BANK FOR INTERNATIONAL SETTLEMENTS

BIS Working Papers

No 921

Firm-specific risk-neutral distributions with options and CDS

by Sirio Aramonte, Mohammad R Jahan-Parvar, Samuel Rosen and John W Schindler

Monetary and Economic Department

January 2021

JEL classification: G12, G13, G14.

Keywords: risk neutral distributions, investor expectations, CDS spreads.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Firm-Specific Risk-Neutral Distributions with Options and CDS*

Sirio Aramonte Mohammad R. Jahan-Parvar
Samuel Rosen John W. Schindler†

December 15, 2020

Abstract

We propose a method to extract the risk-neutral distribution of firm-specific stock returns using both options and credit default swaps (CDS). Options and CDS provide information about the central part and the left tail of the distribution, respectively. Taken together, but not in isolation, options and CDS span the intermediate part of the distribution, which is driven by exposure to the risk of large but not extreme returns. Through a series of asset-pricing tests, we show that this intermediate-return risk carries a premium, particularly at times of heightened market stress.

Keywords: risk neutral distributions; investor expectations; CDS spreads

JEL classification: G12; G13; G14

*This paper was previously circulated under the title "Firm-Specific Risk-Neutral Distributions: The Role of CDS Spreads"

†Aramonte is at the Bank for International Settlements; Jahan-Parvar and Schindler are at the Federal Reserve Board. Rosen is at the Fox School of Business at Temple University. Contact author e-mail and phone number: sirio.aramonte@bis.org, +41 (0) 61 280 8561. We would like to thank Gurdip Bakshi, Jack Bao, Daniel Beltran, Yasser Boualam, John Calarco, Stijn Claessens, Jennifer Conrad, Max Croce, Jesse Davis, Hitesh Doshi, Art Duquette, Bruno Feunou, A. Ronald Gallant, Michael Gordy, Patrick Herb, Christian Skov Jensen, Pawel Szerszen, Andreas Schrimpf, Claudio Tebaldi, and seminar participants at the Bank for International Settlements, Bocconi University, Federal Reserve Board, Kenan-Flagler Business School, West Virginia University, University of Warsaw School of Management, CFE 2016 (Seville), IFABS 2016 (Barcelona), JSM 2016 (Chicago), FMA 2017 (Boston), Georgetown Center for Economic Research Biannual Conference (2017), JSM 2017 (Baltimore), NFA 2017 (Halifax), North American Summer Meeting of the Econometric Society 2017 (St. Louis), and SNDE 2017 (Paris). This article represents the views of the authors, and not necessarily those of the Bank for International Settlements or of the Board of Governors of the Federal Reserve System, or other members of their staff.

1 Introduction

Option-implied risk-neutral distributions of stock returns (RNDs) are a popular instrument for gauging investor expectations and risk preferences. When recovering an RND, one would ideally use data that are informative about the entire distribution. For individual companies, however, equity options trade actively only at strikes around the current price. As a result, they do not provide much information about the tails of the RND, which are particularly interesting for researchers and practitioners as they characterize the risk of large negative returns. Existing methods for extracting RNDs from individual-stock options (such as Bakshi, Kapadia, and Madan, 2003, henceforth BKM) rely on assumptions rather than data to address this issue.

Our main contribution is a new method to extract firm-specific RNDs that incorporate information from credit default swaps (CDS) in addition to information from option prices. CDS are derivatives that pay off in case of default, hence their prices reflect tail risk. Since the early 2000s, the market for CDS has developed quickly. Our approach has two distinct advantages, as illustrated by Figure 1. First, estimated tail risk is driven by actual market prices because information from CDS spreads helps pin down the far left tail of the distribution. Second, options and CDS together span the intermediate part of the distribution, which is driven by exposure to the risk of large but not extreme returns. Considered individually, neither options nor CDS are informative about this intermediate part of the RND, which we document is directly related to skewness. Therefore, we can study the risk of large but not extreme negative returns by focusing on the difference between the skewness based on CDS and options, which captures intermediate risk, and the skewness based only on options, which does not.

We blend together the information from CDS and options using the skew-t distribution, which is widely used in empirical finance research (e.g., Hansen, 1994, Jondeau and Rockinger, 2003, Patton, 2004 and Oh and Patton, 2017). Together with the inclusion of CDS, the flexibility of the skew-t is a crucial element of our method. In particular, blending CDS and traded options using a versatile distribution is instrumental to extracting information that

Figure 1: The contribution of options and CDS to risk-neutral return distributions

The figure provides a stylized illustration of the contribution that options and CDS make to the risk-neutral distribution of returns. The red markers correspond to traded option strikes, which are normally clustered close to the current stock price. The CDS-implied default probability determines the cumulative density of the distribution up to the default threshold, which is the negative equity return at which a company would default. Our method, including the estimation of the default threshold, is described in detail in Section 3.

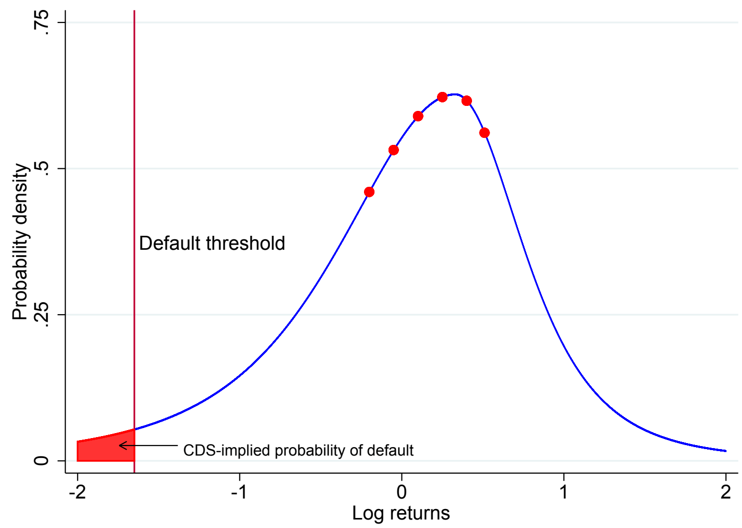

cannot be obtained from options with strikes far below at-the-money because they seldom trade. In this sense, our method amounts to more than simply adding one data point (the CDS) to existing option prices. Instead, the combination of CDS *plus* suitable parametric assumptions allows us to recover information equivalent that what could be implied from many additional options with relatively low strike prices, if only these options traded in practice.

After detailing the technical aspects of our method, the paper studies the asset-pricing characteristics of a portfolio that buys and sells stocks on the basis of differences between options/CDS skewness and options-only skewness. This portfolio earns a 0.5% monthly risk premium, which is robust to different assumptions about key inputs such as CDS tenor and default threshold. This premium is particularly elevated at times of financial distress, since we make efficient use of multiple sources of data and our RNDs are less affected by the price volatility that characterizes such periods. As a whole, these findings show that combining options and CDS yields a more informative skewness measure. The tests we conduct are ancillary to and informed by our methodological contribution. Thus, ours is not a "factor study" that simply proposes a new pricing factor.

Relative to traditional approaches based only on options, the trade-off from using CDS is that the number of firms we can study is reduced because of the limited availability of CDS data. Given that we extract RNDs for each stock and on each day independently, our sample of 275 U.S. companies is considerably larger than the potential set of firms for which one could estimate a structural model based on both options and CDS. For example, Carr and Wu (2010) develop a framework for the joint valuation of options and CDS, where the default rate is affected by stock volatility. They explicitly model the stock and default dynamics, which entails estimating a relatively large set of parameters. As a result, they apply their method to only eight large companies.

Our work is closely related to a variety of other studies of individual-stock risk-neutral distributions. These papers most often build on the popular method of Bakshi and Madan (2000) and Bakshi, Kapadia, and Madan (2003) to calculate higher-order risk-neutral moments from option prices. Examples include Dennis and Mayhew (2002), Rehman and Vilkov (2012), Bali and Murray (2013), Conrad, Dittmar, and Ghysels (2013), DeMiguel, Plyakha, Uppal, and Vilkov (2013), and Stilger, Kostakis, and Poon (2017).[1] Some implementations rely on a potentially small number of options (e.g., Dennis and Mayhew, 2002, among others) or on interpolated option prices (e.g., An, Ang, Bali, and Cakici, 2014). Our empirical analysis, which revolves around the cross-sectional pricing of risk-neutral higher moments, is close to Stilger, Kostakis, and Poon (2017), Conrad, Dittmar, and Ghysels (2013) and Rehman and Vilkov (2012). The key difference, however, is that these papers are focused on the direct contribution of risk-neutral skewness to stock returns, while we are interested in what information is unlocked by including CDS in the calculation of skewness.

[1] Studies of RNDs from single-name options are preceded by the voluminous literature on extracting RNDs from equity index options. A partial list of the wide range of available parametric and non-parametric methods includes Bates (1991), Madan and Milne (1994), Rubinstein (1994), Longstaff (1995), Jackwerth and Rubinstein (1996), Aït-Sahalia and Lo (1998), Bates (2000). Empirical applications underscore the importance of risk-neutral higher-order moments for asset pricing, for instance by computing risk aversion as the wedge between option-implied distributions and statistical estimates of expected returns (Rosenberg and Engle, 2002, Bliss and Panigirtzoglou, 2004, Ross, 2015, and Jensen, Lando, and Pedersen, 2017, among others).

Research on risk-neutral skewness follows earlier work on skewness based on historical returns. Harvey and Siddique (2000) highlight that systematic skewness helps explain the cross-section of returns. Accounting for skewness is also important to identify the sign of the risk-return relation (Feunou, Jahan-Parvar, and Tédongap, 2013). Amaya, Christoffersen, Jacobs, and Vasquez (2015) find that realized skewness generates cross-sectional predictability in stock returns. Recently, Colacito, Ghysels, Meng, and Siwasarit (2016) investigate the effect of skewness in firm-level and macroeconomic fundamentals on stock returns. As documented by Kim and White (2004), measuring historical higher moments is difficult, and researchers have adopted a variety of strategies to tackle this issue. For instance, Neuberger (2012) develops a realized estimator for skewness based on high-frequency data, Kelly and Jiang (2014) exploit information in the cross-section of company-specific price jumps, and Ghysels, Plazzi, and Valkanov (2016) use quantile-based measures of skewness to overcome data constraints in emerging markets. Feunou, Jahan-Parvar, and Tédongap (2016) examine alternative parametric structures for skewness models.

The rest of the paper is organized as follows. Section 2 describes the data. Section 3 describes the method for extracting the options/CDS-implied risk-neutral distributions. Section 4 discusses our empirical investigation and findings, and Section 5 concludes.

2 Data

Our sample includes option prices, CDS spreads, and company stock returns from January 2006 to December 2015. We focus on this period due to CDS data constraints. While CDS spreads are available starting in 2001, the population of firms in the data grows rapidly until 2006.

Option and interest rate data are from OptionMetrics through Wharton Research Data Services (WRDS). We collect American options (with an "A" exercise style flag) written on individual common stocks (CRSP share codes 10 and 11) that trade on AMEX, NASDAQ, or NYSE (CRSP exchange codes 1, 2, and 3). As is customary with option data, we apply a series of filters to discard thinly-traded options and likely data errors. We keep observations with positive volume, positive bid and ask prices, and an ask price higher than the bid price. Following Santa-Clara and Saretto (2009), we drop options with a bid-ask spread smaller than the minimum tick (0.05 if the ask is less than 3, and 0.1 if the bid is more than or equal to 3). Finally, we discard options with missing values for implied volatility.

CDS data are from Markit. They include the term structure of CDS spreads between 6 months and 30 years, in addition to recovery rates and restructuring clauses. Moreover, Markit provides information on the reference obligation, including seniority and country of domicile for the issuer. We focus on U.S. Dollar-denominated CDS contracts on senior unsecured obligations issued by U.S.-based entities and on spreads pertaining to contracts with an XR restructuring clause. Restructuring clauses determine what credit events trigger the payout of the CDS, and the XR clause excludes all debt restructurings and focuses on defaults.

We obtain stock returns from the Center for Research on Security Prices (CRSP) and balance-sheet items through Compustat. We manually match companies in Markit and CRSP by name, and we merge Compustat and OptionMetrics using the the *lpermno* and *cusip* variables.

On a given day, we select the cross section of options with maturity closest to 90 calendar days, as long as the maturity is between 15 and 180 days. In all cases, implied volatilities are compounded to 90-day. We require that the CDS spread and at least five option observations are available for each company/day combination. Table 1 shows selected summary statistics for the resulting 275 companies as of 2006. These companies are large, with median book assets equal to about \$13 billion. The average of book assets is about \$63 billion, which indicates the presence of several very large companies. The remaining summary statistics are financial and balance sheet ratios showing that there is substantial heterogeneity in our sample along several dimensions, including cash flows, sales, investment, research and development expenses, and stock and debt issuance. Options on individual firms are generally less traded than those on the S&P 500 (e.g., Pan and Poteshman, 2006), and this observation holds in our sample as well (see Appendix A for details).

The default threshold that we use to blend option- and CDS-implied probabilities is

Table 1: Selected balance-sheet characteristics

The table presents descriptive summary statistics for the 275 companies for which we estimate the options/CDS risk-neutral return distributions. Book assets are reported in $ million. Ratios are calculated with respect to book assets at the end of the previous year. The data are as of the end of fiscal year 2006.

	Average	75^{th} perc.	50^{th} perc.	25^{th} perc.
Book assets	62,899	31,777	12,864	6,590
Sales ratio	114.59	143.85	95.42	53.56
Cash flow ratio	16.45	21.81	15.97	9.91
Investment ratio	6.38	8.03	4.51	2.27
R&D ratio	4.00	4.87	1.75	0.15
Stock issuance ratio	-3.51	0.03	-1.88	-5.71
Debt issuance ratio	3.38	3.66	0.00	-0.99

computed using a set of bankruptcies compiled from Capital IQ. We consider Chapter 7 and Chapter 11 filings (event code 89) between 1990 and 2015 for companies with common stock trading on large exchanges, with market capitalization in excess of 100 million one year before the filing, and with a delisting payment within 30 days of bankruptcy. When a company defaults multiple times, only the first instance is included unless the bankruptcies are at least five years apart. The final number of company-bankruptcy observations in our sample is 112. This figure is smaller than the typical sample in the relevant literature because we need to ensure that we calculate the thresholds using firms comparable to those with traded options and CDS. For instance, Subrahmanyam, Tang, and Wang (2014) consider 1,628 filings between 1997 and 2009. We start with 1,827 filings and we are left with 367 firms simply by focusing on companies with common stock trading on AMEX, NASDAQ or NYSE.

3 Options/CDS-implied return distributions

Our approach to estimating RNDs combines three ingredients: option prices, CDS spreads, and a default threshold. Option prices provide information about the central part of the distribution, while CDS spreads anchor the left tail. The third ingredient comes into play because the probability of default embedded in the CDS contract is the cumulative density up to the default threshold, which is the transition point between option- and CDS-based information. By definition, a CDS contract with an XR restructuring clause pays upon default, meaning that its "strike price" is the return to default over the horizon of the risk neutral distribution. In terms of horizon, we focus on three-month ahead returns.

The output of our method that combines the three ingredients is a skew-t distribution for expected risk-neutral stock returns. Hansen (1994), Patton (2004), and Jondeau and Rockinger (2003) study financial applications of the skew-t distribution under the physical measure. They find that this distribution fits stock returns well and that it spans a large set of the theoretically admissible values of skewness and kurtosis.

All of the input data that we use are available at a daily frequency, and we obtain a separate RND for each firm-day pair with sufficient data (for a total of 253,790 distributions). Parameters are estimated by minimizing squared deviations between parametric CDFs and cumulative probabilities extracted from options and CDS. In the remainder of this section, we outline the various steps of our method. Details can be found in Appendix B.

3.1 Option-implied return probabilities

We compute option-implied probabilities using established methods that build on the link between risk-neutral probabilities and call-option prices that was highlighted by Breeden and Litzenberger (1978). At a given strike price, the empirical density is the scaled second derivative of the European call price function. We follow the implementation of Figlewski (2010) in terms of both processing the options data and calculating the option-implied cumulative probabilities at the traded strikes. See Appendix B for details.

The empirical density is calculated only at traded strikes, which tend to be close to at-the-money for individual firms. Existing studies often obtain a denser set of option prices by interpolating/extrapolating the cross-section of implied volatilities (as a function of strike prices) and inverting the set of traded and interpolated implied volatilities back to prices. The interpolation is often based on non-parametric techniques, like parabolic functions (Shimko,

1993) or cubic and quartic splines. Crucially, the shape of the tails of the implied distributions depends on the extrapolation method and on the leverage exerted by observations at the extremes of the set of traded strikes. To limit this sensitivity, the literature has proposed to model the tails parametrically (Shimko, 1993, Figlewski, 2010). In our framework, we do not need to interpolate/extrapolate because we use CDS to pin down the left tail of a parametric distribution.

3.2 CDS-implied default probability

We extract risk-neutral default probabilities from CDS spreads using standard formulas that equate expected losses to expected CDS premia (see, for instance, Duffie, 2003), using the appropriately compounded spread on CDS with five years to maturity. The reason for choosing five-year CDS is that trading is concentrated in this tenor (Chen, Fleming, Jackson, Li, and Sarkar, 2011), with the consequence that spreads at other maturities are more likely to have a lower signal-to-noise ratio. Zhang, Zhou, and Zhu (2009) study five-year CDS for similar reasons. In principle, CDS spreads incorporate compensation for counterparty credit risk, which could induce downward bias in the measured default risk. However, Arora, Ghandi, and Longstaff (2012) [pg. 1] find that the effect is "vanishingly small and is consistent with [...] collateralization of swap liabilities by counterparties".

3.3 Combining CDS- and option-implied probabilities

In order for our method to be feasible, stock prices must be positive just before default. If stock prices always approached zero before default, bankruptcy would not span a meaningful set of a firm's equity return space, and the CDS-implied default probability could not be used to anchor the left tail of the risk neutral distribution of returns. There is evidence that debtholders generally have an incentive to strategically force a default before the value of net assets approaches zero, in order to maximize the recovery rate on the debt (see Fan and Sundaresan, 2000, Carey and Gordy, 2016, Garlappi and Yan, 2011, and other references in Carr and Wu, 2011). The post-default value that accrues to shareholders is also increasing in equity-ownership concentration (Alanis, Chava, and Kumar, 2018).

Table 2: Default thresholds

The table shows the default thresholds, in percentage points, that we use when extracting the options/CDS risk-neutral distributions. The thresholds are the average cumulative equity returns over the three-month period leading to bankruptcy calculated after double sorting firms based on size and leverage. Low/high leverage means leverage below/above the sample median. Large firms are those with market capitalization in excess of $1 billion.

	Low leverage	High leverage
Market cap $<$ 1bn	-81%	-70%
Market cap $\geq$ 1bn	-90%	-89%

We define the default threshold (R_D) as the cumulative return over the three-month period prior to bankruptcy. We consider thresholds that vary with two key company characteristics: leverage and market capitalization (henceforth, size). We compute four different thresholds that reflect the double sorting of firms into two size categories and two leverage categories. Table 2 reports these values. The sorting is based on data as of five quarters before the bankruptcy filing. Of the five quarters by which we lag the data, one quarter simply reflects the horizon of the default threshold (three months). The remaining four quarters are meant to limit the effect of the eventual bankruptcy on size and leverage.

3.4 Estimating RND Parameters

In the previous steps, we computed a set of cumulative density values corresponding to the strikes of traded options, a cumulative density value from CDS that quantifies perceived default risk, and a return threshold beyond which a company is expected to default. In the final step, we recover the parameters of a skew-t distribution by minimizing squared deviations between the parametric CDF and the set of cumulative density values mentioned above. Specifically, the returns associated with each option-implied cumulative density value correspond to the returns from the current stock price to the strike price, and the return associated with the CDS-implied cumulative density value is the threshold that corresponds to default (see Appendix B for more details). This procedure is applied to each firm-day pair, for a total of 253,790 sets of distribution parameters.

Figlewski (2010) notes that this approach (selecting a known parametric distribution, solving for the set of parameters that minimize the discrepancy between the fitted function and the empirical density values) is the most common method to model the RND. The key difference in our case is the inclusion of CDS, which characterizes the left tail of the distribution and allows us to assess intermediate risk.

4 Empirical investigation

In this section, we study the performance of our proposed method. We largely focus on skewness because, as shown in detail in Appendix C, it is the risk-neutral moment that is most affected by the inclusion of CDS. Intutively, one can expect CDS to inform higher order moments more than volatility, which is well characterized by options with strike prices around the current stock price. Skewness measures the shape of a distribution, and it is influenced by both the central part and the tails.

In Section 4.1, we discuss the properties of the skew-t moments extracted using, first, options only and, second, options and CDS. We also compare these moments to those based on the non-parametric procedure of Bakshi, Kapadia, and Madan (2003), which is widely used in the literature. In Section 4.2, we assess the economic value of our method with asset pricing tests on portfolios that capitalize on the additional information unlocked by considering options and CDS together rather than options only. In Section 4.3, we interpret our results, discuss potential issues, and provide evidence from robustness tests.

4.1 Evaluating measures of risk-neutral skewness

In Panel A of Table 3, we report summary statistics for skewness measures extracted from options and CDS, from options only, and for the difference between the two. These moments are recovered using our parametric method. In Panel B, we compare the non-parametric skewness of BKM, as implemented by Conrad, Dittmar, and Ghysels (2013), to our options/CDS and options-only skewness. As shown in Panel A, differences between the options/CDS and options-only measures are concentrated in the tails. Panel B highlights that the magnitude

Table 3: Selected summary statistics for skewness and skewness differences

In Panel A, the table reports selected summary statistics for the options/CDS and options-only skewness, and for their difference. Panel B shows selected summary statistics for the skewness measure of Bakshi, Kapadia, and Madan (2003), and the options/CDS and options-only skewness. The figures in this panel are based on observations for which all three measures can be calculated. The sample includes the last three available observations of each month for each company, from 2006 to 2015.

Panel A		Distribution of:		Panel B		Distribution of:	
	Skewness (options/CDS)	Skewness (options only)	Difference in skewness		Skewness (BKM)	Skewness (options/CDS)	Skewness (options only)
1%	-14.402	-12.192	-0.754	1%	-3.294	-11.693	-9.741
5%	-3.405	-3.190	-0.011	5%	-1.862	-2.835	-2.679
10%	-1.868	-1.805	-0.002	10%	-1.414	-1.608	-1.566
25%	-0.793	-0.784	0.000	25%	-0.955	-0.726	-0.719
Average	-0.884	-0.844	-0.040	Average	-0.664	-0.769	-0.735
50%	-0.316	-0.312	0.000	50%	-0.593	-0.316	-0.313
75%	0.033	0.032	0.000	75%	-0.250	0.000	0.000
90%	0.466	0.457	0.001	90%	0.107	0.379	0.371
95%	0.910	0.885	0.004	95%	0.334	0.761	0.732
99%	3.502	3.048	0.183	99%	0.831	3.009	2.467

of BKM skewness is generally comparable to the magnitude of ours. However, there are noticeable differences, particularly in the tails.

In what follows we evaluate whether differences between BKM, options-only, and options/CDS skewness are economically meaningful for investors who trade skewness risk. We do so by forming three portfolios on the basis of the three measures. Each portfolio is long (short) stocks in the top (bottom) 25% of the skewness distribution at the end of the previous month. Figure 2 shows the cumulative returns on the different skewness portfolios. The returns on the options/CDS ($\text{SKEW}_{opt/CDS}$) and options-only (SKEW_{opt}) portfolios are similar, but not identical. The BKM portfolio (SKEW_{BKM}) tracks the other two closely until early 2009. However, the portfolios based on our parametric skewness measures outperform the non-parametric BKM portfolio through 2009. This gap only starts to close in late 2014. Since the options/CDS and options-only portfolios have similar cumulative returns, the outperformance of $\text{SKEW}_{opt/CDS}$ and SKEW_{opt} is not driven by the inclusion of CDS. The substantive difference lies with our use of a parametric method to recover risk-neutral

Figure 2: Cumulative returns of momentum and skewness-related portfolios

The solid blue line is the return on the portfolio based on options/CDS-skewness ($\text{SKEW}_{opt/CDS}$). The portfolio is long stocks in the top 25% of the distribution of average skewness in month t-1 and short those in the bottom 25%. The dashed red line is the return on the portfolio based on options-only skewness (SKEW_{opt}). The black dash-dot line is the return on a third skewness-based portfolio, with skewness computed according to the Bakshi, Kapadia, and Madan (2003) method (SKEW_{BKM}). The grey solid thick line is the return on the UMD momentum factor. The shaded area shows the December 2007–June 2009 recession. The sample period is February 2006 to December 2015.

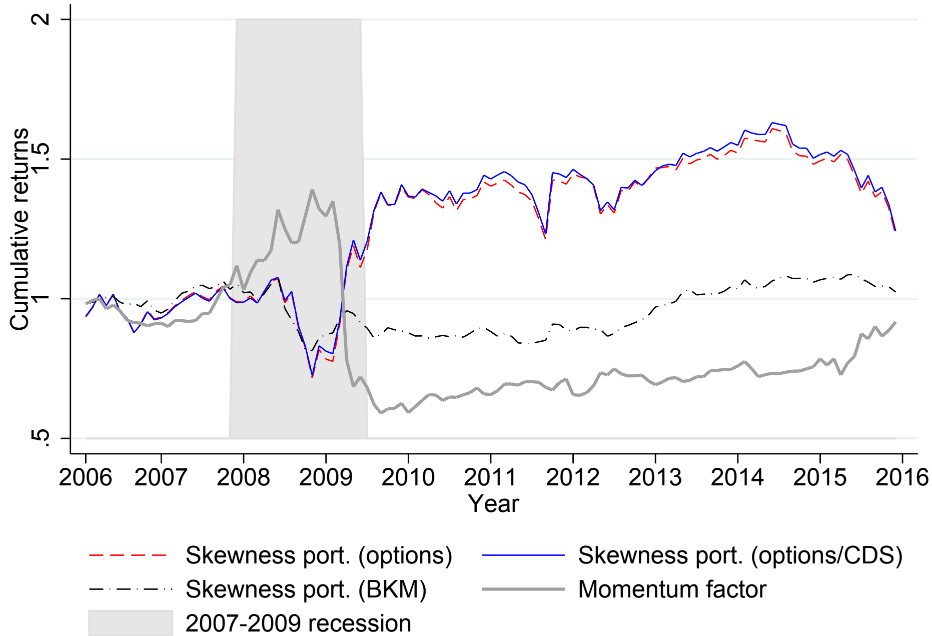

distributions. A plausible explanation is that, by combining multiple sources of data and imposing a parametric structure, we can better extract information at times of large movements in the investment opportunity set. Indeed, early 2009 is when the stock market bottomed out and a sustained recovery commenced. This sudden turn led to large losses for the momentum strategy (also shown in Figure 2; see Barroso and Santa-Clara, 2015 and Daniel and Moskowitz, 2016), which, by construction, is not well suited for capturing rapid changes in trend after a prolonged slump. In Appendix D, we constrast the patterns in momentum returns and in $\text{SKEW}_{opt/CDS}$ and SKEW_{opt} returns to support our interpretation that the parametric nature of our approach can extract signals about future investment opportunities more efficiently.

We test the relationship between the BKM and options/CDS portfolio returns formally, by regressing options/CDS-portfolio returns on BKM-portfolio returns and a set of additional risk factors, including the five Fama and French (2015) factors, changes in the spread of the CDX High Yield index (to capture economy-wide default risk), and a factor-mimicking portfolio for changes in the implied volatility index VIX. Factor-mimicking portfolios allow us to interpret the intercepts of time-series regressions as risk-adjusted average returns.[2] We report our findings in Table 4. The coefficients on the BKM portfolio are statistically significant both in the full sample and in the 2008-2011 sample, which focuses on the global financial crisis and its wake. The statistical significance of the estimated intercept in the 2008-2011 period implies that the returns on the options/CDS portfolio are not fully explained by the other factors included in the regression. This result is consistent with Figure 2, which shows that the options/CDS portfolio outperforms the BKM portfolio precisely in the aftermath the 2008 financial crisis.

Besides the months that immediately follow the 2008 financial crisis, there are other periods in which the return patterns of the skewness portfolios are noteworthy. First, returns are relatively high when the broader stock market performs strongly, as was the case between mid-2012 and mid-2014. The reason is that these portfolios buy high-skewness stocks, for which the probability mass is shifted to the left but the right tail is thicker. Such stocks are expected to earn modest returns but they have the potential to outperform and do particularly well when the broader stock market rises. Second, the skewness portfolios experience negative returns when the stock market moves sideways, for instance in 2015. The reason is that the portfolios sell low-skewness stocks, whose probability mass is shifted to the right and the left tail is thicker. These stocks are expected to earn relatively high returns but they can underperform. At times of sideways market moves, this left-tail risk is unlikely to materialize, and stocks shorted by the skewness portfolios are likely to earn higher returns,

[2] The factor mimicking portfolio is built using stocks for which we can calculate the volatility spreads of Bali and Hovakimian (2009) and the implied-volatility smirk of Xing, Zhang, and Zhao (2010). These stocks are also used to replicate the 25 size/book-to-market portfolios of Fama and French (1993) that we use in the cross-sectional asset pricing tests discussed in Section 4.2. We regress monthly stock returns in excess of the risk-free rate on the three Fama and French (1993) factors plus momentum, and changes in VIX, changes in VIX squared, and the variance risk premium of Bollerslev, Tauchen, and Zhou (2009) over the 2006-2015 time period. The equally-weighted replicating portfolio is long (short) stocks whose factor loadings on VIX changes are in the top (bottom) 25% of the distribution.

Table 4: Risk-adjusted returns of the options/CDS skewness portfolio

This table shows the results of regressing returns on the $\text{SKEW}_{opt/CDS}$ portfolio on selected risk factors and the SKEW_{BKM} portfolio. $\text{SKEW}_{opt/CDS}$ and SKEW_{BKM} are portfolios that buy (sell) stocks in the top (bottom) 25% of the distribution of average skewness in month t-1, with skewness computed with the options/CDS and BKM method, respectively. The sample covers 2006 to 2015.

	Dependent variable: $\text{SKEW}_{opt/CDS}$			
	2006-2015		2008-2011	
MKT	0.263***	0.229*	0.309**	0.314*
	(2.95)	(1.82)	(2.38)	(1.70)
SMB	0.425***	0.403***	0.447**	0.497***
	(3.60)	(3.34)	(2.21)	(2.86)
HML	0.200	0.382**	0.201	0.343
	(1.36)	(2.10)	(0.81)	(1.29)
RMW		-0.074		-0.279
		(-0.39)		(-0.84)
CMA		-0.593**		-0.510
		(-2.56)		(-1.29)
UMD	-0.320***	-0.293***	-0.328***	-0.306***
	(-6.69)	(-4.99)	(-5.50)	(-4.57)
ΔVIX_{repl}	-0.108	-0.117	-0.107	-0.137*
	(-1.44)	(-1.58)	(-1.68)	(-1.86)
ΔCDX_{HY}		-0.001		0.005
		(-0.17)		(0.57)
SKEW_{BKM}	0.884***	0.795***	0.931***	0.881***
	(5.90)	(5.09)	(4.69)	(3.95)
Intercept	0.001	0.002	0.008*	0.011**
	(0.42)	(0.85)	(1.73)	(2.14)
Obs.	119	119	48	48
Adj.R^2	0.680	0.692	0.779	0.775

dragging down the overall performance of the portfolios. Also recall that the long leg of the portfolios includes stocks with modest expected returns and upside risk, with the latter unlikely to materialize when the broad stock market remains flat.

4.2 Asset-pricing tests

Thus far we have compared the different skewness measures using the returns on portfolios based on each measure. We have shown that our parametric approach compares well to the established BKM method, and it outperforms BKM during times of high volatility. We now turn to evaluating the economic significance of including CDS in the computation of risk-neutral skewness. We do so with a variety of time-series and cross-sectional asset pricing tests.

As shown in Figure 2, the cumulative returns on the SKEW$_{opt/CDS}$ and SKEW$_{opt}$ portfolios are fairly close, with the former outperforming the latter only slightly over the sample period. To properly assess the economic significance of the contribution of CDS, however, one should not compare SKEW$_{opt/CDS}$ and SKEW$_{opt}$. Rather, the focus should be on a portfolio that explicitly loads on the differences between the options/CDS and options-only skewness measures. By doing so, the dynamics of portfolio returns are driven specifically by skewness differences that arise from the inclusion of CDS. When comparing portfolios based on skewness levels, the role of CDS is likely to be overshadowed by broad skewness-related dynamics.

The battery of tests that we conduct revolve around a trading strategy that buys and sells stocks on the basis of differences between options/CDS and options-only risk-neutral skewness (henceforth, the DS portfolio). DS is defined as the monthly returns on a portfolio that buys (sells) stocks in the top (bottom) 25% of the distribution of skewness differences in month t-1 (in unreported results, we use top/bottom 10% cutoffs, with similar conclusions). We assign companies to portfolio legs based on moments as of the end of the previous month. In order to reduce potential estimation noise, while also retaining timely information, we average the moments over the last three days of the month (Stilger, Kostakis, and Poon, 2017 also use end-of-month moments). For each company, we discard stock returns on the first day of the month to avoid possible issues with non-synchronous trading between options, CDS, and stocks.

The cumulative return on the DS portfolio is shown in Figure 3, together with the cumulative return on the broad market factor of Fama and French (1993) (henceforth, MKT). As it is clear from the figure, DS returns are distinct from market returns. Figure 3 also shows the cumulative returns on a strategy that is long SKEW$_{opt/CDS}$ and short SKEW$_{opt}$. This strategy tracks the gap between the solid blue line and the dashed red line in Figure 2. The return profile of the DS portfolio, compared to the difference between the skewness-based portfolios, confirms that, in order to assess the contribution of CDS to risk-neutral skewness, the trading strategy needs to focus explicitly on skewness differences as the sorting variable.

We choose to evaluate our method with asset pricing tests for three reasons. First,

Figure 3: DS portfolio and market factor, cumulative returns

In month t, the difference-in-skewness (DS) portfolio buys (sells) stocks in the top (bottom) 25% of the month $t-1$ distribution of skewness differences between the options/CDS and options-only distributions. End-of-month values for skewness are the average of the last three daily observations within each month. The market factor is the return of the Fama-French market portfolio (Fama and French, 1993) over the risk-free rate. Cumulative returns for DS are shown in solid red, and cumulative returns for the market factor are in dashed blue. The dash-dot black line shows the cumulative returns for a trading strategy that buys the portfolio based on options/CDS-skewness and sells the portfolio based on options-only skewness. The sample period is February 2006 to December 2015.

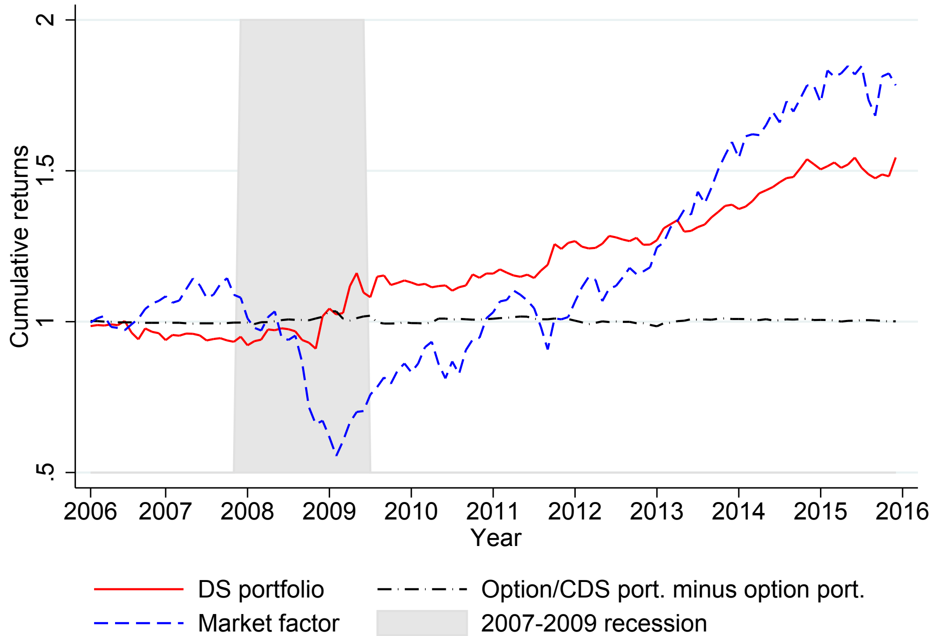

risk-adjusted returns reflect how much investors care about the risk measured by skewness differences. Therefore, studying risk-adjusted returns is a simple way of evaluating the economic significance of our contribution to the measurement of firm-specific risk. Second, due to data limitations, our panel of firm/day risk-neutral moments is unbalanced. For the median company, we can calculate risk-neutral moments on 34% of the business days. At the 10^{th} and 90^{th} percentiles, we can calculate risk-neutral moments on 12% and 68% of the business days, respectively. By forming portfolios and studying returns, we obtain continuous time series that also smooth out estimation noise. Third, we do not have to worry about potential econometric issues from extreme skewness values, as we would have had to if we had decided to assess the predictive power of options/CDS skewness for future risk neutral or historical moments.

4.2.1 Time-series evidence

The tests in this section focus on whether existing asset-pricing factors can explain the returns of the DS portfolio. In the context of asset pricing tests, the intercept of a time series regression represents a risk-adjusted average return, which can be interpreted as a risk premium. If the estimated intercept is statistically different from zero, we conclude that the factors included in the time-series regression cannot fully account for the returns on the DS portfolio.

We fit the following regression model to the data:

$$DS_t = \alpha + \sum_{i=1}^{N} \beta_i f_t^i + \varepsilon_t, \tag{1}$$

where α is the intercept (our parameter of interest), f_t^i are the factors discussed below, and ε_t is the error term. We report heteroskedasticity-consistent standard errors for all estimated parameters, following White (1980).

We use a large number of relevant asset-pricing factors, including the five Fama and French (2015) factors (the market, MKT, size, SMB, book-to-market, HML, profitability, RMW, and investment, CMA), the Carhart (1997) momentum factor (UMD), the Pastor and Stambaugh (2003) liquidity factor (LIQ), and changes in the spread of the CDX high-yield CDS index (ΔCDX_{HY}) to control for aggregate default risk.[3]

The long run and short run reversals factors, LT REV and ST REV respectively, capture the effect of past and recent stock performance on current stock returns (see, among others, Fama and French, 1996 and Novy-Marx, 2012, and references therein). ΔVIX_{repl} and $\Delta\text{VIX}^2_{repl}$ are replicating portfolios for changes and squared changes in the implied volatility index VIX, VRP_{repl} is the replicating portfolio for Bollerslev, Tauchen, and Zhou (2009)'s

[3] Qiu and Yu (2012) document that CDS liquidity affects CDS spreads, especially for smaller companies. The companies in our sample are much larger than in Qiu and Yu (2012)'s. In addition, the liquidity factor LIQ directly controls for broad market liquidity, and ΔCDX_{HY} likely reflects CDS-market liquidity.

variance risk premium.[4] We also include the $\text{SKEW}_{opt/CDS}$ factor to ensure that the results are not driven by skewness.

The DS portfolio has a moderate positive correlation (0.27) with the market factor (MKT) and a negative correlation with ΔVIX (-0.33), UMD (-0.38), and ΔCDX_{HY} (-0.38). The skewness portfolios $\text{SKEW}_{opt/CDS}$ and SKEW_{opt} are highly correlated (0.99), though the correlation of DS with both $\text{SKEW}_{opt/CDS}$ (0.52) and SKEW_{opt} (0.53) is only moderate. The BKM-skewness portfolio SKEW_{BKM} is weakly correlated with DS (0.28) and moderately correlated with $\text{SKEW}_{opt/CDS}$ (0.52) and SKEW_{opt} (0.53).

We report the time series test results in Tables 5 (for the full sample) and 6 (for 2008 to 2011). The first column of Table 5 shows that, when using a small set of factors that includes MKT, SMB, HML, UMD, LIQ, and ΔVIX_{repl}, the intercept – the risk-adjusted average return – is about 0.3% per month and is statistically significant at the 10% level. The intercept remains about the same as we include progressively more factors, and the statistical significance improves to the 5% level. The inclusion of the $\text{SKEW}_{opt/CDS}$ portfolio improves the adjusted R^2 from 39% to 44%. When we substitute this portfolio with SKEW_{BKM}, the R^2 declines back to 39%. A pertinent question is which leg of the DS portfolio is delivering the results. In columns (6) and (7) of Table 5, we regress the returns of the long leg and of the short leg, respectively, on the full set of factors. We observe the following. First, the risk-adjusted average return of the DS portfolio is driven by the long leg. Second, both legs load heavily on the MKT factor, which results in high adjusted R^2s. Finally, both legs are highly correlated with $\Delta\text{VIX}^2_{repl}$ and the $\text{SKEW}_{opt/CDS}$ portfolio.

In Table 6, we restrict our sample to the 2008-2011 period. Focusing on this period of heightened financial uncertainty, we find that the risk-adjusted average return is noticeably higher, ranging between 0.8% to 1.1% per month. The statistical significance is about the same or stronger, with the exception of the specification with the smallest set of risk factors. While larger, the adjusted R^2 pattern is similar to what we observed in Table 5. In contrast

[4] The replicating portfolio for VIX changes is described in footnote 2. The replicating portfolios for squared VIX changes and for the variance risk premium are built in exactly the same manner. The only difference is that stocks are sorted into portfolios using the coefficients on squared VIX changes and on the variance risk premium, respectively, rather than on VIX changes.

Table 5: Time series intercept of the DS portfolio: full sample

The table reports coefficients, intercepts, t-statistics, number of observations, and adjusted R^2s from regressions of the DS portfolio on the listed factors: $DS_t = \alpha + \sum_{i=1}^{N} \beta_i f_t^i + \varepsilon_t$. MKT, SMB, HML, RMW, and CMA are the market, size, book-to-market, profitability, and investment factors of Fama and French (2015). UMD is the Carhart (1997) momentum factor. LT REV and ST REV are the long run and short run reversals factors (see, among others, Fama and French, 1996 and Novy-Marx, 2012, and references therein). LIQ is the Pastor and Stambaugh (2003) liquidity factor. ΔVIX_{repl}, ΔVIX^2_{repl}, and VRP_{repl} are replicating portfolios for changes and squared changes in the implied volatility index VIX, and for the variance risk premium of Bollerslev, Tauchen, and Zhou (2009). ΔCDX_{HY} is changes in the spread of the CDX high-yield CDS index. The t-statistics are based on heteroskedasticity-consistent standard errors. ***, **, and * indicate statistical significance at the 1, 5, and 10 % levels, respectively. The sample includes February 2006 to December 2015.

	(1)	(2)	(3)	(4)	(5)	(6) long leg	(7) short leg
MKT	0.060	0.081	0.078	0.021	0.071	0.943***	0.921***
	(1.23)	(1.62)	(1.09)	(0.35)	(0.99)	(11.48)	(12.12)
SMB	-0.016	0.015	-0.004	-0.056	0.000	-0.022	0.034
	(-0.22)	(0.21)	(-0.05)	(-0.78)	(0.00)	(-0.22)	(0.34)
HML	0.001	0.191	0.202	0.145	0.209	0.114	-0.031
	(0.01)	(1.44)	(1.49)	(1.16)	(1.53)	(0.86)	(-0.24)
UMD	-0.164***	-0.146***	-0.132***	-0.098**	-0.138***	-0.075	0.023
	(-3.42)	(-3.72)	(-2.64)	(-2.10)	(-2.78)	(-1.37)	(0.43)
LIQ	0.045	0.031	0.039	0.030	0.040	-0.000	-0.030
	(0.82)	(0.52)	(0.62)	(0.50)	(0.63)	(-0.00)	(-0.40)
RMW		-0.110	-0.147	-0.089	-0.134	0.167	0.257
		(-0.79)	(-1.14)	(-0.73)	(-0.97)	(1.16)	(1.59)
CMA		-0.177	-0.141	-0.059	-0.145	0.027	0.085
		(-1.02)	(-0.84)	(-0.35)	(-0.86)	(0.13)	(0.40)
LT REV		-0.199**	-0.241***	-0.184*	-0.215**	-0.089	0.095
		(-1.99)	(-2.63)	(-1.83)	(-2.03)	(-0.62)	(0.76)
ST REV		-0.126*	-0.135*	-0.118*	-0.132*	0.029	0.147*
		(-1.68)	(-1.82)	(-1.69)	(-1.78)	(0.37)	(1.78)
ΔVIX_{repl}	-0.235***	-0.248***	-0.212***	-0.219***	-0.216***	-0.195*	0.025
	(-4.15)	(-3.96)	(-3.64)	(-4.22)	(-3.76)	(-1.97)	(0.32)
ΔVIX^2_{repl}			0.132	0.006	0.112	0.330**	0.323***
			(0.99)	(0.06)	(0.91)	(2.27)	(2.67)
VRP_{repl}			-0.061	-0.119	-0.089	-0.062	0.058
			(-0.50)	(-1.43)	(-0.63)	(-0.43)	(0.40)
ΔCDX_{HY}			0.001	-0.000	0.000	0.005	0.006
			(0.14)	(-0.11)	(0.08)	(1.11)	(1.32)
$SKEW_{opt/CDS}$				0.197***		0.372***	0.175**
				(3.00)		(5.63)	(2.43)
$SKEW_{BKM}$					0.070		
					(0.59)		
Intercept	0.003*	0.004**	0.004**	0.003**	0.004**	0.003	-0.000
	(1.86)	(2.07)	(2.17)	(1.98)	(2.17)	(1.46)	(-0.07)
Obs.	119	119	119	119	119	119	119
Adj.R^2	0.329	0.397	0.389	0.440	0.385	0.904	0.865

Table 6: Time series intercept of the DS portfolio: between 2008 and 2011

The table reports coefficients, intercepts, t-statistics, number of observations, and adjusted R^2s from regressions of the DS portfolio on the listed factors: $DS_t = \alpha + \sum_{i=1}^{N} \beta_i f_t^i + \varepsilon_t$. MKT, SMB, HML, RMW, and CMA are the market, size, book-to-market, profitability, and investment factors of Fama and French (2015). UMD is the Carhart (1997) momentum factor. LT REV and ST REV are the long run and short run reversals factors (see, among others, Fama and French, 1996 and Novy-Marx, 2012, and references therein). LIQ is the Pastor and Stambaugh (2003) liquidity factor. ΔVIX_{repl}, ΔVIX^2_{repl}, and VRP_{repl} are replicating portfolios for changes and squared changes in the implied volatility index VIX, and for the variance risk premium of Bollerslev, Tauchen, and Zhou (2009). ΔCDX_{HY} is changes in the spread of the CDX high-yield CDS index. The t-statistics are based on heteroskedasticity-consistent standard errors. ***, **, and * indicate statistical significance at the 1, 5, and 10 % levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
						long leg	short leg
MKT	0.070	0.056	0.132	0.025	0.118	0.959***	0.933***
	(0.88)	(0.69)	(1.18)	(0.24)	(1.19)	(6.48)	(6.22)
SMB	-0.082	-0.043	0.011	-0.110	-0.023	0.066	0.176
	(-0.54)	(-0.29)	(0.08)	(-0.78)	(-0.18)	(0.36)	(1.03)
HML	0.008	0.243	0.338	0.281	0.463*	0.045	-0.236
	(0.03)	(1.03)	(1.47)	(1.52)	(1.95)	(0.19)	(-0.97)
UMD	-0.210***	-0.140**	-0.129*	-0.075	-0.170**	-0.101	-0.026
	(-4.25)	(-2.60)	(-2.00)	(-1.19)	(-2.48)	(-1.18)	(-0.35)
LIQ	0.049	0.055	0.097	0.075	0.136	-0.067	-0.142
	(0.57)	(0.57)	(1.11)	(0.91)	(1.58)	(-0.56)	(-1.25)
RMW		-0.334	-0.440*	-0.311	-0.419*	0.555*	0.867***
		(-1.19)	(-1.86)	(-1.24)	(-1.79)	(1.89)	(3.63)
CMA		-0.508**	-0.565*	-0.433	-0.694**	-0.072	0.362
		(-2.05)	(-1.96)	(-1.51)	(-2.70)	(-0.19)	(1.18)
LT REV		-0.074	-0.183	-0.102	0.022	-0.036	0.066
		(-0.47)	(-1.35)	(-0.61)	(0.14)	(-0.18)	(0.45)
ST REV		-0.169	-0.210*	-0.177	-0.218**	0.171	0.348***
		(-1.55)	(-1.87)	(-1.56)	(-2.05)	(1.32)	(3.60)
ΔVIX_{repl}	-0.235***	-0.263***	-0.201***	-0.215***	-0.249***	-0.146	0.069
	(-3.78)	(-4.33)	(-3.31)	(-3.73)	(-3.70)	(-1.45)	(0.79)
ΔVIX^2_{repl}			0.315**	0.109	0.161	0.370*	0.262*
			(2.31)	(0.70)	(1.09)	(1.90)	(1.73)
VRP_{repl}			-0.191	-0.204**	-0.406***	0.070	0.274*
			(-1.49)	(-2.06)	(-3.01)	(0.47)	(1.94)
ΔCDX_{HY}			0.008	0.004	0.007	0.011*	0.007
			(1.24)	(0.85)	(1.14)	(1.76)	(1.13)
$SKEW_{opt/CDS}$				0.256***		0.423***	0.166
				(3.03)		(3.87)	(1.53)
$SKEW_{BKM}$					0.427**		
					(2.35)		
Intercept	0.004	0.008**	0.008**	0.006*	0.011***	-0.001	-0.008*
	(1.33)	(2.05)	(2.05)	(1.95)	(3.06)	(-0.27)	(-1.98)
Obs.	48	48	48	48	48	48	48
Adj.R^2	0.455	0.513	0.546	0.623	0.590	0.929	0.908

to what we observe in the previous table, columns (6) and (7) imply that the short leg is driving the results. Our estimates of the risk-adjusted average returns for the DS portfolio are comparable to the skewness premium estimated by Rehman and Vilkov (2012) and Stilger, Kostakis, and Poon (2017). Both studies find skewness premia of roughly 0.5% per month.

4.2.2 Cross-sectional evidence

In this section, we carry out a second set of tests where we use the two-stage method of Fama and MacBeth (1973) in order to introduce stock-level characteristics that can proxy for risk exposure (see, for instance, Daniel and Titman, 1997 and Daniel, Titman, and Wei, 2001).

The first stage of the Fama and MacBeth (1973) procedure estimates the sensitivity of portfolio returns to the various factors with a series of portfolio-specific time-series regressions:

$$r_t^j - r_t^f = \alpha_j + \sum_{i=1}^{N} \beta_i^j f_t^i + \varepsilon_t^j, \forall j \tag{2}$$

where r_t^j is the return on portfolio j, r_t^f is the risk-free rate, and f_t^i is one of the N factors included. We consider 35 portfolios: 10 decile portfolios based on the distributions of DS in month t-1, and the replication of the 25 size/book-to-market portfolios of Fama and French (1993) using stocks for which we can calculate the volatility spreads of Bali and Hovakimian (2009) and the implied-volatility smirk of Xing, Zhang, and Zhao (2010).

The second step of the Fama and MacBeth (1973) method uses cross-sectional regressions to evaluate how differences in the estimated factor loadings explain excess returns:

$$r_t^j - r_t^f = \lambda_t^0 + \sum_{i=1}^{N} \lambda_t^i \hat{\beta}_i^j + \sum_{k=1}^{K} \phi_k \gamma_{t-1,k}^j + \epsilon_j, \forall t, \tag{3}$$

where λ_t^0 is the pricing error at time t, λ_t^i is the risk premium on factor i at time t, $\hat{\beta}_t^j$ are the estimates from the first step, $\gamma_{t-1,k}^j$ is characteristic k for portfolio j as of time t-1 (calculated as the average of stock-level characteristics), and ϕ_k is the regression coefficient for characteristic $\gamma_{t-1,k}^j$. The risk premium on factor f_t^i is computed as the average of the coefficients from the T cross-sectional regressions and its statistical significance is assessed with Shanken (1992)-adjusted standard errors with Newey-West correction:

$$\hat{\lambda}^i = \frac{1}{T}\sum_{t=1}^{T}\hat{\lambda}^i_t. \tag{4}$$

The characteristics we consider are related to higher-moment risks in stock and option returns, and have been shown to explain the cross sections of equity and option returns: idiosyncratic volatility (Ang, Hodrick, Xing, and Zhang, 2006), volatility spreads (Bali and Hovakimian, 2009), the implied-volatility smirk of Xing, Zhang, and Zhao (2010), the changes in call and put implied volatility of An, Ang, Bali, and Cakici (2014), and the tail covariance between stock market returns and company-specific returns introduced by Bali, Cakici, and Whitelaw (2014). In addition, we include the 5-year CDS spread to evaluate whether the contribution of the DS portfolio to the cross-section of returns is explained by default risk.

In the Fama-MacBeth procedure, the risk premia are identified from cross-sectional variation in factor sensitivities and our cross section includes 35 portfolios. As such, each specification we discuss includes only a subset of the factors we use in the time series tests, and we study various combinations of factors and characteristics to ensure that our findings are robust. We present the first set of results in Table 7 where we keep the factors constant, but change the set of characteristics included in the regression. In the full sample, the risk premium on DS is statistically significant at the 5% level and equal to 0.5% per month. It is comparable to the intercept in the time-series regressions of DS on other factors.[5] In line with Ang, Hodrick, Xing, and Zhang (2006), we find that stocks with high idiosyncratic volatility have lower returns. In addition, controlling for skewness using either $\text{SKEW}_{opt/CDS}$ or SKEW_{BKM} portfolios does not affect our findings. In the 2008-2011 sub-sample, the DS portfolio commands a higher premium, around 1% per month, which is in line with our time-series results. We find that the estimated coefficients for Δcvol and Δpvol (changes in call or put implied volatilities) have the same signs as in An, Ang, Bali, and Cakici (2014) and

[5] The risk premia on most factors are statistically not different from zero. This is not surprising, since the test portfolios we use are designed to generate cross-sectional dispersion in the sensitivity to the risk expressed by DS. Hence the test is designed to capture whether the risk premium on the DS portfolio is absorbed by other factors. The 25 Fama-French portfolios are meant to generate cross-sectional sensitivity to SMB and HML factors. However, since the mean returns to these two factors are statistically not different from zero over the sample period, we do not expect to observe a risk premium for either factor.

Table 7: Fama-MacBeth regressions (1/2)

The table shows the second-stage coefficients from Fama and MacBeth (1973) regressions. The test assets are 25 book-to-market/size portfolios replicated using stocks with available volatility spreads (Bali and Hovakimian, 2009) and smirks (Xing, Zhang, and Zhao, 2010), and 10 decile portfolios based on the DS measure in month t-1. The characteristics (portfolios averages in month t-1) are idiosyncratic volatility (Ang, Hodrick, Xing, and Zhang, 2006) (ivol), average 5-year CDS log-spread (sprd_{5yr}), hybrid tail covariance (Bali, Cakici, and Whitelaw, 2014) (cov_{tail}), changes in call/put implied volatility of An, Ang, Bali, and Cakici (2014) (Δcvol and Δpvol), and volatility spreads (vspread) and smirks (smirk). Adj.R^2_{cr} is the average adj. R^2 of the second-stage regressions. The sample covers February 2006 to December 2015.

	2006-2015				2008-2011			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
λ^0	0.034	0.023	0.027	0.019	0.003	0.004	0.023	0.020
	(1.37)	(0.94)	(1.61)	(1.16)	(0.07)	(0.10)	(0.87)	(0.63)
ivol	-0.251*	-0.215*	-0.154*	-0.153*	-0.293	-0.168	-0.285**	-0.234
	(-1.83)	(-1.77)	(-1.71)	(-1.70)	(-1.35)	(-0.93)	(-2.17)	(-1.58)
sprd_{5yr}	0.354	0.235	0.248	0.133	-0.015	0.058	0.259	0.170
	(1.04)	(0.72)	(0.94)	(0.52)	(-0.03)	(0.11)	(0.72)	(0.39)
cov_{tail}	-0.007	-0.007	-0.010	-0.011*	-0.012*	-0.012*	-0.010**	-0.008
	(-0.96)	(-0.97)	(-1.55)	(-1.87)	(-1.93)	(-1.99)	(-2.08)	(-1.18)
Δcvol	0.098		0.100	0.121	0.215		0.194*	0.214
	(0.93)		(1.13)	(1.31)	(1.51)		(1.85)	(1.51)
Δpvol	-0.027		-0.049	-0.066	-0.223		-0.200*	-0.206
	(-0.27)		(-0.58)	(-0.77)	(-1.47)		(-1.78)	(-1.39)
vspread	0.034	0.028			0.008	0.007		
	(0.80)	(0.69)			(0.10)	(0.09)		
smirk	0.079	0.080			0.205*	0.139		
	(0.92)	(0.94)			(1.94)	(1.40)		
MKT	0.003	0.007	0.002	0.005	0.012	0.013	0.010	0.005
	(0.37)	(0.78)	(0.25)	(0.56)	(0.98)	(1.14)	(0.79)	(0.39)
SMB	0.002	0.002	0.002	0.002	0.002	-0.001	0.001	-0.001
	(0.46)	(0.64)	(0.46)	(0.59)	(0.43)	(-0.20)	(0.33)	(-0.33)
HML	-0.007	-0.008	-0.005	-0.004	0.006	0.004	0.003	0.005
	(-1.36)	(-1.62)	(-1.11)	(-0.88)	(0.77)	(0.53)	(0.44)	(0.64)
UMD	-0.002	0.004	-0.004	-0.006	-0.006	0.007	-0.002	-0.001
	(-0.20)	(0.31)	(-0.40)	(-0.59)	(-0.27)	(0.32)	(-0.13)	(-0.09)
LT REV	-0.006*	-0.005	-0.003	-0.003	-0.001	0.002	0.001	0.001
	(-1.71)	(-1.32)	(-1.03)	(-0.77)	(-0.18)	(0.26)	(0.21)	(0.10)
ST REV	-0.008	-0.009	-0.008	-0.009	-0.003	-0.005	0.000	-0.002
	(-1.26)	(-1.21)	(-1.11)	(-1.28)	(-0.19)	(-0.31)	(-0.02)	(-0.12)
LIQ	0.000	0.000	0.000	0.000	0.012	0.010	0.009	0.005
	(-0.04)	(0.00)	(0.04)	(0.03)	(0.65)	(0.54)	(0.55)	(0.31)
ΔVIX	0.001	0.001	0.004	-0.002	-0.020	-0.017	-0.015	-0.016
	(0.07)	(0.11)	(0.34)	(-0.19)	(-1.44)	(-1.20)	(-1.06)	(-1.06)
ΔVIX^2	0.003*	0.003*	0.002	0.002	0.003	0.003	0.002	0.003
	(1.85)	(1.84)	(1.63)	(1.65)	(1.27)	(1.25)	(1.00)	(1.24)
$\text{SKEW}_{opt/CDS}$	0.008	0.007	0.007		0.017	0.012	0.013	
	(1.08)	(0.97)	(1.09)		(1.23)	(0.96)	(1.12)	
SKEW_{BKM}				0.002				-0.001
				(0.46)				(-0.16)
DS	0.005**	0.005*	0.004**	0.005**	0.011*	0.010	0.010*	0.010*
	(2.01)	(1.82)	(2.01)	(2.22)	(1.99)	(1.64)	(1.90)	(1.92)
Adj.R^2_{cr}	0.529	0.497	0.498	0.489	0.653	0.644	0.633	0.626

Table 8: Fama-MacBeth regressions (2/2)

The table shows the second-stage coefficients from Fama and MacBeth (1973) regressions. The test assets are 25 book-to-market/size portfolios replicated using stocks with available volatility spreads (Bali and Hovakimian, 2009) and smirks (Xing, Zhang, and Zhao, 2010), and 10 decile portfolios based on the DS measure in month t-1. The characteristics (portfolios averages in month t-1), are idiosyncratic volatility (Ang, Hodrick, Xing, and Zhang, 2006) (ivol) average 5-year CDS log-spread ($sprd_{5yr}$), hybrid tail covariance (Bali, Cakici, and Whitelaw, 2014) (cov_{tail}), and changes in call/put implied volatility of An, Ang, Bali, and Cakici (2014) (Δcvol and Δpvol). Adj.R^2_{cr} is the average adj. R^2 of the second-stage regressions. The sample covers February 2006 to December 2015.

	2006-2015				2008-2011			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
λ^0	0.026	0.017	0.027*	0.018	0.017	0.012	0.015	0.003
	(1.60)	(1.12)	(1.67)	(1.08)	(0.64)	(0.45)	(0.60)	(0.13)
ivol	-0.138	-0.134	-0.155	-0.143	-0.249*	-0.266*	-0.282**	-0.236*
	(-1.54)	(-1.55)	(-1.64)	(-1.53)	(-1.78)	(-1.96)	(-2.02)	(-1.74)
$sprd_{5yr}$	0.243	0.120	0.204	0.053	0.142	0.078	0.122	0.010
	(0.91)	(0.48)	(0.78)	(0.20)	(0.37)	(0.20)	(0.32)	(0.03)
cov_{tail}	-0.009	-0.011*	-0.010	-0.013**	-0.009**	-0.006	-0.008**	-0.006
	(-1.44)	(-1.85)	(-1.55)	(-2.06)	(-2.03)	(-1.08)	(-2.04)	(-1.17)
Δcvol	0.112	0.136	0.118	0.167*	0.156	0.173	0.197	0.134
	(1.22)	(1.49)	(1.37)	(1.83)	(1.22)	(1.23)	(1.55)	(1.02)
Δpvol	-0.066	-0.094	-0.076	-0.128	-0.170	-0.199	-0.239*	-0.149
	(-0.74)	(-1.09)	(-0.89)	(-1.46)	(-1.28)	(-1.33)	(-1.77)	(-1.10)
MKT	0.001	0.004	0.001	0.002	0.008	0.008	0.008	0.011
	(0.18)	(0.50)	(0.07)	(0.23)	(0.64)	(0.60)	(0.58)	(0.77)
SMB	0.000	0.001	0.000	0.001	0.001	0.000	0.001	0.000
	(0.05)	(0.31)	(0.01)	(0.40)	(0.15)	(-0.10)	(0.14)	(0.11)
HML	-0.007*	-0.006	-0.005	-0.004	0.002	0.004	0.003	0.005
	(-1.71)	(-1.45)	(-1.27)	(-1.10)	(0.26)	(0.56)	(0.51)	(0.70)
RMW			0.000	-0.001			-0.002	-0.004
			(0.00)	(-0.23)			(-0.58)	(-0.74)
CMA			0.000	0.000			0.000	-0.002
			(-0.03)	(0.03)			(-0.09)	(-0.30)
UMD	-0.005	-0.008	-0.004	-0.008	-0.005	-0.005	-0.010	-0.010
	(-0.49)	(-0.84)	(-0.42)	(-0.83)	(-0.31)	(-0.28)	(-0.57)	(-0.56)
LIQ	0.003	0.002	0.000	0.000	0.010	0.009	0.008	0.010
	(0.38)	(0.20)	(0.05)	(-0.02)	(0.71)	(0.61)	(0.53)	(0.60)
ΔVIX	0.005	0.000	0.005	0.002	-0.012	-0.016	-0.016	-0.023
	(0.43)	(-0.04)	(0.42)	(0.11)	(-0.78)	(-1.03)	(-1.08)	(-1.47)
ΔVIX^2	0.002	0.001	0.002	0.002	0.002	0.002	0.002	0.001
	(1.17)	(1.07)	(1.16)	(1.15)	(1.05)	(0.98)	(0.87)	(0.57)
$SKEW_{opt/CDS}$	0.006		0.006		0.011		0.013	
	(0.93)		(0.95)		(0.96)		(1.08)	
$SKEW_{BKM}$		0.002		0.003		0.000		-0.002
		(0.54)		(0.59)		(-0.07)		(-0.23)
DS	0.004*	0.005**	0.004	0.005*	0.010*	0.010	0.011**	0.010**
	(1.75)	(2.05)	(1.51)	(1.85)	(1.88)	(1.64)	(2.04)	(2.09)
Adj.R^2_{cr}	0.461	0.453	0.487	0.475	0.568	0.560	0.609	0.613

are statistically significant in several instances during this period.

Studies such as Acharya and Johnson (2007), Ni and Pan (2011), and Han and Zhou (2011) indicate the presence of information flows from the CDS market to the equity market. Our results about the DS portfolio are unlikely to be driven by such flows. We include CDS spreads in our cross-sectional regression and do not observe a statistically significant coefficient for the spreads. Moreover, we skip the first return of the month in building the test portfolios, which eliminates the observation most likely affected by the delayed information flow.

In Table 8 we keep the set of characteristics constant and change the factors. We use either the three or the five Fama-French factors; see Fama and French (1993, 2015). In both the full sample and in the 2008-2011 sub-sample, the estimated DS risk premium remains comparable to the values reported in Table 7. In model (3), the estimated DS premium is statistically zero. However, the asset pricing test is rejected for model (3), since the estimated pricing error is significantly different from zero.

The cross-sectional evidence presented in this section corroborates the time-series evidence reported in Section 4.2.1. In addition, our cross-sectional results demonstrate robustness of DS premia to the inclusion of stock characteristics in asset pricing tests.

4.3 Further discussion and robustness

Our asset-pricing results can be interpreted in terms of intermediate-return risk, that is the risk of large but not extreme returns. Intuitively, options drive the central part of the distribution, while CDS pin down the left tail. Jointly, options and CDS characterize returns in between these two regions, an area of substantially negative but not bankruptcy-inducing returns that cannot be spanned by options or CDS in isolation. With parametric assumptions, options and CDS together also help determine the shape of right part of the distribution.

The sign of the difference between options/CDS and options-only skewness indicates whether a company is more likely to experience positive or negative intermediate returns. Higher (lower) skewness translates into a shift of the probability mass to the left (right), with intermediate negative (positive) returns being more likely than under the options-only estimation. As an illustration using actual data, this shift is visible in Figure 4, which corresponds to a day when the options/CDS skewness is much more negative than the options-only skewness, hence the DS measure is large and negative. The options/CDS distribution indicates a higher likelihood of intermediate positive returns (note that the higher incidence of tail returns is captured by the CDS spread we include as a control in our regressions).

To corroborate our interpretation in terms of intermediate-return risk, we provide additional details on the risk characteristics of stocks with high/low DS in Appendix E. In the same appendix, we also include an in-depth discussion of the statistical properties of differences between the options/CDS skewness and options-only skewness. Since the DS portfolio is long (short) stocks that are more likely to experience negative (positive) intermediate returns, the risk premium on DS should be positive. This is what we find in the tests we detail in Sections 4.2.1 and 4.2.2. In addition, recall from our discussion of Table 3 in Section 4.1 that large wedges between options/CDS and options-only skewness open up sporadically. Similar to Kapadia and Zekhnini (2019), these results confirm that infrequent dynamics in expected returns contribute meaningfully to observed returns.

We now evaluate the robustness of our findings to alternative specifications and possible confounding factors. We start by considering three different default thresholds, which determine the transition point between option-implied and CDS-implied probabilities, and two less liquid shorter-horizon CDS tenors, to better align the maturity of CDS with that of options. As detailed in Appendix F, the three default thresholds are fixed across companies and are combinations of the leverage- and size-specific thresholds. In place of appropriately compounded five-year CDS spreads, where trading liquidity is concentrated, we use one-year and synthetic three-month CDS spreads, with the latter non-traded maturity obtained by extrapolating the CDS curve. Overall, the DS premium remains similar in magnitude and statistical significance to the estimates we presented in the previous section. Importantly, robustness to different default thresholds also addresses robustness to different recovery rates, since changing the threshold and changing the recovery rate are observationally equivalent in terms of the default probability extracted from CDS spreads.

Figure 4: Example of risk neutral density estimated with and without CDS

The left chart shows the risk-neutral distribution of returns for Citigroup Inc. on July 20, 2009. The distribution in solid red is estimated using options and CDS, while the one based on options only is in dashed blue. On the same day and for the same company, the right chart shows the probability density of the options/CDS distribution when this probability density is higher than the probability density of the options-only distribution.

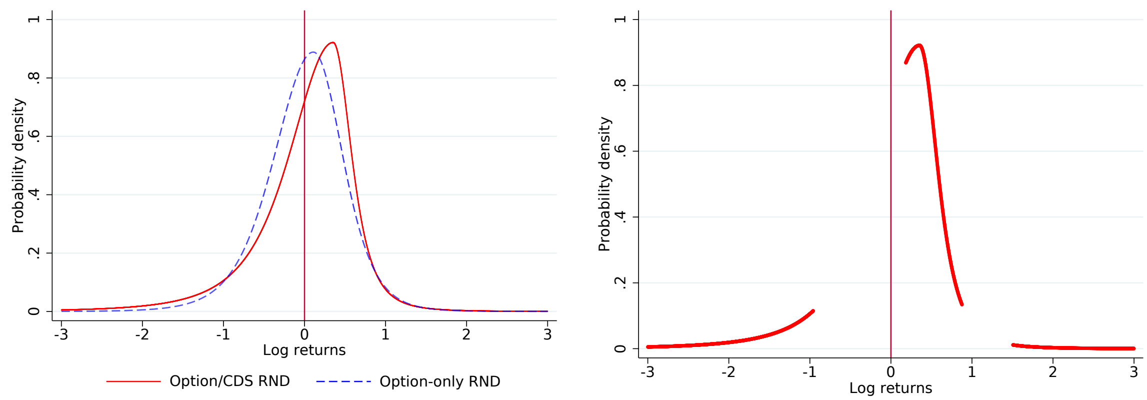

Our decision to use the skew-t distribution among various possible alternatives is firmly rooted in the literature, which indicates that this distribution is well suited to model stock return dynamics (eg, Hansen, 1994, Jondeau and Rockinger, 2003, Patton, 2004 and Oh and Patton, 2017). In Appendix F, we consider the effects of using alternative distributions and interpolation/extrapolation methods. In summary, we find that using distributions other than the skew-t does not produces significant results. Moreover, simply using the skew-t is also insufficient. However, the use of options data interpolated with kernel smoothing (the OptionMetrics volatility surface), together with the skew-t, yields intercepts with magnitudes similar to our baseline results and a p-value of 11% in the crisis sample. Overall, we interpret these findings as suggesting that the skew-t distribution is important, but it is the combination of skew-t plus options/CDS that clearly contributes to our results.

Throughout the paper, portfolio returns are gross of transaction costs, and shorting fees, in particular, might be elevated for risky stocks at times of market volatility. As such, the estimated risk premia should be considered upper bounds. We should emphasize that our asset pricing tests are meant to evaluate whether combining option prices and CDS spreads provides additional information relative to standard methods for extracting risk-neutral distributions for individual stocks. The presence of transaction costs can impact replicability but does not detract from our finding that, jointly, options and CDS unlock information about certain aspects of risk-neutral distributions. Still, based on the estimates for shorting fees at times of market distress and even assuming full porfolio turnover every month, a strategy based on skewness differences remains profitable over the sample we study.[6]

A potential concern is that CDS spreads may embed a liquidity premium, which could distort our implied probabilities of default. Qiu and Yu (2012) find that CDS spreads reflect liquidity risk, especially in the case of smaller firms. Our work covers companies that are, on average, about four times as large as those studied by Qiu and Yu (2012) because we require that both options and CDS data are available.[7] In addition, the CDS market was still developing during the earlier part of the period they study (2001-2008), which is one of the reasons why our sample starts in 2006. As a result, the presence of a liquidity premium is less of a concern in our analysis. Nevertheless, we control for liquidity risk in our asset-pricing tests by including a broad liquidity risk factor and CDS index spreads. Besides credit risk, CDS index spreads also reflect broad liquidity conditions in the CDS market. In general, asset prices are affected by liquidity conditions, and returns tend to be higher (lower) when liquidity improves (deteriorates). As such, CDS index spreads reflect both credit risk and liquidity risk.

We now turn to the possible perception that our results might arise from data mining. Harvey, Liu, and Zhu (2016) recommend setting a high bar for cross-sectional asset pricing studies, with t-statistics in excess of 3 when four factors are included and far more stringent significance criteria as the number of factors increases. However, they limit their recom-

[6] D'Avolio (2002) finds that, betwen April 2000 and September 2001, 91% of stocks have a shorting fee below 1% per year, with a value-weighted average of 0.17%. Using a more comprehensive sample that spans January 2004 through December 2013, Drechsler and Drechsler (2016) report meaningful time variation in shorting fees. In 2008 and 2009, the average fees were 1.34% ad 0.81%, respectively. As shown in Figure 3, the DS factor earns a substantial return at the very end of 2008 and in the first half of 2009. As such, we evaluate the effect of transaction costs using a 1% shorting fee, together with a $5 transaction cost per trade. Under assumptions that are conservative for our results (every stock in the long and short legs is bought and sold each month, and a 1% per year fee applies to the short leg), $10,000,000 invested in the DS strategy yield a net-of-fees risk premium of 0.47% per month with a t-statistic of 1.86.

[7] Table 1 shows selected summary statistics for the 275 companies in our sample as of 2006. Compare the average total assets figure in this table (about $63 billion) against the average total assets figure (about $17 billion) in Table 1 of Qiu and Yu (2012).

mendation to new and purely empirical results. They explicitly state that micro-founded, theoretically-motivated results – like those, in our opinion, that belong to the literature on risk-neutral skewness – need not pass their proposed bar. The link between skewness and equity returns is well established both theoretically and empirically, with a number of papers that investigate the role of risk-neutral firm-specific skewness in the cross-section of returns (e.g., Conrad, Dittmar, and Ghysels, 2013; Rehman and Vilkov, 2012; Stilger, Kostakis, and Poon, 2017). In addition, studies such as Chabi-Yo, Leisen, and Renault (2014), Feunou, Jahan-Parvar, and Tédongap (2013), and Feunou, Jahan-Parvar, and Okou (2018) explicitly derive equilibrium asset pricing models for skewness preferences. As such, our empirical results are in line with a theoretically-founded line of research. In addition, the relation that we uncover between the momentum strategy and the DS portfolio (see Appendix D) is incidental to the central topic of our work, yet it speaks to an important stylized fact in asset pricing (momentum crashes) and it arises naturally within our setup. It is unlikely that both the main difference-in-skewness results and the link between momentum and the DS portfolio are the outcome of data mining.

Additionally, data mining concerns can be addressed by gauging model stability, specifically the likelihood that our results are tied to the factors and sample we use. For example, Lewellen, Nagel, and Shanken (2010) propose to assess three quantities: the average excess returns of the portfolios used as factors, the risk premia from time-series regressions, and the risk premia from cross-sectional regressions (in our case, Fama and MacBeth, 1973 regressions). Large discrepancies in these quantities would indicate that the estimated risk premia could reflect model instability. Based on time-series regressions, the monthly risk premium on the DS factor is about 30-40 basis points in the full sample and 80-110 basis points in the crisis sample. The corresponding estimates from cross-sectional regressions are about 40-50 basis points and 100-110 basis points, respectively. The monthly average return of the DS factor is 39 basis points (statistically significant at the 10% level). The similarity of these figures indicate model stability and further dispels concerns of data mining.

5 Conclusions

Traditionally, one extracts risk neutral distributions using only option prices. For stock indexes, like the S&P 500, the number of actively traded options with out-of-the-money strikes is large enough that risk-neutral distributions are informative about both the intermediate part of the distribution and tail risk. For individual companies, however, only options with strike prices close to the current stock price are actively traded. We develop a new method to extract the risk-neutral distribution of firm-specific stock returns that combines data from options and credit default swaps (CDS). By construction, options/CDS-implied distributions reflect the default risk embedded in CDS spreads. In addition, options and CDS are, jointly but not in isolation, informative about the intermediate part of the distribution. As a result, combining options and CDS can yield a more informative skewness measure.

We evaluate our method with a series of asset pricing tests that revolve around a portfolio that buys (sells) stocks with large differences between options/CDS and options-only skewness. This portfolio is long (short) stocks exposed to negative (positive) intermediate-return risk that is not spanned by options and CDS individually. After controlling for a large set of asset pricing factors and stock characteristics, we find that the portfolio earns a risk premium of 0.5% per month, rising to about 1% during periods of financial stress when efficient extraction of information becomes particularly valuable. The results are robust to different assumptions about key inputs – including, crucially, the tenor of the CDS. However, we find that it is important to choose default thresholds that reflect the size and leverage of a company.

References

Acharya, V., and T. Johnson, 2007, "Insider trading in credit derivatives," *Journal of Financial Economics*, 84(1), 110–141.

Aït-Sahalia, Y., and A. Lo, 1998, "Nonparametric estimation of state-price densities implicit in financial asset prices," *Journal of Finance*, 53(2), 499–547.

Alanis, E., S. Chava, and P. Kumar, 2018, "Shareholder bargaining power, debt overhang, and investment," *Review of Corporate Finance Studies*, 7(2), 276–318.

Amaya, D., P. Christoffersen, K. Jacobs, and A. Vasquez, 2015, "Does realized skewness predict the cross-section of equity returns?," *Journal of Financial Economics*, 118(1), 135–167.

An, B.-J., A. Ang, T. Bali, and N. Cakici, 2014, "The joint cross section of stocks and options," *Journal of Finance*, 69(5), 2279–2337.

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

Arora, N., P. Ghandi, and F. Longstaff, 2012, "Counterparty credit risk and the credit default swap market," *Journal of Financial Economics*, 103(2), 280–293.

Azzalini, A., 2013, *The skew-normal distribution: probability,*Cambridge University Press, pp. 24–56.

Bakshi, G., N. Kapadia, and D. Madan, 2003, "Stock return characteristics, skew laws, and the differential pricing of individual equity options," *Review of Financial Studies*, 16(1), 101–143.

Bakshi, G., and D. Madan, 2000, "Spanning and derivative-security valuation," *Journal of Financial Economics*, 55(2), 205–238.

Bali, T., N. Cakici, and R. Whitelaw, 2014, "Hybrid tail risk and expected stock returns: When does the tail wag the dog?," *Review of Asset Pricing Studies*, 4(2), 206–246.

Bali, T., and A. Hovakimian, 2009, "Volatility spreads and expected stock returns," *Management Science*, 55(11), 1797–1812.

Bali, T., and S. Murray, 2013, "Does risk-neutral skewness predict the cross-section of equity option portfolio returns?," *Journal of Financial and Quantitative Analysis*, 48(4), 1145–1171.

Barroso, P., and P. Santa-Clara, 2015, "Momentum has its moments," *Journal of Financial Economics*, 116(1), 111–120.

Bates, D., 1991, "The crash of '87: was it expected? The evidence from options markets," *Journal of Finance*, 46(3), 1009–1044.

——— , 2000, “Post-’87 crash fears in the S&P 500 futures option market,” *Journal of Financial Economics*, 94(1–2), 181–238.

Bliss, R., and N. Panigirtzoglou, 2002, “Testing the stability of implied probability density functions,” *Journal of Banking & Finance*, 26(2-3), 381–422.

——— , 2004, “Option-implied risk aversion estimates,” *Journal of Finance*, 59(1), 407–446.

Bollerslev, T., G. Tauchen, and H. Zhou, 2009, “Expected stock returns and variance risk premia,” *Review of Financial Studies*, 22(11), 4463–4492.

Breeden, D., and R. Litzenberger, 1978, “Prices of state-contingent claims implicit in option prices,” *Journal of Business*, 51(4), 621–651.

Broadie, M., M. Chernov, and M. Johannes, 2007, “Model specification and risk premia: Evidence from futures options,” *Journal of Finance*, 62(3), 1453–1490.

Carey, M., and M. Gordy, 2016, “The bank as grim reaper: Debt composition and recoveries on defaulted debt,” *Finance and Economics Discussion Series, Federal Reserve Board*.

Carhart, M., 1997, “On persistence in mutual fund performance,” *Journal of Finance*, 52(1), 57–82.

Carr, P., and L. Wu, 2010, “Stock options and credit default swaps: A joint framework for valuation and estimation,” *Journal of Financial Econometrics*, 8(4), 409–449.

——— , 2011, “A simple robust link between American puts and credit protection,” *Review of Financial Studies*, 24(2), 473–505.

Chabi-Yo, F., D. P. Leisen, and E. Renault, 2014, “Aggregation of preferences for skewed asset returns,” *Journal of Economic Theory*, 154, 453 – 489.

Chan, L., N. Jegadeesh, and J. Lakonishok, 1996, “Momentum strategies,” *Journal of Finance*, 51(5), 1681–1713.

Chen, K., M. Fleming, J. Jackson, A. Li, and A. Sarkar, 2011, “An analysis of CDS transactions: Implications for public reporting,” *FRB of New York Staff Report No. 517*.

Colacito, R., E. Ghysels, J. Meng, and W. Siwasarit, 2016, “Skewness in expected macro fundamentals and the predictability of equity returns: evidence and theory,” *Review of Financial Studies*, 29(8), 2069–2109.

Conrad, J., R. Dittmar, and E. Ghysels, 2013, “Ex ante skewness and expected stock returns,” *Journal of Finance*, 68(1), 85–124.

Cox, J., S. Ross, and M. Rubinstein, 1979, “Option pricing: A simplified approach,” *Journal of Financial Economics*, 7(3), 229–263.

Daniel, K., and T. Moskowitz, 2016, “Momentum crashes,” *Journal of Financial Economics*, 122(2), 221–247.

Daniel, K., and S. Titman, 1997, "Evidence on the characteristics of cross sectional variation in stock returns," *Journal of Finance*, 52(1), 1–33.

Daniel, K., S. Titman, and K. Wei, 2001, "Explaining the cross-section of stock returns in Japan: Factors or characteristics," *Journal of Finance*, 56(2), 743–766.

D'Avolio, G., 2002, "The market for borrowing stock," *Journal of Financial Economics*, 66(2–3), 271–306.

DeMiguel, V., Y. Plyakha, R. Uppal, and G. Vilkov, 2013, "Improving portfolio selection using option-implied volatility and skewness," *Journal of Financial and Quantitative Analysis*, 48(6), 1813–1845.

Dennis, P., and S. Mayhew, 2002, "Risk-neutral skewness: Evidence from stock options," *Journal of Financial and Quantitative Analysis*, 37(3), 471–493.

Drechsler, I., and Q. Drechsler, 2016, "The shorting premium and asset pricing anomalies," *Working paper.*

Duffie, D. Singleton, K., 2003, *Credit risk: Pricing, measurement, and management.* Princeton University Press, Princeton, N.J., 1 edn.

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

——— , 1996, "Multifactor explanations of asset pricing anomalies," *Journal of Finance*, 51(1), 55–84.

——— , 2015, "A five-factor asset pricing model," *Journal of Financial Economics*, 116(1), 1–22.

Fama, E., and J. MacBeth, 1973, "Risk, return, and equilibrium: Empirical tests," *Journal of Political Economy*, 81(3), 607–636.

Fan, H., and S. Sundaresan, 2000, "Debt valuation, renegotiation, and optimal dividend policy," *Review of Financial Studies*, 13(4), 1057–1099.

Feunou, B., M. Jahan-Parvar, and C. Okou, 2018, "Downside variance risk premium," *Journal of Financial Econometrics*, 16(3), 341–383.

Feunou, B., M. Jahan-Parvar, and R. Tédongap, 2013, "Modeling market downside volatility," *Review of Finance*, 17(1), 443–481.

——— , 2016, "Which parametric model for conditional skewness?," *European Journal of Finance*, 22(13), 1237–1271.

Figlewski, S., 2010, "Estimating the implied risk-neutral density for the US market portfolio," *in "Volatility and time series econometrics: essays in honor of Robert Engle", edited by T. Bollerslev, J. Russell, and M.Watson. Oxford University Press.*

Friewald, N., C. Wagner, and J. Zechner, 2014, "The cross-section of credit risk premia and equity returns," *Journal of Finance*, 69(6), 2419–2469.

Garlappi, L., and H. Yan, 2011, "Financial distress and the cross-section of equity returns," *Journal of Finance*, 66(3), 789–822.

Ghysels, E., A. Plazzi, and R. Valkanov, 2016, "Why invest in emerging markets? The role of conditional return asymmetry," *Journal of Finance*, 71(5), 2145–2192.

Han, B., and Y. Zhou, 2011, "Term structure of credit default spreads and cross-section of stock returns," *Working paper*.

Hansen, B., 1994, "Autoregressive conditional density estimation," *International Economic Review*, 35(3), 705–730.

Harvey, C., Y. Liu, and H. Zhu, 2016, "... and the cross-section of expected returns," *Review of Financial Studies*, 29(1), 5–68.

Harvey, C., and A. Siddique, 2000, "Conditional skewness in asset pricing tests," *Journal of Finance*, 55, 1263–1295.

Jackwerth, J., and M. Rubinstein, 1996, "Recovering probability distributions from option prices," *Journal of Finance*, 51(5), 1611–1631.

Jensen, C., D. Lando, and L. Pedersen, 2017, "Generalized recovery," *Journal of Financial Economics*, Forthcoming.

Jondeau, E., and M. Rockinger, 2003, "Conditional volatility, skewness, and kurtosis: Existence, persistence, and comovements," *Journal of Economic Dynamics & Control*, 27(10), 1699–1737.

JPMorgan, 2001, "Par credit default swap spread approximation from default probabilities," *JPMorgan Securities Inc. Report*.

Kapadia, N., and M. Zekhnini, 2019, "Do idiosyncratic jumps matter?," *Journal of Financial Economics*, 131(3), 666 – 692.

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

Kim, T., and H. White, 2004, "On more robust estimation of skewness and kurtosis," *Finance Research Letters*, 1(1), 56–73.

Lewellen, J., S. Nagel, and J. Shanken, 2010, "A skeptical appraisal of asset pricing tests," *Journal of Financial Economics*, 96(2), 175–194.

Longstaff, F., 1995, "Option pricing and the martingale restriction," *Review of Financial Studies*, 8(4), 1091–1124.

Madan, D., and F. Milne, 1994, "Contingent claims valued and hedged by pricing and investing in a basis," *Mathematical Finance*, 4(3), 223–245.

Neuberger, A., 2012, "Realized skewness," *Review of Financial Studies*, 25(11), 3423–3455.

Ni, S., and J. Pan, 2011, "Trading puts and CDS on stocks with short sale ban," *Working paper.*

Novy-Marx, R., 2012, "Is momentum really momentum?," *Journal of Financial Economics*, 103(3), 429–453.

Oh, D., and A. Patton, 2017, "Modeling dependence in high dimensions with factor copulas," *Journal of Business & Economic Statistics*, 35(1), 139–154.

Pan, J., and A. Poteshman, 2006, "The information in option volume for future stock prices," *Review of Financial Studies*, 19(3), 871–908.

Pastor, L., and R. Stambaugh, 2003, "Liquidity risk and expected stock returns," *Journal of Political Economy*, 111(3), 642–685.

Patton, A., 2004, "On the out-of-sample importance of skewness and asymmetric dependence for asset allocation," *Journal of Financial Econometrics*, 2(1), 130–168.

Qiu, J., and F. Yu, 2012, "Endogenous liquidity in credit derivatives," *Journal of Financial Economics*, 103(3), 611–631.

Rehman, Z., and G. Vilkov, 2012, "Risk-neutral skewness: Return predictability and its sources," *Working paper, Frankfurt University.*

Rodriguez, R. N., 1977, "A Guide to the Burr Type XII Distributions," *Biometrika*, 64(1), 129–134.

Rosenberg, J., and R. Engle, 2002, "Empirical pricing kernels," *Journal of Financial Economics*, 64(3), 341–372.

Ross, S., 2015, "The recovery theorem," *Journal of Finance*, 70(2), 615–648.

Rubinstein, M., 1994, "Implied binomial trees," *Journal of Finance*, 49(3), 771–818.

Santa-Clara, P., and A. Saretto, 2009, "Option strategies: Good deals and margin calls," *Journal of Financial Markets*, 12(3), 391–417.

Shanken, J., 1992, "On the estimation of beta-pricing models," *Review of Financial Studies*, 5(1), 1–33.

Shimko, D., 1993, "Bounds of probability," *Risk*, 6, 33–37.

Stilger, P., A. Kostakis, and S. Poon, 2017, "What does risk-neutral skewness tell us about future stock returns?," *Management Science*, 63(6), 1814–1834.

Subrahmanyam, M., D. Tang, and S. Wang, 2014, "Does the tail wag the dog?': The effect of credit default swaps on credit risk," *Review of Financial Studies*, 27(10), 2927–2960.

White, H., 1980, "A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity," *Econometrica*, 48(4), 817–838.

Xing, Y., X. Zhang, and R. Zhao, 2010, "What does the individual option volatility smirk tell us about future equity returns?," *Journal of Financial and Quantitative Analysis*, 45(3), 641–662.

Zhang, B., H. Zhou, and H. Zhu, 2009, "Explaining credit default swap spreads with the equity volatility and jump risks of individual firms," *Review of Financial Studies*, 22(12), 5099–5131.

Appendix A: Differences Between Computing Risk-Neutral Distributions for Indexes and for Individual Firms

The companies in our sample are relatively large, yet the strike coverage and traded volume of options written on their stocks are pervasively low. The summary statistics in Table A1 highlight the differences in the availability of strike prices between options on the S&P 500 index and those on the companies we study. We define a call (put) option as out-of-the-money (OTM) if the strike is above (below) the stock price. A put option is considered deep OTM if its strike price is less than 80% of the stock price.

The top panel of Table A1 shows that, for the S&P 500, OTM options trade more often than their in-the-money (ITM) counterparts, with the average number of OTM options equal to 14.69 (puts) and 12.02 (calls) and the corresponding averages for ITM options equal to 4.29 (puts) and 5.40 (calls). The average daily number of deep OTM put options is about 8, or roughly half the number for OTM put options. Turning to individual-stock options in the bottom panel, the average number of option prices is considerably lower across both moneyness and call/put types, and there is little difference in availability between ITM and OTM options. The average number of deep OTM put options is roughly 2, and the median equals 1.

The two rightmost columns of Table A1 report summary statistics for option volume along moneyness and call/put types. We observe the same patterns that we discussed for the number of available options. Trading volumes are one order of magnitude smaller for individual-stock options compared to index options, and the difference is even more pronounced for OTM puts. Overall, the evidence in Table A1 supports the hypothesis that CDS can provide useful information about the left tail of the risk-neutral distribution of individual stocks. In Section 4, we find results in favor of this hypothesis.

To illustrate the contribution of options and CDS to firm-specific RNDs, Table A2 outlines the typical availability and informativeness about default risk of CDS and options with different strikes. Here, the current stock price is \$100 and default happens when the stock price drops to \$20. Options with a strike price close to the underlying price (strikes equal to \$110 and \$90) are actively traded but they are not informative about the probability of default. Instead, these options speak to the probability of moderate price changes ($\pm$10%), and do not provide information about large price drops that are associated with company default. Single-name options with a strike price significantly below the current stock price and close to the default threshold (\$20, in this case) rarely trade. Thus, while potentially informative, they are not readily available. CDS, on the other hand, have an implicit strike price that is equal to the default threshold and they are actively traded, hence they meet both requirements.

Table A1: Summary statistics on option-data availability

The table reports the average and median number of daily option observations, and the associated volumes, after applying the filters discussed in Section 2. The statistics are reported separately for S&P 500 index options and for options on the 275 companies for which we estimate options/CDS risk-neutral distributions. The statistics are shown by call and put option types (C and P, respectively), and for different moneyness levels. A call (put) option is out-of-the-money (OTM) if the strike price is above (below) the stock price. A put option is deep OTM if the strike price is less than 80% of the stock price. The sample period is 2006 to 2015.

		Options on the S&P 500 index			
Call/Put	Moneyness	Average obs.	Median obs.	Av. volume	Med. volume
C	ITM	5.40	4	6,541	788
C	OTM	12.02	11	10,772	5,825
P	ITM	4.29	3	5,023	709
P	OTM	14.69	13	19,702	11,942
P	Deep OTM	8.11	6	4,074	1,312

		Options on the firms we study			
Call/Put	Moneyness	Average obs.	Median obs.	Av. volume	Med. volume
C	ITM	2.91	2	393	66
C	OTM	3.24	3	701	173
P	ITM	2.21	2	256	48
P	OTM	2.51	2	427	92
P	Deep OTM	1.82	1	219	30

Table A2: Illustration of the default-risk informativeness and trading liquidity of options and CDS

The table illustrates the informativeness about default risk and the typical trading liquidity of a CDS and of a set of call and put options with different strike prices. The assumption is that the current stock price is \$100, and that the company defaults if the stock price drops to \$20. A security is useful for characterizing the deep left tail of a risk-neutral distribution of returns only if it is informative about default risk and it is actively traded.

Security	Option								CDS
Option type	Call				Put				
Option strike	110	100	90	20	110	100	90	20	
Typically traded?	yes	yes	yes	**no**	yes	yes	yes	**no**	**yes**
Informative about the deep left tail?	no	no	no	**yes**	no	no	no	**yes**	**yes**

Appendix B: Computational Details

Option-implied return probabilities

Our method for extracting the risk-neutral distribution of stock returns builds on the established literature that uses the first derivative of the European call pricing function, relative to the strike price, to approximate the CDF of stock returns. We follow Figlewski (2010) when calculating the option-implied cumulative probabilities at the traded strikes.

For each company/day combination we have a cross section of call and put options with the same maturity. If both a call and a put option are available for the same strike price i, we compute the volume-weighted implied volatility as follows:

$$IV_i = \frac{v_{i,C} \cdot IV_{i,C} + v_{i,P} \cdot IV_{i,P}}{v_{i,C} + v_{i,P}}, \tag{B1}$$

where $v_{i,C}$ and $v_{i,P}$ are the call and put volumes, and $IV_{i,C}$ and $IV_{i,P}$ are the call and put implied volatilities. The volume associated with the weighted implied volatility is:

$$v_i^{av} = \frac{v_{i,C}^2 + v_{i,P}^2}{v_{i,C} + v_{i,P}}. \tag{B2}$$

Compared with a simple average, this formula ensures that the relative liquidity of the call and put options is appropriately taken into account.[8]

In line with the literature, we refer to implied volatilities as a function of strike prices as the volatility smile. Following Bliss and Panigirtzoglou (2002), we interpolate the volatility smile with a natural smoothing cubic spline before converting the implied volatilities to Black-Scholes call prices. We use weights based on log-volume in the interpolation. Then we obtain the option-implied cumulative probabilities by taking finite differences of the interpolated Black-Scholes European price function. We approximate the first derivative using points 0.1% to the left and right of the traded strikes. Specifically, we compute

$$CP_{si} = e^{0.25r} \frac{C_{i+1} - C_{i-1}}{X_{i+1} - X_{i-1}} + 1$$

where C_i and X_i are the call price and strike for the traded option i, the $i+1$ and $i-1$ values are the points 0.1% to the left and right, and r is the risk-free rate. Since the strikes of the interpolated implied volatilities are very close to the strikes of the traded options, the effect of alternative interpolation methods on our results is minimal. The trade off for achieving this robustness is that

[8] For instance, if the call-option volume is 100 and the put-option volume is 1, the volume-weighted implied volatility mostly reflects the information contained in the call implied volatility. Taking the average of the option volumes would yield 50.5, while Equation B2 yields 99.02, which better represents the liquidity of the option that drives the average implied volatility.

we impose a parametric specification on the risk-neutral distribution.

Note that the risk-neutral cumulative density function (CDF) of returns is the first derivative of the *European* price function, but exchange-traded options on individual companies are *American* options. We convert American prices into their European equivalent with three-month maturity by calculating the Black-Scholes price based on the implied volatility provided by OptionMetrics, which is computed according to a Cox, Ross, and Rubinstein (1979) binomial tree and does not incorporate the early exercise premium. In this respect, we follow Broadie, Chernov, and Johannes (2007).

CDS-implied default probability

We compute CDS-implied default probability by equating the expected value of the payments received by the protection seller (fee leg) and the expected value of the losses incurred by the protection seller (contingent leg). As detailed below, we observe the CDS spread and all variables that determine the value of the fee and contingent legs, with the exception of the hazard rate. This rate is computed numerically by equating the CDS spread to the ratio of the contingent and fee legs.

The value of the fee leg of a CDS with maturity T_N and payment dates $\{T_i\}_{i=1}^N$ can be expressed as a function of the spread, s, the hazard rate, λ, the risk-free discount rate, r_i^f, and of the time between T_{i-1} and T_i (Δ_i):[9]

$$\begin{aligned} V_f(\lambda, T_N) &= s \cdot \Sigma_{i=1}^N \left\{ \Delta_i e^{-\lambda T_{i-1}} \left[e^{-\lambda \Delta_i} + \left(1 - e^{-\lambda \Delta_i}\right) \frac{\lambda^{-1} - e^{-\lambda \Delta_i}\left(\Delta_i + \lambda^{-1}\right)}{1 - e^{-\lambda \Delta_i}} \right] e^{-r_i^f T_i} \right\} \\ &= s \cdot \Sigma_{i=1}^N \left\{ \Delta_i e^{-\lambda T_{i-1}} \left[e^{-\lambda \Delta_i} + \lambda^{-1} - e^{-\lambda \Delta_i}\left(\Delta_i + \lambda^{-1}\right) \right] e^{-r_i^f T_i} \right\} \qquad \text{(B3)} \end{aligned}$$

The first exponential in curly brackets is the survival probability up to time T_{i-1}. The expression in square brackets gives the expected fee between time T_{i-1} and time T_i: the firm survives one more period, in which case the full fee is collected; or the firm defaults, so that only the accrued premium is collected. The accrued premium is given by the spread times the expected time of default, conditional on default taking place in the interval Δ_i, as captured by the fraction next to the closing square bracket. The last term in the formula is the discount factor.

[9] Our specification builds on Duffie (2003) and JPMorgan (2001). See the ISDA "Standard North American Corporate CDS Contract Specification" for details on the pricing and timing conventions for corporate CDS.

The expected time of default, conditional on default taking place in the interval Δ_i, is:

$$
\begin{aligned}
E\left[x|0<x\leq\Delta_i\right] &= \int_0^{\Delta_i} x\frac{f(x)}{\int_0^{\Delta_i} f(x)\mathrm{d}x}\mathrm{d}x = \int_0^{\Delta_i} x\frac{\lambda e^{-\lambda x}}{1-e^{-\lambda\Delta_i}}\mathrm{d}x \\
&= \frac{1}{1-e^{-\lambda\Delta_i}}\lambda\left[\left.\frac{-xe^{-\lambda x}}{\lambda}\right|_0^{\Delta_i} + \left.\frac{-e^{-\lambda x}}{\lambda^2}\right|_0^{\Delta_i}\right] \\
&= \frac{\lambda^{-1}-e^{-\lambda\Delta_i}(\Delta_i+\lambda^{-1})}{1-e^{-\lambda\Delta_i}} \qquad \text{(B4)}
\end{aligned}
$$

Note that we make the expected time of default a function of the hazard rate, rather than assuming it takes place in the middle of the quarter, because companies with high default risk are likely to default earlier. In expectation, the final CDS payment should be discounted for a shorter period of time when default risk is higher.

The value of the contingent leg is:

$$
V_c(\lambda,T_N) = L\cdot\Sigma_{i=1}^N\left[e^{-r_i^f T_i}\left(e^{-\lambda T_{i-1}}-e^{-\lambda T_i}\right)\right] \qquad \text{(B5)}
$$

The par spread sets the value of the contract equal to zero at initiation, which means that the fee and the default legs have the same value. It follows that:

$$
s = \frac{L\cdot\Sigma_{i=1}^N\left[e^{-r_i^f T_i}\left(e^{-\lambda T_{i-1}}-e^{-\lambda T_i}\right)\right]}{\Sigma_{i=1}^N\left\{\Delta_i e^{-\lambda T_{i-1}}\left[e^{-\lambda\Delta_i}+\lambda^{-1}-e^{-\lambda\Delta_i}\left(\Delta_i+\lambda^{-1}\right)\right]e^{-r_i^f T_i}\right\}} \qquad \text{(B6)}
$$

Assuming that the loss given default L is equal to 1 minus the recovery rate provided by Markit, all the variables in the equation except λ are observable. We then recover the hazard rate as the value of λ which verifies equation B6. Note that Markit's recovery rates are, for our sample of companies, virtually constant at 40%, which is an often-used industry convention. Of the company/date combinations for which we extract risk-neutral distributions, 89.17% have a recovery rate of 40%, 8.38% between 35% and 40%, and 0.61% between 40% and 45%. Only in 1.84% of the cases is the recovery rate outside of this narrow range.

In equation (B3), the variable λ is the company's hazard rate expressed on an annualized basis. In the main estimation, we obtain this annualized hazard rate from 5-year credit default swaps. Our risk-neutral densities have a 3-month horizon, which means we need to calculate the CDS-implied default probability over 3-months, which can be obtained as: $CP_{cds} = 1-e^{-\lambda\cdot 0.25}$.

Estimating the parameters of the distribution

The probability density function of the skew-t distribution is given by:

$$f(z^j) = \begin{cases} bc\left(1 + \frac{1}{\eta-2}\left(\frac{bz^j+a}{1-\lambda}\right)^2\right)^{-(\eta+1)/2} & \text{if } z^j < -a/b \\ bc\left(1 + \frac{1}{\eta-2}\left(\frac{bz^j+a}{1+\lambda}\right)^2\right)^{-(\eta+1)/2} & \text{if } z^j \geq -a/b \end{cases} \tag{B7}$$

where λ ($|\lambda| < 1$) is the shape parameter, η ($2 < \eta < \infty$) are the degrees of freedom, and $z^j = \frac{r^j - r^f}{\sigma^j}$ is the standardized return for company j with expected return r^f and volatility σ^j. In addition, $a = 4\lambda c\frac{\eta-2}{\eta-1}$, $b = \sqrt{1+3\lambda^2 - a^2}$, $c = \frac{\Gamma((\eta+1)/2)}{\Gamma(\eta/2)\sqrt{\pi(\eta-2)}}$. Setting the expected return equal to the risk-free rate r^f for all companies enforces the martingale restriction under the risk-neutral measure.

The skew-t CDF with parameter set $\Phi = \{\mu, \sigma, \lambda, \eta\}$ under the risk-neutral measure is represented by $F_{sk}(.;\Phi)$:

$$F_{sk}(z^j, \Phi) = \begin{cases} (1-\lambda)\, t_{cdf}\left(\frac{b\frac{z^j-\mu}{\sigma}+a}{1-\lambda}\sqrt{\frac{\eta}{\eta-2}}, \eta\right) & \text{if } z^j < -a/b \\ \frac{1-\lambda}{2} + (1+\lambda)\left(t_{cdf}\left(\frac{b\frac{z^j-\mu}{\sigma}+a}{1+\lambda}\sqrt{\eta/\left(\eta-2\right)}, \eta\right) - 0.5\right) & \text{if } z^j \geq -a/b \end{cases} \tag{B8}$$

where $t_{cdf}\left(\cdot, \eta\right)$ is the cumulative standard symmetric Student-t distribution.

The CDS-implied cumulative probability up to the default threshold is CP_{cds}, and the option-implied cumulative probability up to strike s_i is CP_{s_i}. For each company/day combination, we estimate the parameters of a skew-t distribution by minimizing the squared deviations between the skew-t CDF and the cumulative probabilities extracted from options and CDS:

$$\hat{\Phi} = \underset{\Phi}{\operatorname{argmin}} \left\{ w_{cds} \cdot \left[F_{sk}\left(R_D, \Phi\right) - CP_{cds}\right]^2 + \Sigma_{i=1}^{N} w_{s_i}\left[F_{sk}\left(s_i, \Phi\right) - CP_{s_i}\right]^2 \right\}, \tag{B9}$$

where w_{cds} is the weight assigned to the CDS-implied cumulative probability and w_{s_i} is the weight assigned to the cumulative probability corresponding to strike i, and R_D is the default threshold corresponding to the latest size and leverage values for the given company. If no option trades with a strike below R_D, each observation is equally weighted. If options trade with strikes below R_D, w_{cds} is equal to the volume of these options divided by the total option volume, and w_{s_i} are equal weights scaled so that $w_{cds} + \Sigma_{i=1}^{N} w_{s_i}$ sums to one. We do not use options with strikes below R_D in computing risk-neutral densities because it is very unusual for single-name options to trade this far from at-the-money, and prices are likely to contain a considerable level of noise. For symmetry, we also discard options with strikes that imply a return above the opposite of R_D.

We recover the skew-t parameters for each stock/day pair by first estimating 10 sets of parameters from randomized starting values. We then carry out a grid search centered around the parameter set $\tilde{\Phi} = \{\tilde{\sigma}, \tilde{\lambda}, \tilde{\eta}\}$ which yields the smallest squared deviations. The skew-t distribution is characterized by three parameters (as discussed above, the mean is set to the risk-free rate to satisfy the martingale restriction). We constrain the volatility, shape, and degrees of freedom parameters to be between [0.05,1.25], [-0.995,0.995], and [2.1,100], respectively. For the volatility and shape parameters, the grid search is focused on ± 0.05 around $\tilde{\sigma}$ and $\tilde{\lambda}$, in 0.005 steps. For the degrees of freedom, we focus on ± 0.1 around $\tilde{\eta}$, in 0.01 steps.

We repeat the above procedure (equation B9) using only option-implied cumulative probabilities, effectively setting the term $[F_{sk}(R_D, \Phi) - CP_{cds}]^2$ equal to zero. As a result, on a given day and for a given company, we recover one risk neutral distribution based on options and CDS, and another distribution based on options only.

Appendix C: Relationship Between CDS Inclusion and Skewness

In order to understand which moment of the risk-neutral distributions is most affected by the inclusion of CDS, we compare the moments of options/CDS and options-only risk-neutral distributions. We expect CDS to affect kurtosis and skewness more than volatility, since the latter is well characterized by options with strike prices around the current stock price. Kurtosis, on the other hand, measures the thickness of tails, which is directly linked to the default probability expressed by CDS spreads. Skewness measures the shape of a distribution, which is influenced by both the central part and the tails.

The skewness sk and kurtosis ku of the skew-t distribution are defined as follows (Feunou, Jahan-Parvar, and Tédongap, 2016):

$$\begin{cases} sk = (m_3 - 3am_2 + 2a^3)/b^3 \\ ku = (m_4 - 4am_3 + 6a^2m_2 - 3a^4)/b^4 \end{cases} \tag{C10}$$

where $a = 4\lambda c\frac{\eta-2}{\eta-1}$, $b = \sqrt{1+3\lambda^2 - a^2}$, $c = \frac{\Gamma((\eta+1)/2)}{\Gamma(\eta/2)\sqrt{\pi(\eta-2)}}$, $m_2 = 1+3\lambda^2$, $m_3 = 16c\lambda(1+\lambda^2)\frac{(\eta-2)^2}{(\eta-1)(\eta-3)}$ (for $\eta > 3$), and $m_4 = 3\frac{\eta-2}{\eta-4}(1+10\lambda^2+5\lambda^4)$ (for $\eta > 4$).

For each moment (volatility, skewness, and kurtosis), we construct the absolute difference between the options/CDS moment and the options-only moment, scaled by the absolute value of the options-only moment ($\frac{|mom_{opt/CDS} - mom_{opt}|}{|mom_{opt}|}$). Panel A of Table C1 shows the average and upper percentiles of the absolute relative differences' distribution. It is immediately clear that, first, the moment most affected by the inclusion of CDS is skewness, followed by kurtosis. Second, we observe that meaningful differences between the options/CDS and options-only skewness are concentrated above the 75^{th} percentile.

The information about default risk embedded in CDS spreads affects kurtosis directly, while it affects skewness indirectly and only in conjunction with options data. As a result, we expect CDS spreads to explain the difference between options/CDS and options-only kurtosis to a larger extent than they explain the difference between options/CDS and options-only skewness. We expect this to be the case even though skewness is the moment that changes the most with the inclusion of CDS, as shown in Panel A. In Panel B of Table C1, we report the results of regressing $sk_{opt/CDS} - sk_{opt}$ or $ku_{opt/CDS} - ku_{opt}$ on 5-year CDS log spreads, with standard errors clustered by date since the data are pooled and innovations to CDS spreads could be correlated in the cross-section. The regression results point to stronger direct influence of CDS on kurtosis than on skewness. When focusing on non-zero differences, the slope parameter for CDS spreads is statistically significant only for kurtosis. This coefficient becomes statistically significant for skewness only when the differences become increasingly larger. The magnitude and statistical significance of the slope coefficient, as well as the R^2s, are always larger for kurtosis than for skewness.

In light of the evidence presented in Panels A and B, one can conclude that the *joint* consideration of options and CDS influences skewness more than kurtosis for two reasons. First, including CDS has a stronger effect on skewness than kurtosis. Second, less of the difference between options/CDS and options-only moments is explained by CDS spreads for skewness than for kurtosis, suggesting that skewness differences reflect more than just the information incorporated with the inclusion of CDS. For this reason, we focus on skewness in our empirical analysis.

Table C1: Differences between options/CDS and options-only moments

In panel A, the table shows selected summary statistics for the relative absolute difference, in percentage points, between options/CDS and options-only moments (volatility, skewness, and kurtosis). Panel B reports the coefficients, t-statistics (with standard errors clustered by date), and R^2s of regressions of skewness differences and kurtosis differences (options/CDS minus options-only) on 5-year log-CDS spreads. The regressions are repeated on three samples: observations with non-zero differences, observations where the absolute differences are greater than 0.01, and observations where the absolute differences are greater than 0.1. The sample includes the last three available observations of each month for each company, from 2006 to 2015.

Panel A

	Distribution of absolute relative difference (in %) between options/CDS and options-only moments ($\frac{\lvert mom_{opt/CDS} - mom_{opt}\rvert}{\lvert mom_{opt}\rvert}$)		
	Volatility	Skewness	Kurtosis
Average	0.31	6.41	2.03
50%	0.00	0.07	0.06
75%	0.00	0.27	0.21
90%	0.00	1.79	1.06
95%	1.96	17.63	3.10
99%	6.67	119.00	66.54

Panel B

		Regression on $\ln(\text{sprd}_5 yr)$	
		Skewness diff.	Kurtosis diff.
Diff.	coeff.	-0.002	0.150
Non-zero	t-stat	(-0.18)	(12.64)
	R^2	0.00	0.50
|Diff.|	coeff.	0.245	0.675
>0.01	t-stat	(2.39)	(13.31)
	R^2	0.11	2.40
|Diff.|	coeff.	1.029	1.445
>0.1	t-stat	(3.37)	(11.90)
	R^2	0.70	4.18

Appendix D: Relationship Between Skewness Measures and Momentum

A possible reason for the options/CDS and options-only portfolios outperforming the BKM portfolio during 2009 is that the parametric nature of our method enables us to more efficiently extract information about future investment opportunities. Indeed, early 2009 is when the stock market bottomed out and a sustained recovery commenced. We can make an interesting comparison of the returns on the options/CDS and options-only portfolios with the momentum strategy. As discussed by Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016), and as shown in Figure 2, the momentum strategy experienced very negative returns through 2009. Momentum is a backward-looking strategy that performs poorly when confronted with sharp changes in the business cycle. At the end of a long recession, like the one that began in December 2007, the momentum strategy would buy stocks that have performed well (e.g., low-beta stocks) and sell stocks that have performed poorly (e.g., high-beta stocks). After a sharp change in the business cycle and in the stock market trend, and for a length of time that depends on the portfolio-formation period, the momentum strategy would keep buying (selling) stocks that have done relatively well during the recession but that will perform relatively poorly with the new upward trend (e.g., the strategy would keep buying low-beta stocks and selling high-beta stocks).

To the extent that our method extracts information about future investment opportunities efficiently, the returns on the portfolios based on the options/CDS and options-only skewness measures should be negatively correlated with the returns on the momentum strategy during the year 2009. The top panel in Figure D1 shows that such is the case. The returns on the options/CDS portfolio are clearly negatively correlated with momentum returns, and the most negative momentum return (April 2009) corresponds to the largest return on the options/CDS portfolio. The bottom panel of this figure shows that the correlation between momentum and BKM skewness portfolio returns is far less pronounced.

We formally test the relationship between the momentum strategy and options/CDS or the BKM skewness, by regressing momentum portfolio returns on options/CDS or the BKM portfolio returns, along with the same risk factors used to evaluate the relationship between the two skewness portfolio returns. Table D1 reports these results. As expected, the estimated slope parameters for options/CDS portfolio returns are statistically different from zero and negative-valued in the full sample and in the 2008-2011 sub-sample. Estimated intercepts for the regression models that include options/CDS portfolio returns as a factor are statistically indistinguishable from zero both for the full sample and the 2008-2011 sub-sample. Conversely, the estimated slope coefficients for the BKM portfolio returns are statistically insignificant regardless of the sample used. The estimated intercept parameter for the model containing the BKM portfolio returns as an independent variable is statistically different from zero at the 10% significance level in the 2008-2011 sub-sample.

These results support the visual evidence presented in Figure 2. The BKM portfolio returns are not significantly correlated with the momentum factor. Moreover, particularly during the global financial crisis and its immediate aftermath, BKM portfolio returns and other risk factors used in these tests do not fully span the momentum factor. On the other hand, the momentum factor is fully spanned, regardless of the sample used, by the options/CDS portfolio returns and the additional risk factors.

Since the philosophies behind the construction of the momentum and options/CDS skewness portfolios are completely different, these findings cannot be attributed to data mining – a serious concern in recent studies of factors aiming to explain the cross-sectional variations of stock returns (see Harvey, Liu, and Zhu, 2016).

Table D1: Risk-adjusted returns of the momentum factor

The table shows the results of regressing the UMD momentum factor of Carhart (1997) on selected risk factors, including either the $\text{SKEW}_{opt/CDS}$ or SKEW_{BKM} portfolio. MKT, SMB, HML, RMW, and CMA are the market, size, book-to-market, profitability, and investment factors of Fama and French (2015). LIQ is the Pastor and Stambaugh (2003) liquidity factor. ΔVIX_{repl} is a factor-mimicking portfolio for changes in the implied volatility index (VIX). ΔCDX_{HY} is changes in the spread of the CDX High-Yield CDS index. The sample covers 2006 to 2015.

	Dependent variable: UMD			
	2006-2015		2008-2011	
MKT	0.195	0.334	0.334	0.554
	(0.81)	(1.46)	(0.89)	(1.50)
SMB	0.201	0.369*	0.423	0.641
	(0.89)	(1.77)	(1.14)	(1.67)
HML	-0.702**	-0.454*	-0.895*	-0.650
	(-2.36)	(-1.72)	(-1.89)	(-1.62)
RMW	0.315	0.235	0.971	0.644
	(0.94)	(0.77)	(1.48)	(1.10)
CMA	0.252	-0.240	0.772	0.047
	(0.66)	(-0.63)	(1.26)	(0.06)
ΔVIX_{repl}	-0.132	-0.189	-0.082	-0.176
	(-0.95)	(-1.64)	(-0.53)	(-1.11)
ΔCDX_{HY}	0.028	0.024*	0.029	0.025
	(1.56)	(1.70)	(1.18)	(1.37)
SKEW_{BKM}	-0.072		-0.126	
	(-0.29)		(-0.33)	
$\text{SKEW}_{opt/CDS}$		-0.488***		-0.611**
		(-2.66)		(-2.05)
Intercept	-0.003	-0.002	-0.019*	-0.010
	(-0.61)	(-0.37)	(-1.92)	(-1.31)
Obs.	119	119	48	48
Adj.R^2	0.285	0.391	0.283	0.425

Fig. D1: Covariation of the momentum and skewness portfolios

The top chart is a scatter plot of the monthly returns on the UMD momentum portfolio and of the returns on the options/CDS skewness-based portfolio. In the bottom chart, skewness is calculated according to the Bakshi, Kapadia, and Madan (2003) method. In both charts, linear fit lines are shown in solid red. The vertical dashed lines center the plots on zero returns for the skewness portfolios. Solid blue markers correspond to 2009, with months as labels. The sample covers February 2006 to December 2015.

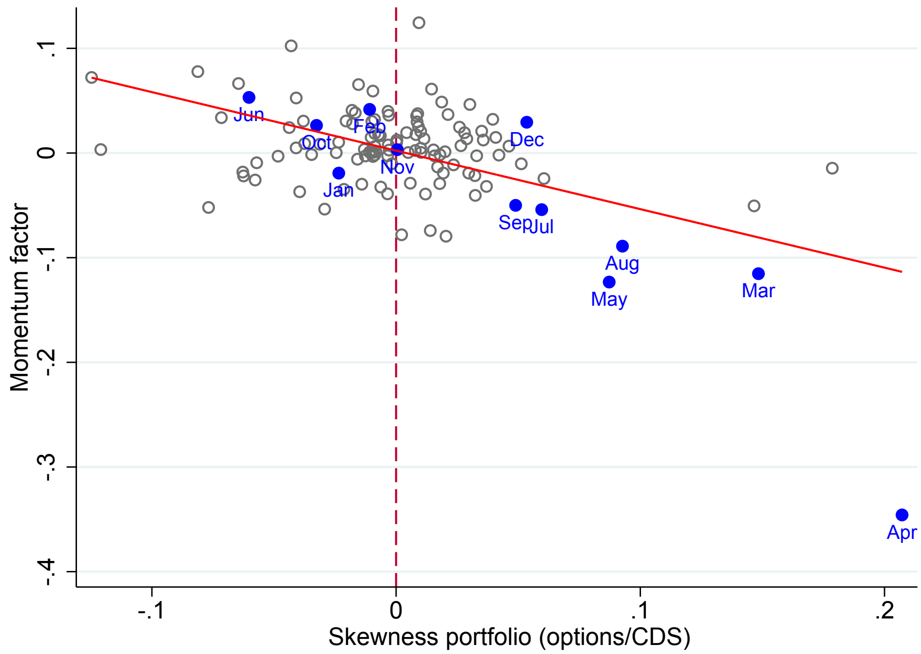

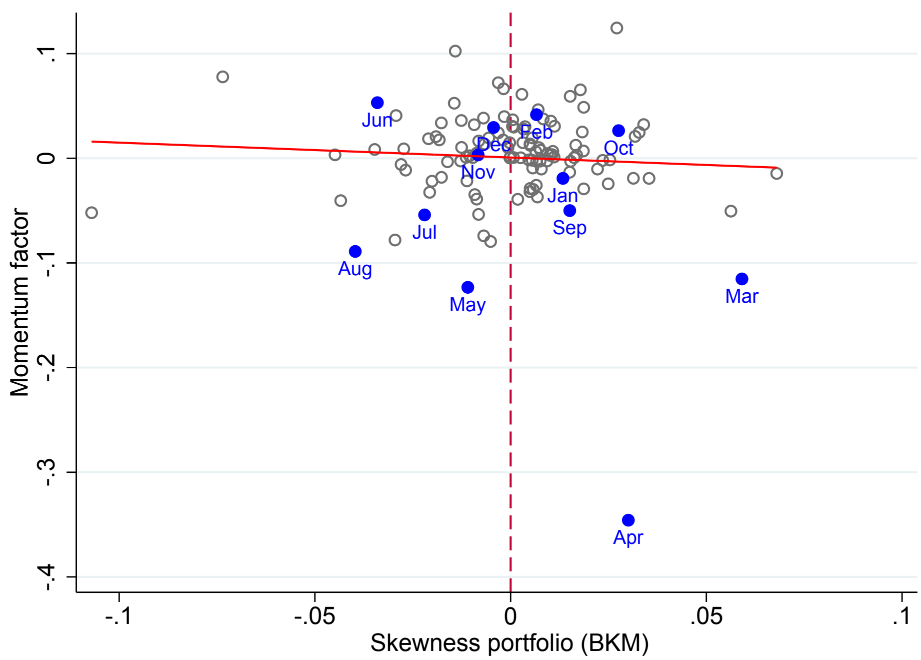

Appendix E: Insights from the Composition of the Difference-in-Skewness Decile Portfolios

In the body of the paper, we discuss differences between options/CDS and options-only skewness in Panel A of Table 3. We highlight that these differences become sizable only in the tails of the distribution of differences, roughly below the 10^{th} percentile and above the 99^{th} percentile. Here, we compare the distributions of options/CDS and options-only skewness, with the aim of understanding whether the two are statistically different, especially in the tails. We start by presenting quantile-quantile (QQ) plots in Figure E1. These plots provide an intuitive graphical comparison of the two distributions. If the distributions were identical, QQ plots would line up perfectly with a 45-degree line. The left chart of Figure E1 indicates that there are noticeable differences in the tails of the options/CDS and options-only skewness distributions, mainly in the left tail. In particular, options/CDS skewness tends to have a heavier left tail, with more negative values than options-only skewness for a given percentile. Consistent with Table 3, differences are less pronounced and typically positive in the right tail, which is thus heavier for options-only skewness.

We test whether tail differences between options/CDS and options-only skewness are statistically significant. We do so with Kolmogorov-Smirnov tests, which are known to be not particularly powerful against the null of no difference, hence our results are conservative. As shown in the right chart of Figure E1, the null of distribution equality is rejected when we consider observations in the left tail, roughly below the 20^{th} percentile, and when we consider observations in the far right tail, approximately above the 90^{th} percentile. These results formalize – and are fully consistent with – those reported in Table 3.

Fig. E1: Statistical properties of options/CDS and options-only skewness

The left chart shows quantile-quantile (QQ) plots for options/CDS and options-only skewness. In a given month, we calculate percentiles from 1 to 99 for the cross-sectional distributions of options/CDS skewness and options-only skewness. Each marker represents the value of options/CDS skewness (x-axis) and options-only skewness (y-axis) for a given percentile and in a given month. Monthly QQ plots from February 2006 to December 2015 are overlaid to to obtain the chart. In the right chart, we show p-values of Kolmogorov-Smirnov tests of the null that the options/CDS and options-only skewness distributions are equal. We consider observations of options/CDS and options-only skewness that are below or above the indicated percentiles of the respective distributions. Blue (red) markers show p-values of tests on observations *above* (*below*) the indicated percentiles.

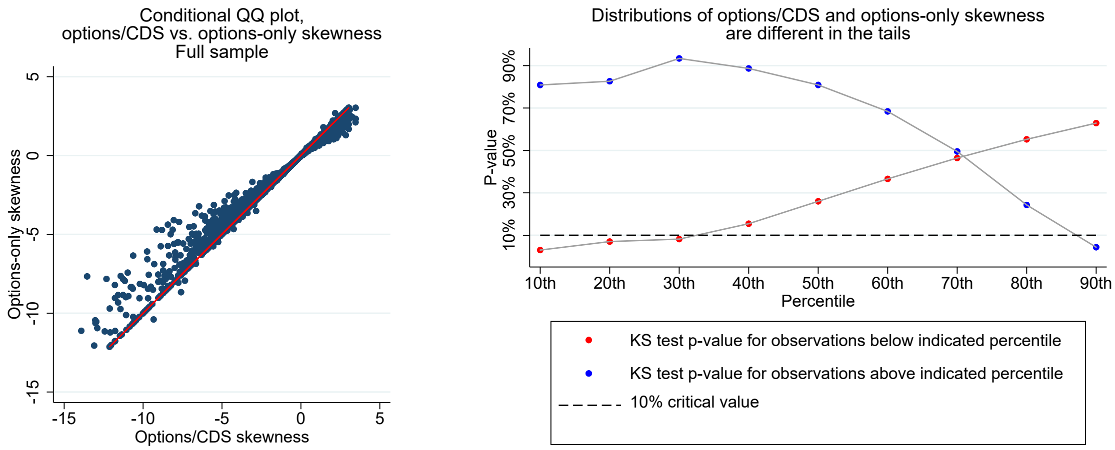

We now turn to investigating in more detail the link between differences in skewness and intermediate risk. To the extent that differences between options/CDS and options-only skewness reflect intermediate risk, a large difference (in absolute value) would foreshadow a higher incidence of intermediate returns in the near future. We compute the change in the probability of intermediate returns under the options/CDS distribution relative to the options-only distribution, and we present the average (top chart) and the 90^{th} percentile (bottom chart) of this change in probabilities across DS deciles.[10] The increase in intermediate-return probabilities should be larger in the top and bottom DS deciles. As shown in Figure E2, this is the case. The probability of intermediate returns is, on average, slightly more than 2% higher in the bottom DS decile, and about 1% higher in the top DS decile. At the 90^{th} percentile, the increase is more than 5% and about 3% in the bottom and top deciles, respectively.

[10] Intermediate returns are defined as those in the intervals $[-2\sigma, -0.5\sigma]$ and $[0.5\sigma, 2\sigma]$, where σ is the volatility of the distribution. The probability of intermediate returns under the options/CDS distribution is $P^{int}_{opt/CDS} = P[r < 2\sigma_{opt/CDS}] - P[r < 0.5\sigma_{opt/CDS}] + P[r < -0.5\sigma_{opt/CDS}] - P[r < -2\sigma_{opt/CDS}]$, where r indicates returns. The probability of intermediate returns under the options-only distribution (P^{int}_{opt}) is defined in a similar way, with σ_{opt} replacing $\sigma_{opt/CDS}$. The increase in the probability of intermediate returns under the options/CDS distribution is defined as $\Delta P = ln\frac{P^{int}_{opt/CDS}}{P^{int}_{opt}}$. When computing the increase in the probability of intermediate positive returns, the interval we consider is $[0.5\sigma, 2\sigma]$. For the increase in the probability of intermediate negative returns, the interval is $[-2\sigma, -0.5\sigma]$. DS deciles and probability increases are as of month t.

The sign of the difference between options/CDS and options-only skewness indicates whether intermediate returns are likely to be positive or negative. As illustrated in Figure E3, lower skewness translates into a shift of the probability mass to the right. Figure E2 illustrates the relation between the sign of the DS measure and the sign of realized intermediate returns by showing the increase in the probability of positive and negative intermediate returns separately. To the extent that stocks in the top (bottom) DS decile are exposed to negative (positive) intermediate-return risk that is not spanned by options and CDS individually, the increase in the probability of intermediate returns should be more pronounced for negative (positive) returns in the top (bottom) DS decile. This is what we find.

We provide additional details on the companies in the short and long legs of the DS factor. In Figure E4, we show two company characteristics, averaged by DS decile. DS is measured in month t, and the two characteristics are measured in the year preceding month t. These variables are realized daily volatility and average standardized earnings surprises (henceforth, SUE. See Chan, Jegadeesh, and Lakonishok, 1996, p. 1685). Companies with high DS, in absolute value, have higher volatility and relatively low profitability. In the first and tenth deciles, realized volatility is about 5 percentage points higher than in the other deciles. In line with higher volatility, profitability also indicates that these companies tend to grow less than the other firms in our sample.

Note that, while companies in the top and bottom DS deciles have noticeably higher volatility and lower SUE than other companies in the sample, the variation across DS deciles is relatively small when compared to full-sample variation. Specifically, realized volatility for high-DS companies is about 5 percentage points higher than it is for the other deciles, and is equal to about 37%. Before averaging by DS decile, the sample median and 75% percentile volatilities are about 28% and 40%, respectively. As for SUE, the averages in the first and tenth DS deciles are about 0 and -0.07. Before averaging by DS decile, the sample median and the 25% percentile are about 0.25 and -0.29. These figures confirm that companies in the top and bottom DS deciles are only moderately riskier than those in the remainder of the sample.

We now focus to the economic content of large DS observations. As shown in Table 3, noticeable gaps between options/CDS and options-only skewness measures are concentrated in the tails of the distribution. Understanding the drivers of these large differences is important, to ensure that they reflect economically meaningful information.

In Table E1, we focus on the company/day observations with the top five largest positive and top five largest negative differences between skewness measures. For each company and day, we report firm-specific events that occurred on the same or previous day as the large DS observation. We identify these events by searching business news and corporate filings. Based on earlier discussion, negative (positive) DS observations should correspond to positive (negative) news, because negative (positive) DS corresponds to a higher probability of positive (negative) intermediate returns. Indeed, we find such correspondence between differences in skewness and corporate events. The fact that we satisfy this theoretically-grounded constraint is further evidence that large DS observations reflect economically meaningful events. In two cases (Commercial Metals and Lexmark International) we were unable to find relevant news stories. However, the companies' stock prices moved by substantial amounts and in the expected direction based on the sign of the respective DS observation. Importantly, there is no mechanical relation between stock prices and skewness measures, because we do not use stock returns when extracting risk-neutral densities.

Fig. E2: Increase in the probability of intermediate returns under the options/CDS distribution, by DS decile

The figure shows the average (top chart) and 90^{th} percentile (bottom chart) increase in the probability of intermediate returns from the options-only distribution to the options/CDS distribution. See footnote 10 for the definition of intermediate returns and for the calculation of the increase in probability. DS deciles and probability increases are as of month t. The sample period is 2006 to 2015.

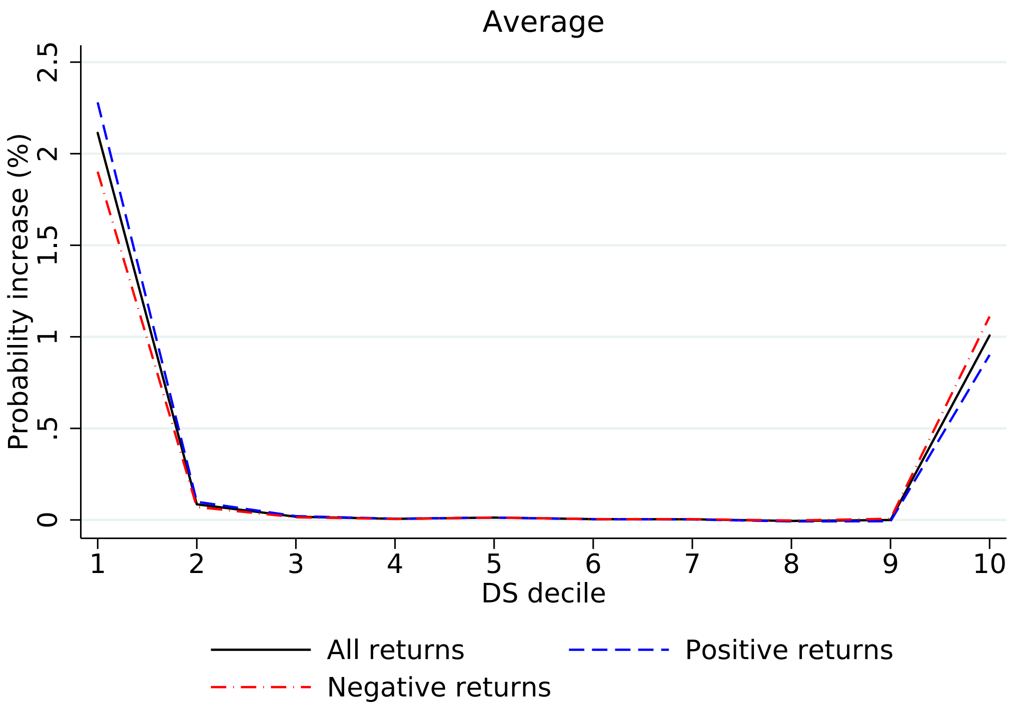

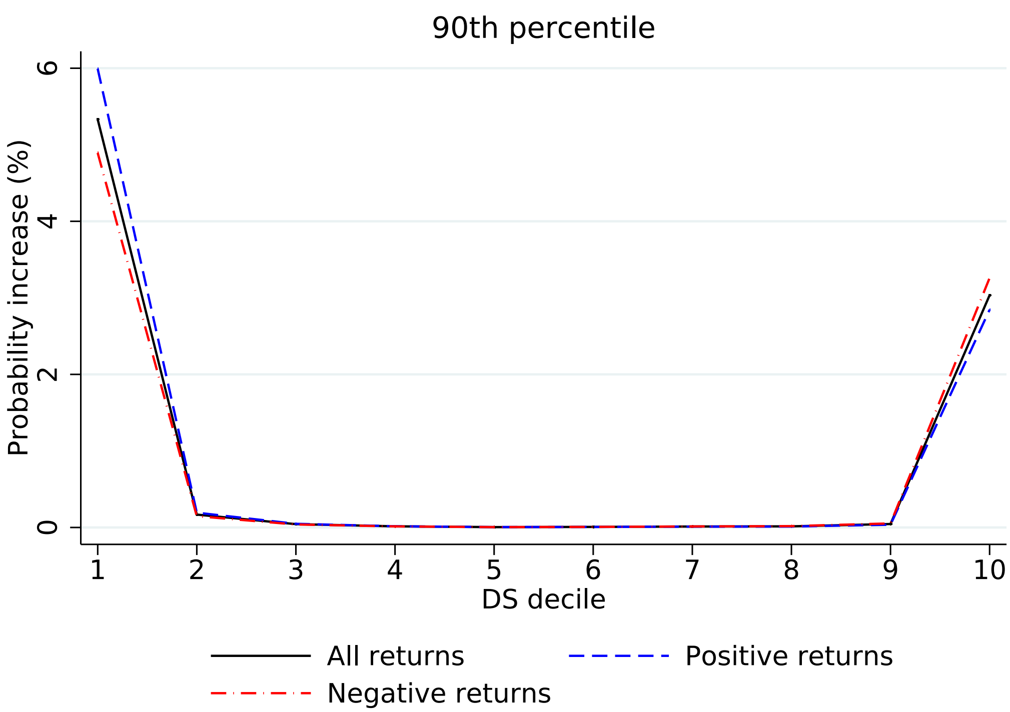

Fig. E3: Skewed-t distribution for different levels of skewness

The figure illustrates the shape of the skew-t distribution for different levels of skewness. Each distribution has a volatility of 15%, 5 degrees of freedom, and the shape parameter is equal to -0.5 (negative skewness, blue dash-dot line), 0 (zero skewness, black solid line), and 0.5 (positive skewness, red dashed line).

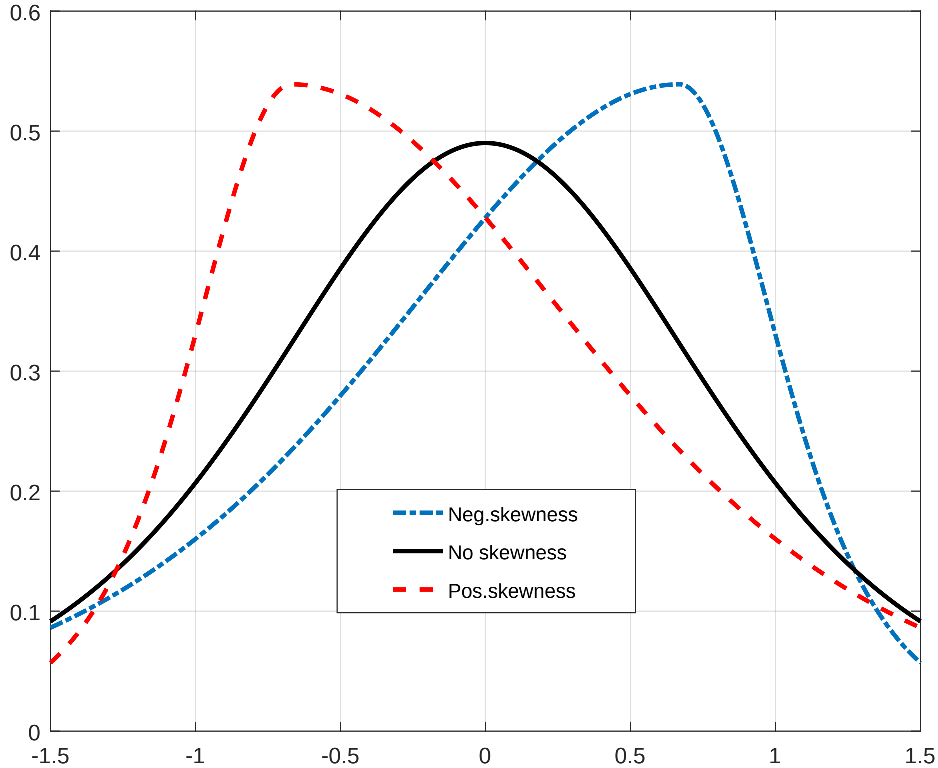

Fig. E4: Realized volatility and standardized earnings surprises across DS deciles

The charts show the DS-decile average of two stock characteristics. DS deciles are recalculated each month t. The two stock characteristics are realized daily stock volatility one year prior to month t (top), and average standardized earnings surprises (SUE) in the year up to month t (Chan, Jegadeesh, and Lakonishok, 1996, p. 1685) (bottom). The sample period is 2006 to 2015.

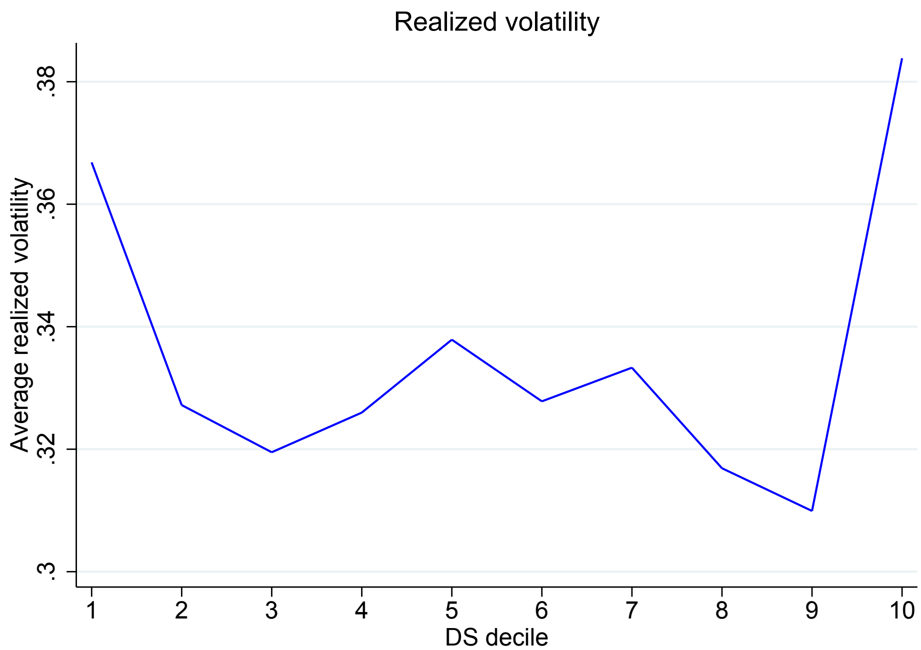

Table E1: Drivers of extreme differences in skewness

We discuss the drivers of extreme differences in skewness, focusing on the top 5 negative DS observations in Panel A and on the top 5 positive DS observations in Panel B. Negative DS implies a higher likelihood of intermediate positive returns, and it should be associated with positive news. On the contrary, positive DS implies a higher likelihood of intermediate negative returns, and it should be associated with negative news.

Panel A – Top 5 negative DS observations

Rank	Ticker	Name	Date	Description
1	GE	General Electric	5/14/2015	- GE announced plans to divest most of its Japanese commercial finance business
2	LPX	Louisiana-Pacific	4/6/2011	- LPX was downgraded by Credit Suisse due to short-term headwinds, with reaffirmed long term guidance
3	RAI	Reynolds American	2/16/2006	- Post-earnings announcement call with analysts
4	FE	FirstEnergy	10/12/2011	- A key step was taken to resolve a dispute between JCP&L (a large New Jersey subsidiary of FE) and New Jersey's utilities regulator
5	CMC	Commercial Metals	1/22/2013	- CMC stock price increased 3.4%

Panel B – Top 5 positive DS observations

Rank	Ticker	Name	Date	Description
1	JBLU	JetBlue Airways	7/24/2006	- Earnings announcement scheduled for the following day. JBLU stock price declined sharply after the announcement
2	ABT	Abbott Laboratories	6/30/2010	- The previous day, ABT filed a lawsuit to prevent the manufacturing of a generic version of an own drug in the lucrative cholesterol-control market. The main active ingredient of the drug in question was off-patent
3	LXK	Lexmark International	7/26/2007	- LXK stock price dropped 4%
4	JCP	J.C. Penney	6/17/2015	- Macy's filed a lawsuit seeking damages from JCP as part of an existing dispute on the licensing of Martha Stewart merchandise
5	FDX	FedEx	5/19/2011	- Important trade conference held on May 19-20, 2011. On the first day of the conference, UPS (a major competitor) presented positive results relative to FDX

Appendix F: Robustness of Asset Pricing Tests

Default thresholds and CDS tenors

We start by evaluating the sensitivity of our findings to alternative default thresholds and to two additional CDS tenors. We report these results in Table F1.

Relative to the default thresholds shown in Table 2, we consider three additional default thresholds. First, we take the average of the thresholds obtained by double sorting companies on the basis of size and leverage. Second, we use a threshold equal to 120% of this average. Third, we set the threshold to 80% of the average. Evaluating the robustness to different default thresholds is also informative about the robustness of our results to different recovery rates, since changing the threshold and changing the recovery rate are observationally equivalent in their impact on the default probability extracted from CDS spreads. In columns 1 to 6 of Table F1, we report results based on these alternative default thresholds. The estimated time-series intercepts are slightly smaller than those reported in Tables 5 and 6. When compared with the findings presented in Sections 4.2.1 and 4.2.2, these results highlight the importance of choosing default thresholds based on company characteristics, such as size and leverage. The intercepts for 120% of the mean for 2008-2011 sub-sample and 80% of the mean in the full sample are statistically significant, and most of the remaining marginally fail the statistical significance test. The adjusted R^2s are comparable to our previously reported values. The size of the DS risk premium is similar to those shown in Tables 7 and 8 in both full sample and in the 2008-2011 sub-sample.

As noted by Chen, Fleming, Jackson, Li, and Sarkar (2011), trading liquidity in the CDS market is concentrated in the five-year tenor. Therefore, we focus on risk-neutral distributions extracted using options and the appropriately compounded five-year CDS spread, although the horizon of the distributions is three months. Even if five-year spreads are the most liquid and thus more likely to have a higher signal-to-noise ratio, Friewald, Wagner, and Zechner (2014) find that the term-structure of CDS spreads is informative about the equity risk premium. We repeat our analysis using either one-year CDS spreads from Markit or three-month CDS spreads, obtained by linearly extrapolating log-CDS spreads with maturities between six months and 10 years. The last four columns of Table F1 show that the results are robust to using alternative CDS tenors. In both time-series and cross-sectional regressions, the size of the estimated intercepts and DS premia are comparable to our reported results. The estimated intercepts have larger t statistics compared with those reported in Tables 5 and 6. In the Fama-MacBeth regressions, we obtain slightly lower t statistics for the DS premia.

Table F1: Results sensitivity to alternative thresholds and CDS maturities

For a variety of default thresholds and CDS spreads with different maturities, the table reports the intercepts, t-statistics, and adjusted R^2s of time series regressions (as in Tables 5 and 6) and second-stage Fama and MacBeth (1973) coefficients, t-statistics, and average adjusted R^2 (as in Tables 7 and 8). The default thresholds are the mean, 120% of the mean, and 80% of the mean of the four thresholds that vary on the basis of size and leverage. The 1-year CDS spreads are provided by Markit, and the 3-month CDS spreads are linearly extrapolated from the log-spreads of CDS with maturity between six months and ten years.

CDS tenor			5 years				1 year		3 months	
Default threshold	mean		120% of mean		80% of mean		original		original	
Sample period	full	crisis	full	crisis	full	crisis	full	crisis	full	crisis
Time series intercept	0.002	0.005	0.002	0.006*	0.003*	0.004	0.005***	0.008**	0.005***	0.008**
t-statistic	(1.08)	(1.66)	(1.44)	(1.73)	(1.92)	(1.39)	(2.89)	(2.54)	(2.78)	(2.57)
Adj.R^2	0.467	0.617	0.385	0.590	0.574	0.721	0.371	0.592	0.416	0.619
FMB risk premium	0.006**	0.010*	0.004	0.010*	0.005*	0.007	0.004	0.010*	0.004	0.012**
t-statistic	(2.09)	(1.72)	(1.63)	(1.69)	(1.95)	(1.27)	(1.41)	(1.93)	(1.35)	(2.13)
Adj.R^2_{cr}	0.499	0.643	0.506	0.620	0.503	0.639	0.499	0.614	0.497	0.628

Interpolated data and alternative parametric distributions

Our work contributes to a large literature on how to extract information from option prices. Researchers have explored a variety of approaches to augment the cross-section of traded prices, using both parametric and non-parametric methods. In our case, we propose to use the information embedded in CDS prices, and rely on a parametric distribution that is fairly common in empirical asset-pricing (the skew-t) to combine information from CDS and options, and also to reduce the effect of noise in option prices.

We now compare our main asset-pricing results with results we obtain when using several alternative distributions for risk-neutral returns: skew-normal, Burr Type XII, and a non-parametric density augmented with tails from generalized extreme value distributions as in Figlewski (2010) (henceforth, hybrid GEV). The skew-normal is an extension of the normal distribution that allows for the presence of skewness (see., e.g., Azzalini, 2013). The Burr Type XII distribution has two scale parameters and a shape parameter to cover a relatively broad set of skewness and kurtosis values.[11] Given that its support is the positive real line, we fit this distribution in the space of gross returns instead of log returns. The hybrid GEV distribution is described in Figlewski (2010). In addition, we re-estimate skew-t distributions using implied volatilities extrapolated according to Stilger, Kostakis, and Poon (2017), and using the volatility surface provided by OptionMetrics.

The results shown in Table F2 indicate that the choice of the distribution is important, but it is not the only driving factor of our findings: using options and CDS, together with the skew-t, is *the* key driver. In particular, building the test assets and the DS factor on the basis of skewness calculated from the skew-normal distribution (first column of the table) generates statistically insignificant time-series intercepts both over the full sample an around the 2008 financial crisis (these and later results correspond to specifications (4) in Tables 5 and 6 in the manuscript). We obtain similar results with the Burr distribution. When we use the hybrid GEV distribution for the long leg of the DS portfolio, and the options-only skew-t for the short leg, the intercepts are not negative any longer but remain clearly statistically insignificant. Only when we consider a skew-t also for the long leg of the DS factor (estimated from the OptionMetrics volatility surface) do the intercepts approach the values reported in the paper. The t-statistics also increase markedly, even if the results remain statistically insignificant.

Overall, choosing the skew-t distribution is important, but it is the combination of skew-t plus options *and* CDS that clearly contributes to our results. Crucially, our decision to use the skew-t distribution among various possible alternatives is firmly rooted in the literature, which indicates that this distribution is well suited to model stock return dynamics (eg, Hansen, 1994, Jondeau and Rockinger, 2003, Patton, 2004 and Oh and Patton, 2017).

[11] See Rodriguez (1977) for a discussion of the Burr family of distributions.

Table F2: Time series intercepts when using alternative distributions

The table shows intercepts from time-series asset-pricing tests conducted under alternative definitions of the DS factor and using data interpolated/extrapolated with different methods. "SKP interp." refers to the interpolation and extrapolation of the volatility smile conducted according to Stilger, Kostakis, and Poon (2017). "Vol surface" refers to the volatility surface provided by OptionMetrics, which is interpolated/extrapolated with kernel smoothing.

	Definition of DS factor				
Long leg:	Skew normal w/ CDS	Burr w/ CDS	Hybrid GEV w/o CDS	Skew-t w/ SKP interp.	Skew-t w/ vol surface
Short leg:	Skew normal w/o CDS	Burr w/o CDS	Skew-t w/o CDS	Skew-t w/o CDS	Skew-t w/o CDS
	Time series asset pricing tests – Full sample				
Intercept	-0.002	-0.000	0.000	0.001	0.002
t-stat	(-0.48)	(-0.11)	(0.16)	(0.36)	(1.28)
Obs.	119	119	119	119	119
Adj.R^2	0.336	0.284	0.396	0.55	0.274
	Time series asset pricing tests – crisis only				
Intercept	-0.005	-0.003	0.002	0.001	0.004
t-stat	(-0.36)	(-0.74)	(0.70)	0.35	(1.61)
Obs.	48	48	48	48	48
Adj.R^2	0.344	0.484	0.485	0.532	0.378

Alternative computation of BKM skewness

The width of the cross-section of option prices used in computing BKM moments can have a significant effect on the values of these moments. We evaluate the implications of calculating BKM skewness using interpolated prices based on the method of Stilger, Kostakis, and Poon (2017).[12] We refer to this variable as SKP-BKM skewness, and we include the corresponding skewness factor (defined similarly to the skewness factors in Tables 5 and 6) in the main time-series asset pricing regressions.

The two leftmost columns of Table F3 show the results of time-series asset pricing regressions based on the full sample; the two rightmost columns show results for the crisis sample (2008-2011). Within each panel, the first colum shows our baseline results from Tables 5 and 6 (specifications (4)). In the second colum, we replace the BKM factor with the SKP-BKM factor, while every other aspect of the analysis stays unchanged. Three results are noteworthy. First, the intercept remains positive and statistically significant in all cases, althought the point estimate is smaller for the crisis sub-sample (0.7% rather than 1.1%). Second, the coefficient on the SKP-BKM factor turns positive and statistically significant. Third, the coefficient on the variance risk premium is now negative and stongly statistically significant, highlighting the link between skewness risk and compensation for variance risk.

In summary, we find that the method for interpolating/extrapolating the available option data has a clear effect on calculating BKM moments. Our main results, however, carry through even with SKP-BKM skewness.

[12] We thank an anonymous referee for suggesting this analysis.

Table F3: Asset-pricing results with the SKP-BKM skewness factor

The table shows the coefficients of time-series regressions of the DS factor on a number of established asset pricing factors. The first and third colums are specifications (4) from Tables 5 and 6, respectively. In the second and fourth columns, the skewness factor is based on SKP-BKM skewness. The full sample period runs from 2006 to 2015. The crisis period includes 2008 through 2011.

	Full sample			Crisis (2008-2011)	
	Baseline	SKP-BKM		Baseline	SKP-BKM
MKT	0.069	0.026		0.118	0.016
	(0.97)	(0.41)		(1.19)	(0.15)
SMB	-0.001	-0.038		-0.023	-0.040
	(-0.01)	(-0.53)		(-0.18)	(-0.32)
HML	0.208	0.183		0.463*	0.265
	(1.52)	(1.43)		(1.95)	(1.42)
UMD	-0.138***	-0.120**		-0.170**	-0.105
	(-2.77)	(-2.48)		(-2.48)	(-1.68)
LIQ	0.039	0.032		0.136	0.071
	(0.63)	(0.51)		(1.58)	(0.85)
RMW	-0.134	-0.127		-0.419*	-0.399*
	(-0.97)	(-0.99)		(-1.79)	(-1.74)
CMA	-0.144	-0.112		-0.694**	-0.544*
	(-0.85)	(-0.66)		(-2.70)	(-2.03)
LT REV	-0.214**	-0.214**		0.022	-0.104
	(-2.02)	(-2.19)		(0.14)	(-0.62)
ST REV	-0.132*	-0.124*		-0.219**	-0.186*
	(-1.78)	(-1.75)		(-2.05)	(-1.76)
ΔVIX_{repl}	-0.216***	-0.217***		-0.249***	-0.206***
	(-3.75)	(-3.89)		(-3.70)	(-3.91)
$\Delta\text{VIX}^2_{repl}$	0.111	0.059		0.161	0.161
	(0.90)	(0.47)		(1.09)	(1.15)
VRP_{repl}	-0.089	-0.082		-0.406***	-0.208**
	(-0.63)	(-0.78)		(-3.01)	(-2.25)
ΔCDX_{HY}	0.000	-0.001		0.007	0.004
	(0.05)	(-0.24)		(1.14)	(0.77)
SKEW_{BKM}	0.070	0.118	$\text{SKEW}_{SKP-BKM}$	0.427**	0.228**
	(0.59)	(1.49)		(2.34)	(2.17)
Intercept	0.004**	0.003**		0.011***	0.007**
	(2.17)	(2.15)		(3.06)	(2.24)
Obs.	119	119		48	48
Adj.R^2	0.385	0.399		0.590	0.592

Previous volumes in this series

920 January 2021	Investing like conglomerates: is diversification a blessing or curse for China's local governments?	Jianchao Fan, Jing Liu and Yinggang Zhou
919 January 2021	Understanding bank and nonbank credit cycles: A structural exploration	C Bora Durdu and Molin Zhong
918 January 2021	Sovereign credit and exchange rate risks: Evidence from Asia-Pacific local currency bonds	Mikhail Chernov, Drew Creal and Peter Hördahl
917 January 2021	Trade sentiment and the stock market: new evidence based on big data textual analysis of Chinese media	Marlene Amstad, Leonardo Gambacorta, Chao He and Dora Xia
916 January 2021	Firm-level R&D after periods of intense technological innovation: The role of investor sentiment	Sirio Aramonte and Matthew Carl
915 December 2020	The globalisation of inflation in the European emerging countries	Horatiu Lovin
914 December 2020	Demographic shifts, macroprudential policies, and house prices	Jieun Lee and Hosung Jung
913 December 2020	Forecasting expected and unexpected losses	Mikael Juselius and Nikola Tarashev
912 December 2020	Regulatory capital, market capital and risk taking in international bank lending	Stefan Avdjiev and Jose Maria Serena
911 December 2020	Bilateral international investments: The Big Sur?	Fernando Broner, Tatiana Didier, Sergio L Schmukler and Goetz von Peter
910 December 2020	Recessions and mortality: a global perspective	Sebastian Doerr and Boris Hofmann
909 December 2020	Dealing with bank distress: Insights from a comprehensive database	Konrad Adler and Frederic Boissay
908 December 2020	Contagion Accounting	Iñaki Aldasoro, Anne-Caroline Hüser and Christoffer Kok
907 December 2020	Low price-to-book ratios and bank dividend payout policies	Leonardo Gambacorta, Tommaso Oliviero and Hyun Song Shin

All volumes are available on our website www.bis.org.