

Low-risk Anomalies?

Schneider, Paul; Wagner, Christian; Zechner, Josef

Document Version

Final published version

Published in:

The Journal of Finance

DOI:

[10.1111/jofi.12910](https://doi.org/10.1111/jofi.12910)

Publication date:

2020

License

CC BY-NC-ND

Citation for published version (APA):

Schneider, P., Wagner, C., & Zechner, J. (2020). Low-risk Anomalies? *The Journal of Finance*, 75(5), 2673-2718. <https://doi.org/10.1111/jofi.12910>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 04. Jul. 2025



Low-Risk Anomalies?

PAUL SCHNEIDER, CHRISTIAN WAGNER, and JOSEF ZECHNER*

ABSTRACT

This paper shows that low-risk anomalies in the capital asset pricing model and in traditional factor models arise when investors require compensation for coskewness risk. Empirically, we find that option-implied ex ante skewness is strongly related to ex post residual coskewness, which allows us to construct coskewness factor-mimicking portfolios. Controlling for skewness renders the alphas of betting-against-beta and betting-against-volatility insignificant. We also show that the returns of beta- and volatility-sorted portfolios are driven largely by a single principal component, which in turn is explained largely by skewness.

*Paul Schneider is at USI Lugano and SFI. Christian Wagner and Josef Zechner are at WU Vienna University of Business and Economics. This paper received the 2015 Jack Treynor Prize sponsored by the Q-Group (The Institute for Quantitative Research in Finance). We are grateful to Kevin Aretz; Turan Bali; Nick Baltas; Michael Brennan; John Campbell; Mikhail Chernov; Peter Christoffersen; Mathijs Cosemans; Andrea Gamba; Patrick Gagliardini; Christopher Hennessey; Christopher Hrdlicka; Leonid Kogan; Lorian Mancini; Miriam Marra; Ian Martin; Yoshio Nozawa; Lasse Pedersen; Paulo Rodrigues; Ivan Shaliastovich; Christian Schlag; Fabio Trojani; Rossen Valkanov; Pietro Veronesi; Arne Westerkamp; Paul Whelan; Liuren Wu; and participants at the American Finance Association Meetings 2017 (Chicago), the European Finance Association Meetings 2016 (Oslo), the UBS Quantitative Investment Conference 2016 (London), the Spring Seminar of the Q Group 2016 (Washington, D.C.), the Annual Conference on Advances in the Analysis of Hedge Fund Strategies 2015 (London), the IFSID Conference on Derivatives 2015 (Montreal), the SAFE Asset Pricing Workshop 2015 (Frankfurt); and seminar participants at Cass Business School, Dauphine, Copenhagen Business School, Hong Kong University of Science and Technology, Imperial College, Singapore Management University, Stockholm School of Economics, University of Geneva, University of Toronto (Rotman), Warwick Business School, and WU Vienna for helpful comments. Paul Schneider acknowledges support from the Swiss National Science Foundation grant “Model-Free Asset Pricing.” Christian Wagner acknowledges support from the Center for Financial Frictions (FRIC), grant no. DNRFF102. We are especially indebted to Stefan Nagel and Kenneth Singleton (the Editors), two anonymous referees, and an anonymous associate editor for their extensive comments that have greatly helped to improve the paper. The authors alone are responsible for any errors and for the views expressed in the paper. We have read *The Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

Correspondence: Josef Zechner, Department of Finance, Accounting and Statistics, Institute for Finance, Banking and Insurance, Vienna University of Economics and Business, Welthandelsplatz 1, A-1020 Vienna, Austria; e-mail: josef.zechner@wu.ac.at.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](#) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

DOI: 10.1111/jofi.12910

© 2020 The Authors. *The Journal of Finance* published by Wiley Periodicals LLC on behalf of American Finance Association

EMPIRICAL FINDINGS THAT LOW-BETA stocks outperform high-beta stocks and that (idiosyncratic) volatility negatively predicts equity returns have spurred a large literature on “low-risk anomalies” (LRAs; e.g. Haugen and Heins (1975), Ang et al. (2006), Baker, Bradley, and Wurgler (2011), Frazzini and Pedersen (2014)). In this paper, we show that the returns to trading such LRAs can be explained by the skewness of equity returns, which is ignored by standard measures of market and idiosyncratic risk.

Our theoretical analysis starts with a stochastic discount factor (SDF) that is quadratic in the market excess return (e.g., Kraus and Litzenberger (1976), Harvey and Siddique (2000)). This SDF implies that investors demand compensation for covariance risk as in the capital asset pricing model (CAPM), and that investors accept lower (demand higher) expected returns on assets with positive (negative) coskewness, as expressed by the covariance between the asset and squared market excess returns. If excess returns are determined in this way, then some trading strategies will generate alphas and thus appear as anomalies in the CAPM. Such alphas, however, merely reflect compensation for coskewness risk ignored by the CAPM. We show that these alphas are indeed driven by the correlation between CAPM residual returns and squared market returns.

Next, we demonstrate the link between coskewness and LRA alphas in a Merton-type model of firm risk, where (co)skewness arises endogenously from leverage and stochastic asset volatility. In this model, we simulate three worlds. First, we simulate a world with a standard CAPM pricing kernel. In this case, we find no LRAs. Second, we simulate a skew-aware world in which the pricing kernel also depends on the squared market return. In this case, we find that LRAs appear and are driven by coskewness. Third, we simulate a world in which all moments higher than skewness are also taken into account. We find that these higher moments contribute much less toward explaining LRAs. Taken together, the simulation results suggest that (co)skewness is the main determinant of LRAs. We also use the model to demonstrate that there is a direct link between a firm’s ex ante skewness, its realized coskewness, and its alpha. This motivates our empirical approach of using equity option-implied ex ante skewness to construct coskewness factor-mimicking portfolios in our study of LRAs.

We establish our main empirical results for a cross-section of approximately 5,000 U.S. firms over the period January 1996 to August 2014. This sample covers all CRSP firms for which data on common stock and equity options are available. To comprehensively capture asymmetries in the return distribution, we compute three measures of ex ante skewness from portfolios of out-of-the-money (OTM) options: upper skewness from OTM call options, which covers the right part of the distribution, lower skewness from a portfolio that is short OTM put options, which by definition is negative, to account for the left part of the distribution, and total skewness, which is the sum of upper and lower skewness. Thus, total skewness becomes more negative, as the cost of put options relative to call options rises, that is, as the premium that investors are willing to pay for protection against downside risk rises.

Our empirical analysis starts by showing that in the data, *ex ante* skewness is related to residual coskewness and alphas in the same way as in our simulated skew-aware world: the more extreme a firm's *ex ante* skewness, the higher its residual coskewness and the lower its CAPM alpha. The results are virtually unchanged when we compute alphas and residual coskewness relative to the Fama-French three-factor model (FF3, Fama and French (1993)). When we additionally control for momentum (FF4, Carhart (1997)), the results become quantitatively less pronounced but the qualitative patterns remain the same for lower and upper *ex ante* skewness. These findings suggest that *ex ante* skewness is linked to residual coskewness and alphas in a way that is consistent with skew-aware asset pricing, but is not captured empirically by standard risk factors.

Having established that *ex ante* skewness is a forward-looking proxy for residual coskewness, we study the main prediction of our model, namely, that controlling for skewness should eliminate positive alphas and negative residual coskewness of beta- and volatility-related LRAs. To do so, we construct coskewness factor-mimicking portfolios from decile portfolios sorted by measures of firms' *ex ante* skewness. Using several alternative specifications of skewness factors, we find that LRA alphas decrease substantially and become statistically insignificant when we control for skewness. For instance, using our most flexible specification, we find for betting-against-beta (BaB) that the CAPM alpha drops from 125 to 33 basis points per month, the FF3 alpha from 109 to 21 basis points, and the FF4 alpha from 73 to 21 basis points. For all anomalies, the reduction in alphas is in lockstep with a reduction in the strategies' negative coskewness. These results suggest that controlling for *ex ante* skewness does indeed render alphas insignificant because it captures coskewness risk. This is confirmed in cross-sectional regressions of alphas on residual coskewness betas for 80 CAPM beta- and volatility-sorted portfolios. Without controlling for skewness, the regression R^2 s are 73%, 73%, and 48% using CAPM, FF3, and FF4 alphas and residual coskewness, respectively. When we control for skewness, the R^2 s drop to 22% for the CAPM-based regression and to less than 4% for the FF3- and FF4-based regressions.

Given the skew-aware SDF, the different anomalies that prior literature has established as mostly unrelated asset pricing puzzles, should all be exposed to a common factor. To explore this conjecture, we proceed in three steps. First, using principal component analyses (PCAs) of anomaly returns, we show that the anomalies based on CAPM betas, idiosyncratic volatility, and option-implied variance do indeed have a common determinant. We find that the first principal component (PC) explains more than 90% of the variation in anomaly excess returns and more than 70% of the variation in FF4 residual returns. Second, we show that the first PC is related to the returns of skewness factors. When we regress the first PC on skew factor returns, we find R^2 s of up to 95% for excess returns and up to 80% for FF4 residual returns. Furthermore, we show that the skew exposures of beta- and volatility-sorted portfolios monotonically decrease in beta and volatility. These results provide strong evidence that LRAs have a common determinant related to (co)skewness.

In short, our results imply that empirical patterns referred to as “low-risk anomalies” may not in fact pose asset pricing puzzles. Taking into account the fact that stock returns exhibit (co)skewness, our findings suggest that the CAPM beta may not be a sufficient metric to capture a firm’s market risk, and hence equity returns, which reflect the firm’s true market risk, may appear anomalous only when benchmarked against the CAPM. Our findings also help explain the seemingly anomalous relations between (idiosyncratic) volatility and stock returns, as these empirical patterns can be directly connected to skewness.

Various robustness checks corroborate our results. We first find that controlling for skewness also leads to a substantial decrease in the alpha of the BaB factor returns of Frazzini and Pedersen (2014), even though their factor is constructed from a broader cross-section that does not require options data. Next, we show that our results are robust to considering alternative portfolio-weighting schemes, including additional control factors, and examining subsample periods. We further show that the model-implied relations between ex ante skewness and credit spreads continue to hold when we use credit ratings and CDS spreads. Finally, using the methodology of Nagel and Singleton (2011) to estimate conditional SDFs, we show that a conditional skew-aware SDF fits the data better than the conditional CAPM.

Related Literature. While the CAPM (see Sharpe (1964), Lintner (1965), Mossin (1966)) postulates a positive relation between beta and returns, a large body of research documents that the empirical relation is flatter than that implied by the CAPM or even negative. Early studies that provide such evidence and attempt to explain the empirical failure of the CAPM include Brennan (1971), Black (1972), Black, Jensen, and Scholes (1972), and Haugen and Heins (1975). Recent research confirms these patterns. Ang et al. (2006, 2009) show that (idiosyncratic) volatility negatively predicts equity returns and that stocks with high sensitivity to aggregate volatility risk earn low returns. While Fu (2009) finds that the sign on the relation between idiosyncratic risk and returns depends on the specific risk measure employed, other papers argue that a negative relation can arise when accounting for leverage (e.g. Johnson (2004)) or differences in beliefs and short-selling constraints (e.g. Boehme et al. (2009)). Related, Stambaugh, Yu, and Yuan (2015) argue that the sign on the relation between idiosyncratic risk and returns depends on whether stocks are over- or underpriced and that arbitrage asymmetry explains why the overall relation is negative. Campbell et al. (2018) show that the low returns of stocks with high sensitivity to aggregate volatility risk are consistent with the intertemporal CAPM (Campbell (1993)), which allows for stochastic volatility.

To rationalize the profitability of BaB strategies, Frazzini and Pedersen (2014) build on the idea of Black (1972) that restrictions on borrowing affect the shape of the security market line (SML). Frazzini and Pedersen (2014) present a model in which leverage-constrained investors bid up high-beta assets and in turn generate low risk-adjusted returns. Jylhä (2018) provides further evidence on the role of leverage constraints by showing that the slope of the SML is connected to margin requirements. Baker, Bradley, and Wurgler (2011)

argue that institutional investors' mandates to beat a fixed benchmark discourages arbitrage activity and thereby contributes to the anomaly. Hong and Sraer (2016) present a model with short-sale-constrained investors in which high-beta assets are more prone to speculative overpricing because they are more sensitive to macro-disagreement. Bali, Cakici, and Whitelaw (2011) find that accounting for the lottery characteristics of stocks reverses the relation between idiosyncratic volatility and equity returns, and Bali et al. (2017) argue that the BaB anomaly is consistent with investors' preference for lottery stocks. Other papers study the properties of BaB returns. For example, Baker, Bradley, and Taliaferro (2014) decompose BaB returns into micro and macro components. The results of Novy-Marx (2016) suggest that the performance of LRA strategies is linked to firms' size, profitability, and book-to-market. Recent evidence suggests that similar patterns hold in other asset classes (e.g., Frazzini and Pedersen (2014)) and in international markets (Walkshäusl (2014)). Huang, Lou, and Polk (2016) find that BaB activity itself affects the profitability of the strategy. Cederburg and O'Doherty (2016) study BaB returns through the lens of the conditional CAPM and argue that the positive alphas reported in prior research are due to biases in unconditional performance measures.

This paper takes a different approach by directly linking LRAs to return skewness. Specifically, we build on the insight of Rubinstein (1973) and Kraus and Litzenberger (1976) that the empirical failure of the CAPM may be due to ignoring the effect of skewness on asset returns. Friend and Westerfield (1980) also find that coskewness with the market contains information about stock returns that is incremental to covariance, Sears and Wei (1985) discuss the interaction between skewness and the market risk premium in asset pricing tests, and Harvey and Siddique (2000) show that conditional skewness helps explain the cross-section of equity returns.¹ With the widespread availability of equity options data, recent papers explore the effect of option-implied ex ante skewness on subsequent equity returns but find mixed evidence (e.g. Bali and Hovakimian (2009), Xing, Zhang, and Zhao (2010), Rehman and Vilkov (2012), Bali and Murray (2013), Conrad, Dittmar, and Ghysels (2013)), with differences in results driven by differences in the construction of skew measures.² Bali, Hu, and Murray (2015) provide complementary evidence by showing that ex ante skewness is positively related to ex ante stock returns estimated from analyst price targets. Other recent papers suggesting that skewness matters

¹ Our approach builds on the idea that skewness matters for asset prices through their coskewness component, which is distinct from prior work that focuses on how idiosyncratic skewness can be priced in stock returns (e.g. Brunnermeier, Gollier, and Parker (2007), Mitton and Vorkink (2007), Barberis and Huang (2008), Boyer, Mitton, and Vorkink (2010)). The pricing of coskewness is also conceptually and empirically distinct from the pricing of downside risk (e.g. Ang, Chen, and Xing (2006), Lettau, Maggiori, and Weber (2014)) as Ang, Chen, and Xing (2006) discuss in detail.

² For instance, Rehman and Vilkov (2012) and Conrad, Dittmar, and Ghysels (2013) both use the ex ante skew measure of Bakshi, Kapadia, and Madan (2003) but find a positive relation and a negative relation with subsequent returns, respectively. This difference in results stems from Rehman and Vilkov (2012) measuring ex ante skew from the latest option data only whereas Conrad, Dittmar, and Ghysels (2013) compute ex ante skew as the average across each day over the past quarter, which smooths out recent changes in skewness.

for the cross-section of equity returns include Amaya et al. (2015), who find a negative relation between realized skewness and subsequent equity returns, and Chang, Christoffersen, and Jacobs (2013), who show that those stocks that are most sensitive to changes in the market's ex ante skewness exhibit the lowest returns. Buss and Vilkov (2012) apply the measure of Chang, Christoffersen, and Jacobs (2013) to individual stocks, but do not find a pronounced relation with equity returns, whereas they do find that betas constructed from option-implied correlations exhibit a positive relation with subsequent stock returns. This paper differs from the literature by focusing on the relation between skewness and LRAs. We show that option-implied ex ante skewness of LRA portfolios is informative about their future residual coskewness and thus about CAPM mispricing.

The remainder of this paper is organized as follows. Section I presents the theoretical framework that guides our empirical analysis. We describe the data and variable construction in Section II and present our empirical results in Section III. In Section IV, we provide results of additional analyses and robustness checks, and we discuss potential extensions. Section V concludes. The Appendix contains technical details and the Internet Appendix reports results not reported here for brevity.³


I. Motivation and Theoretical Framework

In this section, we present the theory that guides our analysis of LRAs, such as the finding that high-CAPM-beta stocks underperform relative to low beta stocks. We work with a skewness-aware SDF and show that alphas compensate for (co)skewness risk ignored by the CAPM and other standard factor models (FMs).

Our choice of pricing kernel follows Kraus and Litzenberger (1976), Harvey and Siddique (2000), and others who study the role of (co)skewness risk for stock returns. Kraus and Litzenberger (1976) are the first to note that the lack of empirical support for the CAPM may be due to the model ignoring the effect of skewness on asset prices. Specifically, they argue that investors demand compensation for accepting negative coskewness. Building on this insight, Harvey and Siddique (2000) show that conditional skewness does indeed help explain the cross-section of equity returns. Therefore, skewness itself is a plausible candidate to provide insights for beta- and volatility-based LRAs that receive considerable attention in the recent literature (e.g. Ang et al. (2006), Frazzini and Pedersen (2014)).

To motivate our analysis, we document that the positive alphas of LRA strategies come with negative residual coskewness, in that their residual returns exhibit negative covariance with squared market excess returns. As can be seen in Figure 1, CAPM alphas of beta-sorted portfolios are negatively related to the portfolios' residual coskewness betas obtained from regressing the CAPM residual returns on squared market excess returns. The results are

³ The Internet Appendix may be found in the online version of this article.



Panel B. Beta-Sorted Portfolios from 1996 to 2014

Figure 1. CAPM alphas and residual coskewness of beta-sorted portfolios. This figure plots results for equal-weighted (EW) and value-weighted (VW) decile portfolios sorted by CAPM beta at the end of each month. We plot CAPM alphas (in basis points per month) against their corresponding residual coskewness betas, measured from regressing CAPM residual returns on squared market excess returns. Blue circles mark the low-beta portfolios that a betting-against-beta (BaB) strategy goes long. Red diamonds mark the high-beta portfolios that a BaB strategy goes short. In the legends, we also report results for a cross-sectional regression of portfolio alphas on residual coskewness betas, with t -statistics based on White (1980) standard errors, and regression R^2 s. Panel A presents results for the CRSP sample from June 1963 to December 2014, which covers a total of 2,017,271 monthly observations across 15,843 U.S. firms. Panel B presents results for our main sample that additionally requires stock option data from OptionMetrics and comprises a total of 400,449 monthly observations across 4,967 U.S. firms over the period January 1996 to August 2014. In each panel, the left plot presents results for standard portfolio sorts and the right plot presents results for portfolios (de-)levered to a CAPM beta of one at portfolio formation, that is, initially, all portfolios are equally risky as measured by CAPM beta. (Color figure can be viewed at wileyonlinelibrary.com)

similar in both the CRSP sample from 1963 to 2014 (Panel A) and in our main sample, which also requires options data, from 1996 to 2014 (Panel B). Similar results also obtain when we construct portfolios that are identical in terms of CAPM risk by (de-)levering each portfolio such that it exhibits a beta of one at portfolio formation.

These findings extend to other LRA portfolios and to alphas that additionally control for size, value, and momentum factors. Table I shows that the alphas of LRAs, computed as low-minus-high returns of decile portfolios sorted by CAPM

Table I

Low-Risk Anomalies (LRAs): Alphas and Residual Coskewness

This table reports excess returns, factor model alphas, and residual coskewness of LRAs using value- and equal-weighted decile portfolios (Panels A and B, respectively). At the end of each month, we sort firms into decile portfolios based on their CAPM beta, idiosyncratic volatility (measured from the residual variance of CAPM and Fama-French three-factor model regressions), or equity option-implied ex ante variance. From these portfolios, we compute low-minus-high returns generated by betting-against-beta/volatility strategies and report raw excess returns as well as alphas of CAPM, Fama-French three-factor (FF3), and Fama-French-Carhart four-factor (FF4) regressions. Excess returns and alphas are reported in basis points per month. Values in square brackets are *t*-statistics based on standard errors following Newey and West (1987), where we choose the optimal truncation lag as suggested by Andrews (1991). Additionally, we report annualized Sharpe ratios associated with excess returns and annualized information ratios associated with the alphas, measured as the average residual return divided by its standard deviation. Finally, we report (residual) coskewness, measured as the coskewness beta obtained from regressing (residual) returns on squared market excess returns. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations.

	CAPM Beta	Ivol (CAPM)	IVol (FF3)	ExAnte Var
Panel A: Low-Minus-High Returns of Value-Weighted Portfolios				
Excess return	20.43	30.18	33.44	31.12
<i>t</i> -statistic	[0.27]	[0.40]	[0.54]	[0.37]
Sharpe ratio	0.07	0.09	0.12	0.09
Coskewness	3.55	2.45	2.07	3.24
CAPM alpha	125.06	123.79	109.02	133.17
<i>t</i> -statistic	[2.87]	[2.52]	[2.43]	[2.49]
Information ratio	0.63	0.47	0.51	0.51
Residual coskewness	-2.86	-3.29	-2.56	-3.01
FF3 alpha	109.07	121.31	112.05	132.34
<i>t</i> -statistic	[2.82]	[2.91]	[2.53]	[3.00]
Information ratio	0.61	0.57	0.57	0.61
Residual coskewness	-2.78	-3.45	-2.71	-3.19
FF4 alpha	72.98	60.61	52.04	68.11
<i>t</i> -statistic	[1.58]	[1.50]	[1.78]	[1.69]
Information ratio	0.45	0.35	0.35	0.40
Residual coskewness	-2.06	-2.24	-1.51	-1.91
Panel B: Low-Minus-High Returns of Equal-Weighted Portfolios				
Excess return	-7.58	32.21	28.65	5.83
<i>t</i> -statistic	[-0.10]	[0.46]	[0.49]	[0.08]
Sharpe ratio	-0.02	0.10	0.11	0.02
Coskewness	2.18	3.39	2.35	2.74
CAPM alpha	91.80	121.56	102.27	96.97
<i>t</i> -statistic	[1.87]	[2.47]	[2.27]	[1.79]
Information ratio	0.44	0.51	0.50	0.39
Residual coskewness	-3.91	-2.08	-2.16	-2.85
FF3 alpha	79.47	116.71	98.29	91.24
<i>t</i> -statistic	[2.23]	[4.31]	[3.72]	[2.81]
Information ratio	0.47	0.80	0.74	0.54
Residual coskewness	-3.93	-2.26	-2.31	-3.02

(Continued)

Table I—Continued

	CAPM Beta	Ivol (CAPM)	IVol (FF3)	ExAnte Var
FF4 alpha	39.04	78.55	56.77	38.91
<i>t</i> -statistic	[0.87]	[2.33]	[2.26]	[1.09]
Information ratio	0.27	0.64	0.56	0.31
Residual coskewness	−3.12	−1.50	−1.48	−1.97

beta, idiosyncratic volatility, and option-implied variance, are associated with negative residual coskewness.⁴ Figure 2 plots the alphas of the decile portfolios against their residual coskewness betas and illustrates a strong negative relation across the 80 decile portfolios, with cross-sectional R^2 s ranging from 45% to 73%. The significant slope estimates from the cross-sectional regressions imply an economically important relation between FM alphas and coskewness risk: per unit of residual coskewness beta, on average portfolios earn a monthly CAPM alpha of -0.38% (-0.35% when all portfolios are levered to an initial beta of one), FF3 alpha of -0.33% (-0.29%), and FF4 alpha of -0.23% (-0.21%). This negative relation is in line with Kraus and Litzenberger (1976), Harvey and Siddique (2000), and others who argue that investors demand compensation for accepting negative coskewness.

Below, we show that alphas that appear “anomalous” from the perspective of the CAPM reflect compensation for residual coskewness. We provide evidence for the link between coskewness and LRA alphas by simulating a skew-aware economy populated by Merton-type firms, where (co)skewness arises endogenously from leverage and stochastic asset volatility. Within this framework, we also show that moments beyond (co)skewness matter much less. Finally, we demonstrate that there is a direct link between a firm’s ex ante skewness, its realized residual coskewness, and its alpha. These insights guide our empirical approach in Section III, where we construct factors based on option-implied skewness.

A. Skew-Aware Asset Pricing

Previous research estimates LRA alphas relative to the CAPM and other (multi)factor models. We focus our discussion on the CAPM as the simplest possible benchmark, but the same arguments apply to other FMs as well. To estimate a stock’s or a portfolio’s CAPM alpha (α_i), we regress its excess returns ($R_{i,t+1}$) on the market’s excess returns ($R_{m,t+1}$),


$$R_{i,t+1} = \alpha_i + \beta_i R_{m,t+1} + \epsilon_{i,t+1}. \quad (1)$$

⁴ In our main sample period, the FF4 alphas of some LRAs are not statistically significant, which suggests there may be some link between coskewness and momentum as discussed by Harvey and Siddique (2000). To account for this possibility, we control for momentum throughout our empirical analysis.

Since the regression imposes $\mathbb{E}[\epsilon_{i,t+1}] = \mathbb{E}[\epsilon_{i,t+1}R_{m,t+1}] = 0$, the alpha measures the time-series average of CAPM pricing errors from

$$R_{i,t+1} - \beta_i R_{m,t+1} = \alpha_i + \epsilon_{i,t+1}. \tag{2}$$

Panel A. CAPM Residual Returns



Panel B. FF3 Residual Returns




Figure 2. Low-risk anomalies (LRAs): Alphas and residual coskewness. This figure plots results for the equal-weighted and value-weighted decile portfolios used to compute the LRA returns in Table I. At the end of each month, we sort firms into decile portfolios based on their CAPM beta, idiosyncratic volatility (measured from the residual variance of CAPM and Fama-French three-factor model regressions), or equity option-implied ex ante variance. In total, we have 80 portfolios: 10 portfolios for each of the four LRAs, both, equally- and value-weighted. We plot CAPM-, Fama-French (FF) three-, and four-factor-alphas (in basis points per month) against their corresponding residual coskewness betas, measured from regressing factor model residual returns on squared market excess returns. In each panel, the left plot presents results for standard portfolio sorts and the right plot presents results for portfolios (de-)levered to a CAPM beta of one at portfolio formation, that is, initially, all portfolios are equally risky as measured by CAPM beta. Blue circles mark the low-risk portfolios that a betting-against-beta (BaB)/volatility strategy goes long. Red diamonds mark the high-risk portfolios that a BaB/volatility strategy goes short. In the legends, we also report results for a cross-sectional regression of portfolio alphas on residual coskewness betas, with *t*-statistics based on White (1980) standard errors, and regression *R*²s. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations. (Color figure can be viewed at wileyonlinelibrary.com)

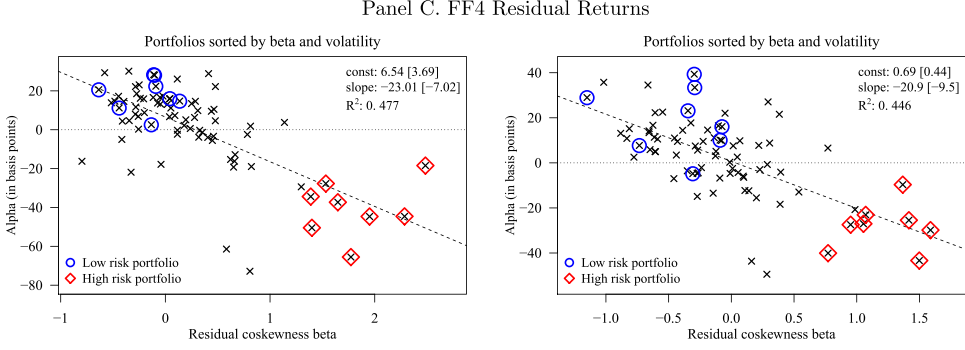


Figure 2. *Continued*

Should exposure to a factor ignored by the CAPM be correlated with the CAPM's pricing errors, this could provide a natural explanation for alphas relative to the CAPM. In other words, any “anomaly” relative to the CAPM really corresponds to residual returns after adjusting for the CAPM. In this paper, we provide evidence that LRA alphas do not represent anomalous returns, but instead reflect compensation for coskewness risk ignored by the CAPM and other FMs.

To show this, we work with an SDF that is quadratic in the market return, as suggested by Kraus and Litzenberger (1976), Harvey and Siddique (2000), and others who study the role of (co)skewness risk for stock returns. With the SDF specified as

$$m_{t+1} = 1 - b_{1,t}(R_{m,t+1} - \mathbb{E}_t[R_{m,t+1}]) - b_{2,t}(R_{m,t+1}^2 - \mathbb{E}_t[R_{m,t+1}^2]), \quad (3)$$

the expected excess return on stock i , given by $\mathbb{E}_t[R_{i,t+1}] = -\text{Cov}_t[m_{t+1}, R_{i,t+1}]$, is

$$\mathbb{E}_t[R_{i,t+1}] = b_{1,t}\text{Cov}_t[R_{i,t+1}, R_{m,t+1}] + b_{2,t}\text{Cov}_t[R_{i,t+1}, R_{m,t+1}^2], \quad (4)$$

where $b_{1,t} > 0$ and $b_{2,t} < 0$. This specification implies that investors demand compensation for covariance risk as in the CAPM. Moreover, it implies that investors accept lower (demand higher) expected return on assets with positive (negative) coskewness.

In this skew-aware world, we identify a stock's alpha as its expected CAPM pricing error and show that alpha arises from coskewness risk. Using the fact that $\mathbb{E}_t[R_{m,t+1}] = -\text{Cov}_t[m_{t+1}, R_{m,t+1}]$, we rewrite equation (4) as

$$\mathbb{E}_t[R_{i,t+1}] = \frac{\text{Cov}_t[m_{t+1}, R_{i,t+1}]}{\text{Cov}_t[m_{t+1}, R_{m,t+1}]} \mathbb{E}_t[R_{m,t+1}],$$

which we further rewrite as

$$\mathbb{E}_t[R_{i,t+1}] = \left(\underbrace{\frac{b_{1,t} \sigma_{i,m,t}}{b_{1,t} \sigma_{m,t}^2 + b_{2,t} \sigma_{m,m^2,t}}}_{\text{covariance risk}} + \underbrace{\frac{b_{2,t} \sigma_{i,m^2,t}}{b_{1,t} \sigma_{m,t}^2 + b_{2,t} \sigma_{m,m^2,t}}}_{\text{coskewness risk}} \right) \mathbb{E}_t[R_{m,t+1}], \quad (5)$$

where $\sigma_{i,m,t} = \text{Cov}_t[R_{i,t+1}, R_{m,t+1}]$, $\sigma_{i,m^2,t} = \text{Cov}_t[R_{i,t+1}, R_{m,t+1}^2]$, $\sigma_{m,t}^2 = \text{Var}_t[R_{m,t+1}]$, and $\sigma_{m,m^2,t} = \text{Cov}_t[R_{m,t+1}, R_{m,t+1}^2]$. Equation (5) shows that the stock's expected excess return depends on its covariance risk and its coskewness risk. Next, we decompose the expected excess return into a CAPM expected return component, $\beta_{i,t}^{\text{CAPM}} \mathbb{E}_t[R_{m,t+1}]$, where $\beta_{i,t}^{\text{CAPM}} = \sigma_{i,m,t} / \sigma_{m,t}^2$, and the expected CAPM pricing error,

$$\mathbb{E}_t[R_{i,t+1}] = \beta_{i,t}^{\text{CAPM}} \mathbb{E}_t[R_{m,t+1}] + \mathbb{E}_t[\epsilon_{i,t+1}^{\text{CAPM}}]. \quad (6)$$

Combining equations (5) and (6), and defining alpha as the time- t expected return of stock i in excess of the CAPM expected return, alpha is given by

$$\alpha_{i,t} := \mathbb{E}_t[\epsilon_{i,t+1}^{\text{CAPM}}] = \underbrace{(\sigma_{i,m^2,t} - \beta_{i,t}^{\text{CAPM}} \sigma_{m,m^2,t})}_{\sigma_{\alpha_i,m^2,t}} \underbrace{\left(\frac{b_{2,t}}{b_{1,t} \sigma_{m,t}^2 + b_{2,t} \sigma_{m,m^2,t}} \right)}_{B_t \mathbb{E}_t[R_{m,t+1}]} \mathbb{E}_t[R_{m,t+1}]. \quad (7)$$

The first term on the right-hand side of equation (7) measures the stock's residual coskewness, that is, the stock's coskewness beyond the coskewness (implicitly) accounted for by the CAPM.⁵ This residual coskewness can be measured as the covariance of CAPM residual returns with squared market excess returns,

$$\sigma_{\alpha_i,m^2,t} = \text{Cov}_t[\epsilon_{i,t+1}^{\text{CAPM}}, R_{m,t+1}^2] = \sigma_{i,m^2,t} - \beta_{i,t}^{\text{CAPM}} \sigma_{m,m^2,t}. \quad (8)$$

The second term on the right-hand side of equation (7) quantifies the pricing of residual coskewness, with the price of risk being negative due to $\mathbb{E}_t[R_{m,t+1}] > 0$ and $B_t < 0$.⁶ Hence, negative residual coskewness is associated with positive alpha and vice versa.

The portfolio results reflected in Figures 1 and 2 are consistent with the negative relation between alpha and residual coskewness suggested by equation (7). To see this more explicitly, rewrite equation (7) as

$$\alpha_{i,t} = \beta_{i,t}^{\text{RESCOS}} \lambda_t^{\text{RESCOS}}, \quad (9)$$

⁵ Using the CAPM to forecast a stock's excess return in a skew-aware world (implicitly) presumes that the stock's coskewness is $\text{Cov}_t[\beta_{i,t}^{\text{CAPM}} R_{m,t+1}, R_{m,t+1}^2] = \beta_{i,t}^{\text{CAPM}} \text{Cov}_t[R_{m,t+1}, R_{m,t+1}^2]$.

⁶ We have $B_t < 0$ because $b_{1,t} > 0$ and $b_{2,t} < 0$, and because $|b_{1,t} \sigma_{m,t}^2|$ is larger than $|b_{2,t} \sigma_{m,m^2,t}|$ by several orders of magnitude except under pathological market return distributions.

where $\beta_{i,t}^{RESCOS} = \sigma_{\alpha_i, m^2, t} / \sigma_{m^2, t}^2$ and $\lambda_t^{RESCOS} = B_t \sigma_{m^2, t}^2 \mathbb{E}_t[R_{m, t+1}] < 0$. Thus, the CAPM alpha can be expressed as the residual return's coskewness beta times the price of coskewness risk.

Empirically, we obtain the (unconditional) sample estimates as follows. First, we use the CAPM regression specified in equation (1) to estimate alphas, $\hat{\alpha}_i$. Second, we use the residual returns from this regression, that is, the CAPM pricing errors $\hat{\alpha}_i + \hat{\epsilon}_{i, t+1}$ as in equation (2), to estimate residual coskewness betas, $\hat{\beta}_i^{RESCOS}$, from time-series regressions of portfolios' residual returns on squared market excess returns. Finally, we run cross-sectional regressions of alphas on residual coskewness betas, where the negative slope can be interpreted as our sample estimate for the price of coskewness risk, $\hat{\lambda}^{RESCOS}$. We also estimate alphas and residual coskewness betas from FMs that additionally control for size, value, and momentum. The results depicted in Figures 1 and 2 show that the relation between alphas and residual coskewness is negative and that this relation is both statistically significant and economically large: the alpha per unit of residual coskewness beta ranges from -21 to -42 basis points per month, depending on the period studied, the LRAs considered, the portfolio formation procedure, and the benchmark FM used.

Our theory implies that LRAs should disappear when we control for (residual) coskewness. We test this prediction empirically in Section III. While this prediction appears straightforward, measuring ex ante coskewness is empirically challenging, as stressed by, for example, Harvey and Siddique (2000) and Christoffersen et al. (2019). Previous research has developed methods to measure individual firms' ex ante moments from equity options data, but it is not possible to measure comoments following a similar approach, because derivatives do not have both the market and the stock as the underlying. However, to the extent that a firm's skewness reflects information about its coskewness, such measures of ex ante skewness are informative about the coskewness of the firm's future returns. In the next subsection, we present a simulation study that uses our framework to illustrate that (i) coskewness is indeed the main determinant of LRAs and (ii) firms' ex ante skewness is informative about their future residual coskewness and alphas.

B. Simulation of a Skew-Aware World with Merton-Type Firms

To shed additional light on the interaction between skewness and LRAs, we simulate an economy that is populated by levered Merton type firms. In this economy, skewness and coskewness arise endogenously from firms' leverage, stochastic asset volatility, and market skewness. For details on the setup of the simulation study, see the Appendix. There we also discuss alternative model specifications, such as including jumps in the asset process or allowing a firm to default before debt reaches its maturity, for which we present results in the Internet Appendix.

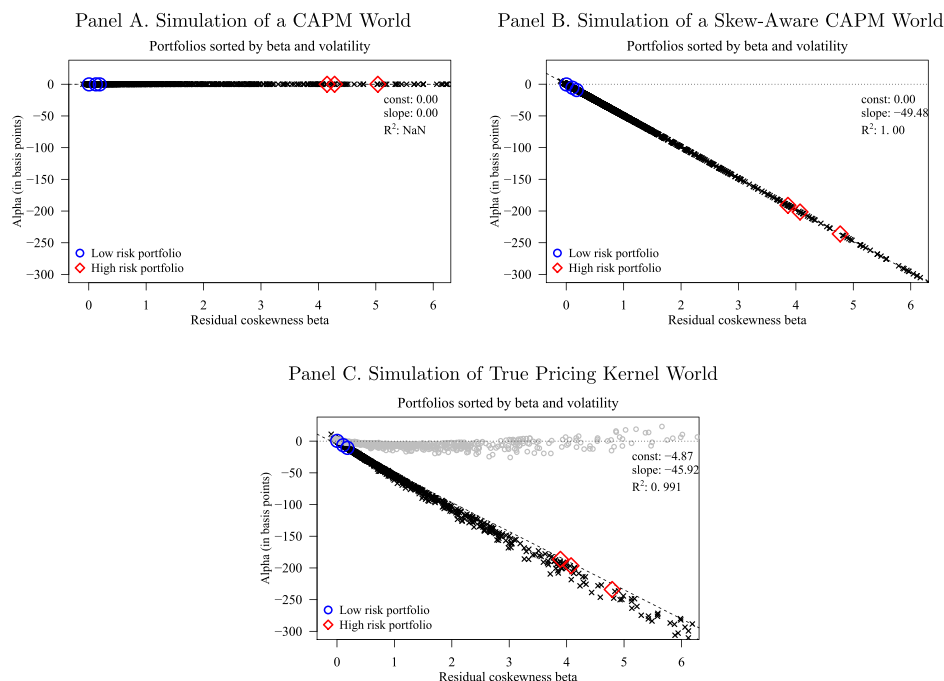


Figure 3. Alphas and residual coskewness: Simulation evidence. This figure plots results for the relation between CAPM alphas and residual coskewness betas in simulated worlds with 2,000 Merton-type firms. The three plots are based on different pricing kernel specifications: a CAPM world (Panel A), a skew-aware world (Panel B), and a world in which moments higher than skewness are also accounted for (Panel C). The figure also presents results for LRA strategies based on equal-weighted decile portfolios, betting against CAPM beta, idiosyncratic volatility, and implied volatility. Blue circles mark the low-risk portfolios that a betting-against-beta (BaB)/volatility strategy goes long. Red diamonds mark the high-risk portfolios that a BaB/volatility strategy goes short. In the legends, we also report results for a cross-sectional regression of alphas on residual coskewness betas. In Panel C, the gray circles mark the alphas after additionally controlling for skewness. (Color figure can be viewed at wileyonlinelibrary.com)

The simulation evidence in Figure 3 illustrates the relation between alphas and residual coskewness for a population of 2,000 (levered) firms for three different pricing kernel specifications. Similar to Figure 2, the panels of Figure 3 also illustrate the long (i.e., low risk) and short (i.e., high risk) positions of an investor who bets against beta, idiosyncratic volatility, or total volatility. Also, as in Figure 2, these positions are marked by blue circles and red diamonds, respectively, in the various panels of the figure. Panel A shows that in a world governed by a CAPM pricing kernel, there would be no alpha to BaB or betting-against-volatility, investors do not care about coskewness, and all CAPM alphas are zero. Panel B plots the simulation results for a skew-aware CAPM world and reports the relation between alpha and residual coskewness discussed above. These patterns are qualitatively identical to those in the empirical data, that is, low-beta stocks generate higher alphas than high-beta stocks

and hence BaB delivers positive alpha.⁷ Finally, Panel C shows the results for a world in which investors also care about moments higher than (co)skewness. The cross-sectional patterns in CAPM alphas and residual coskewness in Panel C are very similar to those in the skew-aware CAPM world (Panel B), which suggests that coskewness is the main determinant of LRA alphas. To make this point more clear, we also plot (in gray circles) coskew-adjusted alphas, that is, firms' expected excess returns due to exposures to moments higher than skewness. These are close to zero, indicating that moments beyond skewness matter much less. The results are essentially the same when we allow for jumps in firm asset value and for default prior to debt maturity (see Figures IA.1 and IA.2 in the Internet Appendix).

We next use the simulation framework to study the relation between a firm's skewness and coskewness. The results, which are depicted in Figure 4, guide our empirical analysis in which we construct forward-looking factors based on ex ante skewness implied by firms' stock option prices. Specifically, we study how measures of implied skewness (under the risk-neutral \mathbb{Q} -measure) as well as measures of expected realized skewness (under the physical \mathbb{P} -measure) are related to each other as well as to residual coskewness and alphas.

We evaluate three dimensions of skewness to comprehensively capture asymmetries in the return distribution. First, we measure the implied skewness that originates from the upper part of the distribution. This *upper skewness* is defined to be positive and can be measured from prices of OTM call options. Second, we measure the implied skewness from the lower part of the distribution, which is defined to be negative. This *lower skewness* can be measured as the price of a portfolio that is short OTM put options. Third, we measure the firm's overall skewness, defined as the sum of upper and lower skewness, which can be positive or negative and quantifies the overall asymmetry of the distribution.⁸

⁷ In the simulations, BaB is driven mostly by the short position in high beta stocks and less by the long position in low-beta stocks. Empirically, the relative contribution of the long position can be somewhat higher. This can be seen from the fitted lines of regressing alphas on residual coskewness betas being on a higher level in the empirical data illustrated in Figure 2 compared to the simulated data in Figure 3. We experimented with different parameter values for the stochastic volatility market model and found that we can reduce this level-difference when we make the market variance less persistent (explosive). Given that our main goal is to show the negative relation between LRA alphas and residual coskewness, we show this relation only for standard model parametrizations established in the literature (e.g., Aït-Sahalia and Kimmel (2007)), as we discuss in the Appendix. In future research, it would be interesting to explore the quantitative implications in more detail, building on ongoing work that aims to model more realistic interactions between the cross-section of firms and the pricing kernel (e.g., Gouriéroux (2016), Boloorforoosh et al. (2017)).

⁸ Considering upper and lower skewness separately in addition to total skewness can be informative, because identical values of total skewness can arise from different combinations of values for lower and upper skewness. This can be easily illustrated by two firms with total skewness close to zero, where for one firm both upper and lower skewness may be close to zero, while for the other firm upper skewness takes a high positive value (e.g., due to growth options) that is offset by a large negative value of lower skewness (e.g., due to leverage).

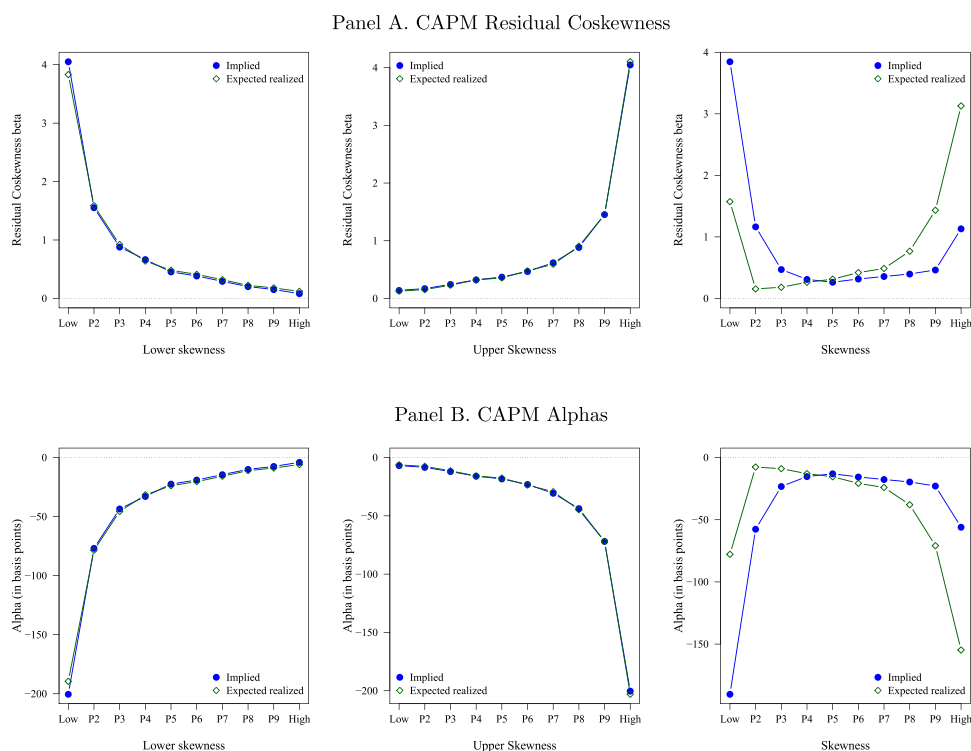


Figure 4. Implied skewness, residual coskewness, and alphas: Simulation evidence.

This figure plots results for the relation between firms' skewness, CAPM residual coskewness, and CAPM alphas in a simulated skew-aware world with 2,000 Merton-type firms. For each firm, we compute measures of implied skewness (under the Q-measure) and expected realized skewness (under the P-measure). We compute lower skewness (by definition always negative), upper skewness (by definition always positive), and total skewness, which is the sum of lower and upper skewness. We sort firms into decile portfolios based on the three Q- and the three P-skew measures, and we compute the portfolios' CAPM alphas and residual coskewness betas. Portfolios P1 and P10 contain firms with the lowest and highest values of the sort variables, respectively. Panel A reports the portfolios' CAPM residual coskewness betas when using measures of Q-skew (blue line with bullets) and P-skew (green line with diamonds) using lower skewness (left), upper skewness (middle), and total skewness (right). Panel B plots the portfolios' corresponding CAPM alphas. (Color figure can be viewed at wileyonlinelibrary.com)

Bearing in mind that our empirical objective is to construct factors based on equity option-implied skew measures, Figure 4 shows the residual coskewness and the alphas of decile portfolios sorted by implied and expected realized lower skewness, upper skewness, and skewness. Portfolios P1 and P10 contain the firms with the lowest and highest values of the conditioning variables, respectively. The left column shows that firms with low (i.e., most negative) lower skewness have the highest residual coskewness and lowest alphas, and that residual coskewness decreases whereas alphas increase when moving toward portfolio P10, which contains firms with the highest lower skewness

(i.e., close to zero). Conversely, in the middle column we see that residual coskewness increases from firms with the lowest (i.e., closest to zero) upper skewness in P1 to firms with the highest (i.e., most positive) upper skewness in P10. Accordingly, the alphas decrease from P1 to P10. In other words, the higher the implied lower or upper skewness in absolute terms, the higher the portfolio's residual coskewness and the more negative its alpha. In line with these patterns, we find that portfolios sorted by implied skewness (i.e., the sum of lower and upper semiskew) exhibit a U-shaped relation toward coskewness and an inversely U-shaped relation to alphas. These findings remain unchanged when we allow for jumps in firms' assets and for default prior to debt maturity (Figures IA.3 and IA.4 in the Internet Appendix).

In our simulations, (co)skewness arises from firms' leverage and stochastic asset volatility. The simulation results suggest a strong link between implied skewness measures, residual coskewness, and alphas. This direct link between a firm's ex ante skewness, its realized coskewness, and its alpha guides our empirical approach to constructing mimicking factors based on portfolios sorted by equity option-implied skewness. The empirical results in Section III are qualitatively similar to the simulation evidence presented above.

II. Setup of Empirical Analysis

In this section, we discuss the data used in the empirical analysis, we describe the construction of beta, volatility, and coskewness measures from stock and market returns, and we provide details on the computation of ex ante skewness and ex ante variance from equity option data. Summary statistics for the variables used in our empirical analysis are reported in Table IA.I in the Internet Appendix.

A. Data

To construct our sample, we first identify U.S. firms for which daily stock and options data are available. We start with volatility surface data from Option-Metrics. We keep those firms for which we find corresponding stock returns in CRSP (common stocks with share code 10 or 11) and firm data in COMPUSTAT to compute market capitalization and book-to-market. The merged data set contains 400,449 monthly observations across 4,967 firms from January 1996 to August 2014. Additionally, we obtain daily return data for the CRSP value-weighted market index as well as daily and monthly market, size, value, and momentum factors and risk-free returns from Ken French's data library. Other data used in additional analyses and in robustness checks are described in the corresponding sections.

For empirical analyses of beta-sorted portfolios using all CRSP stocks over the period June 1963 to December 2014, we drop the requirement of options data. This data set contains 2,017,271 monthly observations across 15,843 firms.

B. Construction of Variables

In our empirical analysis of beta- and idiosyncratic volatility-related LRAs, we apply the measures that were used in the studies originally documenting the anomalies. Below, we describe how we estimate (i) CAPM betas and measures of (idiosyncratic) volatility that we use as sort variables in constructing the LRA portfolios, (ii) measures of ex ante skewness that we use to construct skewness factors, and (iii) residual coskewness.

CAPM betas. First, we estimate CAPM betas as described in Frazzini and Pedersen (2014) using the CRSP value-weighted market index. For security i , the beta estimate is given by

$$\hat{\beta}_i^{TS} = \hat{\rho}_i \frac{\hat{\sigma}_i}{\hat{\sigma}_m},$$

where $\hat{\sigma}_m$ and $\hat{\sigma}_i$ denote the volatility of stock i and of the market excess returns, respectively, and $\hat{\rho}_i$ denotes their correlation with the market. We estimate volatilities as one-year rolling standard deviations of one-day log returns and correlations using a five-year rolling window of overlapping three-day log returns. We require at least 120 and 750 trading days of nonmissing data, respectively. To reduce the influence of outliers, Frazzini and Pedersen (2014) shrink the time-series estimate $\hat{\beta}_i^{TS}$ to the cross-sectional beta mean ($\hat{\beta}^{XS}$),

$$\hat{\beta}_i = w \times \hat{\beta}_i^{TS} + (1 - w) \times \hat{\beta}^{XS},$$

where they set $w = 0.6$ and $\hat{\beta}^{XS} = 1$. Following this procedure, we generate end-of-month estimates of CAPM betas for the sample periods described above.

Idiosyncratic volatility. We use the residuals of this CAPM estimation to construct our measure of idiosyncratic volatility relative to the CAPM. We estimate idiosyncratic volatility following Ang et al. (2006) as the square root of the residual variance from regressing daily equity excess returns of firm i on the daily returns of the three Fama-French factors (market, size, and value) over the previous month.

Ex ante variance and ex ante skewness. We use OTM stock options data to measure firms' ex ante variance (VAR) and ex ante skewness (SKEW), following the approach of Schneider and Trojani (2014).⁹ Let the price of a

⁹ Building on the concepts of Breeden and Litzenberger (1978) and Neuberger (1994), recent research assesses ex ante moments of the equity return distribution based on stock option prices (e.g., Bakshi, Kapadia, and Madan (2003), Kozhan, Neuberger, and Schneider (2013), Martin (2013), Schneider and Trojani (2014)). The common theme is to measure ex ante moments from portfolios of OTM options, but differences arise from the associated portfolio weights. We choose to follow Schneider and Trojani (2014) because their portfolio weights comply with put-call symmetry as in Carr and Lee (2009) and their measures are well defined when the stock price reaches zero, a feature that is essential for individual firms that can default. Also note that the exposition below rests on the assumption that options markets are complete, but only for notational convenience. In our empirical analysis, we use the "tradable" counterparts that are computed from available option data only; see Schneider and Trojani (2014). In other words, they account for market incompleteness and they do not require interpolation or extrapolation schemes to satisfy the assumption that a continuum of option prices is available.

zero-coupon bond with maturity T be denoted by $p_{t,T}$, the forward price of the stock (contracted at time t for delivery at time T) by $F_{t,T}$, and the prices of European put and call options with strike price K by $P_{t,T}(K)$ and $C_{t,T}(K)$, respectively. In our empirical analysis, we measure ex ante moments from equity options with a maturity of 30 days, thereby matching the monthly horizon of equity returns. We compute a firm's options-implied ex ante variance ($VAR_{t,T}^Q$) and ex ante skewness ($SKEW_{t,T}^Q$) as

$$VAR_{t,T}^Q = \frac{2}{p_{t,T}} \left(\int_0^{F_{t,T}} \frac{\sqrt{\frac{K}{F_{t,T}}} P_{t,T}(K)}{K^2} dK + \int_{F_{t,T}}^{\infty} \frac{\sqrt{\frac{K}{F_{t,T}}} C_{t,T}(K)}{K^2} dK \right),$$

$$SKEW_{t,T}^Q = \frac{6}{p_{t,T}} \left(\int_{F_{t,T}}^{\infty} \log\left(\frac{K}{F_{t,T}}\right) \frac{\sqrt{\frac{K}{F_{t,T}}} C_{t,T}(K)}{K^2} dK \right. \\ \left. - \int_0^{F_{t,T}} \left(\log \frac{F_{t,T}}{K} \right) \frac{\sqrt{\frac{K}{F_{t,T}}} P_{t,T}(K)}{K^2} dK \right).$$

To capture asymmetries in the return distribution, we decompose $SKEW_{t,T}^Q$ into upper and lower skewness, to separately account for the left and the right parts of the distribution,

$$upperSKEW_{t,T}^Q = \frac{6}{p_{t,T}} \left(\int_{F_{t,T}}^{\infty} \log\left(\frac{K}{F_{t,T}}\right) \frac{\sqrt{\frac{K}{F_{t,T}}} C_{t,T}(K)}{K^2} dK \right),$$

$$lowerSKEW_{t,T}^Q = -\frac{6}{p_{t,T}} \left(\int_0^{F_{t,T}} \left(\log \frac{F_{t,T}}{K} \right) \frac{\sqrt{\frac{K}{F_{t,T}}} P_{t,T}(K)}{K^2} dK \right),$$

that is, we have $SKEW_{t,T}^Q = upperSKEW_{t,T}^Q + lowerSKEW_{t,T}^Q$. While upper skew is by definition positive and lower skew is by definition negative, the sign of overall ex ante skew depends on the relative prices of OTM put and OTM call options. As mentioned above (in footnote), considering upper and lower skew separately is informative, because identical values of ex ante skew can arise from different combinations of values for lower and upper skew.

Residual coskewness. Finally, to measure (residual) coskewness, we compute the covariance of portfolio (residual) returns with squared market excess returns. More specifically, to measure return coskewness, we compute the sample covariance of the portfolios' raw excess returns with squared market excess returns. To measure FM residual coskewness, we regress the portfolios' raw excess returns on the factors suggested by the corresponding FM, obtain the residual returns unexplained by the factors, and compute the sample covariance of these residual returns with squared market excess returns,

that is,

$$\text{FM residual coskewness} = \text{Cov}[\hat{\alpha}_i^{FM} + \hat{\epsilon}_{i,t+1}^{FM}, R_{m,t+1}^2].$$

When the FM is specified to be the CAPM, this corresponds to the sample estimate of residual coskewness as defined in equation (8). To facilitate interpretation, we also compute residual coskewness betas, by regressing FM residual returns on squared market excess returns, that is, residual coskewness scaled by the variance of squared market excess returns,

$$\text{FM residual coskewness beta: } \hat{\beta}_i^{FM-RESCOS} = \frac{\text{Cov}[\hat{\alpha}_{i,t}^{FM} + \hat{\epsilon}_{i,t+1}^{FM}, R_{m,t+1}^2]}{\text{Var}[R_{m,t+1}^2]}.$$

Consistent with the market choice in the CAPM beta estimation, we use excess returns of the CRSP value-weighted market index for the computation of $R_{m,t+1}^2$, but results are virtually identical when we use the Fama-French market factor instead.

III. Empirical Analysis

Our theoretical results in Section I suggest that the positive alphas of BaB and betting-against-volatility strategies may be driven by compensation required by skew-aware investors. In this case, the (residual) returns of LRAs should have a common determinant that is related to coskewness. Controlling for coskewness should reduce the negative residual coskewness of LRA strategies and, as a consequence, render LRA alphas insignificant. In this section, we provide strong support for these predictions.

A. Implied Skewness, Residual Coskewness, and Alphas

To study the importance of skewness for LRAs, we use stock option-implied measures of a firm's ex ante skewness. Figure 5 shows that the empirical links between skewness, residual coskewness, and alpha are qualitatively identical to those in the simulations presented in Figure 4 above.

The empirical results in Figure 5 are based on decile portfolios sorted by three measures of option-implied skewness, as defined in Section II.B: lower skewness (left column), upper skewness (middle column), and skewness computed as the sum of lower and upper skewness (right column). Panel A plots the portfolios' CAPM residual coskewness betas, obtained from regressing the residual returns on squared market excess returns, and Panel B plots the corresponding CAPM alphas. In all plots, blue lines with bullets correspond to value-weighted portfolios and green lines with diamonds correspond to equal-weighted portfolios. The results accord well with those from the simulation study: firms with low (i.e., the most negative) lower skewness have the highest residual coskewness and the lowest alphas. The residual coskewness decreases whereas alpha increases when moving toward portfolio P10, which contains firms with the highest lower skewness (i.e., close to zero). Conversely, in the middle column, we see that residual coskewness increases from firms with the




Figure 5. Option-implied ex ante skewness, residual coskewness, and alphas. This figure plots results for the relation between firms' equity-option implied ex ante skewness, CAPM residual coskewness, and CAPM alphas. For each firm, we compute option-implied lower skewness (by definition always negative), upper skewness (by definition always positive), and total skewness, which is the sum of upper and lower skewness. We sort firms into decile portfolios based on the three option-implied skew measures and compute the portfolios' CAPM alphas and residual coskewness betas. Portfolios P1 and P10 contain firms with the lowest and highest values of the sort variables, respectively. Panel A plots the portfolios' CAPM residual coskewness betas for equal-weighted portfolios (blue line with bullets) and value-weighted portfolios (green line with diamonds) using lower skewness (left), upper skewness (middle), and total skewness (right). Panel B plots the portfolios' corresponding CAPM alphas. Alphas are reported in basis points per month. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations. (Color figure can be viewed at wileyonlinelibrary.com)

lowest (i.e., closest to zero) upper skewness in P1 to firms with the highest (i.e., most positive) upper skewness in P10. Accordingly, alphas decrease from P1 to P10. In other words, the higher the implied lower skewness or the implied upper skewness in absolute terms, the higher the portfolio's residual coskewness and the more negative its alpha.¹⁰ In line with these patterns, we find that portfolios sorted by implied skewness, computed as the sum of lower and

¹⁰ In the model in Section I.B, the linkages between ex ante skewness, residual coskewness, and alpha arise endogenously from firm leverage and stochastic asset volatility. In the empirical analysis, we do not take a stand on the specific source of (co)skewness. Identifying the determinants of coskewness would be in its own right. In Section IV, we extend the empirical analysis to

upper skewness, exhibit a U-shaped relation with residual coskewness and an inversely U-shaped relation with alpha.

Encouragingly, these empirical results show that in the data ex ante skewness is as informative for future CAPM residual coskewness and alpha as it is in our simulated world. The results are similar when we also control for the small-minus-big (SMB) and high-minus-low (HML) factors of Fama and French (1993, FF3); see Figure IA.5 in the Internet Appendix. When we additionally control for momentum (FF4, following Carhart, 1997), we find the same patterns in the relations between residual coskewness and alpha on the one hand and to upper and lower skewness on the other hand (see Figure IA.6). We also find the U-shaped relation between skewness and residual coskewness whereas the relation between skewness and alpha is mostly increasing and the inverse U-shape is less pronounced. In a robustness check, we repeat the analysis with the equity option-implied skewness measure of Kozhan, Neuberger, and Schneider (2013) and obtain similar results (see Figure IA.7).¹¹ All of these results are consistent with the notion that alpha reflects compensation for residual coskewness.

B. Skew-Adjusted Returns of BaB and Betting-Against-Volatility

The results above suggest that measures of ex ante skewness are indeed informative for stocks' future residual coskewness and alpha. We now construct skew factors and study whether controlling for skewness explains LRAs.

We construct the skew factors from the portfolios sorted by lower, upper, and total ex ante skewness such that they generate positive alpha and negative residual coskewness. In total, we consider four skew factor specifications. First, we compute HML returns of portfolios sorted by lower skewness (*LSK* factor). Second, we compute low-minus-high returns of portfolios sorted by upper skewness (*USK* factor). Third, we construct a skew factor from portfolios sorted by ex ante skewness (shown in the right column of Figure 5). To capture the U-shaped relation between ex ante skewness and residual coskewness, we compute the *SK* factor from the returns of going short the extreme portfolios (P1 and P10) and going long the middle portfolios (P5 and P6). Similarly, we construct the *LUSK* factor as the sum of *LSK* and *USK*.

Table II presents summary statistics for the skew factors.¹² In line with our results on LRAs that we study in this paper, we find that the raw excess return

credit ratings and CDS spreads and find that the leverage-induced predictions of our model find support in the data, which suggest that firms' (co)skewness is indeed related to their credit risk.

¹¹ We use the skewness measure of Kozhan, Neuberger, and Schneider (2013) in the robustness check because, similar to our measure, it is a measure of skewness that is not standardized by variance, whereas the measure by Bakshi, Kapadia, and Madan (2003), for instance, is a standardized measure. Since we explicitly study the interaction between (co)skewness and (co)variance, our analysis focuses on non standardized skewness measures. In addition, we decompose the skewness measure of Kozhan, Neuberger, and Schneider (2013) to separate upper skewness and lower skewness components and find results similar to those reported above.

¹² For detailed summary statistics for the underlying decile portfolios' returns, alphas, and other characteristics, see Tables IA.II and IA.III for value-weighted portfolios and Tables IA.IV and IA.V for equal-weighted portfolios.

Table II
Summary Statistics for Ex Ante Skewness Factors

This table reports excess returns, factor model alphas, and residual coskewness of ex ante skewness factors constructed from value- and equal-weighted decile portfolios (Panels A and B, respectively). At the end of each month, we sort firms into decile portfolios based on their option-implied lower skewness, upper skewness, or total skewness. From these portfolios, we construct four factors. First, we compute the high-minus-low returns of portfolios sorted by lower skewness (*LSK* factor). Second, we compute the low-minus-high returns of portfolios sorted by upper skewness (*USK* factor). Third, we compute the *SK* factor from the portfolios sorted by skewness by going short the extreme portfolios (P1 and P10) and going long the middle portfolios (P5 and P6). Finally, we construct the *LUSK* factor as the sum of *LSK* and *USK*. We report raw excess returns as well as alphas of CAPM, Fama-French three-factor (FF3), and Fama-French-Carhart four-factor (FF4) regressions. Excess returns and alphas are reported in basis points per month. Values in square brackets are *t*-statistics based on standard errors following Newey and West (1987), where we choose the optimal truncation lag as suggested by Andrews (1991). Additionally, we report the annualized information ratio associated with the alphas, measured as the average residual return divided by its standard deviation. Finally, we report (residual) coskewness, measured as the coskewness beta obtained from regressing (residual) returns on squared market excess returns. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations.

	LSK	USK	SK	LUSK
Panel A: Skewness Factors Constructed from Value-Weighted Portfolios				
Excess return	58.04	−8.43	32.54	49.61
<i>t</i> -statistic	[0.73]	[−0.10]	[0.34]	[0.30]
Sharpe ratio	0.17	−0.02	0.07	0.07
Coskewness	3.21	3.45	2.83	6.66
CAPM alpha	158.94	92.35	156.45	251.29
<i>t</i> -statistic	[2.93]	[1.66]	[2.53]	[2.32]
Information ratio	0.63	0.34	0.45	0.48
Residual coskewness	−2.97	−2.73	−4.76	−5.69
FF3 alpha	156.02	93.95	178.68	249.97
<i>t</i> -statistic	[3.10]	[1.95]	[3.17]	[2.65]
Information ratio	0.74	0.41	0.56	0.57
Residual coskewness	−3.11	−2.92	−5.16	−6.03
FF4 alpha	94.31	25.99	70.86	120.30
<i>t</i> -statistic	[2.39]	[0.60]	[1.17]	[1.39]
Information ratio	0.57	0.14	0.31	0.35
Residual coskewness	−1.88	−1.57	−3.01	−3.45
Panel B: Skewness Factors Constructed from Equal-Weighted Portfolios				
Excess return	35.59	−24.75	55.41	10.84
<i>t</i> -statistic	[0.52]	[−0.36]	[0.75]	[0.08]
Sharpe ratio	0.11	−0.08	0.17	0.02
Coskewness	2.69	2.58	2.19	5.27
CAPM alpha	123.87	64.57	144.52	188.44
<i>t</i> -statistic	[2.36]	[1.22]	[2.47]	[1.79]
Information ratio	0.52	0.27	0.55	0.39
Residual coskewness	−2.72	−2.89	−3.27	−5.61

(Continued)

Table II—Continued

	LSK	USK	SK	LUSK
FF3 alpha	117.95	59.08	147.53	177.03
<i>t</i> -statistic	[3.79]	[1.80]	[3.50]	[2.79]
Information ratio	0.74	0.36	0.74	0.55
Residual coskewness	−2.88	−3.05	−3.51	−5.93
FF4 alpha	69.08	8.61	83.14	77.69
<i>t</i> -statistic	[1.99]	[0.22]	[2.11]	[1.06]
Information ratio	0.56	0.07	0.57	0.31
Residual coskewness	−1.91	−2.04	−2.23	−3.95

of the skew factors are not significantly different from zero. However, once we control for the CAPM, the skew factors deliver large positive alphas (in the range of 0.65%–2.51%) that are associated with negative residual coskewness betas (between −2.72 and −5.69). The level of significance depends on the particular measure of ex ante skewness, with results being most significant and associated with the highest information ratios (0.63 value-weighted, 0.52 equal-weighted) for *LSK* and least significant with the lowest information ratios (0.34 value-weighted, 0.27 equal-weighted) for *USK*. The results are similar for the skew factors’ FF3 alphas and residual coskewness betas, but become less pronounced once we also control for momentum. The latter finding may indicate that coskewness and momentum are related, as discussed by Harvey and Siddique (2000).

Next, we show that controlling for skewness reduces both the positive alphas and the negative coskewness of LRAs, consistent with the view that investors require compensation for negative coskewness. We use the skew factors constructed above and estimate the LRAs’ skew-adjusted alphas by running the regressions

$$\begin{aligned} R_{i,t+1} &= \alpha_i + \sum_j \beta_i^j FF_{t+1}^j + \gamma_i^{SK} SK_{t+1} + \varepsilon_{i,t+1}, \\ R_{i,t+1} &= \alpha_i + \sum_j \beta_i^j FF_{t+1}^j + \gamma_i^{LU} LUSK_{t+1} + \varepsilon_{i,t+1}, \\ R_{i,t+1} &= \alpha_i + \sum_j \beta_i^j FF_{t+1}^j + \gamma_i^L LSK_{t+1} + \gamma_i^U USK_{t+1} + \varepsilon_{i,t+1}, \end{aligned}$$

where FF_{t+1}^j denotes the corresponding benchmark factor returns included in the CAPM, FF3, and FF4 regressions. Figure 6 shows that any of the skew adjustments leads to a substantial decrease in alphas compared to those without skew adjustments. The figure also shows that the reductions in alphas are associated with the residual coskewness of the LRA strategies becoming much less negative and closer to zero. These results are consistent with the predictions developed in Section I. First, measures of ex ante skewness contain information about future residual coskewness, which can be seen from the fact

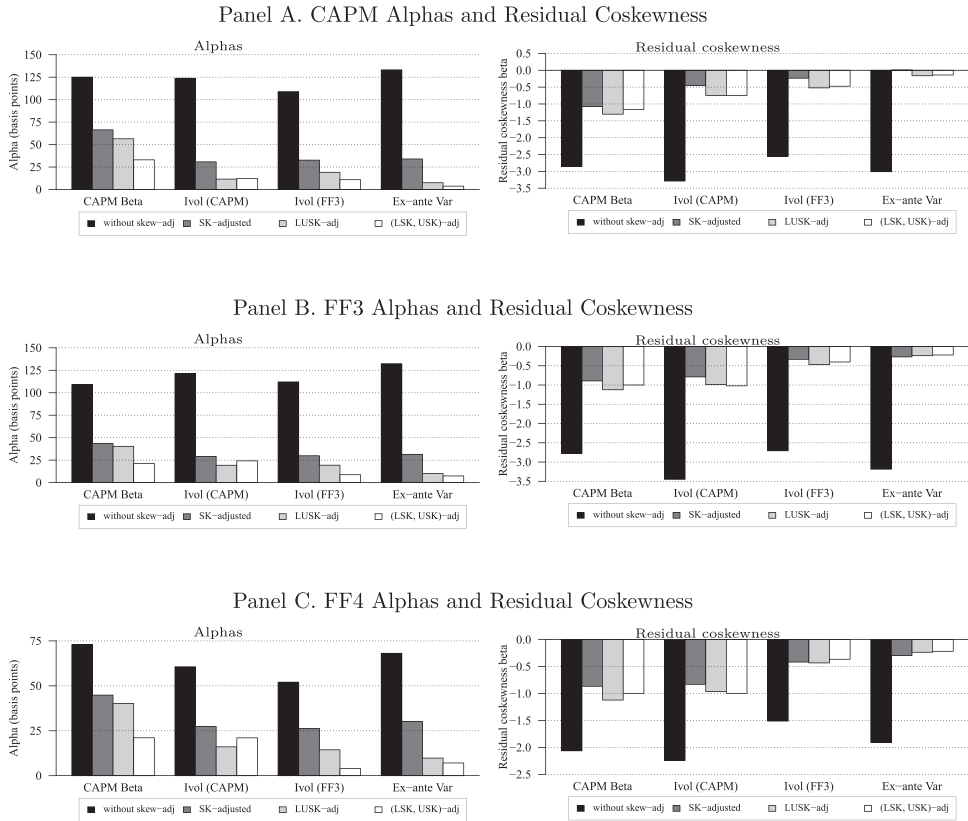


Figure 6. Skew-adjusted alphas and residual coskewness. This figure plots alphas and skew-adjusted alphas of low-risk anomalies (LRAs). At the end of each month, we sort firms into value-weighted decile portfolios based on their CAPM beta, idiosyncratic volatility (measured from the residual variance of CAPM and Fama-French three-factor model regressions), or equity option-implied exante variance. From these portfolios, we compute low-minus-high returns generated by betting-against-beta (BaB)/volatility strategies. On the left, we present alphas (in basis points per month) of CAPM-, Fama-French three-, and four-factor regressions (black bars in Panels A, B, and C, respectively) as well as alphas that also include controls for skewness. To adjust for skewness, we use the *SK* factor (dark gray bars), the *LUSK* factor (light gray bars), or both the *LSK* factor and the *USK* factor (white bars). On the right, we present the residual coskewness betas associated with the LRAs. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations.

that skew-adjusted LRAs have less negative residual coskewness. Second, controlling for coskewness largely eliminates the alphas.

The reductions in alphas are economically large and render most alphas of BaB and betting-against-volatility statistically insignificant, as we show in more detail in Table III. The only alphas that remain borderline significant are the CAPM alphas of BaB when we use the *SK* factor or the *LUSK* factor as a control variable, with the *t*-statistics being 1.87 and 1.70,

Table III

Low-Risk Anomalies (LRAs): Controlling for Ex Ante Skew Factors

This table reports skew-adjusted factor model alphas and residual coskewness of LRAs using value-weighted decile portfolios. At the end of each month, we sort firms into decile portfolios based on their CAPM beta, idiosyncratic volatility (measured from the residual variance of CAPM and Fama-French three-factor model regressions), or equity option-implied ex ante variance. From these portfolios, we compute low-minus-high returns generated by betting-against-beta/volatility strategies and report alphas of CAPM-, Fama-French three-, and four-factor regressions that additionally include controls for skewness. To adjust for skewness, we use either the *SK* factor (Panel A), the *LUSK* factor (Panel B), or both the LSK- and the USK-factor (Panel C). Alphas are reported in basis points per month. Values in square brackets are *t*-statistics based on standard errors following Newey and West (1987), where we choose the optimal truncation lag as suggested by Andrews (1991). Additionally, we report the annualized information ratio associated with the alphas, measured as the average residual return divided by its standard deviation. Finally, we report residual coskewness, measured as the coskewness beta obtained from regressing residual returns on squared market excess returns. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations.

	CAPM Beta	Ivol (CAPM)	IVol (FF3)	ExAnte Var
Panel A: Controlling for Skewness using the <i>SK</i> Factor				
CAPM alpha	66.45	30.80	32.70	34.01
<i>t</i> -statistic	[1.86]	[0.84]	[1.06]	[1.22]
Information ratio	0.45	0.19	0.24	0.24
Residual coskewness	−1.07	−0.46	−0.24	0.01
FF3 alpha	43.70	29.17	29.91	31.37
<i>t</i> -statistic	[1.34]	[0.83]	[0.92]	[1.09]
Information ratio	0.32	0.22	0.23	0.26
Residual coskewness	−0.89	−0.79	−0.33	−0.27
FF4 alpha	44.85	27.41	26.32	30.23
<i>t</i> -statistic	[1.40]	[0.82]	[0.94]	[0.94]
Information ratio	0.33	0.20	0.21	0.25
Residual coskewness	−0.87	−0.83	−0.42	−0.30
Panel B: Controlling for Skewness using the <i>LUSK</i> Factor				
CAPM alpha	56.47	11.60	19.18	7.53
<i>t</i> -statistic	[1.74]	[0.39]	[0.61]	[0.82]
Information ratio	0.41	0.10	0.17	0.18
Residual coskewness	−1.30	−0.75	−0.52	−0.16
FF3 alpha	40.35	19.29	19.38	10.16
<i>t</i> -statistic	[1.22]	[0.69]	[0.65]	[1.10]
Information ratio	0.30	0.17	0.18	0.25
Residual coskewness	−1.12	−0.99	−0.47	−0.24
FF4 alpha	40.16	16.06	14.42	9.80
<i>t</i> -statistic	[1.12]	[0.57]	[0.56]	[1.04]
Information ratio	0.30	0.14	0.14	0.24
Residual coskewness	−1.12	−0.96	−0.43	−0.23

(Continued)

Table III—Continued

	CAPM Beta	Ivol (CAPM)	Ivol (FF3)	ExAnte Var
Panel C: Controlling for Skewness using the LSK and the USK Factor				
CAPM alpha	33.02	12.33	10.93	3.72
<i>t</i> -statistic	[1.04]	[0.36]	[0.35]	[0.32]
Information ratio	0.24	0.10	0.10	0.09
Residual coskewness	−1.16	−0.75	−0.47	−0.14
FF3 alpha	21.27	24.21	8.80	7.43
<i>t</i> -statistic	[0.63]	[0.75]	[0.31]	[0.61]
Information ratio	0.16	0.21	0.08	0.18
Residual coskewness	−1.00	−1.02	−0.40	−0.22
FF4 alpha	21.11	21.06	3.99	7.09
<i>t</i> -statistic	[0.58]	[0.65]	[0.15]	[0.62]
Information ratio	0.16	0.18	0.04	0.17
Residual coskewness	−1.00	−1.00	−0.36	−0.22

respectively. But even for those alphas there is a large reduction in the magnitude, from an unadjusted 125 basis points (with *t*-statistic of 2.87) to 66 basis points and 56 basis points, respectively. All other alphas are insignificant, with all three skew adjustments producing quite similar results; in most cases the skew adjustment using both *LSK* and *USK* leads to the largest reduction in alphas and residual coskewness. The results remain qualitatively unchanged when we use equal and rank-weighted portfolios (Figures IA.8 and IA.9 in the Internet Appendix).

Finally, we revisit the link between alphas and residual coskewness at the portfolio level, as we did at the outset of Section I in Figure 2. In Figure 7, we illustrate the relation between alphas and residual coskewness when controlling for skewness (using *LSK* and *USK*) at the portfolio level. The results show that after accounting for skewness, there is much less (systematic) dispersion in the 80 portfolios' alphas and residual coskewness. Compared to the results without skew adjustment in Figure 2, we continue to find a negative relation between CAPM alphas and residual coskewness, but the cross-sectional regression yields a less negative slope estimate and the R^2 decreases from 73% to 22%. Hence, ex ante skewness is not a perfect predictor of future realized coskewness but it accounts for a large share of its variation and thereby renders LRAs insignificant. For the FF3 and FF4 residual returns, we find that the slope estimates from regressing alphas on residual coskewness are zero and that R^2 decreases from 73% to 4% and from 48% to 3%, respectively. Our findings are similar when we construct portfolios that are identical in terms of CAPM risk, that is, when we (de-)lever all portfolios to a beta of one at portfolio formation.

Taken together, the results above support the view that the positive alphas of BaB and betting-against-volatility represent compensation for coskewness risk, rather than anomalous returns.

C. Skewness as a Common Driver of LRAs

The results in the previous subsection show that controlling for skew factors, constructed from portfolios sorted by ex ante skewness, reduces the positive alphas as well as the negative coskewness of LRA strategies. To corroborate our interpretation that LRA alphas represent compensation for coskewness risk, we now show that LRAs have a common determinant that is related to coskewness risk. We conduct PCAs of the excess and FM residual returns of LRA strategies, and we use all portfolios sorted by beta and volatility. We show that the first principal component of LRAs is connected to coskewness risk as measured by the skew factors.¹³

We start by performing a PCA on the LRA excess returns as well as on their CAPM, FF3, and FFC4 residual returns. In the main text, we present results for low-minus-high returns of value-weighted decile portfolios sorted by beta or volatility in Table IV. Panel A shows that the first PC (PC1) explains around 91% of the variation in LRA excess returns. After controlling for the market and other risk factors (size, value, and momentum), PC1 still explains around 72% of variation in the LRA residual returns. In turn, when we regress the individual LRA's (residual) returns on PC1, we find that all LRAs load on PC1 with coefficients in the range of 0.37 to 0.58 (across all LRAs and all specifications). The explanatory power of PC1 for the individual LRAs is high, with R^2 s in the range of 85% to 95% for excess returns and 61% to 80% for FF4 residual returns. For equal- and rank-weighted portfolios, the same analysis suggests an even higher degree of comovement among LRAs. The results are similar when additionally controlling for liquidity, profitability, and investment factors.¹⁴

Next, we extend the analysis to the portfolios underlying the LRA strategies. In Panel B of Table IV, we first report the variation explained by the common components from a PCA of the eight high and low beta/volatility portfolios

¹³ Our approach to studying whether LRAs are driven by a common component that is linked to coskewness risk can be connected to recent research in asset pricing that uses PCs of returns to reduce the dimensionality of SDFs and to shrink the crosssection. Kozak, Nagel, and Santosh (2018) show that the crosssection of expected returns can be priced with a low-dimensional SDF that uses the first few return PCs as factors. Kozak, Nagel, and Santosh (2020) show that parsimonious SDF representations with a small number of PCs of characteristic-based factor returns achieve good out-of-sample performance. They conclude that the PCs distill the characteristics' SDF contributions, which allows one to focus on the determinants of the PCs (rather than of all characteristics). This is exactly our intent here, by showing that the common component in excess and residual returns of beta- and volatility-sorted portfolios is related to coskewness risk.

¹⁴ In the Internet Appendix, Tables IA.VI and IA.VII report results for equal- and rank-weighted portfolios, respectively. Table IA.VIII presents results when we augment the FF4 specification with the liquidity factor of Pástor and Stambaugh (2003, FF4 LIQU) and when we compute residuals relative to the five-factor model of Fama and French (2015, FF5). That is, when we augment the FF3 specification with profitability and investment factors. Since all of these results are qualitatively identical to those reported in the paper, and the LRA literature typically reports CAPM, FF3, and FF4 alphas, we focus on these specifications in our empirical analysis. Similar to the results reported in the paper, we find that controlling for skewness also reduces the FF4, LIQU, and FF5 alphas and the negative residual coskewness of LRAs.

Table IV
Principal Components of Low-Risk Anomalies (LRAs)

This table presents evidence on a common determinant of beta- and volatility-based LRAs. In Panel A, we compute LRA returns as low-minus-high returns of value-weighted decile portfolios that we sort by their CAPM beta, idiosyncratic volatility (measured from the residual variance of CAPM and Fama-French three-factor model regressions), or equity option-implied ex ante variance at the end of each month. We compute the LRAs' excess returns as well as alphas of CAPM, Fama-French three-factor (FF3), and Fama-French-Carhart four-factor (FF4) regressions and present results for principal component analyses of the corresponding residual returns. The left part of Panel A reports the return variation explained by each of the four principal components (PCs). The right part reports the coefficients as well as the associated R^2 s of regressing LRA returns on PC1. Values in square brackets are t -statistics based on standard errors following Newey and West (1987), where we choose the optimal truncation lag as suggested by Andrews (1991). Panel B presents results from principal component analyses using the portfolios underlying the computation of the LRAs. On the left, we report the variation explained by the first four principal components based on the returns of the eight high and low beta/volatility-portfolios (rather than the four low-minus-high returns). On the right, we report the variation explained by the first four principal components based on the returns of all 40 portfolios sorted by beta and volatility. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations.

Panel A: Principal Components of Low-Risk Anomalies								
	Variation Explained by PCs				Anomaly Return Loadings on PC1			
	PC1	PC2	PC3	PC4	CAPM Beta	Ivol (CAPM)	Ivol (FF3)	Ex Ante Var
Excess returns	91.05	4.67	2.24	2.04	0.47 [36.70]	0.54 [40.67]	0.43 [22.88]	0.56 [46.74]
R^2 (%)					84.60	94.08	88.76	94.68
CAPM residuals	85.41	6.94	3.96	3.69	0.37 [20.05]	0.57 [31.62]	0.45 [17.96]	0.58 [36.68]
R^2 (%)					66.94	91.55	83.09	91.29
FF3 residuals	81.34	8.80	5.08	4.77	0.39 [16.69]	0.55 [27.76]	0.49 [19.92]	0.56 [34.17]
R^2 (%)					64.06	87.34	81.50	87.21
FF4 residuals	72.43	12.64	7.71	7.21	0.46 [14.01]	0.55 [23.62]	0.44 [17.26]	0.54 [25.00]
R^2 (%)					61.09	80.46	67.54	78.42
Panel B: Principal Components of Portfolios Sorted by Beta and Volatility								
	Low and high portfolios Variation explained by PCs				All portfolios Variation explained by PCs			
	PC1	PC2	PC3	PC4	PC1	PC2	PC3	PC4
Excess returns	89.17	5.25	2.11	1.46	84.06	5.94	1.73	0.94
CAPM residuals	78.47	7.14	4.97	4.30	51.83	9.36	6.84	3.45
FF3 residuals	77.11	7.57	5.65	5.16	48.61	9.11	5.15	4.00
FF4 residuals	65.45	11.28	8.64	7.76	31.94	12.19	6.20	5.22

underlying the LRAs. We find that the common variation in portfolio excess and residual returns explained by the first PC is of a similar magnitude as in the previous exercise that uses the four low-minus-high returns. When we repeat the PCA using all 40 beta- and volatility-sorted portfolios, we find that PC1 explains around 84% of the variation in excess returns, approximately

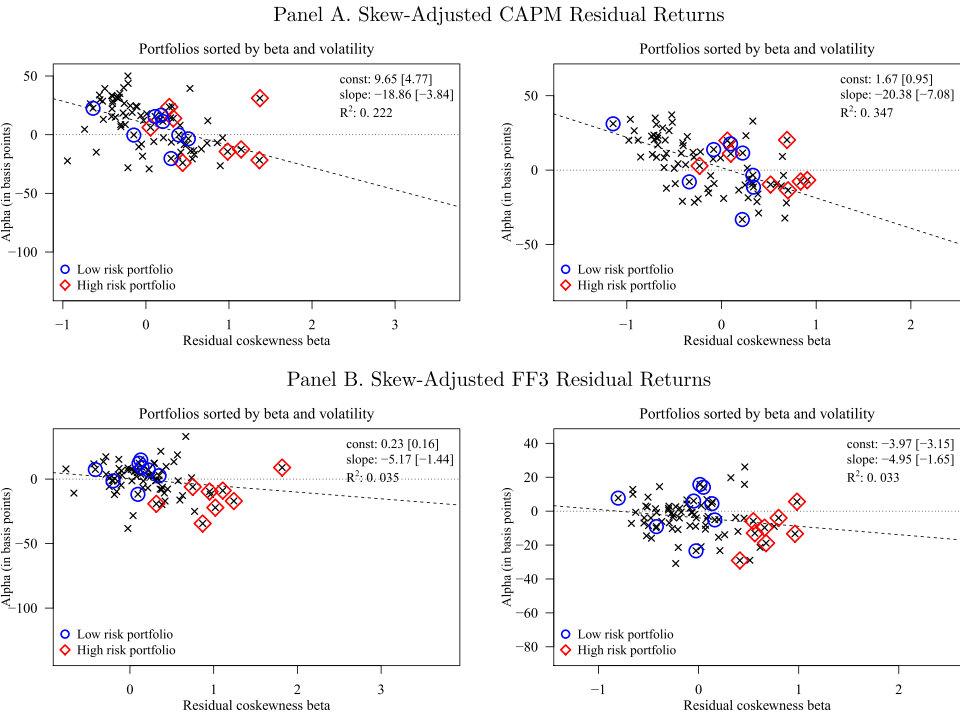


Figure 7. Low-risk anomalies (LRAs): Skew-adjusted alphas and residual coskewness. This figure plots results for the equal-weighted and value-weighted decile portfolios used to compute the skew-adjusted LRA returns in Table III. At the end of each month, we sort firms into decile portfolios based on their CAPM beta, idiosyncratic volatility (measured from the residual variance of CAPM and Fama-French three-factor model regressions), or equity option-implied ex-ante variance. In total, we have 80 portfolios: 10 portfolios for each of the four LRAs, both equal- and value-weighted. We plot skew-adjusted CAPM, Fama-French three-factor, and Fama-French-Carhart four-factor alphas (in basis points per month) against their corresponding coskewness betas, measured from regressing factor model residual returns on squared market excess returns. To adjust for skewness, we add the LSK and the USK factors to the factor model regressions. In each panel, the left plot presents results for standard portfolio sorts and the right plot presents results for portfolios (de-)levered to a CAPM beta of one at portfolio formation, that is, initially, all portfolios are equally risky as measured by CAPM beta. Blue circles mark the low-risk portfolios that a betting-against-beta (BaB)/volatility strategy goes long. Red diamonds mark the high-risk portfolios that a BaB/volatility strategy goes short. In the legends, we also report results for a cross-sectional regression of portfolio alphas on coskew betas; we report the slope coefficient, its t -statistic (based on White (1980) standard errors), and regression R^2 s. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations. (Color figure can be viewed at wileyonlinelibrary.com)

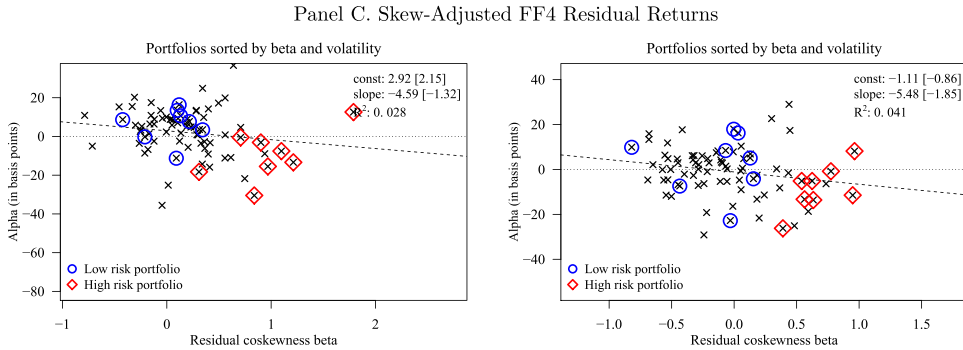


Figure 7. *Continued*

50% of the variation in CAPM and FF3 residual returns, and 32% of the variation in FF4 residual returns. To better understand the structure in these returns, we regress portfolio (residual) returns on PC1. Figure 8 presents the results for CAPM residual returns. In line with the low-minus-high strategies' positive PC1 loadings (Panel A), we find that the portfolio loadings decrease as beta and volatility increase, with coefficients being significantly positive for low beta/volatility stocks and significantly negative for high beta/volatility stocks. The regression R^2 s exhibit a U-shape, consistent with our results above (e.g., in Figure 2) that coskewness matters most for extreme beta/volatility portfolios, particularly for high compared to low beta/volatility stocks. Our findings are similar when we use FF3 and FF4 residual returns, as we show in Figures IA.10 and IA.11 in the Internet Appendix.

These results indicate that LRA strategies as well as their underlying beta- and volatility-sorted portfolios have a common determinant that explains a large part of the variation in their returns. To explore whether this common determinant is related to (co)skewness, we use the ex ante skew factors proposed in the previous subsection. Table V presents the results for regressions of PC1 on skew factor returns.

Panel A reports results for regressions that use the PC1 estimated from the four LRA strategies as the common determinant. We find that *LSK* and *USK* each explain more than 93% of the variation in PC1 estimated from LRA excess returns and more than 75% when PC1 is estimated from FF4 residual returns. Using the *SK* factor, the R^2 s of these regressions are substantial as well, at 83% in the first specification and 52% in the second. Next, we use the *LUSK* factor and find that the explanatory power is between 95% for the PC1 estimated from excess returns and 80% for the PC1 estimated from FF4 residual returns. Finally, the results are similar when we regress PC1 on both *LSK* and *USK*. In all specifications, the regression coefficients are significantly positive.

Panel B reports results for regressions that use the PC1 estimated from all 40 beta- and volatility-sorted portfolios. We find that the explanatory power of skew factor returns for PC1 is between 67% and 71% when using excess

returns, between 76% and 86% when using CAPM residual returns, between 75% and 79% when using FF3 residual returns, and between 54% and 69% when using FF4 residual returns. Hence, coskewness risk, as captured by ex ante skew factors, appears to be strongly related to the common return determinant of beta- and volatility-sorted portfolios as well. To shed light on the portfolios' exposures to the skewness factors, we regress the portfolios on skewness factor returns. Figure 9 presents results for CAPM residual returns using the *LUSK* factor, which show that the exposures decrease from significantly positive estimates for low beta/volatility stocks to significantly negative estimates for high beta/volatility stocks. The regression R^2 s exhibit a U-shape,

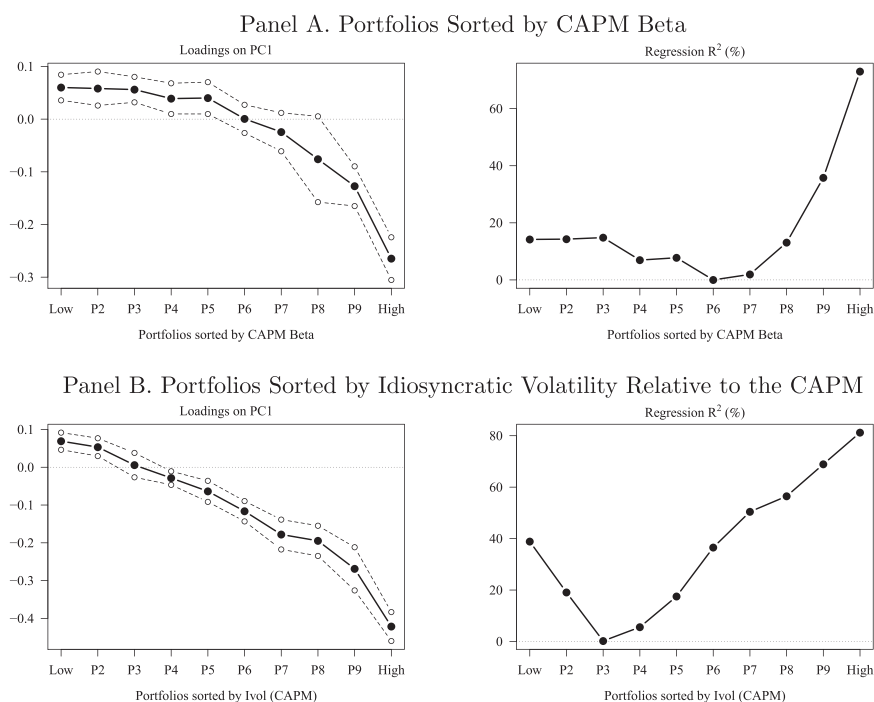


Figure 8. Common return component in beta- and volatility-sorted portfolios. This figure plots evidence for a common return determinant of portfolios sorted by beta and volatility. We present results for value-weighted decile portfolios that we sort by their CAPM beta, idiosyncratic volatility (measured from the residual variance of CAPM and Fama-French three-factor model regressions), or equity option-implied ex ante variance at the end of each month. For each of the 40 portfolios, we compute the portfolios' CAPM residual returns. We then use the 40 residual return series in principal component analysis (PCA) to extract the first principal component. For each of the four sort variables we present results (in Panels A to D, respectively) from regressing portfolio residual returns on PC1. On the left, we present the regression slope estimates, along with 95% confidence intervals based on standard errors following Newey and West (1987), where we choose the optimal truncation lag as suggested by Andrews (1991). On the right, we plot the regression R^2 s. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations.

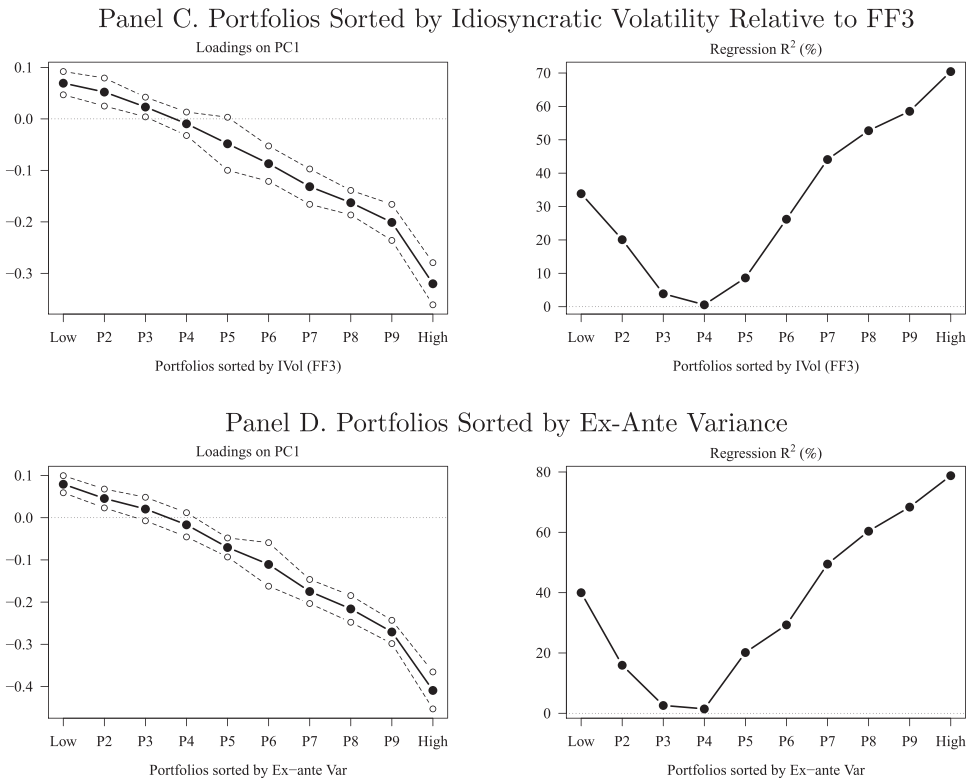


Figure 8. *Continued*

with values in the range of 15% to 40% for low beta/volatility stocks and 55% to 90% for high beta/volatility stocks. The results for other return and skew factor specifications, presented in the Internet Appendix, corroborate our findings.^{15,16}

Summarizing, 72% of the variation in LRA excess returns after controlling for the market, SMB, HML, and momentum are driven by a single component. In turn, up to 80% of this common component can be explained by the returns of factors constructed from portfolios sorted by ex ante skewness. Together with the significantly positive (negative) loadings of low (high) beta and low (high) volatility stocks on the common determinant as well as on the skew

¹⁵ For the *LUSK* factor, we present results for FF3 and FF4 residual returns in Figures IA.12 and IA.13. Additionally, we present results for CAPM, FF3, and FF4 residual returns for the *LSK* factor in Figures IA.14 to IA.16, for the *USK* factor in Figures IA.17 to IA.19, and for the *SK* factor in Figures IA.20 to IA.22.

¹⁶ In the Internet Appendix, we also show that the portfolio composition of beta- and volatility-sorted portfolios is similar. Table IA.IX reports rank correlations of portfolio deciles and standard deviations of firms' portfolio decile allocations across the beta- and volatility-sorted variables.

Table V

Skewness as a Common Determinant of Low-Risk Anomalies (LRAs)

This table presents evidence that skewness is a common determinant of beta- and volatility-based LRAs. In Panel A, we compute LRA returns as low-minus-high returns of value-weighted decile portfolios that we sort by their CAPM beta, idiosyncratic volatility (measured from the residual variance of CAPM and Fama-French three-factor model regressions), or equity option-implied ex ante variance at the end of each month. We compute the LRAs' excess returns as well as alphas of CAPM, Fama-French three-factor (FF3), and Fama-French-Carhart four-factor (FF4) regressions and the panel presents results for principal component analyses of the corresponding residual returns. We show that the LRAs' PC1 is related to skewness by reporting results from regressing PC1 of excess returns as well as CAPM, FF3, and FF4 residual returns on skew factor (residual) returns. We construct four skew factors from value-weighted decile portfolios sorted by firms' option-implied lower skewness, upper skewness, or total skewness. First, we compute the high-minus-low returns of portfolios sorted by lower skewness (*LSK* factor). Second, we compute the low-minus-high returns of portfolios sorted by upper skewness (*USK* factor). Third, we compute the *SK* factor from portfolios sorted by skewness by going short the extreme portfolios (P1 and P10) and going long the middle portfolios (P5 and P6). Finally, we construct the *LUSK* factor as the sum of *LSK* and *USK*. We report results for regressions using these skew factors and additionally for a regression in which we include both *USK* and *LSK*. Values in square brackets are *t*-statistics based on standard errors following Newey and West (1987), where we choose the optimal truncation lag as suggested by Andrews (1991). In Panel B, we repeat this exercise for the first PC obtained from principal component analysis based on the returns of all 40 portfolios sorted by beta and volatility (rather than the four low-minus-high returns). The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations.

	Excess Returns	CAPM Residuals	FF3 Residuals	FF4 Residuals
Panel A: Skew Factor Returns and the First Principal Component of Low-Risk Anomalies				
LSK	1.74 [44.09]	1.65 [37.57]	1.61 [23.51]	1.48 [16.55]
<i>R</i> ² (%)	94.66	90.76	86.61	77.77
USK	1.65 [42.57]	1.52 [38.81]	1.45 [24.78]	1.32 [18.25]
<i>R</i> ² (%)	93.13	89.35	85.18	75.35
SK	1.24 [16.87]	1.06 [15.66]	0.96 [22.71]	0.89 [12.36]
<i>R</i> ² (%)	82.82	71.92	71.40	51.83
LUSK	0.86 [47.74]	0.81 [44.18]	0.78 [28.08]	0.73 [21.28]
<i>R</i> ² (%)	94.98	91.85	88.25	79.96
LSK	1.18 [8.65]	1.00 [7.45]	0.95 [6.31]	0.90 [6.58]
USK	0.55 [4.27]	0.63 [5.55]	0.64 [5.19]	0.58 [4.98]
<i>R</i> ² (%)	95.13	91.95	88.34	80.14

(Continued)

Table V—Continued

	Excess Returns	CAPM Residuals	FF3 Residuals	FF4 Residuals
Panel B: Skew Factor Returns and the First Principal Component of Portfolios Sorted by Beta and Volatility				
LSK	3.08	1.79	1.81	1.48
	[12.98]	[18.73]	[32.72]	[16.98]
R^2 (%)	70.89	84.56	77.93	67.46
USK	2.88	1.64	1.63	1.31
	[13.45]	[16.86]	[24.23]	[15.05]
R^2 (%)	67.78	82.90	76.39	63.70
SK	2.28	1.23	1.17	0.98
	[15.12]	[12.21]	[19.08]	[12.09]
R^2 (%)	66.86	76.40	75.09	54.15
LUSK	1.51	0.87	0.88	0.73
	[13.43]	[18.62]	[37.19]	[18.49]
R^2 (%)	70.11	85.39	79.27	68.45
LSK	3.00	1.13	1.09	1.00
	[4.90]	[4.81]	[4.92]	[4.91]
USK	0.07	0.64	0.69	0.48
	[0.13]	[2.97]	[3.48]	[2.36]
R^2 (%)	70.89	85.53	79.38	68.86

factors, these results corroborate our interpretation that investors require positive LRA alphas as compensation for coskewness risk.

IV. Discussion and Additional Results

In this section, we discuss two additional robustness tests as well as additional analyses that lend further support to our main conclusions. We also discuss potential extensions represent interesting directions for future research.

Robustness checks. In the analyses above, we refer to several robustness tests reported in the Internet Appendix. For instance, we replicate the value-weighted portfolio analysis using equal- or rank-weighted portfolios and we employ alternative skewness factors. We find that our main results continue to hold. In additional robustness checks, we further show that our empirical results are not driven by particular subsample periods (see Internet Appendix Section IA.I) and that skew adjustments also render the alpha of the BaB factor by Frazzini and Pedersen (2014) insignificant (see Internet Appendix Section IA.II).

Default risk as a driver of (co)skewness risk and LRA alphas. The construction of our skewness factors is guided by the model results in Section I.B, where we rely on corporate default risk as a determinant of firms' (co)skewness risk and LRA alphas. We follow Merton (1974) to model levered firms, which implies

that credit risk and equity option-implied skewness are related (e.g., Geske (1979), Hull, Nelken, and White (2005)). We show that the link between credit risk and (co)skewness in the data is qualitatively the same as in our model.

First, we illustrate the link between skewness and credit risk in our simulated skew-aware world using the Merton-implied credit spreads, discussed in more detail in the Appendix. Panel A of Figure 10 shows that credit spreads increase as absolute lower and upper skewness increase, resulting in a U-shaped relation with total skewness. The patterns are qualitatively the same in

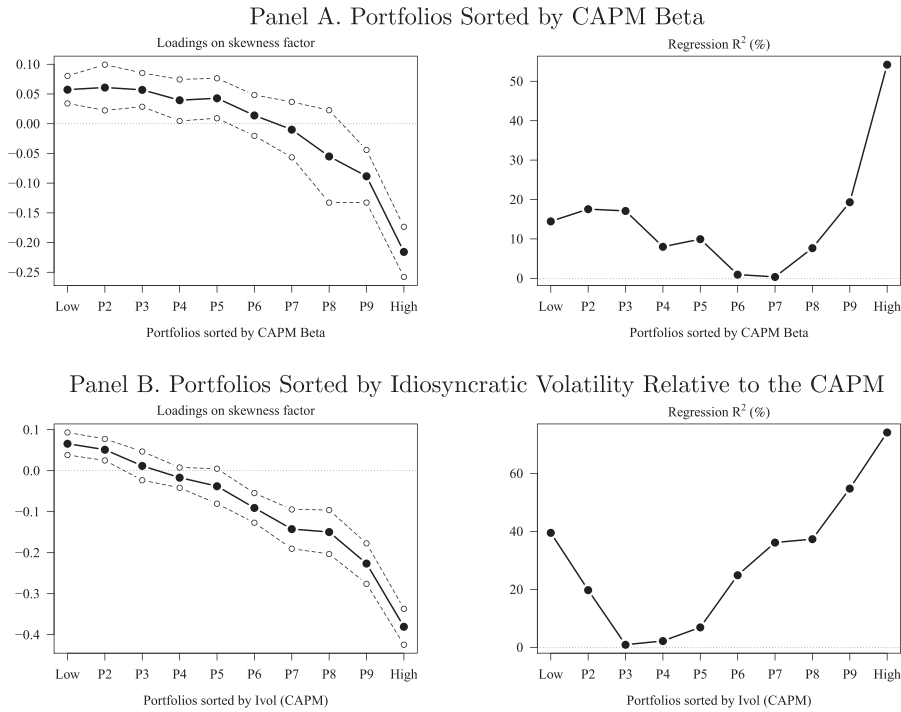


Figure 9. Exposures of beta- and volatility-sorted portfolios to skewness factor. This figure plots *LUSK* skewness factor exposures of portfolios sorted by beta and volatility. We present results for value-weighted decile portfolios that we sort by the CAPM beta, idiosyncratic volatility (measured from the residual variance of CAPM and Fama-French three-factor model regressions), or equity option-implied ex ante variance at the end of each month, that is, in total we have 40 portfolios. We construct the *LUSK* skewness factor as follows. First, we compute the high-minus-low returns of portfolios sorted by lower skewness (*LSK* factor). Second, we compute the low-minus-high returns of portfolios sorted by upper skewness (*USK* factor). The *LUSK* factor is the sum of *LSK* and *USK*. For each of the four beta and volatility sort variables we present results (in Panels A to D, respectively) from regressing the portfolios' CAPM residual returns on the skewness factor's residual returns. On the left, we present the regression slope estimates, along with 95% confidence intervals based on standard errors following Newey and West (1987), where we choose the optimal truncation lag as suggested by Andrews (1991). On the right, we plot the regression R^2 s. The data cover 4,967 U.S. firms, are sampled at a monthly frequency over the period January 1996 to August 2014, and contain a total of 400,449 observations.

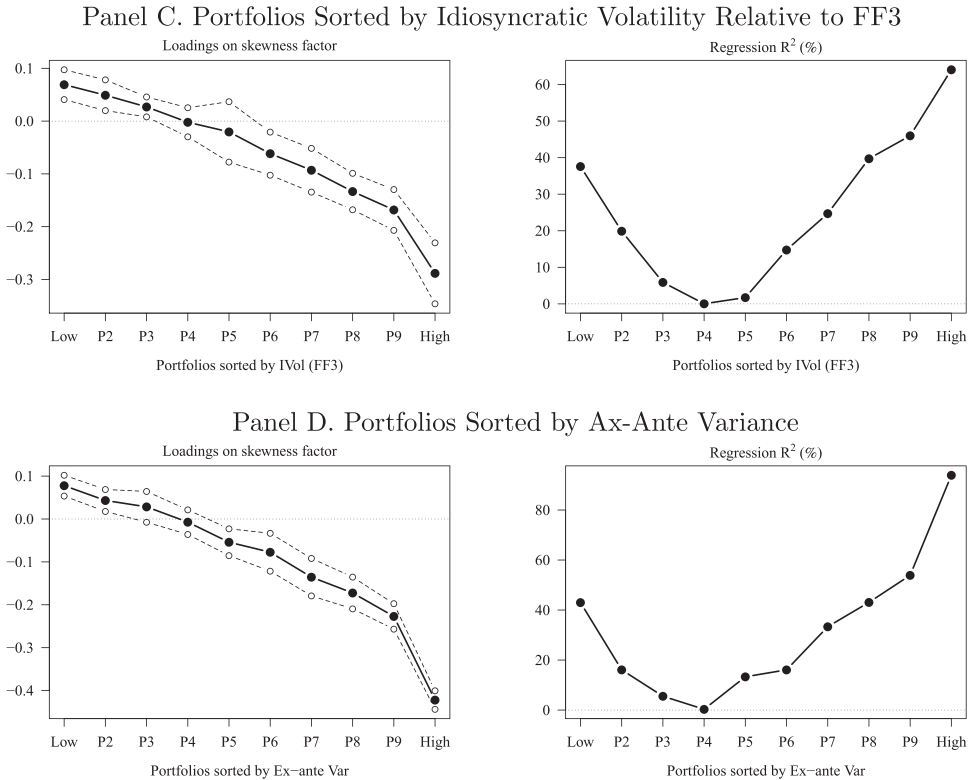


Figure 9. *Continued*

empirical data for CDS spreads (Panel B) and credit ratings (Panel C).¹⁷ Thus, the firms with high absolute values of lower and upper skewness are the firms with the highest CDS spreads and the worst ratings. The closer firms' lower and upper skewness are to zero, the lower are their CDS spreads and the better are their ratings. Accordingly, we find a U-shaped relation between CDS spreads and ratings on the one hand and total skewness on the other. Our model therefore captures not only the relation between skewness and equity returns, but also the link to credit spreads and ratings.

Combining the results for equity risk premia with the results above for credit risk, our paper may also shed light on another LRA, namely, the distress puzzle. Previous research finds that firms with high distress risk underperform

¹⁷ The analyses involving CDS data and credit ratings are conducted on subsamples of our original data set due to data availability. For the analysis of CDS spreads, we use the data set compiled by Friewald, Wagner, and Zechner (2014), which contains Markit CDS data for 573 firms from January 2001 to March 2010 with a total of 37,514 observations. For the analysis of credit ratings, we obtain S&P long-term credit ratings from Compustat when available for firms in our sample. This results in a subsample of 2,066 firms with a total of 179,816 observations over our full sample period.

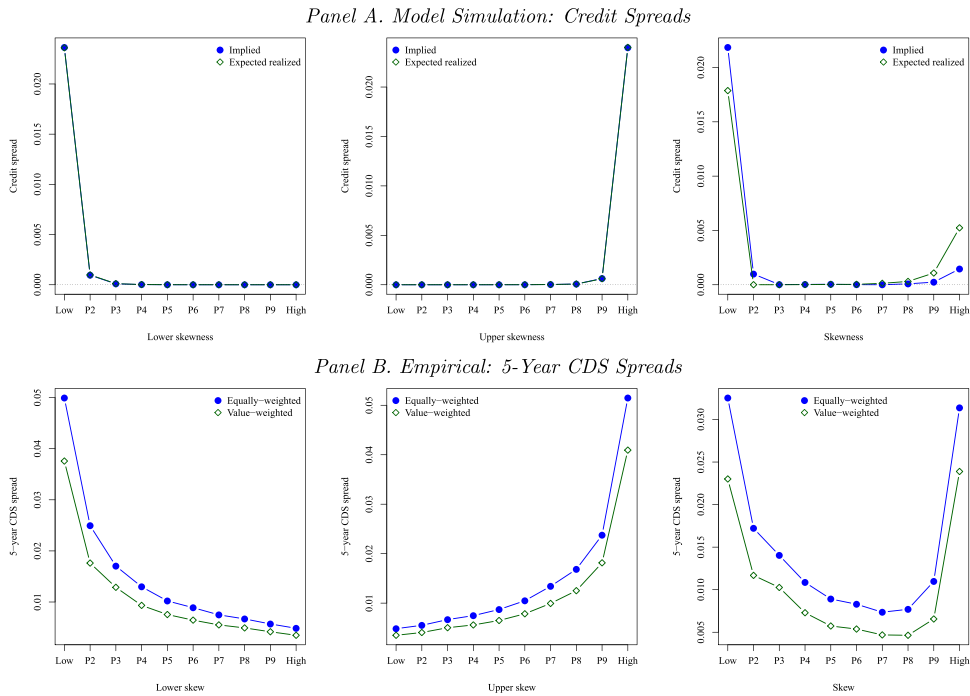


Figure 10. Implied skewness and credit risk. This figure presents results for the relation between firms’ implied skewness and measures of credit risk. Panel A presents simulation evidence from a skew-aware world with 2,000 Merton firms and their model-implied credit spreads. Panels B and C present empirical evidence using option-implied measures of ex ante skewness and CDS spreads or credit ratings respectively. We compute lower skewness (by definition always negative), upper skewness (by definition always positive), and total skewness which is the sum of upper and lower skewness. We sort firms into decile portfolios based on the three option-implied skew measures and compute the portfolios average measure of credit risk, where high (low) values generally imply high (low) credit risk. Portfolios P1 and P10 contain firms with the lowest and highest values of the sort variables, respectively. In the simulated data in Panel A we report results for measures of Q-skew (blue line with bullets) and P-skew (green line with diamonds) using lower skewness (left), upper skewness (middle), and total skewness (right). In the empirical data in Panels B and C, we present results for equal-weighted portfolios (blue line with bullets) and value-weighted portfolios (green line with diamonds). For the analysis of CDS spreads, our data set contains Markit CDS data for 573 firms from January 2001 to March 2010 with a total of 37,514 observations. For the analysis of credit ratings, we obtain S&P long-term credit ratings via Compustat when available for firms in our sample. We convert ratings into numerical data, with lower numbers indicating better ratings, that is, we assign a value of 1 to AAA-rated firms, 2 to firms with a rating of AA+, 3 to firms with a rating of AA, etc. This results in a subsample of 2,066 firms with a total of 179,816 observations from January 1996 to August 2014. (Color figure can be viewed at wileyonlinelibrary.com)




Figure 10. *Continued*

relative to firms with low distress risk. For instance, Campbell, Hilscher, and Szilagyi (2008) show that distressed stocks have low returns, high loadings on risk factors, and negative alphas. Within our framework, these results appear to be consistent with skew-aware asset pricing. Our findings suggest that firms with high credit spreads (bad ratings) are firms with high residual coskewness and hence should earn negative alphas.

Conditional SDF estimation. The goal of our paper is to show that the alphas of LRAs are related to coskewness risk. Since the papers that establish these anomalies focus on *unconditional* alphas, we follow this focus in our empirical analysis as well.¹⁸ In Internet Appendix Section IA.III, we test whether our theory taken to a *conditional* asset pricing model can generate unconditional patterns similar to those we observed in the data. To do so, we use the framework of Nagel and Singleton (2011), who derive an optimal estimator for an SDF that is conditionally affine in a set of priced risk factors and show how to test a null model against a more general SDF. The empirical results of our skew-aware SDF estimation suggest that low- (high-) beta stocks earn positive (negative) conditional alphas that are associated with negative (positive) conditional residual coskewness, in line with our theory. Overall, the conditional skew-aware SDF prices assets significantly better than the conditional CAPM.

Potential extensions. In summary, our results show that coskewness risk can explain the alphas of LRAs. The extent to which exposure to coskewness risk can also explain other cross-sectional asset pricing patterns depends on whether these are driven by coskewness risk in the first place, which ultimately is an empirical question. As a first step in this direction, Internet Appendix Sections IA.III and IA.IV provide conditional and

¹⁸ Few papers study LRAs in a conditional setup. Those papers that do use a conditional setup provide mixed evidence (see, for example, Cederburg and O'Doherty (2016) and Liu, Stambaugh, and Yuan (2018)), due in part to differences in their variable definitions relative to papers that establish the LRAs.

unconditional evidence for portfolios sorted by size and book-to-market. Our findings suggest that coskewness risk may account for (some of) the difference in alphas of small compared to big firms. Coskewness risk does not appear to matter much, however, for portfolios sorted by book-to-market, and the value premium appears unrelated to our skew factors. Since we have shown that our LRA results are robust to controlling for the HML factor, this evidence suggests that LRAs and the value premium are driven by different state variables.

Conceptually, extending our framework to include additional state variables, for example, to start from a Fama-French pricing kernel rather than from the CAPM kernel, should be straightforward. We have not worked this out in our theory section above, because the goal of this paper is to understand LRAs and, empirically, starting from the CAPM and accounting for coskewness is enough to render LRA alphas insignificant. In other words, for our purpose it is not necessary to add more state variables to the SDF, because we find that controlling for size, value, and momentum does not change our results qualitatively. Similarly, we have not extended our SDF to account for (co)moments beyond (co)skewness. Nonetheless, extending our work to alternative SDF specifications and applying it to shed light on other cross-sectional and possibly also time-series patterns in asset prices is likely to be a fruitful avenue for future research. It would be particularly interesting to study whether accounting for coskewness can contribute to recent efforts to organize the factor “zoo” (e.g., Cochrane (2011), Freyberger, Neuhierl, and Weber (2020), Feng, Giglio, and Xiu (2020)).

V. Conclusion

In this paper, we provide a novel perspective on beta- and volatility-based LRAs established in previous research. We show that these apparently anomalous empirical patterns do not necessarily pose asset pricing puzzles after taking (co)skewness of equity returns into account.

Our theory starts from an SDF that is quadratic in the market excess return, and that implies that investors demand compensation for negative coskewness. We show that CAPM pricing errors exhibit residual coskewness and that this residual coskewness is directly linked to a stock’s CAPM beta. CAPM alphas therefore reflect compensation for residual coskewness risk. To study the cross-sectional implications for LRAs, we use a Merton-type model of firm risk and show that the difference in alphas between low- and high-beta/volatility stocks is commensurate with the difference in their residual coskewness.

Our empirical results confirm the model predictions for BaB and betting-against-volatility. We document that these strategies’ positive alphas within the CAPM and other FMs are indeed associated with negative residual coskewness. After we control for skewness, the LRAs disappear: The strategies no longer deliver positive alphas and their residual coskewness becomes substantially less negative.

Initial submission: February 27, 2016; Accepted: March 4, 2019
 Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

Appendix: Simulation Study

This appendix presents details on the simulation study for which we report results in Section I.B of the paper. In particular, we describe (i) how we model the market, (ii) how we model firms, and (iii) the general setup of the simulation analysis.

A. Market Model

To assess the extent to which higher moments such as skewness matter for asset pricing, consider a representative power-utility investor who is exposed to stochastic volatility. We model the dynamics of the forward market price $M_{t,T}$, contracted at time t for delivery at T , as¹⁹

$$\begin{aligned}\frac{dM_{t,T}}{M_{t,T}} &= \eta_t dt + \kappa_t (\xi dW_t^{1\mathbb{P}} + \sqrt{1 - \xi^2} dW_t^{2\mathbb{P}}), \\ d\kappa_t^2 &= (\nu_0 + \nu_1 \kappa_t^2) dt + \kappa_t \vartheta dW_t^{1\mathbb{P}}.\end{aligned}\tag{A1}$$

With γ denoting the coefficient of constant relative risk aversion, $\eta_t = \gamma \kappa_t^2$ is the instantaneous market return in excess of the risk-free rate and κ_t is the associated market volatility. Campbell et al. (2018) develop an empirically successful asset pricing model with stochastic volatility in a similar way.²⁰ We define the discrete forward market excess return $R_{t,T} := \frac{M_{T,T}}{M_{t,T}} - 1$, where we suppress time subscripts hereafter for notational convenience, and we set $M_{0,T} = 1$. Given the agent's local risk aversion γ , we obtain the forward pricing kernel as

$$\mathcal{M}_{0,T} := \frac{(R_{0,T} + 1)^{-\gamma}}{e^{1/2(\gamma - \gamma^2) \int_0^T \kappa_s^2 ds}}.\tag{A2}$$

This kernel is subject to stochastic volatility, such that the kernel is not measurable with respect to the market alone. Stochastic volatility greatly impacts the signs and magnitudes of the coefficients in the projections of the pricing kernel on the market.

We next introduce a crosssection of firms into the economy that exhibits skewness in returns through both stochastic volatility and default risk.

¹⁹ We choose to specify the dynamics of the forward price (rather than the spot price) because this naturally accounts for interest rates and ensures that the forward price is a martingale under the forward measure (\mathbb{Q}_T) with the T -period zero-coupon bond as numeraire. With zero interest rates, the forward price is equal to the spot price and $\mathbb{Q} = \mathbb{Q}_T$.

²⁰ The less realistic but more parsimonious case of modeling the market by a geometric Brownian motion leads qualitatively to the same asset pricing implications as the stochastic volatility dynamics in equation (A1). In other words, higher moments of the return distribution matter for asset prices even if the market does not exhibit skewness. This point is also stressed by Kraus and Litzenberger (1976).

B. Model for Levered Firms

Previous research shows that skewness of stock returns may originate from various sources such as credit risk (Merton (1974)), sentiment (e.g., Han (2008)), demand pressure in option markets (e.g., Gârleanu, Pedersen, and Poteshman (2009)), or differences in beliefs (Buraschi, Trojani, and Vedolin (2014)). Buraschi, Trojani, and Vedolin (2014) also discuss the interaction of disagreement and credit risk. In contrast to the aggregate market, the skewness of individual firms' stock returns is often positive. Recent studies providing evidence on the properties of skewness across firms and in the aggregate market include Albuquerque (2012) and Engle and Mistry (2014).

To parsimoniously model both positive and negative skewness, we specify a firm's asset process A to incorporate jumps and stochastic volatility,

$$\begin{aligned} d \log A_t &= \left(\mu - \frac{\sigma_t^2}{2} \right) dt + \sigma_t \left(\rho dW_t^{\mathbb{P}} + \sqrt{1 - \rho^2} dB_t^{\mathbb{P}} \right) + \eta dJ_t^{\mathbb{P}}, \\ d\sigma_t^2 &= (v_2 + v_3\sigma_t^2 + v_4\kappa_t^2)dt + \psi\sigma_t dB_t^{\mathbb{P}}, \end{aligned} \quad (\text{B1})$$

where J is a pure jump process with intensity ω and η is a constant, $W_t^{\mathbb{P}} = \xi W_t^{1\mathbb{P}} + \sqrt{1 - \xi^2} W_t^{2\mathbb{P}}$, and κ_t^2 is the stochastic market variance from equation (A1). We consider two different default specifications for this setup for a level of debt $D_0 \leq A_0$. The first is a Merton (1974)-style default at maturity T if $A_T < D_0$. The second allows for default prior to maturity if $A_t < D_0$, for at least one $t \in [0, T]$. Equity (E) represents the forward price of a European call option on the firm's assets with strike equal to D_0 . The corresponding forward price $F_{0,T} := \mathbb{E}_0^{\mathbb{P}}[\mathcal{M}_{0,T}(A_T - D_0)^+]$, so that the forward gross return on equity is

$$\frac{(A_T - D_0)^+}{\mathbb{E}_0^{\mathbb{P}}[\mathcal{M}_{0,T}(A_T - D_0)^+]}.$$

The model-implied credit spreads for Merton-style default, that is, when $A_T < D_0$, are given by

$$cs_T := \mathbb{E}_0^{\mathbb{P}} \left[\underbrace{\mathcal{M}_{0,T} \frac{D_0 - A_T}{D_0}}_{\text{loss rate conditional on default}} \underbrace{\mathbb{1}(D_0 \geq A_T)}_{\text{default probability}} \right], \quad (\text{B2})$$

and when we allow the firm to default before maturity, that is, as soon as $A_t < D_0$, by

$$cs_T^{\text{early def.}} := \mathbb{E}_0^{\mathbb{P}} \left[\underbrace{\mathcal{M}_{0,T} \frac{D_0 - A_T}{D_0}}_{\text{loss rate conditional on def.}} \underbrace{\mathbb{1}(D_0 \geq A_t; 0 \leq t \leq T)}_{\text{def. probability}} \right]. \quad (\text{B3})$$

With the asset value dynamics accounting for systematic and idiosyncratic shocks, we explore the impact of higher moments on expected equity returns within the market framework discussed above. In the paper, we report results for the baseline specification without jumps (i.e., we set the jump intensity $\omega = 0$) and for Merton-style default at maturity. In the Internet Appendix, we also present results for specifications that include jumps and early defaults.

C. Simulation Analysis

Our simulation study is designed to generate data that match the properties of our empirical data along several dimensions. In what follows, we sketch the most important steps of this procedure.

To simulate an economy according to the joint model for the market and asset prices from the appendix subsections above, we first generate sets of parameters with plausible values. To model the dynamics of the market, we fix the coefficient of relative risk aversion γ at 2, the instantaneous correlation between forward market returns and stochastic variance ξ to -0.85 , the unconditional mean of index variance $-\nu_0/\nu_1$ to 0.048 , and the mean reversion of market variance ν_1 to -1 .²¹ From these parameters we discretize the stochastic differential equation (A1) and simulate a market time series of 320 months from daily increments.

In a second step, we generate 2,000 firms for which we draw the parameters from distributions reflecting the observed cross section. We draw $\rho \sim \mathcal{U}(0, 1)$, a uniform distribution on the unit interval, leverage $D \sim \mathcal{B}(2, 5)$, a Beta distribution, and asset drift $\mu \sim \Gamma(2, 0.01)$, a Gamma distribution, that is the same as the volatility of asset variance ψ , which is taken as the square root of a $\Gamma(2, 0.01)$ random variable. The parameters ν_2 , ν_3 , and ν_4 are drawn from Gamma distributions $\Gamma(2, 0.01)$, $-\Gamma(2, 0.5)$, and $\Gamma(2, 0.25)$, respectively, so that the unconditional mean $\mathbb{E}^{\mathbb{P}}[\sigma_t^2]$ exists. To better reflect the cross section of U.S. corporations, we additionally set 25% of the population's leverage to zero (see, for example, Strebulaev and Yang (2013)). When simulating the asset value processes, we keep the trajectory of the forward market fixed to ensure it is identical for all assets. Given these sample paths for firm assets, we then compute the sample paths of corporate equity values, expected equity returns, implied and expected realized skewness, CAPM betas, etc.

REFERENCES

- Aït-Sahalia, Yacine, and Robert Kimmel, 2007, Maximum likelihood estimation of stochastic volatility models, *Journal of Financial Economics* 83, 413–452.
- Albuquerque, Rui, 2012, Skewness in stock returns: Reconciling the evidence on firm versus aggregate returns, *Review of Financial Studies* 25, 1630–1673.
- Amaya, Diego, Peter Christoffersen, Kris Jacobs, and Aurelio Vasquez, 2015, Does realized skewness predict the cross-section of equity returns? *Journal of Financial Economics* 118, 135–167.
- Andrews, Donald W.K., 1991, Heteroskedasticity and autocorrelation consistent covariance matrix estimation, *Econometrica* 59, 817–858.

²¹ These values are similar to those in Aït-Sahalia and Kimmel (2007).

- Ang, Andrew, Joseph Chen, and Yuhang Xing, 2006, Downside risk, *Review of Financial Studies* 19, 1191–1239.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2009, High idiosyncratic volatility and low returns: International and further U.S. evidence, *Journal of Financial Economics* 91, 1–23.
- Baker, Malcolm, Brendan Bradley, and Ryan Taliaferro, 2014, The low-risk anomaly: A decomposition into micro and macro effects, *Financial Analysts Journal* 70, 43–58.
- Baker, Malcolm, Brendan Bradley, and Jeffrey Wurgler, 2011, Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly, *Financial Analysts Journal* 67, 40–54.
- Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.
- Bali, Turan G., Stephen Brown, Scott Murray, and Yi Tang, 2017, A lottery demand-based explanation of the beta anomaly, *Journal of Financial & Quantitative Analysis* 52, 2369–2397.
- Bali, Turan G., Nusret Cakici, and Robert F. Whitelaw, 2011, Maxing out: Stocks as lotteries and the cross-section of expected returns, *Journal of Financial Economics* 99, 427–446.
- Bali, Turan G., and Armen Hovakimian, 2009, Volatility spreads and expected stock returns, *Management Science* 55, 1797–1812.
- Bali, Turan G., Jianfeng Hu, and Scott Murray, 2015, Option implied volatility, skewness, and kurtosis and the cross-section of expected stock returns, Working paper, Georgetown University, Singapore Management University, and Georgia State University.
- Bali, Turan G., and Scott Murray, 2013, Does risk-neutral skewness predict the cross-section of equity option portfolio returns? *Journal of Financial and Quantitative Analysis* 48, 1145–1171.
- Barberis, Nicholas, and Ming Huang, 2008, Stocks as lotteries: The implications of probability weighting for security prices, *American Economic Review* 98, 2066–2100.
- Black, Fischer, 1972, Capital market equilibrium with restricted borrowing, *Journal of Business* 45, 444–455.
- Black, Fisher, Michael C. Jensen, and Myron S. Scholes, 1972, The capital asset pricing model: Some empirical tests, in M.C. Jensen, ed., *Studies in the Theory of Capital Markets* (Praeger, New York).
- Boehme, Rodney D., Bartley R. Danielsen, Praveen Kumar, and Sorin M. Sorescu, 2009, Idiosyncratic risk and the cross-section of stock returns: Merton (1987) meets Miller (1977), *Journal of Financial Markets* 12, 438–468.
- Bloorforoosh, Ali, Peter Christoffersen, Mathieu Fournier, and Christian Gouriéroux, 2017, Beta risk in the cross-section of equities, Working paper, Concordia University, University of Toronto, and HEC Montreal.
- Boyer, Brian, Todd Mitton, and Keith Vorkink, 2010, Expected idiosyncratic skewness, *Review of Financial Studies* 23, 169–202.
- Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621–651.
- Brennan, Michael J., 1971, Capital market equilibrium with divergent borrowing and lending rates, *Journal of Financial and Quantitative Analysis* 6, 1197–1205.
- Brunnermeier, Markus K., Christian Gollier, and Jonathan A. Parker, 2007, Optimal beliefs, asset prices, and the preference for skewed returns, *American Economic Review* 97, 159–165.
- Buraschi, Andrea, Fabio Trojani, and Andrea Vedolin, 2014, Economic uncertainty, disagreement, and credit markets, *Management Science* 60, 1281–1296.
- Buss, Adrian, and Grigory Vilkov, 2012, Measuring equity risk with option-implied correlations, *Review of Financial Studies* 25, 3113–3140.
- Campbell, John Y., 1993, Intertemporal asset pricing without consumption data, *American Economic Review* 83, 487–512.
- Campbell, John Y., Stefano Giglio, Christopher Polk, and Robert Turley, 2018, An intertemporal CAPM with stochastic volatility, *Journal of Financial Economics* 128, 207–233.

- Campbell, John Y., Jens Hilscher, and Jan Szilagyi, 2008, In search of distress risk, *Journal of Finance* 63, 2899–2939.
- Carhart, Mark, 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.
- Carr, Peter, and Roger Lee, 2009, Put-call symmetry: Extensions and applications, *Mathematical Finance* 19, 523–560.
- Cederburg, Scott, and Michael S. O'Doherty, 2016, Does it pay to bet against beta? On the conditional performance of the beta anomaly, *Journal of Finance* 71, 737–774.
- Chang, Bo Young, Peter Christoffersen, and Kris Jacobs, 2013, Market skewness risk and the cross section of stock returns, *Journal of Financial Economics* 107, 46–68.
- Christoffersen, Peter, Mathieu Fournier, Kris Jacobs, and Mehdi Karoui, 2019, Option-based estimation of the price of co-skewness and co-kurtosis risk, *Journal of Financial and Quantitative Analysis* (forthcoming) 1–51. <https://doi.org/10.1017/S002210902000023X>.
- Cochrane, John H., 2011, Presidential address: Discount rates, *Journal of finance* 66, 1047–1108.
- Conrad, Jennifer, Robert F. Dittmar, and Eric Ghysels, 2013, Ex ante skewness and expected stock returns, *Journal of Finance* 68, 85–124.
- Engle, Robert, and Abhishek Mishra, 2014, Priced risk and asymmetric volatility in the cross section of skewness, *Journal of Econometrics* 182, 135–144.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.
- Feng, Guanhao, Stefano Giglio, and Dacheng Xiu, 2020, Taming the factor zoo: A test of new factors, *Journal of Finance* 75, 1327–1370.
- Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.
- Freyberger, Joachim, Andreas Neuhierl, and Michael Weber, 2020, Dissecting characteristics non-parametrically, *Review of Financial Studies* 33, 2326–2377. <https://doi.org/10.1093/rfs/hhz123>
- Friend, Irwin, and Randolph Westerfield, 1980, Co-skewness and capital asset pricing, *Journal of Finance* 35, 897–913.
- Friedwald, Nils, Christian Wagner, and Josef Zechner, 2014, The cross-section of credit risk premia and equity returns, *Journal of Finance* 69, 2419–2469.
- Fu, Fangjian, 2009, Idiosyncratic risk and the cross-section of expected stock returns, *Journal of Financial Economics* 91, 24–37.
- Gârleanu, Nicolae, Lasse H. Pedersen, and Allen M. Poteshman, 2009, Demand-based option pricing, *Review of Financial Studies* 22, 4259–4299.
- Geske, Robert, 1979, The valuation of compound options, *Journal of Financial Economics* 7, 63–81.
- Gourier, Elise, 2016, Pricing of idiosyncratic equity and variance risks, Working paper, Queen Mary University London.
- Han, Bing, 2008, Investor sentiment and option prices, *Review of Financial Studies* 21, 387–414.
- Harvey, Campbell R., and Akhtar Siddique, 2000, Conditional skewness in asset pricing tests, *Journal of Finance* 55, 1263–1295.
- Haugen, Robert A., and A. James Heins, 1975, Risk and the rate of return on financial assets: Some old wine in new bottles, *Journal of Financial and Quantitative Analysis* 10, 775–784.
- Hong, Harrison, and David A. Sraer, 2016, Speculative betas, *Journal of Finance* 71, 2095–2144.
- Huang, Shiyang, Dong Lou, and Christopher Polk, 2016, The booms and busts of beta arbitrage, Working paper, University of Hong Kong and London School of Economics.
- Hull, John C., Izzy Nelken, and Alan D. White, 2005, Merton's model, credit risk and volatility skews, *Journal of Credit Risk* 1, 3–26.
- Johnson, Timothy C., 2004, Forecast dispersion and the cross section of expected returns, *Journal of Finance* 59, 1957–1978.
- Jylhä, Petri, 2018, Margin constraints and the security market line, *Journal of Finance* 73, 1281–1321.
- Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2018, Interpreting factor models, *Journal of Finance* 73, 1183–1223.

- Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2020, Shrinking the cross-section, *Journal of Financial Economics* 135, 271–292.
- Kozhan, Roman, Anthony Neuberger, and Paul Schneider, 2013, The skew risk premium in the equity index market, *Review of Financial Studies* 26, 2174–2203.
- Kraus, Alan, and Robert H. Litzenberger, 1976, Skewness preference and the valuation of risk assets, *Journal of Finance* 31, 1085–1100.
- Lettau, Martin, Matteo Maggiori, and Michael Weber, 2014, Conditional risk premia in currency markets and other asset classes, *Journal of Financial Economics* 114, 197–225.
- Lintner, John, 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* 13–37.
- Liu, Jianan, Robert F. Stambaugh, and Yu Yuan, 2018, Absolving beta of volatility's effects, *Journal of Financial Economics* 128, 1–15.
- Martin, Ian, 2013, Simple variance swaps, Working paper, Stanford University.
- Merton, Robert, 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449–470.
- Mittton, Todd, and Keith Vorkink, 2007, Equilibrium underdiversification and the preference for skewness, *Review of Financial Studies* 20, 1255–1288.
- Mossin, Jan, 1966, Equilibrium in a capital asset market, *Econometrica* 34, 768–783.
- Nagel, Stefan, and Kenneth J. Singleton, 2011, Estimation and evaluation of conditional asset pricing models, *Journal of Finance* 66, 873–909.
- Neuberger, Anthony, 1994, The log contract, *Journal of Portfolio Management* 20, 74–80.
- Newey, Whitney, and Kenneth West, 1987, A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Novy-Marx, Robert, 2016, Understanding defensive equity, Working paper, University of Rochester.
- Pástor, Luboš, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.
- Rehman, Zahid, and Grigory Vilkov, 2012, Risk-neutral skewness: Return predictability and its sources, Working paper, Goethe University Frankfurt.
- Rubinstein, Mark E., 1973, The fundamental theorem of parameter-preference security valuation, *Journal of Financial and Quantitative Analysis* 8, 61–69.
- Schneider, Paul, and Fabio Trojani, 2014, Fear trading, Working paper, University of Lugano and SFI.
- Stephen Sears, R., and K. C. John Wei, 1985, Asset pricing, higher moments, and the market risk premium: A note, *Journal of Finance* 40, 1251–1253.
- Sharpe, William F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425–442.
- Stambaugh, Robert F., Jianfeng Yu, and Yu Yuan, 2015, Arbitrage asymmetry and the idiosyncratic volatility puzzle, *Journal of Finance* 70, 1903–1948.
- Strebulaev, Ilya A., and Baozhong Yang, 2013, The mystery of zero-leverage firms, *Journal of Financial Economics* 109, 1–23.
- Walkshäusl, Christian, 2014, International low-risk investing, *Journal of Portfolio Management* 41, 45–56.
- White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817–838.
- Xing, Yuhang, Xiaoyan Zhang, and Rui Zhao, 2010, What does the individual option volatility smirk tell us about future equity returns? *Journal of Financial and Quantitative Analysis* 45, 641–662.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.

Replication code.