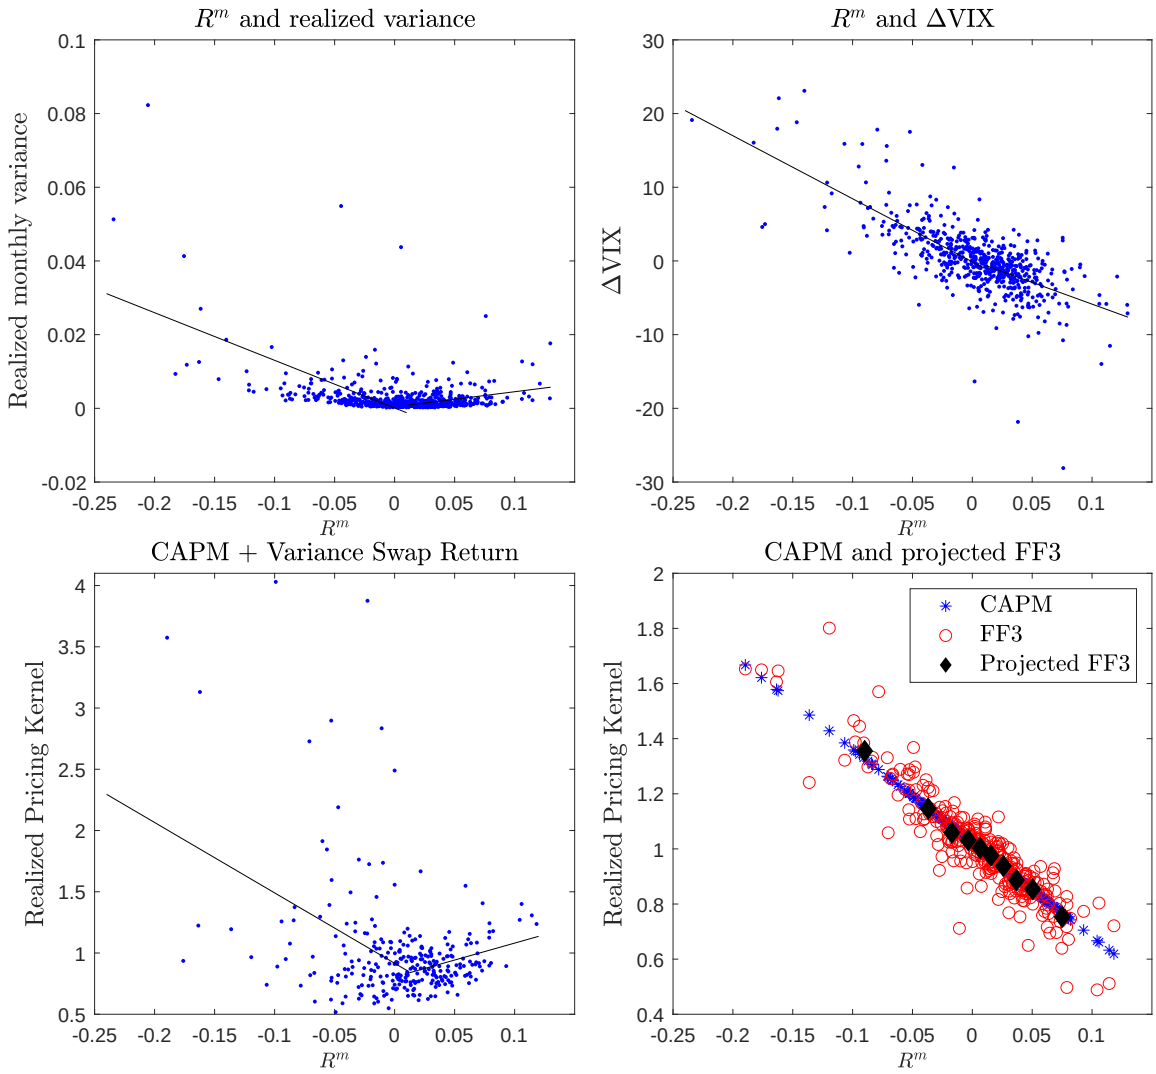

Figure 5: **Relationship between market return, factors, and pricing kernels.** The top left plot shows the relationship between monthly market return R^m and contemporaneous realized monthly variance. The top right plot shows the relationship between R^m and the contemporaneous change in the VIX index, ΔVIX. The bottom left plot shows the realized pricing kernel for the CAPM + variance swap return model from (19). The bottom right plot shows the CAPM and the Fama-French three factor model projected onto market returns. The solid black lines are regression lines fitted in the domain of positive and negative returns separately. The projected FF3 pricing kernel is calculated by taking the average pricing in each decile of R^m.

Ang et al. (2006b) use a stock's beta with respect to daily changes in the VIX index ($\beta_{\Delta\text{VIX}}$) as their measure of a stock's exposure to changes in future expected variance. The top right plot shows the relationship between R^m and the contemporaneous ΔVIX over the same time

lastly, there could be other risk premia embedded in the option-implied M such as for example jumps or higher moment risk.

period. It is evident that this relationship is monotonically decreasing. Furthermore, when considering that γ values of the $\beta_{\Delta\text{VIX}}$-sorted portfolios in Panel F of Table 1 are much lower than for all other sorts, these findings strongly suggest that a systematic variance risk factor based on $\beta_{\Delta\text{VIX}}$ captures neither curvature nor upside risk.

Taken together, this appears to be the reason why the $\beta_{\Delta\text{VIX}}$-factor in Ang et al. (2006b) cannot explain the Vol and Ivol anomaly, while my M and M^{VS} can. Lastly, the bottom right plot shows that the FF3 model projected onto index returns is also monotonically decreasing, and therefore does not capture any upside risk premium.

VI. Upside Risk in an Equilibrium Model

This section derives the pricing implications of a long-run risks type model with recursive preferences where consumption growth and its variance feature GARCH type dynamics. Due to the feedback process of high positive consumption shocks to variance, both large positive consumption shocks and large positive stock returns are associated with large variance. In equilibrium, the representative agent dislikes this increase in variance more than she likes the associated positive consumption shock. Hence, these states are associated with high marginal utility, the projected M is U-shaped, and there is upside risk.

A. Setup

The model setup follows the long-run risks framework of Bansal & Yaron (2004). The representative agent has recursive Epstein & Zin (1989) preferences over total consumption. The dynamics of fundamentals are:

$$x_{t+1} = \rho x_t + \rho_e \sigma_{c,t+1} \eta_{x,t+1} \tag{20}$$

$$g_{c,t+1} = \mu + x_t + \sigma_{c,t+1} \eta_{c,t+1} \tag{21}$$

$$g_{d,t+1} = \mu + \phi x_t + \pi \sigma_{c,t+1} \eta_{c,t+1} + \rho_d \sigma_{c,t+1} \eta_{d,t+1} \tag{22}$$

$$\sigma^2_{c,t+1} = w + b\sigma^2_{c,t} + a\sigma^2_{c,t}(\eta_{c,t} - h)^2 \tag{23}$$

$$\eta_{c,t+1}, \eta_{x,t+1}, \eta_{d,t+1} \sim i.i.d.\ N(0,1). \tag{24}$$

The dynamics of the long-run growth component x and log consumption growth g_c are identical to those in Bansal & Yaron (2004). Dividends are modeled with an exposure π to consumption shocks, which is absent in the original model, but included in many follow-up specifications, as for example Bansal et al. (2012) and Schorfheide et al. (2018). The key innovation is to replace the stochastic variance of the original model with the GARCH dynamics of Engle & Ng (1993), where the shocks to consumption growth have a feedback effect on future variance. This effect is symmetric around the parameter h, and hence both large positive and negative η_c lead to an increase in future σ^2_c (given the usual assumption of $w, a, b > 0$). This relationship is in contrast to the usual modeling frameworks of either pure stochastic variance or Markov chains, where

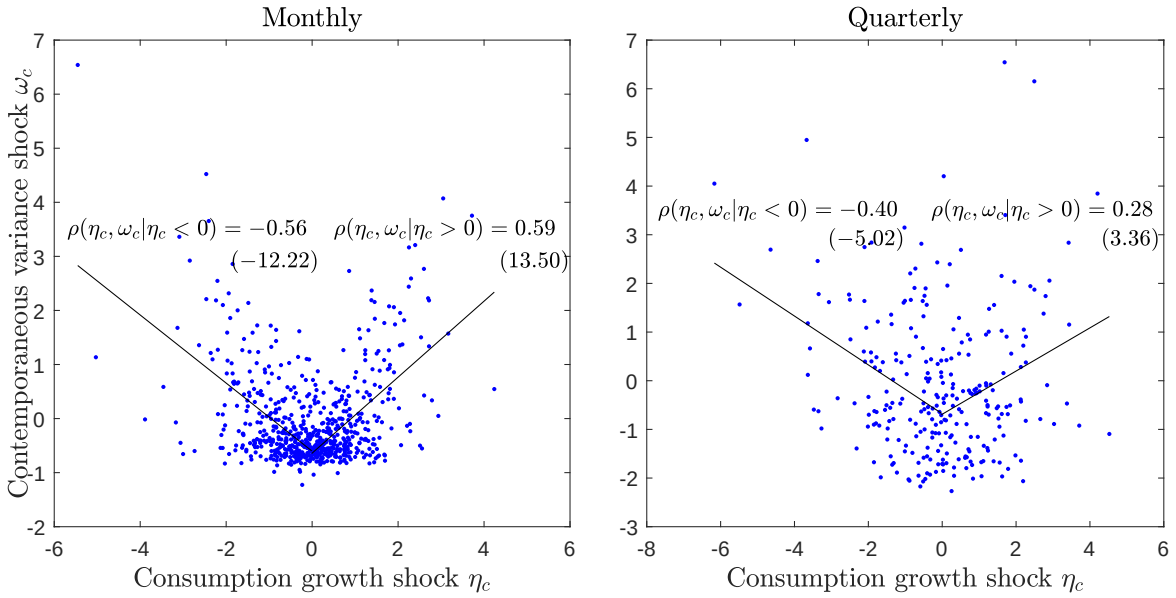

Figure 6: **Relationship between filtered consumption and variance shocks.** The figure plots contemporaneous shocks η_c to the consumption growth process and shocks ω_c to the variance process in the model of Schorfheide et al. (2018). Shocks are calculated using data and filtered state variables provided by Schorfheide et al. (2018). The left plot is based on monthly consumption data from Jan 1959 - Sept 2014, and the right plot on quarterly data from Jan 1947 - Sept 2014. $\rho(\eta_c, \omega_c|\eta_c <> 0)$ denotes the correlation between the shocks conditional on η_c being above or below 0, respectively. t-statistics are in parenthesis. Black lines are regressions lines for positive and negative η_c, respectively.

the shocks to g_c and σ_c^2 are assumed to be independent. The following section as well as Section IA.6 provide empirical support for the GARCH modeling assumption.

B. Evidence supporting GARCH dynamics

In order to study the relationship between shocks to consumption growth and consumption variance, one needs an estimate for the variance, since it is usually not observed. Schorfheide et al. (2018) perform a Bayesian estimation of an augmented version of the Bansal & Yaron (2004) long-run risks model, and one of their outputs is a filtered series of consumption variance. Their model features stochastic variance, which resembles the dynamics of $\sigma_{c,t+1}^2$ in (23), but with an extra variance shock $\omega_{c,t+1}$ instead of the $\sigma_{c,t}^2(\eta_{c,t} - h)^2$-term. By rearranging the dynamics in Equation (9) of Schorfheide et al. (2018), one can back out the contemporaneous shocks η_c to consumption growth and shocks ω_c to consumption variance.[16]

[16]I downloaded the code and data from the website from Dongho Song, whom I thank for sharing. Further details are provided in the IA.6.1.

Figure 6 plots the backed out η_c against the contemporaneous ω_c, together with regression lines and correlation coefficients conditional on positive and negative η_c, both for monthly and quarterly data. It is evident that η_c and ω_c have a stylized U-shaped relationship. This strongly supports the GARCH dynamics assumed above, but is at odds with stochastic variance models, where the two shocks are typically assumed to be independent.

Moreover, as a second source of evidence, I estimate the alternative GARCH specification

$$\sigma^2_{c,t+1} = w + b\sigma^2_{c,t} + a\sigma^2_{c,t}(a^+\eta^2_{c,t}\mathbb{I}_{\eta_{c,t}>0} + a^-\eta^2_{c,t}\mathbb{I}^2_{\eta_{c,t}<0}) \tag{25}$$

that allows for different effects of positive and negative η_c innovations. The estimated coefficient a^+ is significant and larger than the coefficient a^- (details in IA.6.2), which supports the assumed dynamics.

C. Estimation and calibration

The parameters of the GARCH process for σ^2_c in (23) are estimated via maximum likelihood. The data are real monthly consumption growth per capita in the US from Jan 1959 - Dec 2019.[17] The estimated parameters imply a relatively high persistence of 0.9966, which leads to problems with the (numerical) solutions for some parameter combinations. I therefore slightly lower the estimated b parameter, such that the persistence is 0.992, which is between the 0.987 used in (Bansal & Yaron 2004) and the 0.999 used in (Bansal et al. 2012).[18] Furthermore, I set $\pi = 2.6$ as in Bansal et al. (2012), and set $\phi = 1$. All other parameters are the same as in Bansal & Yaron (2004). Section IA.7 present an analysis of the effect of changes in parameters on the main result.

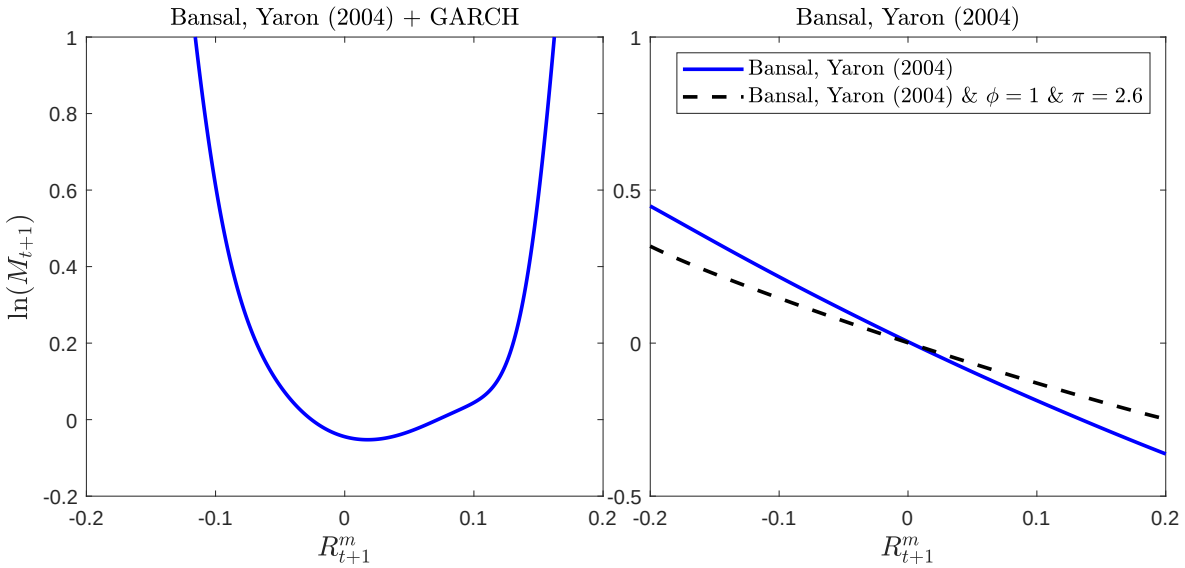

Figure 7: M **in the model.** The figure plots M, i.e., the pricing kernel projected on to market returns R^m, in the modified Bansal & Yaron (2004) model with GARCH dynamics (left), and the original model (right). In the right plot, the blue line represent the original model, while the dashed black line represents a modified version that uses the same values for the parameters ϕ and π as the model with GARCH dynamics in this paper.

D. Pricing implications

I solve the model numerically by iterating on the Euler equation, using a fine grid for the two state variables, x and σ_c^2.[19] For clarity, I denote the model's unprojected pricing kernel as M^{unproj}, and the projection on the return on the aggregate dividend claim R^m again as M. In the model, M^{unproj} is given by:

$$\ln(M_{t+1}^{unproj}) = \ln(\delta) - \gamma g_{c,t+1} - (\gamma - 1/\psi) \ln\left(\frac{V_{t+1}/C_{t+1}}{\mu_t(V_{t+1})/C_t}\right) \tag{26}$$

$$= \underbrace{\ln(\delta) - \gamma(\mu + x_t)}_{\text{constant}} \underbrace{-\gamma\sigma_{c,t+1}\eta_{c,t+1}}_{\text{linearly decreasing in } \eta_c} + \underbrace{(\gamma - 1/\psi)}_{>0} \underbrace{\ln\left(\frac{\mu_t(V_{t+1})/C_t}{V_{t+1}/C_{t+1}}\right)}_{\text{U-shaped in } \eta_c}, \tag{27}$$

[17] Further details on the estimation are provided in the IA.6.2, as well as alternative GARCH dynamic specifications.

[18] In addition, I adjust w such that the average σ_c^2 remains unchanged. Changing the persistence does not affect the results qualitatively. I provide a detailed analysis and discussion of the sensitivity of the main result with respect to all chosen parameters in IA.7.

[19] My solution builds on the code provided by Beason & Schreindorfer (2022), whom I thank for sharing.

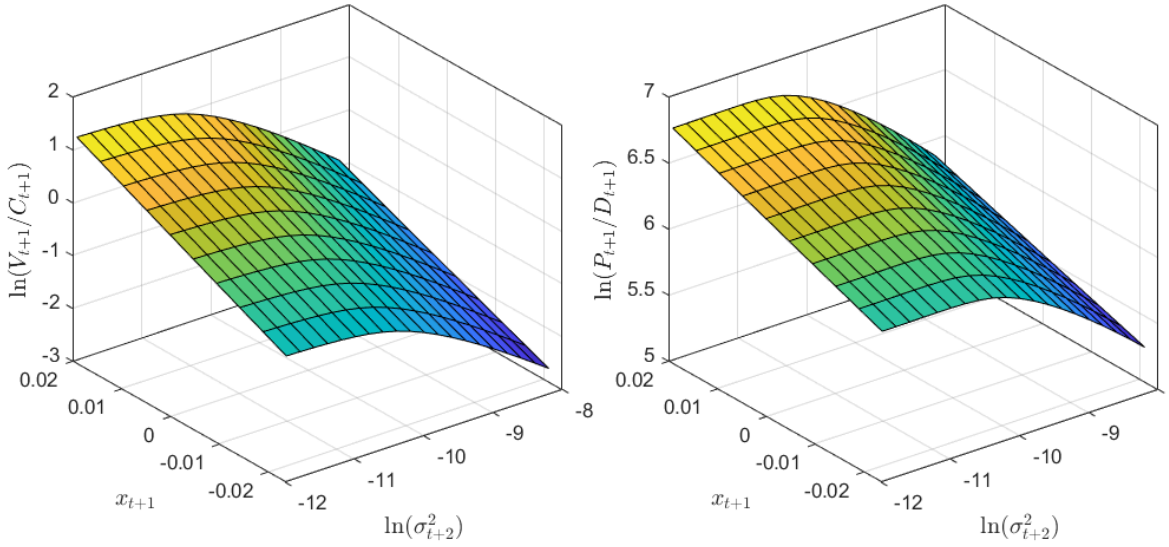

Figure 8: **Utility level and price-dividend ratio in the model.** The figure plots the natural logarithm of the ratio of utility level to consumption, $\ln(V_{t+1}/C_{t+1})$, and the natural logarithm of the price-dividend ratio, $\ln(P_{t+1}/D_{t+1})$ in the solved model against the two state variables.

where V_t is the continuation value of the investor's lifetime utility, $\mu_t(V_{t+1})$ is the certainty equivalent of the next period continuation value, γ is the coefficient of relative risk aversion, ψ is the elasticity of intertemporal substitution, and δ is the time discount factor. The left plot in Figure 7 shows M, which is a U-shaped function of R^m. In contrast, in the original model M is monotonically decreasing, as shown in the right plot. In short, the U-shape stems from large consumption growth shocks η_c leading to both high returns and high marginal utility.[20]

For a more detail understanding of origins of the U-shaped M, I will next go through all relevant effects on M^{unproj} and R^m. To begin with, note that the model has three shocks, but only η_c is important for the result, because η_d is not priced, and η_x affects returns positively and M^{unproj} negatively, i.e., leads to a strictly downward sloping component in M. Hence one only needs to understand the effects of η_c.

Next, consider the effect of η_c on M^{unproj}, which can partially be studied analytically as annotated in (27). While the first term on the right-hand-side is constant, the second term is the standard consumption shock term which implies that $\ln(M^{unproj})$ decreases linearly in η_c.

[20]Technically, M depends on the two state variables x_t and $\sigma_{c,t+1}$ (which is known at time t). For ease of exposition, I fix x_t and $\sigma_{c,t+1}$ at their median for the graphs. However, all studied relationships are qualitatively unchanged for other $\sigma_{c,t+1}$ values, and for other x_t values they are even quantitatively the same (details in IA.7).

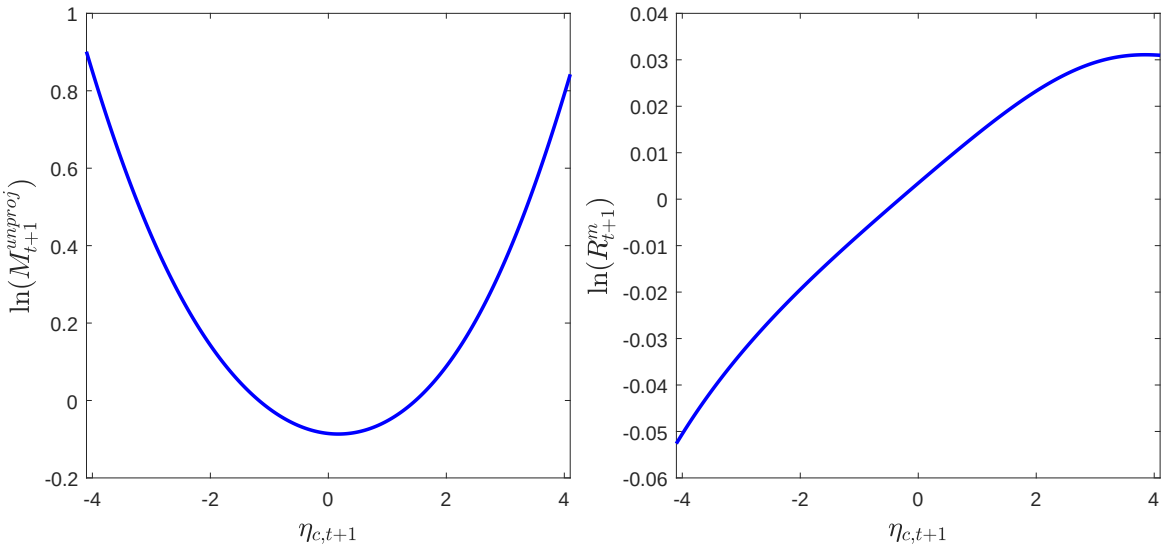

Figure 9: **Relationship between consumption shocks, pricing kernel and returns.** The figure plots the relationship between consumption growth shocks η_c and the average realized M^{unproj} in the left plot, and the relationship between η_c and the realized market return R^m in the right plot.

The third term implies a U-shaped relationship between η_c and M^{unproj}, but this can only be seen from the numerical solution. To understand the properties of the third term, one needs to analyze the relationship between η_c and V/C. The left plot in Figure 8 shows that in equilibrium the V_{t+1}/C_{t+1} ratio is decreasing in $\sigma^2_{c,t+2}$ (that is known at time t), which is consistent with the original model.

In addition, recall that the variance dynamics imply that shocks to σ_c^2 have a U-shaped relationship to η_c. Hence the future V_{t+1}/C_{t+1} is inversely U-shaped in η_c, and the third term in (27) is U-shaped in $\eta_{c,t+1}$. For negative η_c, all effects work in the same direction. For positive η_c, the increasing part of the U-shaped effect from the third term dominates the decreasing linear second term, and as a result, the overall relationship between η_c and M^{unproj} is U-shaped, as depicted in the left plot of Figure 9.

Next, turning to returns, consider the effect of η_c on R^m. For this, note that

$$\ln(R^m_{t+1}) = g_{d,t+1} + \ln\left(\frac{P_{t+1}/D_{t+1}+1}{P_t/D_t}\right) \tag{28}$$

$$= \underbrace{\mu + \phi x_t}_{\text{constant}} + \underbrace{\rho_d \sigma_{c,t+1}\eta_{d,t+1}}_{\text{uncorrelated with } \eta_c} + \underbrace{\underbrace{\pi\sigma_{c,t+1}\eta_{c,t+1}}_{\text{linearly increasing in } \eta_c} + \underbrace{\ln\left(\frac{P_{t+1}/D_{t+1}+1}{P_t/D_t}\right)}_{\text{ambiguous in } \eta_c}}_{\text{"largely increasing in" } \eta_c}, \tag{29}$$

where the third term implies a linearly increasing relationship, and the fourth term is only available numerically. The aggregate relationship between η_c and R^m is depicted in the right plot of Figure 9. For most of the (relevant) range of values of η_c, the relationship is monotonically increasing.

Taking all effects together, a high positive consumption shock η_c on average leads to a positive R^m due to a high positive dividend growth, but also to an increase in σ_c^2, and hence a high realized M^{unproj}, which generates the U-shaped M depicted in Figure 7. In other words, a high positive R^m is on average associated with an increase in the variance of fundamentals, which the agent dislikes more than she likes the positive consumption growth associated with it. Hence she wants to hedge against these states of the world, and assets that pay off well in these states (positive curvature relative to R^m) will have low average returns and negative CAPM alphas, and vice versa.

VII. Conclusion

This paper provides a novel approach to understanding several prominent stock return anomalies by using a pricing model extracted from option prices. The results show that a non-linear market model can jointly explain cross-sectional stock return anomalies related to beta, (idiosyncratic) volatility, liquidity, and momentum. A key part of the explanation is a new, negative upside risk premium, which is estimated to be 4.0% p.a., or on average almost half of the CAPM alphas of the anomalies. The same results are documented for European stock returns. Moreover, upside risk is also prevalent in other asset classes, such as options and commodities.

The paper further documents that many cross-sectional stock return anomalies have a negative relation between their CAPM alpha and the curvature of their returns relative to the market. This pattern is found for virtually all anomalies where the sorting is done based on return-based characteristics. Qualitatively this pattern resembles the return profile of a short straddle, and hence the large positive alphas on straddles and stock returns anomalies are consistent. In contrast, anomalies based on accounting data have no noticeably curvature patterns.

The paper further provides a structural model that generates upside risk as a result of the exposure to variance risk, and provides supporting evidence for the main mechanism of the model. Empirically, the paper shows that the CAPM augmented by a second variance factor can also explain the studied anomalies. After controlling for the two factors, the strategies have only insignificant alphas left. Finally, this explanation is distinct from a coskewness based explanation.

Appendix

A. Details on the Pricing Kernel Estimation

A.1. Options data

The empirical analysis uses out-of-the-money S&P 500 call and put options that are traded in the period from 02/01/1990 - 31/12/2019. I focus on the most liquid contract, namely the standard third Friday contract (with AM settlement), in order to obtain a series of monthly pricing kernels with a 30-day horizon. Next, only options with positive trading volume are considered and the standard filters proposed by Bakshi et al. (1997) are applied. The options data is cleaned further following the standard approach in the literature:

1. For each month, find the third Friday option series that has a remaining time to maturity closest to 30 days. This follows the standard approach in the literature to use data from Wednesdays, if available.
2. Remove all quotes that either have zero trading volume on the day of the price quote, have best bid below 0.50, are more than 20 index points in-the-money, or violate the standard no-arbitrage bounds considered by Bakshi et al. (1997).

Table A.1 presents descriptive statistics for the cleaned options data by standardized moneyness, i.e., $\tilde{K} = \ln(K/S_t)/\sigma_t$. The implied volatility (IV) mildly shows the typical volatility smirk pattern. The volume and open interest data show that the option contracts are highly liquid and traded for a large range of moneyness levels.

A.2. Estimation of the risk-neutral density

The steps for the estimation of the risk-neutral density from options prices are:

1. Clean the data as described above.
2. Get risk-free rate from OptionMetrics, and interpolate linearly for the correct maturity.
3. Get the dividend yield data from OptionMetrics, and interpolate linearly for the correct maturity (using the implied dividend yield from at the money call and put pair leads to similar results, but the dividend yield estimates are more noisy).

Table A.1: **Options Data Descriptive Statistics**

The table presents descriptive statistics for the cleaned S&P 500 options data. The details of the data filters are listed in Appendix A.1. The first column shows the total number of contracts over the sample. The second column shows the average implied volatility (IV) in % p.a. The third column shows the average volume per contract, and the last column shows the average open interest per contract. The data is all 30 day to maturity, out-of-the money, and 3rd Friday AM settled options contracts from 02/01/1996 - 31/12/2019.

$\tilde{K} = \frac{\ln(K/S_0)}{\sigma_t}$	Contracts	IV (% p.a.)	Volume	Open Interest
$\tilde{K} < -4$	4,860	32.1	1,561	9,712
$-4 < \tilde{K} < -3$	1,861	27.3	2,304	14,466
$-3 < \tilde{K} < -2$	2,123	24.7	2,533	16,756
$-2 < \tilde{K} < -1$	2,328	22.2	3,011	18,658
$-1 < \tilde{K} < 0$	3,594	18.6	3,684	17,953
$0 < \tilde{K} < 1$	3,687	16.0	3,229	15,971
$1 < \tilde{K} < 2$	2,250	14.4	2,119	14,703
$2 < \tilde{K}$	358	15.9	1,231	9,820

4. Transform mid-prices into implied volatilities using Black and Scholes (1973). In the region of +/- 20 points from at-the-money, take a weighted average (by volume) of put and call implied volatilities. The results remain unaltered if the implied volatility provided by OptionMetrics is used.

5. Fit a fourth-order polynomial to the implied volatilities over a dense set of strike prices, and convert back into call option prices using Black-Scholes.

6. Finally, numerically differentiate the call prices using (30) and (31) to recover the risk-neutral return distribution:

$$1 - F^*(S_{t,t+\tau}) = -\exp(r\tau)\Big[\frac{\partial C_{BS}(S_t, X, \tau, r, \hat{\sigma}(S_t, X))}{\partial X}\Big]_{|X=S_{t,t+\tau}} \tag{30}$$

$$f^*(S_{t,t+\tau}) = \exp(r\tau)\Big[\frac{\partial^2 C_{BS}(S_t, X, \tau, r, \hat{\sigma}(S_t, X))}{\partial X^2}\Big]_{|X=S_{t,t+\tau}} \tag{31}$$

A.3. Estimation of the physical return density

The represented approach for obtaining the conditional physical density $f(R_{t+1})$ is semi-parametric and the most common methodology in the literature on pricing kernel estimation (see, e.g., Rosenberg & Engle (2002), Barone-Adesi et al. (2008), Christoffersen et al. (2013), Faias &

Santa-Clara (2017), or Christoffersen et al. (2022)). This method is often referred to as "filtered historical innovations" and it has several distinct advantages. First and foremost, it is only semi-parametric as it only requires a minimum of parametric assumptions, preserves the empirical patterns for moments higher than two and, last but not least, has a good fit to the empirical distribution. Moreover, the only parametric input is a forecast of conditional volatility (details in Appendix B), which can easily come from alternative models, and hence robustness analysis is straightforward.

The starting point is a long daily time series of the natural logarithm of monthly returns on the S&P 500 from 02/01/1990 - 31/12/2019, obtained from CRSP. The monthly return series is then standardized by subtracting the sample mean return $\bar{R}$ and afterwards dividing by the conditional one month volatility σ_t. This yields a series of monthly return shocks Z:

$$Z_{t+1} = (R_{t+1} - \bar{R})/\sigma_t. \tag{32}$$

The conditional distribution f_{t+1} is then constructed by multiplying the standardized return shock series Z with the conditional monthly volatility expectation σ_t:

$$f_t(R_{+1}) = f\left(\bar{R} + \sigma_t Z\right). \tag{33}$$

In this paper $f_t(R_{t+1})$ is calculated using only index return data available up to date t, i.e., f is obtained fully out-of-sample. To find a trade-off between between using more recent and potentially more relevant data on the one hand, and having as much data as possible for less noise in f on the other hand, I use a rolling window of the past ten years of return data as a compromise. The results are, however, fully robust to using a different window length or an expanding window.

B. Realized Variance Modeling

B.1. Overview

The goal is to forecast monthly (21 trading days) volatility. Bekaert & Hoerova (2014) perform a large-scale, out-of-sample performance comparison of numerous state-of-the art heterogeneous

autoregressive (HAR) realized variances models for the S&P 500 at the one month horizon. They find that two specifications perform best: the HAR model combined with either the VIX index or jumps as additional predictors. Building on these results, I focus on these predictors. For completeness, I also include the leverage effect considered in Bekaert & Hoerova (2014), and I confirm their finding that it is the least important predictor. In addition, I employ two recent advances from the literature that improve model performance. First, Bollerslev et al. (2016) show both theoretically and empirically that including realized quarticity corrects for measurement error in realized variance. Second, Clements & Preve (2021) show that estimating the model with weighted-least squares (WLS) instead of OLS further improves out-of-sample forecasts.

The following section briefly lays out the econometric framework, followed by the empirical results. A more detailed discussion of each variable and its foundation in the literature can, for example, be found in Bekaert & Hoerova (2014).

B.2. Econometric framework

In the following, daily realized variance $RV_t^{(1)}$ is defined as the sum of N squared five-minute log returns $r_{t,i}$ of day t:

$$RV_t^{(1)} = \sum_{i=1}^{N} r_{t,i}^2. \tag{34}$$

Weekly ($h = 5$) and monthly ($h = 21$) RV is calculated as $RV_t^{(h)} = \frac{1}{h}\sum_{j=0}^{h-1} RV_{t-j}^{(1)}$. Furthermore, the daily jump $J_t^{(1)}$ is defined as:

$$J_t^{(1)} = \max(RV_t - BV_t^{(1)} \cdot \pi/2, 0), \tag{35}$$

where $BV_t^{(1)}$ denotes the daily bipower variation:

$$BV_t^{(1)} = \sum_{i=1}^{N-1} |r_{t,i}||r_{t,i+1}|. \tag{36}$$

Jumps for other frequencies are calculated as $J_t^{(h)} = \frac{1}{h}\sum_{j=0}^{h-1} J_{t-j}^{(1)}$. The leverage effect at different frequencies is defined as $r_t^{(h)-} = \min\left(r_t^{(h)}, 0\right)$, where $r_t^{(h)} = \frac{1}{h}\sum_{j=0}^{h-1} r_{t-j}$. Lastly,

following Bollerslev et al. (2016), daily realized quarticity $RQ_t^{(1)}$ is defined as

$$RQ_t^{(1)} = \frac{N}{3} \sum_{i=1}^{N} r_{t,i}^4, \tag{37}$$

and its multi-period analogue as $RQ_t^{(h)} = \frac{1}{h} \sum_{j=0}^{h-1} RQ_{t-j}^{(1)}$.

Taking all variables together, the most general model specification considered in the following (as well as in Bekaert & Hoerova 2014) is:

$$\begin{aligned} RV_{t+21}^{(21)} =& c + \alpha VIX_t^2 + \underbrace{(\beta_d + \beta_{d,Q}(RQ_t^{(1)})^{1/2})}_{\beta_{d,t}} RV_t^{(1)} \\ &+ \underbrace{(\beta_w + \beta_{w,Q}(RQ_t^{(5)})^{1/2})}_{\beta_{w,t}} RV_t^{(5)} + \underbrace{(\beta_m + \beta_{m,Q}(RQ_t^{(21)})^{1/2})}_{\beta_{m,t}} RV_t^{(21)} \\ &+ \gamma_d J_t^{(1)} + \gamma_w J_t^{(5)} + \gamma_m J_t^{(21)} + \delta^d r_t^{(1)-} + \delta^w r_t^{(5)-} + \delta^m r_t^{(21)-} + \epsilon_{t+21}. \end{aligned} \tag{38}$$

Let $\widehat{RV_{t+21}^{(21)}}$ denote the out-of-sample forecast from the model estimated at day t using information available up to that day. The paper then refers to the expected variance over the next month, i.e., from day $t+1$ to $t+21$, as $\sigma_t^2 = \widehat{RV_{t+21}^{(21)}}$.

B.3. Model selection results

Next, I follow Bekaert & Hoerova (2014) and perform an out-of-sample model comparison of different model specifications. The keep the analysis simple and not to be overwhelmed by the large number of possible combinations of model specification and estimation method, I build on previous results showing that RV combined with jumps J and the VIX usually perform best. From this specification, I shut down either the VIX, J, or the RQ correction, or use OLS instead of WLS. Furthermore, motivated by the results in Bekaert & Hoerova (2014) that using single variables from a group of variables performs worse, all lags of one group of variables are included or excluded together. Lastly, I add the leverage effect to several model specifications.

The data is S&P 500 spot prices obtained from Tickdata.com as in Bekaert & Hoerova (2014) from 02/01/1990 - 31/12/2019. I use an initial burn-in period of six years, such that the first out-of-sample forecast is available at the beginning of 1996. I then extend the estimation window forward day by day. To measure model performance, I follow Bekaert & Hoerova (2014)

and use root mean squared error (RMSE), mean percentage error (MPE) and out-of-sample R^2 (R^2_{OOS}) as metrics.

Table A.2 presents the results. Overall, one can see that WLS always leads to large improvements over OLS, except for the specifications that include the leverage effect. Furthermore, the specification without any adjustment to the estimation methodology (no RQ and OLS) has the worst accuracy.

When it comes to model specification, I find that no model clearly outperforms in all metrics, which is analogous to Bekaert & Hoerova (2014). Therefore, I follow their approach and employ a ranking for each criterion, and calculate the average rank for each model in the last column. The ranking produces a clear winner: the model which includes both the VIX and jumps J via BV, estimated with WLS. This model ranks second in all metrics, while numerically being close to the best value for each metric, which is reflected in its overall best average rank. Therefore, I use this model as the benchmark model. The model with only the VIX is second best, with some distance to the next best one, which is the winning model estimated with OLS, tied with the model that only includes J. Similar to Bekaert & Hoerova (2014), I find that the leverage effect is the least important predictor variable. The full model including all variables performs worst (last two rows). Also when adding the leverage effect to the second best model the performance ranks low (above the last two rows).

C. Pricing Kernel Calculation

C.1. Anchoring the Pricing Kernel

It is necessary to "anchor" (adjust) M for the pricing analysis, since otherwise $E(M^{mon}) < 1/R_f$, i.e., it is not a valid pricing kernel. Therefore, the raw, unadjusted pricing kernel $\tilde{M}$ is anchored such that it prices the CRSP value-weighted market index (R^I in the following) correctly. The corrected M is obtained from adding a constant a to $\tilde{M}$, i.e.,

$$M = \tilde{M} + a, \tag{39}$$

Table A.2: **RV Model Statistics and Ranking**

The table shows statistics of the out-of-sample volatility forecasting model performance. Parameters are estimated using data from 02/01/1990 - 31/12/2019 using an extending window. Out-of-sample forecasts start on Jan 02, 1996. The first three columns show out-of-sample RMSE, MPE, and R^2_{OOS}. The next three column show the rank of each model for each criterion. The last column averages over the ranks. Bold font indicates the best value for each criterion.

Model	Method	RMSE	MPE	R^2_{OOS}	Rank	Rank	Rank	$\overline{\text{Rank}}$
VIX	OLS	0.0123	0.500	0.335	6	9	7	7.33
VIX	WLS	0.0093	0.466	**0.368**	3	4	1	2.67
J	OLS	0.0130	0.457	0.341	10	3	3	5.33
J	WLS	**0.0068**	0.470	0.292	1	5	10	5.33
VIX, J	OLS	0.0130	**0.452**	0.336	9	1	6	5.33
VIX, J	WLS	0.0085	0.455	0.366	2	2	2	**2.00**
VIX, J, no RQ	OLS	0.0127	0.502	0.281	8	10	12	10.00
VIX, J, no RQ	WLS	0.0104	0.472	0.316	4	6	9	6.33
VIX, r^-	OLS	0.0131	0.483	0.341	11	7	4	7.33
VIX, r^-	WLS	0.0118	0.559	0.340	5	12	5	7.33
VIX, J, r^-	OLS	0.0126	0.503	0.323	7	11	8	8.67
VIX, J, r^-	WLS	0.0136	0.493	0.291	12	8	11	10.33

where

$$a = 1 - \frac{1}{T}\sum_{t=1}^{T}\left(\tilde{M}_{t+1}R^I_{t+1}\right). \tag{40}$$

For the benchmark M, CAPM and FF3 models a is negligible, but for M^{mon} the value of a is around 0.05.

For some of the robustness checks, the ratio of densities in (13) becomes explosive in the tails. I therefore impose a maximum on the realized M of 4.53, which corresponds to the highest realized M of the benchmark method. This is still a large value, compared to the highest realized pricing kernels in the CAPM and FF3 model of 1.67 and 1.80, respectively.

C.2. Pricing kernel in linear factor models

One can represent the pricing kernel in a linear factor model as (see, e.g.,Cochrane (2005))

$$M_t = \bar{M} - bf_t, \tag{41}$$

where

$$b = E(f_t f_t')^{-1}\lambda, \tag{42}$$

and λ denotes the market price of risk for each factor. For the CAPM, I use the CRSP value weighted index and λ is calculated using the sample mean of excess returns. Analogously, for the Fama-French 3-factor model, I add the size and value factor provided by Kenneth French on his website.

D. Construction of Equity Portfolios

Following Bali et al. (2016), beta is estimated using a rolling window of one year of daily returns. Volatility (Vol) and idiosyncratic volatility (Ivol) are calculated over the most recent 21 trading days as in Ang et al. (2006b). The Max sort is based on the maximum daily return over the past 21 trading days as in Bali et al. (2011). Momentum (Mom) is the return over the past 12 months, excluding the last month as in Jegadeesh & Titman (1993). Turnover is calculated as number of shares traded as a fraction of the number of shares outstanding over the past year, as in Datar et al. (1998). The standardized turnover-adjusted number of zero daily trading volumes is calculated over the past year, as in Liu (2006). Beta with respect to changes in the VIX index ($\beta_{\Delta \text{VIX}}$) is calculated as in Ang et al. (2006a).

The ten industry portfolios follow the classification provided by Kenneth French on his website. For the calculation of size, book-to-market (B/M), investment (Inv), and accruals (Acc), I follow the methodology used by Kenneth French.

The stock returns are obtained from CRSP, the risk-free rate and factors returns are from Kenneth French's website. Accounting data are from Compustat. I incorporate delisting returns based on the CRSP daily delisting file into the last return observation for the calculation of portfolio returns. Stocks are excluded if they have less then 180 valid return observations per year, or less than 12 for monthly variables. Lastly, stocks with a price below 1$ (5$ as robustness), exchange code other than 1, 2, 3 or 31 and share code other than 10 or 11 are excluded.

The SPX options expire at the market open ("AM") on the third Friday of each month. To accommodate this, I calculate "open returns" using the approach suggested by Polk et al. (2019). In particular, for each stock, I use the closing return on the settlement day and divide this by the intraday return of that day (open to close). The results remain qualitatively unchanged if the closing returns either from the expiry date or the preceding day are used. For commodities

and European stock returns (ESTX options expire at 12:00) only closing prices exist, but the previous robustness check shows that this is a good approximation.

References

Aït-Sahalia, Y. & Lo, A. W. (2000), 'Nonparametric risk management and implied risk aversion', *Journal of Econometrics* **94**(1), 9–51.

Ang, A., Chen, J. & Xing, Y. (2006a), 'Downside risk', *Review of Financial Studies* **19**(4), 1191–1239.

Ang, A., Hodrick, R. J., Xing, Y. & Zhang, X. (2006b), 'The cross-section of volatility and expected returns', *Journal of Finance* **61**(1), 259–299.

Azzalini, A. & Capitanio, A. (2003), 'Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**(2), 367–389.

Bakshi, G., Cao, C. & Chen, Z. (1997), 'Empirical performance of alternative option pricing models', *Journal of Finance* **52**(5), 2003–2049.

Bakshi, G., Crosby, J. & Gao Bakshi, X. (2022), 'Dark matter in (volatility and) equity option risk premiums', *Operations Research* **70**(6), 3035–3628.

Bakshi, G., Madan, D. & Panayotov, G. (2010), 'Returns of claims on the upside and the viability of u-shaped pricing kernels', *Journal of Financial Economics* **97**(1), 130–154.

Bali, T. G., Cakici, N. & Whitelaw, R. F. (2011), 'Maxing out: Stocks as lotteries and the cross-section of expected returns', *Journal of Financial Economics* **99**(2), 427–446.

Bali, T. G., Engle, R. F. & Murray, S. (2016), *Empirical asset pricing: The cross section of stock returns*, John Wiley & Sons.

Bansal, R., Kiku, D., Shaliastovich, I. & Yaron, A. (2014), 'Volatility, the macroeconomy, and asset prices', *Journal of Finance* **69**(6), 2471–2511.

Bansal, R., Kiku, D. & Yaron, A. (2012), 'An empirical evaluation of the long-run risks model for asset prices', *Critical Finance Review* **1**(1), 183–221.

Bansal, R. & Yaron, A. (2004), 'Risks for the long run: A potential resolution of asset pricing puzzles', *Journal of Finance* **59**(4), 1481–1509.

Banz, R. W. & Miller, M. H. (1978), 'Prices for state-contingent claims: Some estimates and applications', *Journal of Business* **51**(4), 653–672.

Barberis, N. & Huang, M. (2008), 'Stocks as lotteries: The implications of probability weighting for security prices', *American Economic Review* **98**(5), 2066–2100.

Barone-Adesi, G., Engle, R. F. & Mancini, L. (2008), 'A garch option pricing model with filtered historical simulation', *Review of Financial Studies* **21**(3), 1223–1258.

Beason, T. & Schreindorfer, D. (2022), 'Dissecting the equity premium', *Journal of Political Economy* **130**(8), 2203–2222.

Bekaert, G. & Hoerova, M. (2014), 'The vix, the variance premium and stock market volatility', *Journal of Econometrics* **183**(2), 181–192.

Boguth, O., Carlson, M., Fisher, A. & Simutin, M. (2011), 'Conditional risk and performance evaluation: Volatility timing, overconditioning, and new estimates of momentum alphas', *Journal of Financial Economics* **102**(2), 363–389.

Bollen, N. P. & Whaley, R. E. (2004), 'Does net buying pressure affect the shape of implied volatility functions?', *Journal of Finance* **59**(2), 711–753.

Bollerslev, T., Patton, A. J. & Quaedvlieg, R. (2016), 'Exploiting the errors: A simple approach for improved volatility forecasting', *Journal of Econometrics* **192**(1), 1–18.

Breeden, D. T. & Litzenberger, R. H. (1978), 'Prices of state-contingent claims implicit in option prices', *Journal of Business* **51**(4), 621–651.

Broadie, M., Chernov, M. & Johannes, M. (2009), 'Understanding index option returns', *Review of Financial Studies* **22**(11), 4493–4529.

Brunnermeier, M. K., Gollier, C. & Parker, J. A. (2007), 'Optimal beliefs, asset prices, and the preference for skewed returns', *American Economic Review* **97**(2), 159–165.

Burnside, C. (2011), 'The cross section of foreign currency risk premia and consumption growth risk: Comment', *American Economic Review* **101**(7), 3456–76.

Campbell, J. Y., Giglio, S., Polk, C. & Turley, R. (2018), 'An intertemporal capm with stochastic volatility', *Journal of Financial Economics* **128**(2), 207–233.

Chabi-Yo, F. (2012), 'Pricing kernels with stochastic skewness and volatility risk', *Management Science* **58**(3), 624–640.

Chabi-Yo, F., Ruenzi, S. & Weigert, F. (2018), 'Crash sensitivity and the cross section of expected stock returns', *Journal of Financial and Quantitative Analysis* **53**(3), 1059–1100.

Cheng, I.-H. (2019), 'The vix premium', *Review of Financial Studies* **32**(1), 180–227.

Christoffersen, P., Heston, S. & Jacobs, K. (2013), 'Capturing option anomalies with a variance-dependent pricing kernel', *Review of Financial Studies* **26**(8), 1963–2006.

Christoffersen, P., Jacobs, K. & Pan, X. (2022), 'The state price density implied by crude oil futures and option prices', *Review of Financial Studies* **35**(2), 1064–1103.

Clements, A. & Preve, D. P. (2021), 'A practical guide to harnessing the har volatility model', *Journal of Banking & Finance* **133**, 106285.

Cochrane, J. H. (2005), *Asset pricing: Revised edition*, Princeton University Press.

Dahlquist, M. & Söderlind, P. (1999), 'Evaluating portfolio performance with stochastic discount factors', *Journal of Business* **72**(3), 347–383.

Datar, V. T., Naik, N. Y. & Radcliffe, R. (1998), 'Liquidity and stock returns: An alternative test', *Journal of Financial Markets* **1**(2), 203–219.

Dew-Becker, I. & Giglio, S. (2022), 'Risk preferences implied by synthetic options', *Working Paper* .

Dittmar, R. F. (2002), 'Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns', *Journal of Finance* **57**(1), 369–403.

Engle, R. F. & Ng, V. K. (1993), 'Measuring and testing the impact of news on volatility', *Journal of Finance* **48**(5), 1749–1778.

Epstein, L. G. & Zin, S. E. (1989), 'Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework', *Econometrica* **57**(4), 937–969.

Faias, J. A. & Santa-Clara, P. (2017), 'Optimal option portfolio strategies: Deepening the puzzle of index option mispricing', *Journal of Financial and Quantitative Analysis* **52**(1), 277–303.

Fama, E. F. & French, K. R. (1993), 'Common risk factors in the returns on stocks and bonds', *Journal of Financial Economics* **33**(1), 3–56.

Farago, A. & Tédongap, R. (2018), 'Downside risks and the cross-section of asset returns', *Journal of Financial Economics* **129**(1), 69–86.

Figlewski, S. (2010), 'Estimating the implied risk neutral density for the us market portfolio', *Volatility and Time Series Econometrics: Essays in Honor of Robert Engle* pp. 323–353.

Frazzini, A. & Pedersen, L. H. (2022), 'Embedded leverage', *Review of Asset Pricing Studies* **12**(1), 1–52.

Garleanu, N., Pedersen, L. H. & Poteshman, A. M. (2008), 'Demand-based option pricing', *Review of Financial Studies* **22**(10), 4259–4299.

Gormsen, N. J. & Jensen, C. S. (2022), 'Higher-moment risk', *Available at SSRN 3069617* .

Gu, S., Kelly, B. & Xiu, D. (2020), 'Empirical asset pricing via machine learning', *Review of Financial Studies* **33**(5), 2223–2273.

Harvey, C. R. & Siddique, A. (2000), 'Conditional skewness in asset pricing tests', *Journal of Finance* **55**(3), 1263–1295.

Jackwerth, J. C. (2000), 'Recovering risk aversion from option prices and realized returns', *Review of Financial Studies* **13**(2), 433–451.

Jegadeesh, N. & Titman, S. (1993), 'Returns to buying winners and selling losers: Implications for stock market efficiency', *Journal of Finance* **48**(1), 65–91.

Kelly, B. & Jiang, H. (2014), 'Tail risk and asset prices', *Review of Financial Studies* **27**(10), 2841–2871.

Kraus, A. & Litzenberger, R. H. (1976), 'Skewness preference and the valuation of risk assets', *Journal of Finance* **31**(4), 1085–1100.

Lettau, M., Maggiori, M. & Weber, M. (2014), 'Conditional risk premia in currency markets and other asset classes', *Journal of Financial Economics* **114**(2), 197–225.

Liu, W. (2006), 'A liquidity-augmented capital asset pricing model', *Journal of Financial Economics* **82**(3), 631–671.

Mitton, T. & Vorkink, K. (2007), 'Equilibrium underdiversification and the preference for skewness', *Review of Financial Studies* **20**(4), 1255–1288.

Newey, W. K. & West, K. D. (1987), 'A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix', *Econometrica* **55**(3), 703–708.

Polk, C., Lou, D., Skouras, S., Lou, D., Polk, C. & Skouras, S. (2019), 'A tug of war: Overnight versus intraday expected returns', *Journal of Financial Economics* **134**(1), 192–213.

Rosenberg, J. V. & Engle, R. F. (2002), 'Empirical pricing kernels', *Journal of Financial Economics* **64**(3), 341–372.

Schneider, P., Wagner, C. & Zechner, J. (2020), 'Low-risk anomalies?', *Journal of Finance* **75**(5), 2673–2718.

Schorfheide, F., Song, D. & Yaron, A. (2018), 'Identifying long-run risks: A bayesian mixed-frequency approach', *Econometrica* **86**(2), 617–654.

Sichert, T. (2023), 'The pricing kernel is u-shaped', *Available at SSRN 3095551* .

Song, Z. & Xiu, D. (2016), 'A tale of two option markets: Pricing kernels and volatility risk', *Journal of Econometrics* **190**(1), 176–196.

Tédongap, R. (2015), 'Consumption volatility and the cross-section of stock returns', *Review of Finance* **19**(1), 367–405.