Crash Risk in Individual Stocks

JOB MARKET PAPER

Paola Pederzoli*

*University of Geneva and Swiss Finance Institute, paola.pederzoli@unige.ch

January 3, 2018

Abstract

In this study, I develop a novel methodology to extract crash risk premia from options and stock markets. I document a dramatic increase in crash risk premia after the 2008/2009 financial crisis, indicating that investors are willing to pay high insurance to hedge against crashes in individual stocks. My results apply to all sectors but are most pronounced for the financial sector. At the same time, crash risk premia on the market index remained at pre-crisis levels. I theoretically explain this puzzling feature in an economy where investors face short-sale constraints. Under short-sale constraints, prices are less informationally efficient which can explain the increase in downside risk in individual stocks. In the data, I document a strong link between proxies of short-sale constraints and crash risk premia.

Keywords: Skewness risk premium, financial crisis, short-selling constraints.

JEL classification codes: G01, G12, G13.

I thank David Bates, Pierre Collin-Dufresne, Antonio Cosma, Jérôme Detemple, Elise Gourier, Alfredo Ibanez, Ian Martin, Ilaria Piatti, Olivier Scaillet, Paul Schneider, Fabio Trojani, Roberto Tubaldi, Andrea Vedolin, Gyuri Venter and seminar and conference participants at the 2017 SoFiE conference at NYU, the 2017 FMA conference in Lisbon, the 2017 SoFiE summer school in Chicago, the 2017 SFI Research days in Gerzensee, the 2016 CFE conference in Seville, the University of Geneva, the Queen Mary University of London and the University of Lugano for the valuable comments and insights. This paper was written while I was visiting the Financial Markets Group at the London School of Economics. All errors are my own. Email: `paola.pederzoli@unige.ch`

1 Introduction

Classic asset pricing theory posits that an investor who holds a portfolio of equities should only be compensated for systematic risk. The idea is that in *perfect* capital markets, rational investors should hold a diversified portfolio where the impact of idiosyncratic risk is zero, on average. In this study, I examine the idiosyncratic risk in the equity market from a new perspective: I explore the downside risk of individual stocks by implementing a trading strategy with which an investor can hedge himself/herself against individual downside risk. The return of this strategy measures how much an investor is willing to lose (or gain) for this hedge. The results that I find are surprising: after the financial crisis of 2008/2009 investors price individual crashes with a higher probability, whereas the market crash risk does not increase. I also develop a new methodology to measure coskewness risk, which measures the downside correlation among stocks, and I find that the coskewness risk does not increase after the crisis. Thus, all the new crash risks in individual stocks are due to idiosyncratic risk. This result is surprising because it goes against the conventional idea that idiosyncratic risk should not matter, and what is more surprising is that the idiosyncratic crash risk and not the market crash risk increases after the systemic crisis of 2008/2009. I then propose, empirically and theoretically, a *friction-based* explanation that connects the increase in individual crash risk to short-selling constraints, a friction that became especially relevant during the crisis.

A plethora of literature has studied the pricing of crash risk, mainly for the stock market index. Investors who are concerned about crash risk can purchase insurance via a portfolio of out-of-the-money put options. Option markets thus provide an ideal laboratory to measure investor willingness to hedge against tail risk. Empirical evidence suggests that investors are willing to pay extraordinary high premia to insure against market crashes. This has become particularly topical after the financial crisis of 2008, when many investors experienced unprecedented losses. Much less is known, however, for individual stocks, and this is what the present study aims to elucidate.

The contribution of this work is threefold. The first contribution is methodological, where I show how to measure crash risk and coskewness risk in the equity market via a trading strategy. The second and main contribution is empirical where I document the puzzling increase in downside risk in individual stocks but not in the market. Finally, I empirically and theoretically rationalize the empirical facts in an economy with frictions, where investors face short-selling constraints and asymmetric information.

The novel methodology I implement to measure the crash risk premia in single stocks takes the form of a simple trading strategy, a skewness swap. In line with recent literature on crash risk in individual stocks, I construct measures of crash risk from the prices of out-of-the-money (OTM) puts and calls (see, e.g., Kelly and Jiang (2014) or Kelly et al. (2016), among others). The difference between the prices of OTM puts and calls measures the slope of the implied volatility smile and provides a measure of the priced tail risk, which is econometrically linked to the risk-neutral skewness of the asset, the third moment of the return distribution. However, different from this literature whose main focus is the price of skewness risk, the goal of my study is to quantify the size of the *skewness risk premium*, i.e., the difference between the physical and risk-neutral skewness through a tradable strategy. Intuitively, the skewness risk premium tells us how much an investor wants to be rewarded for bearing skewness risk. To this end, I implement the swap trading strategy of Schneider and Trojani (2015) who provide a novel methodology for trading skewness by taking positions in options. This strategy takes the form of a swap, where an investor exchanges at maturity the fixed leg of the swap with the floating leg. The fixed leg is settled at the start date of the swap and is given by the price of a portfolio of options while the floating leg is realized at maturity and is given by the payoff of the same option portfolio plus a delta hedge in the underlying stock. The key theory behind the skewness swap is that the fixed leg measures the *risk-neutral skewness* of the asset, whereas the floating leg measures the *realized skewness* of the asset. In this way, the payoff of the skewness swap, given by the difference between the floating leg and the fixed leg, offers a direct and tradable measure of the *realized skewness risk premium*. In the same way in which the literature (see e.g, Carr and Wu (2009) or Martin (2017), among others) has used the variance swap to study the variance risk premium, I use the skewness swap to study the skewness risk premium. I extend the skewness swap of Schneider and Trojani (2015) along two key dimensions. First, options on individual stocks are typically American, so I show how to build a skewness swap with American options when investors exercise their options optimally. Second, I use a different weight function in the construction of the option portfolio, and I demostrate in a numerical study that the fixed leg of my swap offers a more precise measure of the risk-neutral skewness.

In addition to studying individual skewness risk premia, I also analyze skewness risk in the index. More specifically, by taking a long position in the index skewness swap and a short position in the basket of individual skewness swaps, I show how an investor can trade market-wide *coskewness*. Coskewness measures the correlation in the tails of the distributions. Intuitively, if two random variables exhibit positive (negative) coskewness, they tend to move in extreme positive (negative) directions simultaneously. The coskewness swap allows me to study the *coskewness risk premium*.

Equipped with this new methodology, I implement the skewness swap strategy on all constituents of the S&P500, the S&P500 index itself, and the sector ETFs in the time period 2003-2014. The average monthly gain of the strategy amounts to 80% for the index, 70%–80% for the sector ETFs, and 50% for the individual stocks. These results are economically and statistically significant for almost the entire cross-section of stocks and are robust even in the presence of transaction costs. Most importantly, they provide strong evidence of a *positive* skewness risk premium, revealing that investors like skewness because positive skewness implies a higher probability of having high returns. Hence, risk-averse investors want to hedge against a drop in skewness, and they accept very negative returns for this hedge.

This result supports the theoretical model of Bakshi et al. (2003) who show that within a power utility economy in which returns are leptokurtic, the risk-neutral implied skew is greater in magnitude than the physical $\mathbb{P}$ skew, which implies a positive skewness risk premium. I also document that the monthly gain of the coskewness swap amounts to 80%, and I show that the index skewness risk premium is mainly due to the market-wide coskewness, and that investors have a preference for positive coskewness.

To dissect the risk premia in more detail, I then split my sample period into two subsamples: the first spanning the years before the financial crisis from 2003 to 2007, and the second covering the years after the crisis from 2009 to 2014. I find that while the index and coskewness risk premia do not change after the crisis, the risk premium of the individual basket increases by 50%. The change is statistically significant and I show that it is driven by a peculiar decrease in the risk-neutral skewness of single stocks, a decrease that is not present in the index. Indeed, I document that the implied volatility smile for single stocks becomes *steeper* after the financial crisis, indicating that the difference between the price of OTM puts and OTM calls has widened. The result is robust for different measures of skewness and it is not driven by the difference in data availability between the index and single stocks options, e.g., number of options and moneyness range.

Two questions emerge as to why investors are willing to pay higher premia for insuring themselves against individual stock crashes: Is it because downside risk has increased? Or is it simply because investors have become more risk averse since 2008? I rule out both hypotheses in the data. First, the increase in the skewness risk premium means that the priced skewness risk is higher in absolute value than the physical skewness risk, and, therefore, the increase in price does not reflect an increase in realized risk. Second, if the increase in the individual skewness risk premia after the financial crisis were due to an increase in the risk aversion of investors, one would also expect a similar increase in the risk premium for the market index, if we believe that the option market is integrated across markets.

I empirically link these idiosyncratic skewness risk premia to frictions in the stock market, and this constitutes the third contribution of my study. To this end, I first decompose each individual skewness risk premium into a systematic part, which is common to all stocks, and an idiosyncratic part, which is different for each stock. I exploit the cross-sectional heterogeneity of the idiosyncratic skewness risk premia and I find that short-selling constraints are the key drivers of the idiosyncratic skewness risk premium. I then propose two economic channels that could rationalize why stocks that are more difficult to short-sell command a higher idiosyncratic crash risk premium. First, short-selling constraints prevent negative information to be fully incorporated in the prices, leading to overpriced stocks with a higher disaster risk. Empirically, I find that stocks that are difficult to short-sell are more likely to be overpriced, and they command a higher idiosyncratic skewness risk premium. Second, when a stock is difficult to sell short in the equity market investors might resort to the option market as an alternative by buying puts or by creating a synthetic short position in the stock with a combination of puts and calls. Empirically, I find evidence of this market substitution mechanism; I find that stocks that are difficult to sell short have higher volumes in put options, which, in turn, pushes up the skewness risk premium. By using two proxies for short-selling costs, I find that the stocks that were most affected by short-selling constraints during the crisis have a subsequent higher increase in the idiosyncratic skewness risk premium.

Finally, I rationalize these empirical findings in a parsimonious rational expectations equilibrium economy where investors face short-selling constraints and asymmetric information similar to the model of Venter (2016). The model has three types of investors: informed investors, uninformed investors, and noise traders. Only informed investors know the final payoff of the stock, whereas uninformed investors form rational expectations of the expected final payoff from prices. Noise traders prevent the trading activity of informed investors from being fully revealing. If informed investors are short-sale constrained, the distribution of the final payoff given the price is left-skewed in the equilibrium because the short-selling constraints lead to a higher uncertainty in the left tail of the distribution. Intuitively, this implies that the stock crash risk that can be inferred from the prices increases with short-selling constraints.

Literature Review: My work contributes to different strands of the literature.

First, it contributes to the literature that explores the risk premia of different moments of asset return distribution. In the equity market, a plethora of studies has documented the existence of a positive first-moment risk premium, the equity risk premium, within a CAPM framework and in factor models (see, e.g., Fama and French (1993), Carhart (1997), Pastor and Stambaugh (2003), Fama and French (2015), among many others). A more recent stream of literature investigates the risk premia of higher-order moments of the return distribution, particularly the variance risk premium (see, e.g., Bakshi and Kapadia (2003), Bollen and Whaley (2004), Carr and Wu (2009), Ang et al. (2006), Bollerslev et al. (2009), Buraschi et al. (2014), Choi et al. (2017), and Martin (2017), who derives a lower bound on the equity premium in terms of a volatility index, SVIX). These studies provide evidence of a negative market volatility risk premium. Investors dislike volatility because higher volatility represents a deterioration in investment opportunities. Risk-averse agents want to hedge against a rise in volatility, so the risk-neutral price of volatility is higher than the average realized volatility, leading to a negative volatility risk premium. Bollerslev and Todorov (2011), Bollerslev et al. (2015), and Piatti (2015) refine these results by showing that the compensation for rare events and jump tail risk accounts for a large fraction of the average equity and variance risk premia in the index. The tails of the distribution have an important role in the determination of prices (Bates (1991)), and the skewness is a natural measure of the asymmetry of the tails of the distribution. Other important studies that examine variance and jump risk premia are those of Cremers et al. (2015), who examine jump and volatility risk in the cross-section of stock returns, and Andersen et al. (2017), who study the volatility and jump risk implicit in S&P500 weekly options.

Bakshi et al. (2003) have initiated the study of risk-neutral skewness and the skewness risk premium. In their research, the authors develop a methodology to compute the risk-neutral skewness of an asset via option portfolios. They document that the risk-neutral skewness of the S&P500 index and of 30 single stocks is negative, and it is, in absolute value, higher for the index than for the individual stocks. Since then, the empirical literature on the sign of the skewness risk premium has yielded mixed results: Kozhan et al. (2013) and Schneider and Trojani (2015) find a positive skewness risk premium for the S&P500 index, whereas Chang et al. (2013) document that stocks with high exposure to innovations in implied market skewness exhibit low returns, on average (negative market price of skewness risk). Chang et al. (2013) rationalize their result by showing the negative correlation between changes in the market skewness and market returns; therefore, an increase in the skewness leads to a deterioration in the investment opportunity set. Conrad et al. (2013) and Boyer et al. (2010) find that stocks with the lowest expected idiosyncratic skewness outperform stocks with the highest idiosyncratic skewness, whereas Schneider et al. (2016) and Stilger et al. (2016) find the opposite: the more ex-ante negatively skewed returns yield subsequent lower returns.

I contribute to the literature on the skewness risk premium by showing that the realized skewness risk premium measured by the return of a skewness swap is positive, both for the market and for single stocks. This result supports the theoretical model of Bakshi et al. (2003) and connects my work with the theoretical work of Eeckhoudt and Schlesinger (2006), who show that while risk-averse investors dislike volatility, prudent investors have preferences for lotteries with higher skewness. A positive skewness risk premium uncovers the prudent attitude of investors. I also contribute to the literature on the coskewness risk premium by showing that the coskewness risk premium is positive. My results on coskewness are consistent with the findings of Christoffersen et al. (2016) and Harvey and Siddique (2000). However, while their empirical methodology is nested within the classic two-step Fama and MacBeth (1973) regression, I use a trading strategy in which an investor can directly buy marketwide coskewness. My methodology is more linked to that of Driessen et al. (2009), which extracts marketwide correlation from the S&P500 variance, with the key difference that my coskewness risk premium is the gain of a trading strategy. To my knowledge, I am the first to empirically study the skewness risk premium and the coskewness risk premium in the single stock equity market via a trading strategy.

The results on the financial crisis connect my study with that of Kelly et al. (2016) on put option prices. Carr and Wu (2011) show that American put options are theoretically similar to credit insurance contracts linked to the default of the company. Kelly et al. (2016) document that the difference in costs between OTM put options for individual banks and for the financial sector index increased during the 2007–2009 crisis, revealing that the idiosyncratic risk is priced more heavily than systematic risk for the financial sector only. The authors rationalize these results with a bailout guarantee that lowers the systematic crash risk priced in the index puts, compared with the idiosyncratic risk priced in the puts of the individual stocks. In my work, I find a post-crisis increase in the skewness risk premium for all stocks across different sectors, and not only for financials. Indeed, the most striking result is achieved by the industrial sector, where the basket skewness risk premium increases from 13% to 64% after the crisis. Furthermore, different from Kelly et al. (2016), I rationalize the results in an economy with short-selling constraints.

The decomposition of the individual skewness risk premium into a systematic and an idiosyncratic part connects my work with those of Gourier (2016) and Begin et al. (2017), who show that idiosyncratic variance risk is priced in the cross-section of stocks, and my work extends their results to skewness. In addition, while Gourier (2016) and Begin et al. (2017) use a parametric approach, my methodology is completely model free.

My study is also related to the literature on short-selling constraints. The seminal work of Miller (1977) argues that if short selling is costly and investor beliefs are heterogeneous, a stock can be overvalued and generate low subsequent returns. The economic intuition behind this hypothesis is that short-selling constraints prevent negative information or opinions from being revealed in stock prices. Subsequent empirical literature confirmed this hypothesis in the data (see, e.g., Jones and Lamont (2002), Chen et al. (2002), Desai et al. (2002), D'Avolio (2002), Nagel (2005), Asquith et al. (2005), among others). Another strand of literature shows that short-selling costs and bans have significant effects on option prices (see, e.g., Stilger et al. (2016), Lin and Lu (2015), Battalio and Schultz (2011), Atmaz and Basak (2017)). These results indicate that on the one hand, short-selling constraints generate low returns, and, on the other hand, that stocks that are more difficult to short-sell have more negative risk-neutral skewness. The impact of short-selling constraints on the skewness risk premium, however, has been unexplored. I contribute to this literature by showing that short-selling constraints are associated with a higher skewness risk premium, which reveals that these constraints have a higher impact on the risk-neutral distribution than on the physical distribution. Economically, this means that stocks that are difficult to sell short are perceived as idiosyncratically riskier.

The rest of the paper is organized as follows. Section 2.1 introduces the skewness swap, and Section 2.2 explains the details of the empirical application. Section 3 presents the empirical results of the study: Section 3.1 introduces and analyzes the puzzle on the skewness and coskewness risk premiums and Section 3.2.1 proposes a friction-based explanation based on short-selling constraints. Alternative explanations are discussed in Section 3.2.2. Section 4 presents the simple rational expectation equilibrium model, and Section 5 concludes. In the paper, I use the notations $\mathbb{Q}$ skewness, priced skewness, and implied skewness as synonyms for risk-neutral skewness. Analogously, I use the notation $\mathbb{P}$ skewness for the realized skewness.

2 Methodology

2.1 Theoretical motivation: the skewness swap

The skewness swap is a contract through which an investor can *buy the skewness of an asset* by taking positions in options. At the start date of the contract, the investor buys the portfolio of options, and when the options expire, he or she receives the payoff of the option portfolio. The key theory behind this contract is that the price of the option portfolio measures the risk-neutral skewness of the asset and the payoff of the option portfolio measures the realized skewness of the asset. All trades are done in the forward market, so the contract is more similar to a forward contract on options, but I call it a skewness swap to align my work with the literature on higher-moment swaps. We can think of it as a swap contract whereby two counterparts agree to exchange at maturity a fixed leg, given by the price of the option portfolio, with a floating leg, given by the payoff of the option portfolio.

I build on the general divergence trading strategies of Schneider and Trojani (2015) to construct my skewness swap. Schneider and Trojani (2015) introduce a new class of swap trading strategies with which an investor can take a position in the generalized Bregman (1967) divergence of the asset. The skewness can be seen as a special type of divergence, and Schneider and Trojani (2015) propose a Hellinger skew swap for trading skewness. In this section, I construct a new skewness swap with the following enhancements over the Hellinger skew swap of Schneider and Trojani (2015): (a) my skewness swap is a pure bet on the third moment of the stock returns while being independent of the first, second, and fourth moments, and (b) the swap can be applied directly to American options.

The swap defined by Schneider and Trojani (2015) has the following general form:

$$\overline{S_E} = \frac{1}{B_{0,T}} \left(\int_0^{F_{0,T}} \Phi^{''}(K) P_{E,0,T} dK + \int_{F_{0,T}}^{\infty} \Phi^{''}(K) C_{E,0,T} dK \right) \tag{1}$$

$$\begin{aligned}\widetilde{S_E} &= \left(\int_0^{F_{0,T}} \Phi^{''}(K) P_{E,T,T} dK + \int_{F_{0,T}}^{\infty} \Phi^{''}(K) C_{E,T,T} dK \right) \\ &+ \sum_{i=1}^{n-1} \left(\Phi^{'}(F_{i-1,T}) - \Phi^{'}(F_{i,T}) \right) (F_{T,T} - F_{i,T})\end{aligned} \tag{2}$$

where $\overline{S_E}$ is the fixed leg of the swap and $\widetilde{S_E}$ is the floating leg. $P_{E,0,T}$ and $C_{E,0,T}$ are the prices of a European put and call option, respectively, at time 0 and $C_{E,T,T} = (S_T - K)^+$ and $P_{E,T,T} = (K - S_T)^+$ are the payoffs at maturity of the European call and put options. The subscript E indicates that the prices are European options prices. $F_{0,T}$ is the forward price at time 0 for delivery at time T, and $B_{0,T}$ is the price of a zero-coupon bond at time 0 with maturity T. The function $\Phi : \mathbb{R} \longrightarrow \mathbb{R}$ is a twice-differentiable generating function that defines the moment of the distribution we want to trade. For example, if $\Phi(x) = \Phi_2(x) = -4((x/F_{0,T})^{0.5} - 1)$, then $\overline{S_E} = E_0^{\mathbb{Q}} \left[\log(F_{T,T}/F_{0,T})^2 + O(\log(F_{T,T}/F_{0,T})^3) \right]$. In this example, Φ_2 captures the second-order variation of the returns. $\overline{S_E}$ measures the risk-neutral moment while $\widetilde{S_E}$ measures the realization of the moment. Equation 2 shows that the floating leg is composed of two parts: the payoff of the option portfolio at maturity plus a delta hedge in the forward market, which is rebalanced at the intermediate dates i. All the payments of the swap are made at maturity, when the investors exchange the fixed leg with the floating leg. The value of the swap at time 0 is zero, as $E_0^{\mathbb{Q}}[\widetilde{S_E}] = \overline{S_E}$.

The main shortcoming of this trading strategy is that its implementation is limited to assets that have European options available. In the equity market, only indexes have European options because the options on single stocks are American. I modify the swap defined by Equations 1 and 2 in order to deal directly with American options.

I start by defining the American call option payoff at time T:

$$C_{A,T,T} = \frac{(S_{t^*} - K)}{B_{t^*,T}}$$

where $t^* = \min\{0 \leq t \leq T : (S_{t^*} - K) > C(t^*, S_{t^*}, K, T - t^*)\}$, and analogously, the American put option payoff:

$$P_{A,T,T} = \frac{(K - S_{t^*})}{B_{t^*,T}}$$

where $t^* = \min\{0 \leq t \leq T : (K - S_{t^*}) > P(t^*, S_{t^*}, K, T - t^*)\}$. The idea is that the investor exercises the American options optimally and the final payoff at maturity is given by the compounded optimal exercise proceeds.[1]

[1] Many studies show that investors actually do not optimally exercise their stock options, and in particular they miss most of the advantageous exercise opportunities (see, e.g., Pool et al. (2008); Barraclough and Whaley (2012); Cosma et al. (2017)). This issue is important for in-the-money options, while here the skewness swap is constructed using out-of-the-money options for which the early exercise is less relevant. In the empirical section, I will show that in my analysis, the early exercise value is very small and hence an alteration of the early exercise proceeds given by a suboptimal behavior would not alter the main results of the paper.

I define a new swap whose floating leg is given by

$$\begin{aligned}\widetilde{S_A} = &\left(\int_0^{F_{0,T}} \Phi^{''}(K)P_{A,T,T}dK + \int_{F_{0,T}}^{\infty} \Phi^{''}(K)C_{A,T,T}dK\right) \\ &+ \sum_{i=1}^{n-1}\left(\Phi^{'}(F_{i-1,T}) - \Phi^{'}(F_{i,T})\right)(F_{T,T} - F_{i,T})\end{aligned} \tag{3}$$

and fixed leg is given by the expectation of the floating leg

$$\overline{S_A} = E_0^{\mathbb{Q}}[\widetilde{S_A}] = \frac{1}{B_{0,T}}\left(\int_0^{F_{0,T}} \Phi^{''}(K)P_{A,0,T}dK + \int_{F_{0,T}}^{\infty} \Phi^{''}(K)C_{A,0,T}dK\right). \tag{4}$$

The subscript A indicates that the prices are American option prices. The next proposition shows that $\overline{S_A}$ equals $\overline{S_E}$ plus the price of the early exercise and that $\widetilde{S_A}$ equals $\widetilde{S_E}$ plus the realization of the early exercise.

Proposition 1. *The swap with fixed leg given by Equation 4 and floating leg given by Equation 3 verifies the following properties:*

$$\overline{S_A} = \overline{S_E} + \frac{1}{B_{0,T}}\left(\int_0^{F_{0,T}} \Phi^{''}(K)(P_{A,0,T} - P_{E,0,T})dK\right) + \frac{1}{B_{0,T}}\left(\int_{F_{0,T}}^{\infty} \Phi^{''}(K)(C_{A,0,T} - C_{E,0,T})dK\right)$$

$$\widetilde{S_A} = \widetilde{S_E} + \left(\int_0^{F_{0,T}} \Phi^{''}(K)(P_{A,T,T} - P_{E,T,T})dK\right) + \left(\int_{F_{0,T}}^{\infty} \Phi^{''}(K)(C_{A,T,T} - C_{E,T,T})dK\right)$$

The difference between the American and European prices $(P_{A,0,T} - P_{E,0,T})$ and $(C_{A,0,T} - C_{E,0,T})$ measures the price of the early exercise. The difference between the payoff of American and European options $(P_{A,T,T} - P_{E,T,T})$ and $(C_{A,T,T} - C_{E,T,T})$ measures the realization of the early exercise.

Proof. See Appendix C. □

Proposition 1 shows that when I use American options instead of European options, I have an additional component given by the early exercise. While I cannot avoid trading it (because I can trade only American options), I can measure this early exercise component in order to disentangle it from the main skewness leg of the swap.

If the function Φ is

$$\Phi(x) := \Phi_3\left(\frac{x}{F_{0,T}}\right) = -4\left(\frac{x}{F_{0,T}}\right)^{1/2}\log\left(\frac{x}{F_{0,T}}\right) \tag{5}$$

then $\overline{S_{E,\Phi_3}} = E_0^{\mathbb{Q}}\left[\frac{1}{6}y^3 + \frac{1}{12}y^4 + O(y^5)\right]$, where $y = \log(F_{T,T}/F_{0,T})$. This is the Hellinger skewness swap proposed by Schneider and Trojani (2015) to study the third moment of the returns. However, the formula shows that the swap depends theoretically on the fourth moment as well, and I show numerically in Appendix D that this dependence leads to a biased measure of the third moment.

To overcome this problem and better isolate the third moment, I define a new skewness swap S, whose function Φ_S is a combination of Φ_3 and Φ_4, where Φ_4 is the function that defines the kurtosis swap of Schneider and Trojani (2015). In detail, $\Phi_4\left(\frac{x}{F_{0,T}}\right) = -4\left[(x/F_{0,T})^{1/2}(\log(x/F_{0,T})^2 + 8) - 8\right]$

and verifies $\overline{S_{E,\Phi_4}} = E_0^{\mathbb{Q}}[\frac{1}{12}y^4 + O(y^5)]$, where $y = \log(F_{T,T}/F_{0,T})$. By taking a long position in 6 times Φ_3 and a short position in 6 times Φ_4, I can isolate the third moment from the fourth. The result is formally stated in the following proposition.

Proposition 2. *The skewness swap S with floating leg (2) and fixed leg (1) with*

$$\Phi(x) = \Phi_S\left(\frac{x}{F_{0,T}}\right) = -24\left(\frac{x}{F_{0,T}}\right)^{1/2}\log\left(\frac{x}{F_{0,T}}\right) + 24\left[\left(\frac{x}{F_{0,T}}\right)^{1/2}\left(\log\left(\frac{x}{F_{0,T}}\right)^2 + 8\right) - 8\right] \quad (6)$$

verifies the following property:

$$\overline{S_{E,\Phi_S}} = E_0^{\mathbb{Q}}\left[\left(\log\left(\frac{F_{T,T}}{F_{0,T}}\right)\right)^3 + O\left(\log\left(\frac{F_{T,T}}{F_{0,T}}\right)^5\right)\right]$$

Proof. See Appendix C. □

Proposition 2 shows that the swap S isolates the contribution of the third moment from that of the fourth moment, and Proposition 1 shows how to construct a swap using American options. It is worth noting that dividends do not affect the methodology because the modeled return is the forward return $y = \log(F_{T,T}/F_{0,T})$ in which the dividends are included in the calculation of $F_{0,T}$.

2.2 Data and empirical proxies

I apply the skewness swap introduced in Section 2.1 to all the components of the S&P500 separately in the time period 2003–2014. I fix a monthly horizon for the skewness swaps, starting and ending on the third Friday of each month, consistent with the maturity structure of option data. Because the issue of new options sometimes happens on the Monday after the expiration Friday, I take as the starting day of the swaps the Monday after the third Friday of each month. I consider only the periods in which stocks do not distribute special dividends in order to avoid special behavior of stocks.

Security Data: The list of the actual components of the S&P500 is taken from the Compustat database as of December 2014. I exclude the stocks for which there is not an exact match between the daily close price reported by Optionmetrics, Center for Research in Security Prices (CRSP), and Compustat. After this selection, 489 stocks remain. The data on the security prices and returns are taken from CRSP. My methodology requires the calculation of the forward price at time 0 for delivery of the asset at time T. I calculated it as $F_{0,T} = S_0 e^{rT} - PVD$ according to standard no-arbitrage arguments, where r is the risk-free interest rate, S_0 is the stock price at time 0, and PVD is the present value of the dividends paid by the stock between time 0 and time T. The risk-free rate is taken from the Zero Coupon Yield Curve provided by Optionmetrics. The data on the short interest are taken from Compustat.

The S&P500 index and sector indexes: The S&P500 index is a capitalization-weighted index, defined as $P_{S\&P500} = (\sum_i S_i Sh_i)/(Divisor)$, where S_i is the price of each stock in the index, Sh_i is the number of shares publicly available for each stock, and $Divisor$ is a factor proprietary to Standard &

Poor. The *Divisor* is adjusted whenever a stock is added to or deleted from the index or after corporate actions such as share issuance, spinoffs, etc., to ensure that the level of the index is not affected. The simple return of the index from time t to time T can be approximated with the weighted average of the returns of the components. The approximation holds true if the *Divisor* and the number of shares outstanding of each stock do not change value from time t to time T:

$$R_{t,T,S\&P500} \simeq \sum_i R_{t,T,i} w_i,$$

where $w_i = (S_i Sh_i)/(\sum_j S_j Sh_j)$. I will use this approximation to study the coskewness risk premium. The data on the security prices and shares outstanding are taken from CRSP.

The stocks of the S&P500 are also divided into 11 sector indexes, which are capitalization-weighted indexes like the S&P500. There are no options issued directly on the sector indexes, but there are options issued on the SPDR sector ETFs, which are funds that track the S&P500 sector indexes. There are 10 SPDR sector ETFs, because the information technology sector index is merged with the telecom service index into the technology ETF with ticker 'XLK'. The real estate sector ETF started to be traded in 2015 and hence, it is discarded in this analysis. The historical components of the S&P500 and sector indexes are taken from Compustat.

Options data: The data on the option prices and option attributes on single stocks, S&P500, and SPDR sector ETFs are taken from Optionmetrics. The following data filtering is applied: I consider only options with positive open interest, and I exclude options with negative bid-ask spreads, with negative implied volatility and with bid price equal to 0. I implement the swap only if there are at least two call options and two put options to build the fixed leg of the swap. After the selection, I have on average 80 returns for each stock, covering the full data sample period. On average, 8–10 options are used in the implementation of the swap.

Credit default swap data: The data on the credit default swap (CDS) spreads are taken from Markit. I use the CDS spread of the 5-year contract because they are the most liquid CDS contracts.

Book leverage: The book leverage of each stock is computed as $\log(Asset/Equity)$, where *Asset* is the book value of the total assets and *Equity* is the book value of equity. The data are taken from Compustat.

Start and end of the financial crisis 2007/2009: As in Kelly et al. (2016), I consider the start date of the crisis as August 2007 (the asset-backed commercial paper crisis) and June 2009 as the end date of the crisis.

The empirical swap: The fixed leg of the swap is computed at the start date of the swap by building the portfolio of options described in Equation 4 with $\Phi = \Phi_S$ defined in Equation 6. Equation 4 is written for an options market in which a continuum of options is available covering all the strikes in the range $[0, +\infty]$. In practice, I have only a finite number of strikes for each date. I thus implement a discrete approximation of Equation 4. Suppose that at time 0 there are N calls and N puts traded in the market. I order the strikes of the calls such that $K_1 < ... < K_{Mc} \leq F_{0,T} < K_{Mc+1} < ... < K_N$ and the strikes of the puts such that $K_1 < ... < K_{Mp} \leq F_{0,T} < K_{Mp+1} < ... < K_N$. I approximate the fixed leg with the following quadrature formula:

$$D\left(\overline{S_A}\right) = \frac{1}{B_{0,T}}\left(\sum_{i=1}^{Mp} \Phi''(K_i) P_{A,0,T}(K_i)\Delta K_i + \sum_{i=Mc+1}^{N} \Phi''(K_i) C_{A,0,T}(K_i)\Delta K_i\right) \tag{7}$$

where

$$\Delta K_i = \begin{cases} (K_{i+1} - K_{i-1})/2 & \text{if } 1 < i < N, \\ (K_2 - K_1) & \text{if } i = 1, \\ (K_N - K_{N-1}) & \text{if } i = N. \end{cases}$$

I standardize the fixed leg by variance in order to have a scale-invariant measure of the risk-neutral skewness:

$$SK^{\mathbb{Q}}_{0,T} = \frac{D(\overline{S_A})}{(VAR^{\mathbb{Q}}_{0,T})^{3/2}} \tag{8}$$

where $VAR^{\mathbb{Q}}_{0,T}$ is calculated with Equation 4 and $\Phi_2(x/F_{0,T}) = -4((x/F_{0,T})^{0.5} - 1)$.

The floating leg given by Equation 3 is composed of two parts: the payoff of the option portfolio plus the delta hedge. I calculate the payoff of the option portfolio by checking the optimal exercise of the options each day with the market-based rule proposed by Pool et al. (2008). I approximate the integral of the floating leg with the same quadrature approximation I use for the fixed leg, which is given by Equation 7. I implement the delta hedge each day t_i, starting from day t_1 (the day after the start date of the swap) until day t_{n-1} (the day before the maturity of the swap). I also standardize the floating leg by variance in order to have a scale-invariant measure of the realized skewness:

$$SK^{\mathbb{P}}_{0,T} = \frac{D(\widetilde{S_A})}{(VAR^{\mathbb{Q}}_{0,T})^{3/2}} \tag{9}$$

The realized risk premium of each strategy is calculated at maturity as the difference between the floating leg (realized $\mathbb{P}$ skewness) and the fixed leg of the swap ($\mathbb{Q}$ skewness):

$$RP_{0,T} = D(\widetilde{S_A}) - D(\overline{S_A}) \simeq r^3_{0,T} - E^{\mathbb{Q}}\left[r^3_{0,T}\right]. \tag{10}$$

In the analysis that follows, I consider the risk premium expressed as a percentage gain of the skewness swap strategy:

$$RP_{0,T,\%} = \frac{D(\widetilde{S_A}) - D(\overline{S_A})}{C(\overline{S_A})}, \tag{11}$$

where $C(\overline{S_A})$ is the capital needed to purchase the fixed leg of the swap. It is computed as $C(\overline{S_A}) = \frac{1}{B_{0,T}}\left(\sum_{i=1}^{Mp} \mid \Phi''(K_i) \mid P_{A,0,T}(K_i)\Delta K_i + \sum_{i=Mc+1}^{N} \mid \Phi''(K_i) \mid C_{A,0,T}(K_i)\Delta K_i\right)$.[2]

[2]The formula is different than $D\left(\overline{S_A}\right)$ because in $C(\overline{S_A})$ all the weights, given by $\Phi''(K_i)$, are taken as absolute values.

In Appendix D, I numerically show that the fixed leg of my skewness swap constructed with Equation 7 converges to the true risk-neutral skewness when the number of options grows to infinity and when the range of moneyness available grows to $[-\infty, +\infty]$.

3 Empirical results

3.1 The skewness risk premium puzzle after the financial crisis

In this section, I analyze the returns of the skewness swaps on individual stocks, on the S&P500 index, and the coskewness swap returns. I show that the returns of the skewness swaps on individual stocks increased substantially after the financial crisis and that the increase is economically and statistically significant. Moreover, I show that the return of the skewness swap on the S&P500 index and the return of the coskewness swap did not increase after the crisis, indicating that the increase in the individual skewness risk premia are due to idiosyncratic risk. The results survive a number of robustness checks controlling for transaction costs, early exercise bias, and measurement errors that might come from the methodology employed.

3.1.1 The skewness risk premium in individual stocks: preliminary analysis

I implement the monthly skew swap strategy independently for each stock. Thus, each stock of my sample will have a time series of realized skewness swap returns expressed in percentage gain according to Equation 11.

[Table 1 here]

Table 1 reports the average skewness swap returns across the stocks before and after the 2007–2009 financial crisis. In the individual stock analysis of Panel A, I consider only the stocks that have a good options data coverage since the beginning of the sample period, which reduces the sample to 209 stocks. This filter is relaxed in the portfolio analysis of Panels B and C, where I consider the return of an equally weighted and value-weighted portfolio of skewness swaps. The results are very strong: Panel A shows that the return of the skewness swap is high and positive in all sample periods and it is statistically significant for almost all the stocks. The average Sharpe ratio of the skewness swap strategy is high, 1.5. Moreover, the average swap return goes from 43% to 52% after the financial crisis. The results are even stronger in the portfolio analysis of Panels B and C: the average return of the portfolio of skewness swaps goes from 33% before the crisis to 46% after the crisis. At first sight, these numbers may seem too big given that they are *monthly* returns, but they are comparable with other studies on variance swaps and variance risk premia. For example Carr and Wu (2009) estimate that the average return of a variance swap on the S&P500 is -60% monthly and Bakshi and Kapadia (2003) show that the average monthly gain of selling a delta-hedged put option on the S&P500 is 55% and it goes up to 80% if the option is

The reason for this standardization is that the portfolio of options contains both long and short positions. I thus follow the general practice where the return of a short position is computed based on the initial proceeds from the sale, which is taken with positive sign. In this way $C(\overline{S_A})$ measures the total capital exposure of long and short positions, and the absolute values of the weights ensure that $C(\overline{S_A})$ is always positive.

deep out-of-the-money (Table 2 of Bakshi and Kapadia (2003)). A positive skewness swap return implies that the priced skewness is less than the realized skewness, but because the priced skewness is generally negative, a positive skewness swap return implies that the priced skewness is more negative than the realized skewness. An investor who buys skewness will on average make profit while bearing the risk of a sudden decrease in the realized skewness, that is, a crash of the asset. From a hedging perspective, an investor who wants to hedge against a drop in the skewness will sell the skewness swap, and my results show that investors are willing to accept deep negative returns for this hedge. All this supports the idea that investors have a strong preference towards positive skewness and that the sign of the skewness risk premium is positive.

Panel A of Table 1 also displays the results of two robustness checks. First, I compute the return of a skewness swap that uses synthetic European option prices instead of the actual American option prices, where the synthetic European option prices are recovered with the Black-Scholes formula applied to the implied volatility of the American option prices provided by Optionmetrics. This European swap is not tradable, but its return is useful to measure the importance of the early exercise component in the tradable American swap. Table 1 shows that the return of the synthetic European swap is almost identical to the return of the tradable American swap, meaning that the early exercise (both in the priced skewness and realized skewness) is, in this context, negligible. This is not surprising given that the options used in the implementation of the swap are all out-of-the-money and hence have a low probability of expiring in-the-money. As a second check, I compute the return of the swap considering the bid/ask prices instead of mid-quote prices both for the option portfolio and for the delta hedge. Including transaction costs cuts the average swap return by 40%, which shows that the option bid-ask is an important friction to consider. However, even after considering transaction costs, the average swap return goes from 26% to 39% after the financial crisis, and it becomes significant for the majority of the stocks.

Panel A of Table 1 also shows that the increase in the skewness swap returns after the crisis is due to a massive decrease in the fixed leg of the swap, which measures the risk-neutral skewness of the stock. Indeed, the average risk-neutral skewness of the stocks, measured by the fixed leg of the swap, goes from -0.65 to -0.99 after the crisis. I go deeper into this result in the next section, where I compare the results on single stocks with the results for the S&P500 index and the sector indices.

[Figure 1 here]

Figure 1 shows the time series of the monthly return of the value-weighted portfolio of skewness swaps and the cumulative return of investing one US dollar in the portfolio of skewness swaps at the beginning of the period. We can see that on average the return is positive, but the dispersion is huge: there are months in which the return is below -100% or above $+100\%$. The cumulative return shares some similarities with the return on selling insurance: on average it makes money until the trigger event happens (a crash in this case), at which point it loses a lot. Here the risk is even higher because the return can be less than -100% due to the short positions in the portfolio of options. These graphs show that the skewness swap is a highly profitable and highly risky strategy.

[Table 2 here]

I conclude this preliminary analysis by comparing the returns of the skewness swaps with the returns of the variance swaps. The variance swap is a trading strategy with the same structure as the skewness swap but with the Φ function equal to $\Phi_2(x/F_{0,T}) = -4((x/F_{0,T})^{0.5} - 1)$. Table 2 investigates the relation between variance swap returns and skewness swap returns through a portfolio sort analysis. In Panel A of Table 2, for each month I sort the stocks into three tercile portfolios according to their skewness swap return, where ptf1 (ptf3) denotes the portfolio with the lowest (highest) skewness swap return. Then, for each portfolio I compute the average variance swap return. The variance and skewness swap returns are contemporaneous monthly gains of the variance and skewness swaps, respectively. In Panel B, the sorting variable is the variance swap return and the table reports the statistics for the skewness swap return of each variance-sorted portfolio. When I sort according to the skewness swap return (Panel A of Table 2), the variance swap return has a U-shaped pattern: when the skewness swap return is very low or very high, the variance swap return is high. This is intuitive, because the variance swap return, given by ($\textit{realized } Var^{\mathbb{P}}$)$-$($Var^{\mathbb{Q}}$), is high when the realized variance is high. And the realized variance is high when the stock rises substantially (positive skewness swap return) or falls substantially (negative skewness swap return). While the variance swap return does not distinguish whether the stock increases or decreases, the skewness swap return changes sign according to growth direction. When I sort according to the variance swap return (Panel B of Table 2), I find that the stocks with a positive variance swap return have a lower skewness swap return. This is consistent with the well-known leverage effect, according to which there is a negative correlation between stock return and volatility: when the realized volatility is high, and hence the variance swap return is high, the stock is generally in a downturn, and hence the skewness swap return is smaller. This analysis shows that variance and skewness swap returns convey different information.[3]

3.1.2 The index skewness risk premium and the coskewness risk premium

The S&P500 index is a capitalization-weighted index. The simple return of the index from time t to time T can be approximated with the weighted average of the returns of the components:

$$R_{t,T,S\&P500} \simeq \sum_i R_{t,T,i} w_i,$$

where $w_i = (S_i Sh_i)/(\sum_j S_j Sh_j)$, S_i is the price of stock i, and Sh_i is the number of shares outstanding for stock i. I then use the following mathematical identity:

$$\left(\sum_i w_i R_i\right)^3 = \sum_i w_i^3 R_i^3 + 3\sum_{i,j} w_i w_j^2 R_i R_j^2 + 6 \sum_{i,j,k} w_i w_j w_k R_i R_j R_k. \tag{12}$$

[3]Indeed, in the expected utility theory framework of Eeckhoudt and Schlesinger (2006), the attitude of investors toward variance is linked with risk-aversion while the attitude of investors toward skewness is linked with prudence.

Equation 12 allows me to disentangle the index skewness risk premium in the individual stock skewness risk premium and the coskewness risk premium:

$$\underbrace{E^{\mathbb{P}}\left[\left(\sum_i w_i R_i\right)^3\right] - E^{\mathbb{Q}}\left[\left(\sum_i w_i R_i\right)^3\right]}_{\text{Index risk premium}} = \underbrace{E^{\mathbb{P}}\left[\sum_i w_i^3 R_i^3\right] - E^{\mathbb{Q}}\left[\sum_i w_i^3 R_i^3\right]}_{\text{Basket of individual risk premiums}} \tag{13}$$

$$\underbrace{+\, E^{\mathbb{P}}\left[3\sum_{i,j} w_i w_j^2 R_i R_j^2 + 6\sum_{i,j,k} w_i w_j w_k R_i R_j R_k\right] - E^{\mathbb{Q}}\left[3\sum_{i,j} w_i w_j^2 R_i R_j^2 + 6\sum_{i,j,k} w_i w_j w_k R_i R_j R_k\right]}_{\text{Coskewness risk premium}}$$

I name the second term of the right hand side of Equation 13 'Coskewness risk premium' because it is a factor that measures the correlation between the stocks in the tails of their distributions. The factor is composed of two terms, $\sum_{i,j} w_i w_j^2 R_i R_j^2$ and $\sum_{i,j,k} w_i w_j w_k R_i R_j R_k$. The first term is simply the definition of non-standardized coskewness, which measures the correlation of the return of each stock i with the variance of the other stocks j. The second term is a signed correlation, which is positive if the stocks tend to go up together and negative if the stocks tend to go down together.

The coskewness risk premium defined by Equation 13 is *tradable* with the following two strategies, which together constitute my *coskewness swap*:

1. I go long the skewness swap of Section 2.1 on the index to obtain the realized index skewness risk premium (i.e., the term $(\sum_i w_i R_i)^3$);

2. I go short a basket of skewness swaps on the stocks which are part of the index, each with weight w_i^3, to obtain the basket of individual realized risk premiums (i.e., the term $\sum_i w_i^3 R_i^3$).

The mathematical identity of Equation 13 shows that the payoff of the strategy that goes long the index skewness swap and short the basket of skewness swaps measures the realized marketwide coskewness risk premium. The return captured by the skewness swap trading strategy is the continuously compounded return (see Proposition 2) while the return of Equation 13 is the simple return. However, the simple return is just the first-order Taylor approximation of the continuously compounded return because $\log(x) = (x-1) + O\left((x-1)^2\right)$ in the neighborhood of $x = 1$.

I analyze the returns of the three skewness swap strategies, namely the index skewness swap, the basket of individual swaps, and the coskewness swap, separately for the S&P500 index and for the nine sector indexes of the S&P500. I calculate the return of each swap strategy by normalizing the final payoff for the capital needed to purchase the portfolio of options as explained in Equation 11 of Section 2.2. In detail, the return of the skewness swap on the index is:

$$RP_{0,T,Index} = \frac{D(\widetilde{S_{A,Index}}) - D(\overline{S_{A,Index}})}{C(\overline{S_{A,Index}})}.$$

The return of the basket of individual swaps is:

$$RP_{0,T,Bskt} = \frac{\sum_i w_i^3 \left(D(\widetilde{S_{A,i}}) - D(\overline{S_{A,i}}) \right)}{\sum_i w_i^3 C(\overline{S_{A,i}})},$$

and the return of the coskewness swap is

$$RP_{0,T,Cosk} = \frac{D(\widetilde{S_{A,Index}}) - D(\overline{S_{A,Index}}) - \sum_i w_i^3 \left(D(\widetilde{S_{A,i}}) - D(\overline{S_{A,i}}) \right)}{C(\overline{S_{A,Index}}) + \sum_i w_i^3 C(\overline{S_{A,i}})}.$$

[Figure 2 here]

Figure 2 displays the results for the S&P500. The monthly return of the skewness swap on the S&P500 index is positive and very high; it amounts on average to 79% before the crisis and 84% after the financial crisis. The return of the basket of individual skewness swaps is much lower (32% before the crisis and 49% after the crisis), leading to a very high and positive return for the coskewness swap. These results provide evidence of a positive skewness risk premium for both the index and for individual stocks. They also show that the S&P500 skewness risk premium is mainly due to the marketwide coskewnesses and that investors have a strong preference for positive coskewness. Moreover, while the skewness risk premium of the basket increases by more than 50% after the financial crisis, the S&P500 skewness risk premium increases by only 5%. This result is surprising, because after a big systemic crisis such as the one of 2007–2009, understanding why investors require a higher crash risk premium in individual stocks but not in the market is not straightforward. Investors are pricing individual crashes with a higher probability than before, but this new crash risk is not reflected in the market crash risk or in the coskewnesses, that is, it is new idiosyncratic crash risk. In the next section, I propose a friction-based explanation based on short-selling constraints that can explain the rise in the individual skewness risk premiums but not the market skewness risk premium.

[Table 3 here]

Table 3 presents the complete results for all sector indices. The table shows that for the sector indices, the results are qualitatively the same as those for the S&P500: for all sectors I find that while the skewness risk premiums for the indices increase on average by 25% after the financial crisis, the skewness risk premiums of the individual baskets increase on average by more than 100%. The most striking result is achieved by the industrial sector where the basket skewness risk premium goes from 13% to 65%. After the financial crisis, all sectors except finance have an index skewness risk premium between 70% and 90%, which mainly reflects a coskewness risk premium. The basket of individual skewness risk premiums is significant but lower than the coskewness risk premium. Finance is an exception because it is the only sector for which the skewness risk premium of the basket after the crisis is higher than the corresponding index skewness risk premium (61.62% versus 60.39%). Moreover, the financial sector is the only one for which the index skewness risk premium and the coskewness risk premium slightly decrease after the crisis. This links my work with the paper of Kelly et al. (2016) in which the authors show that during the financial crisis, a bailout sector guarantee is priced in the financial sector index only. The guarantee lowers the systematic crash risk priced in the index compared to the idiosyncratic crash risk priced in the basket.

[Table 4 here]

To complete the analysis, Table 4 assesses the statistical significance of the post-crisis increases in the skewness risk premiums documented in Table 3. This is a difficult task because the skewness risk premiums are gains from strategies that are bets on extreme events and hence the returns are very volatile and include many outliers. In addition, I am considering monthly non-overlapping returns and hence the size of my sample is limited. Taking into account all these limitations, I compute the following statistical exercise. For each skewness swap return time series, I compute the bootstrap confidence interval for the median return in the pre-crisis and post-crisis samples separately, and then I calculate the confidence level α at which the two confidence intervals do not overlap. This α is the confidence level at which I can reject the null hypothesis that the median skewness risk premium in the pre-crisis period (2003–2007) is equal to the median skewness risk premium in the post-crisis period (2009–2014). The exercise is performed for each index skewness swap return time series, basket skewness swap return time series, and coskewness swap return time series of Table 3. Table 4 displays the results. With the exception of the finance and utility sectors, the increase in the basket skewness risk premium is statistically significant at standard confidence levels for all the sectors. The increase is also marginally significant for the S&P500 basket ($\alpha = 14\%$). The confidence level α found for each index skewness risk premium is always higher than the one found for the corresponding basket, and for some sectors the difference is striking. For example, for the industrial sector the increase in the basket skewness risk premium is statistically significant at the 1% level while the increase in the index skewness risk premium is statistically significant at the 41% level, which basically means that I do not find statistical evidence that the index skewness risk premium increased. These results show that the individual skewness risk premiums increased more than the index skewness risk premium not only in terms of economic size but also in terms of statistical significance, which is higher for single stocks than for the index.

In summary, this section documents that the skewness risk premium is positive and economically important, showing that investors fear crashes. The skewness risk premium of the index is higher than that for individual stocks, in line with the idea that investors fear market crashes more than individual stock crashes. However, while the index and coskewness risk premium did not change much through time, I document a massive increase in the individual skewness risk premiums after the financial crisis of 2007–2009. The result is quite surprising, because according to standard asset pricing theory, the idiosyncratic risk should not be priced at all in a frictionless market. However, in practice the market is not frictionless. Investors and market makers are often constrained in their trading activity, they have limited bearing capacity, and information is not fully revealed in prices. As a consequence, the idiosyncratic risk might be priced in this imperfect capital market. In Section 3.2.1, I link the idiosyncratic skewness risk premiums of individual stocks with short-selling constraints.

3.1.3 Robustness checks

In this section, I provide three additional pieces of evidence that show that the financial crisis is a structural break for single stocks but not for the market index. First, I show that the implied volatility smile for single stocks becomes steeper after the financial crisis, whereas for the market the change is less significant. Second, I analyze the data with a different measure of skewness, namely the risk-neutral skewness of Bakshi and Kapadia (2003). Finally, I show that the results are not driven by the difference in data availability between the index and stock options.

[Figure 3 here]

Figure 3 plots the average implied volatility smile for single stocks and for the S&P500 before and after the financial crisis. The graph is constructed following Bollen and Whaley (2004): for each stock and each day I divide the put and call options with maturity up to one year in five moneyness categories according to their deltas, and I average the implied volatilities of the options in each category. Finally, I average the results across stocks and across days in each of the sample periods. I define the slope of the implied volatility smile simply as the difference between the implied volatility of the options of category 1 (deep out-of-the-money puts) and the implied volatility of the options of category 5 (deep out-of-the-money calls). Before the financial crisis, the average slope for the S&P500 is 11% and for single stocks is 5%. After the financial crisis, the average slope for single stocks increases to 12% (more than 100% increase) while that for the index increases to 14%. We can see that the increase is more important for deep out-of-the-money puts (options in moneyness category 1). This empirical exercise does not rely on any measure of skewness and it shows that the structural break found in this section on the skewness of single stocks is not driven by a potential bias of the new methodology used, nor is it driven by the tenor of the strategies (1 month).

[Table 5 here]

As a second robustness check, I compute the risk-neutral skewness with the classic methodology of Bakshi and Kapadia (2003). Appendix E explains the details of the implementation. Table 5 shows that while the risk-neutral skewness of the S&P500 decreased by 34% after the financial crisis, the average risk-neutral skewness of single stocks decreased by almost 60%.

[Figure 4 here]

Finally, it is important to assess to what extent the difference in the liquidity and data availability between the index and stock options is responsible for the structural break. Indeed, while the S&P500 options have always been well traded, in particular out-of-the-money puts, the trading in the single stock options market is more recent. Figure 4 shows the time series of the range of options traded for the cross-section of stocks. We can see that the range of the options available increased after the 2007–2009 financial crisis, in particular the average moneyness of the most out-of-the money put traded on single stocks went from -1.8 standard deviations before the crisis to -2.9 standard deviations. In order to insulate the change in the skewness risk premium from the change in the range of moneyness available, I compute as an additional robustness check the return of a model-based skewness swap with constant moneyness range and constant number of options. More precisely, for each stock and index, I fit at each start day of the swap the empirical implied volatility smile with the Merton jump-diffusion implied volatility smile. Then, I consider an equispaced grid of 40 synthetic option prices, 20 calls and 20 puts, calculated with the Merton jump-diffusion model previously calibrated, which spans the moneyness range $[-4SD, +4SD]$, and I construct my skewness swap with these synthetic options. I carry out this computation separately for each stock and index and for each monthly skewness swap, that is, the model is recalibrated each month for each stock or index. In this way, the range of moneyness and the number of options are constant through time and they are the same between the stocks and the index. I compute the return of this strategy in the pre-crisis and post-crisis period. The details of the implementation of this model-based skewness swap are provided in Appendix F.

[Table 6 here]

Table 6 shows that the average return of the S&P500 model-based skewness swap is 70% before the financial crisis and 73% after the financial crisis. However, for single stocks the change is more important: the average model-based skewness risk premium goes from 41% to 49% after the crisis. The statistical significance of this change is discussed in Appendix F. It is interesting to note that the magnitude of the model-based risk premium is comparable to the tradable risk premium for single stocks and for the S&P500 displayed in Tables 1 and 3. These results confirm that the change in the skewness risk premium for single stocks is higher than the change in the risk premium for the index and that this result is not only due to an increase in the moneyness range of the options available.

3.2 Possible explanations for the skewness puzzle

What happened during the financial crisis and why did the idiosyncratic crash risk premiums increase? In this section, I discuss different hypothesis. Section 3.2.1 outlines a possible friction-based explanation based on short-selling constraints, and in Section 3.2.2 I discuss and rule out alternative hypotheses.

I extract the idiosyncratic skewness risk premium from each individual skewness risk premium by taking away the part of the risk premium that can be explained by the covariation of the stock with the market. In details I run the following time-series regression separately for each stock:

$$RP_{i,t} = \alpha_i + \beta_i RP_{S\&P500,t} + \epsilon_{i,t} \tag{14}$$

where $RP_{i,t}$ ($RP_{S\&P500,t}$) is the time series of the realized skewness risk premium of stock i (of the S&P500) measured with the gain of the monthly swap strategies given by Equation 10. I then extract the realized idiosyncratic risk premium for each stock as follows:

$$\text{Idiosyncratic risk premium}_{i,t} = RP_{i,t} - \widehat{\beta}_i RP_{S\&P500,t}. \tag{15}$$

The idiosyncratic risk premium can also be expressed as percentage gain of the swap strategy:

$$\text{Idiosyncratic risk premium}_{i,t,\%} = \frac{RP_{i,t} - \widehat{\beta}_i RP_{S\&P500,t}}{C_{i,t}(\overline{S_A})}, \tag{16}$$

where $C_{i,t}(\overline{S_A})$ is the capital needed to purchase the fixed leg of the swap of stock i on month t defined in Equation 11.

[Table 7 here]

Table 7 shows the average idiosyncratic risk premium before and after the 2007–2009 financial crisis. The table further supports the main finding of the paper: since the crisis, the idiosyncratic part of the skewness risk premium of individual stocks has increased substantially.[4]

3.2.1 A friction-based explanation: short-selling constraints

In this section, I investigate the relation between short-selling constraints and idiosyncratic crash risk premium. First, I compute a cross-sectional analysis of the short-selling constraints over the *whole sample period*, and then I discuss in more details the short-selling constraints during the *financial crisis*.

Cross-sectional analysis over the whole sample period.

By exploiting the cross-sectional heterogeneity of the idiosyncratic crash risk premiums, I show that the higher the short-selling constraints, the higher the idiosyncratic crash risk premium. I outline two economic channels that lead to this relation: (a) stocks that are hard to short-sell are more likely to be overvalued and hence have a higher crash risk premium; and (b) stocks that are hard to short-sell have higher volumes in put options, which shows that investors resort to the options market as an alternative to shorting the stock. The higher volumes in put options in turn push up the crash risk premium of the stock.

I use two variables as proxies for short-selling costs: the short interest ratio (SI) and the option-implied lending fee of Muravyev et al. (2016). The short interest ratio is defined for each stock as the ratio of the number of shares held short against the total number of shares outstanding; this proxies the demand for short selling (see Asquith et al. (2005) and Stilger et al. (2016)). The idea is that the higher the short interest ratio, that is, the higher the demand for shorting, the more difficult and costly it is to short-sell the stock. The option-implied lending fee h_{imp} defined by Muravyev et al. (2016) is based on the idea that an investor can alternatively short the stock in the options market by purchasing a put option and selling a call option with the same strike and same maturity. h_{imp} measures the costs of shorting the stock in the options market. In this section, I report the analysis only for the short interest ratio, while the results for the option-implied lending fee are reported in Appendix G because they are qualitatively identical. Appendix G also reports the details of the calculation of the option-implied lending fee.

[4] The detailed results of the quantile regressions can be found in Table 20 of Appendix H; Figure 10 of Appendix I breaks down the results by sector.

[Table 8 here]

The cross-sectional relation between the short interest ratio and the idiosyncratic skewness risk premium is investigated via a portfolio sort and a Fama MacBeth regression. In Panel A of Table 8, I form three portfolios sorted according to the idiosyncratic skewness risk premium, with ptf1 (ptf3) denoting the portfolio of stocks with the lowest (highest) risk premium. I then compute the average ex-ante short interest ratio of each portfolio, which is computed at the start date of the swaps. Compustat reports the short interest of each stock every two weeks; hence, I take as ex-ante short interest the closest short interest reported by Compustat before the start date of each swap. The average short interest ratio of the portfolios ranges from 3.62% of portfolio 1 to 4.46% of portfolio 3 with a monotonic increasing pattern, and the t-statistic of 6.59 shows that the short interest ratio of portfolio 3 is statistically higher than the short interest ratio of portfolio 1. In Panel B of Table 8, I compute the same univariate portfolio sort but in the reverse order: first I sort the stocks according to their short interest ratio and then I compute the average ex-post idiosyncratic skewness risk premium for each portfolio. The results show that the portfolio of stocks that are harder to short-sell commands a higher idiosyncratic crash risk premium, which is both economically and statistically higher than the risk premium of the portfolio of stocks that are comparably easier to short-sell. The two portfolio sorts show that the relation between short-selling constraints and crash risk premium is necessary and sufficient. The results are confirmed by the Fama MacBeth regression prsented in Panel C of Table 8. The loading on the short interest ratio is positive and statistically significant, confirming that the higher the short-sale constraints, the higher the idiosyncratic skewness risk premium. This connection reveals that there are short-sale constraints in the market and investors require compensation for the associated risk.

Why do stocks that are harder to short-sell command a higher idiosyncratic crash risk premium? I identify two economic channels: first, the stocks with higher short-sale constraints are more likely to be overvalued, and overvalued stocks command higher crash risk premiums. And second, the stocks that are harder to short-sell have higher volumes in put options, which in turn decreases the risk-neutral skewness of the stock and increases the skewness risk premium.

[Table 9 here]

As a measure of overvaluation of stocks, I use the maximum daily return over the previous month (MAX) as suggested by Bali et al. (2011), which was also used by Stilger et al. (2016). In their work Bali et al. (2011) show that the stocks with the highest MAX value in the past month significantly underperform the stocks with the lowest MAX value in the following month. This relation shows that the stocks with the highest MAX value are more likely to be overpriced. I compute a univariate portfolio sort on the basis of the short interest ratio and I calculate the average MAX value for each SI-sorted portfolio. Panel A of Table 9 shows that the stocks in the upper SI tercile portfolio have a higher value of MAX and that the relation between SI and MAX is monotonic and highly significant (t-statistic of the difference portfolio is 9.62). The stocks that are harder to short-sell are more likely to be overpriced. I then investigate the joint relation between SI, MAX, and the idiosyncratic skewness risk premium with a double independent portfolio sort. At each start date of the swaps t, I independently sort the stocks in ascending order according to their SI estimate and MAX estimate, and I assign each stock to one of the three tercile portfolios for each sorting variable. The intersection of these classifications yields nine portfolios. I then calculate the average idiosyncratic skewness risk premium of these nine portfolios at the end of the following month $t+1$ (i.e., ex-post risk premium). The results are displayed in Panel B of Table 9. When I control for the level of MAX, the relation between SI and the idiosyncratic skewness risk premium is significant only for the stocks that have a medium/high level of MAX. Indeed, the portfolio with the highest idiosyncratic skewness risk premium (6.95%) is the portfolio in the upper SI tercile and in the upper MAX tercile. The stocks that are harder to short-sell *and* that are overpriced command higher idiosyncratic skewness risk premiums.

[Table 10 here]

The second mechanism is a very simple market substitution mechanism: if an investor wants to short-sell a stock but it is too expensive to do it in the equity market, he or she might resort to the options market as a cheaper alternative by simply buying puts or a combination of long puts and short calls, which replicates a synthetic short position in the stock. This in turn creates short-selling pressure in the options market and, according to the demand-based option pricing of Garleanu et al. (2009), this pressure will push up the prices of the put options and push down the prices of the call options. Ceteris paribus, the resulting risk-neutral skewness is therefore more negative and the resulting skewness risk premium is higher. In order to study the connection between short-selling constraints and option trading, I define a new variable, the volume ratio VR, which is calculated as the ratio of the volume of out-of-the-money calls against the volume of out-of-the-money puts traded on that day on the stock. I first compute a univariate portfolio analysis wherein for each month I sort the stocks into three tercile portfolios according to their short interest ratio, and then I calculate the average VR of each portfolio. The VR and the short interest ratio are calculated at each start date of the swaps, that is, the two measures are contemporaneous. Panel A of Table 10 shows that the portfolio of stocks with the highest short interest ratios has the lowest volume ratio, the pattern is monotonic through the portfolios, and the t-statistic of the difference between the volume ratio of portfolio 1 and portfolio 3 is -3.62. This pattern reveals that for the stocks with high short interest ratios, the put options are more frequently traded than for the stocks with low short interest ratios. This supports the idea that investors resort to the options market as an alternative for shorting a stock when shorting the stock directly in the equity market is difficult. The joint relation between SI, VR, and the idiosyncratic skewness risk premium is studied via a double independent portfolio sort. The results displayed in Panel B of Table 10 show that the relation between SI and the idiosyncratic crash risk premium is stronger when VR is low (i.e., portfolio 1 of VR). The idiosyncratic skewness risk premium is the highest for the portfolio of stocks that are more difficult to short sell *and* have a high volume in put options.

Short-selling constraints during the financial crisis.

Short-selling constraints played a dominant role in the financial crisis. In October 2008, naked short-selling was banned, which implies that to short-sell a stock, investors first have to borrow it. However, at the same time, the work of Porras Prado et al. (2016) and the Securities and Exchange Commission (SEC) reports of Baklanova et al. (2015, 2016) document that securities lending activity has decreased substantially since the 2007–2009 crisis.

[Figure 5 here]

Figure 5 is taken from Baklanova et al. (2015); it shows that the volume of securities on loan dropped dramatically after 2009. There are several reasons for this drop. As Baklanova et al. (2015) argue, the decline in securities lending can be partly attributed to policymakers efforts to reduce short selling. During the crisis, the SEC issued several emergency orders that tightened borrowing requirements for shares in the largest financial firms. These efforts were aimed at minimizing the possibility of abusive short selling and preventing potential sudden and excessive fluctuations in security prices that could impair markets. In addition, many security lenders restricted their participation in securities lending activity due to concerns about weak performance and outright losses on cash collateral reinvestment portfolios.[5] All of these elements suggest that since the financial crisis, short selling is more difficult because of increased regulation. The index is less affected by the new regulation, because borrowing ETF shares is easier and requires less quality collateral. Indeed, Li and Zhu (2017) report that the ETF short interest increased massively through time and the authors even provide evidence that investors decide to short ETFs instead of single stocks when shorting single stocks is too costly.

[Table 11 here]

I test whether the stocks that were most strongly hit by the short-selling restrictions during the financial crisis had a subsequent greater increase in idiosyncratic skewness risk premiums. To this end, I compute a cross-sectional regression of the average change in the idiosyncratic skewness risk premium on three measures that quantify how severely each stock was hit by the frictions during the financial crisis. These three measures are the maximum of (a) the option-implied lending fee of Muravyev et al. (2016), (b) the short interest ratio, and (c) the overvaluation proxy of Bali et al. (2011), achieved by each stock during the financial crisis (2007–2009). The results are displayed in Table 11 and show that, except for the short interest ratio, the loadings on the other two measures are positive and statistically significant: the stocks with the highest values in the friction measures during the financial crisis show subsequently higher increases in the idiosyncratic skewness risk premium.

3.2.2 Alternative explanations

In this section, I discuss the following alternative explanations, which could potentially explain the skewness risk premium puzzle: (a) increase in investor *risk-aversion*, (b) increase in *options data availability*, (c) increase in the riskiness of firms in terms of *closeness to default*, and (d) increase in the *correlations* among stocks.

[5] An anecdotal case is the one of AIG, which used its securities lending program as a mechanism to raise cash and generate leverage. AIG almost collapsed in 2008 and risk management inefficiencies and misuse of the cash collateral played an important negative role in its distress.

Risk-aversion.

I can confidently rule out the hypothesis that investors became more risk averse (see the option-implied risk-aversion methodology of Bliss and Panigirtzoglou (2004)). Indeed, if the increase in the risk premiums after the financial crisis were due to an increase in the risk aversion of investors, I should have found an increase in the risk premium for the market index as well, if we believe that the options market is integrated across securities.

Options data availability.

I can also rule out the hypothesis that the risk premiums increased because the moneyness range of available options increased after the financial crisis (see Figure 4). Indeed, in Section 3.1.2, I showed that if I construct model-based skewness swaps with constant moneyness range $[-4SD, +4SD]$ in all sample periods, the results still hold, that is, the individual model-based skewness risk premiums increase more than the market model-based skewness risk premiums after the financial crisis.

Closeness to default.

Another possible explanation could be that firms became idiosyncratically riskier after the financial crisis in terms of their closeness to default. To test this hypothesis, I compute a univariate portfolio sort of the stocks based on their idiosyncratic skewness risk premium and then I compute the average book-leverage of each portfolio. The results are displayed in Table 21 of Appendix H and show that the relation is actually negative: the stocks with the highest idiosyncratic skewness risk premium are those with the lowest book-leverage. Therefore, the increase in book-leverage cannot be responsible for the increase in idiosyncratic skewness risk premiums.

Correlations.

Finally, another possible explanation could be that individual skewness risk premiums increased because the correlation between the stock and the market increased after the financial crisis, and therefore, individual crash risk premiums are catching up with the market crash risk premium. However, the focus of this paper is to understand the increase in the idiosyncratic crash risk premium, and in Equation 14, the part of the risk premium due to the covariation between the stock and the market is removed. Indeed, Table 22 of Appendix H fails to detect a relation between the time-varying CAPM-beta of the stocks and the idiosyncratic skewness risk premiums.

4 The rational expectation equilibrium model

In this section, I show how skewness risk can endogenously arise in a very simple rational expectation equilibrium model with asymmetric information and short-selling constraints. The model is very basic, but it is rich enough to provide an economic sense on how the interplay between short-selling costs and demand and supply frictions can generate a negatively skewed conditional distribution of the final payoff.

The model is a simplified version of the model of Venter (2016), characterized by a two-period economy with dates $t = 0$ and 1. At date 0 investors trade, and at date 1 assets pay off. There are two securities traded in a competitive market, a risk-free bond and a risky stock. The bond is in perfectly elastic supply and is used as the numeraire, with the risk-free rate normalized to 0. The risky asset is assumed to be in net supply of $S > 0$ and has a final dividend payoff d at date 1 that is the sum of two random components: $d = f + n$. The first risky component of the dividend payoff, f, can be regarded as the fundamental value of the asset. The second component, n, is thought of as additional noise, preventing informed agents from knowing the exact dividend payoff. The price of the stock at date 0 is denoted by p. f has an improper uniform distribution, while $n \sim N(0, \sigma_n^2)$.

The market is populated by a unit mass of risk-averse rational traders with mean-variance preferences and risk-aversion ρ. A fraction $[0, \lambda)$ of them are informed and know the final payoff f. The others $[\lambda, 1]$ are uninformed and they form their expectations based only on the price p. A fraction w of the informed traders are subject to short-sale constraints, which in this economy are modeled as the inability to sell the stock. This market also contains noise traders whose trading activity is not derived from utility maximization. Noise traders simply buy u shares, where $u \sim N(0, \sigma_u^2)$.

The definition of the equilibrium is standard in this economy.

Definition 1. *A rational expectation equilibrium (REE) of the asset market is a collection of a price function $P(f, u)$ and individual strategies for constrained informed, unconstrained informed, and uninformed traders $x^c(f, p)$, $x^{uc}(f, p)$, and $x^{ui}(p)$ such that:*

1. *The demand x^k is optimal $\forall k$, i.e.,*

$$x^k \in argmax \quad E\left[W_k \mid \mathcal{I}_k\right] - \frac{\rho}{2} Var\left[W_k \mid \mathcal{I}_k\right],$$

where W_k is the terminal wealth of agent k: $W_k = W_0 + x_k(d - p)$;

2. *market clears $\forall u$, i.e.,*

$$w\lambda x^c + (1 - w)\lambda x^{uc} + (1 - \lambda)x^{ui} + u = S.$$

The following theorem shows that this economy admits a unique equilibrium.

Theorem 1. *A unique piecewise linear REE of the model exists in the form*

$$P = f + \begin{cases} A(u - C) & u \leq C \\ B(u - C) & u > C \end{cases} \tag{17}$$

where $A = (\rho\sigma_n^2)/(\lambda)$, $B = (\rho\sigma_n^2)((1 - w)\lambda)$, and C is the solution of an implicit function $0 = L(C)$. The constant C is well defined and unique.

In equilibrium the quantities traded are, respectively, $x^{uc}(f, p) = (f - p)/(\rho\sigma_n^2)$, $x^c(f, p) = \max\{(f - p)/(\rho\sigma_n^2), 0\}$, $x^{ui}(p) = (D)/(\rho(E - D^2 + \sigma_n^2))$, where $D = -A\int_{-\infty}^{C}(u - C)\phi(u)du - B\int_{C}^{\infty}(u - C)\phi(u)du$ and $E = A^2\int_{-\infty}^{C}(u - C)^2\phi(u)du + B^2\int_{C}^{\infty}(u - C)^2\phi(u)du$.

Proof. See Appendix C. ∎

Intuitively, the equilibrium price function P has a different sensitivity to u (the demand of the uninformed agents) if the stock is overvalued or undervalued, because there are short-sale constraints. If the stock is overvalued, that is, $P > f$ or $u > C$, not all the informed investors can sell the stocks, and hence the relative importance of u in the market clearing condition is higher. This is why $B > A$ and in the case of no short-sale constraints, that is, $w = 0$, we obtain the equality $A = B$.

[Figure 6 here]

The main prediction of the model is that with short-sale constraints, the distribution of the final payoff that can be inferred from the prices, that is, $f|P$, is *left skewed.* Figure 6 shows the result graphically. Mathematically, from the equilibrium price function of Equation 17, the distribution of $f|P$ is Gaussian without short-sale constraints because if $w = 0$, then $A = B$ and the price function is simply linear. If $w > 0$, then the distribution of $f|P$ is instead piecewise Gaussian, the left tail has a higher variance because $B > A$, and the resulting skewness is negative. Economically, the interpretation is very intuitive: the short-sale constraints prevent the bad news from being fully incorporated in the prices, which leads to larger realized losses when the final payoff is realized. Short-sale constraints thus generate skewness risk.

In this equilibrium economy, there is also a positive equity risk premium because the price of the stock is less than the conditional expectation of its future payoff:

$$E\left[f \mid P\right] = P + D,$$

where $D = E[\mathbb{I}_{u\leq C}(-A(u-C))]+E[\mathbb{I}_{u>C}(-B(u-C))] = -A\int_{-\infty}^{C}(u-C)\phi(u)du - B\int_{C}^{\infty}(u-C)\phi(u)du \geq 0$ and ϕ is the Gaussian distribution with mean 0 and variance σ_u^2. Mathematically, the risk premium arises from the constant C in the price function, which is solved in order to clear the market and meet the positive supply S. If $w = 0$ (no short-sale constraints), then

$$D = \frac{\rho\sigma_n^2(\rho^2\sigma_n^2\sigma_u^2 + \lambda^2)}{\lambda\rho^2\sigma_n^2\sigma_u^2 + \lambda^2}S.$$

Economically, the risk premium is required by uninformed traders in order to trade, as $x^{ui} = D/(\rho(E - D^2 + \sigma_n^2))$. If the price of the stock is less than its expected payoff, the price of a put option on the stock will be higher than the expected payoff of the option, and this generates a positive crash risk premium. The simple linear structure of the risk premium $E[f \mid P] = P + D$ shows that the more the prices reveal, the lower the risk premium.

[Figure 7 here]

Figure 7 shows that D is monotonic in the friction variables w and λ. w, the proportion of constraint investors, quantifies the short-selling friction in the economy, while λ, the proportion of informed investors, quantifies the information friction in the economy. When short-sale constraints are high (i.e., w is high) or when the information is more asymmetric (i.e., λ is low) prices are less revealing and the risk premium is higher.

5 Conclusion

In this work I implement a trading strategy to investigate the risk premium associated with the third moment of a stocks' return distribution. The strategy involves buying and selling out-of-the-money put and call options in order to take a position in the underlying skewness and subsequently hedge in the forward market. In this way, I develop a strategy that is independent of the first, second, and fourth moments, and it is a pure bet on the skewness. The return of the strategy measures the realized skewness risk premium.

I apply this strategy to the 500 constituents of the S&P500 in the period 2003–2014 and I find that the skewness risk premium is positive and statistically significant for the majority of the stocks, in particular after the financial crisis of 2007–2009. A positive skewness risk premium implies that the price of the skewness ($\mathbb{Q}$ skewness) is lower than the realized skewness ($\mathbb{P}$ skewness). These results are consistent with the theoretical model of Bakshi et al. (2003), which shows that because investors' preferences are towards positive skewness, the price of the skewness should be lower than the realized skewness.

I decompose the individual skewness risk premiums both systematically, common to all stocks, and idiosyncratically, peculiar to each stock. Although the systematic part is the largest component of the risk premium, I find that the idiosyncratic part also carries a positive loading. Based on this result, I reject the null hypothesis of a frictionless capital market in which the idiosyncratic risk is not priced because it can be diversified away. I therefore study which violations of the perfect capital market hypothesis can generate the idiosyncratic skewness risk premium. My results reveal that short selling constraints are associated with higher idiosyncratic crash risk premiums.

References

T. G. Andersen, N. Fusari, and V. Todorov. Short-term market risks implied by weekly options. *The Journal of Finance*, forthcoming, 2017.

D. W. K. Andrews and J. C. Monahan. An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica*, 60(4):953–966, 1992.

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

P. Asquith, P. A. Pathak, and J. R. Ritter. Short interest, institutional ownership, and stock returns. *Journal of Financial Economics*, 78(2):243–276, 2005.

A. Atmaz and S. Basak. Option prices and costly short-selling. *Working paper*, 2017.

V. Baklanova, A. Copeland, and McCaughrin R. Reference guide to u.s. repo and securities lending markets. *SEC Report*, 2015.

G. Bakshi and N. Kapadia. Delta-hedged gains and the negative market volatility premium. *Review of Financial Studies*, 16(2):527–566, 2003.

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

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

T. G. Bali, N. Cakici, and R. F. Whitelaw. Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, 99(2):427–446, 2011.

K. Barraclough and R. E. Whaley. Early exercise of put options on stocks. *The Journal of Finance*, 67 (4):1423–1456, 2012.

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

R. Battalio and P. Schultz. Regulatory uncertainty and market liquidity: The 2008 short sale ban's impact on equity option markets. *The Journal of Finance*, 66(6):2013–2053, 2011.

J.-F. Begin, C. Dorion, and G. Gauthier. Idiosyncratic jump risk matters: Evidence from equity returns and options. *Working paper*, 2017.

R. R. Bliss and N. Panigirtzoglou. Option-implied risk aversion estimates. *The journal of finance*, 59 (1):407–446, 2004.

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

T. Bollerslev and V. Todorov. Tails, fears, and risk premia. *The Journal of Finance*, 66(6):2165–2211, 2011.

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

T. Bollerslev, V. Todorov, and L. Xu. Tail risk premia and return predictability. *Journal of Financial Economics*, 118(1):113–134, 2015.

B. Boyer, T. Mitton, and K. Vorkink. Expected idiosyncratic skewness. *Review of Financial Studies*, 23 (1):169–202, 2010.

L. M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217, 1967.

A. Buraschi, F. Trojani, and A. Vedolin. When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia. *The Journal of Finance*, 69(1):101–137, 2014.

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

P. Carr and D. Madan. Optimal positioning in derivative securities. *Quantitative Finance*, 2001.

P. Carr and L. Wu. Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341, 2009.

P. Carr and L. Wu. A simple robust link between american puts and credit protection. *Review of Financial Studies*, 24(2):473–505, 2011.

B. Y. Chang, P. Christoffersen, and K. Jacobs. Market skewness risk and the cross section of stock returns. *Journal of Financial Economics*, 107(1):46–68, 2013.

J. Chen, H. Hong, and J. C. Stein. Breadth of ownership and stock returns. *Journal of financial Economics*, 66(2):171–205, 2002.

H. Choi, P. Mueller, and A. Vedolin. Bond variance risk premiums. *Review of Finance*, 21(3):987–1022, 2017.

P. Christoffersen and K. Jacobs. The importance of the loss function in option valuation. *Journal of Financial Economics*, 72(2):291–318, 2004.

P. Christoffersen, M. Fournier, K. Jacobs, and M. Karoui. Option-based estimation of the price of co-skewness and co-kurtosis risk. *Working paper*, 2016.

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

A. Cosma, S. Galluccio, P. Pederzoli, and O. Scaillet. Early exercise decision in american options with dividends, stochastic volatility and jumps. *Working paper*, 2017.

M. Cremers, M. Halling, and D. Weinbaum. Aggregate jump and volatility risk in the cross-section of stock returns. *The Journal of Finance*, 70(2):577–614, 2015.

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

H. Desai, K. Ramesh, S. R. Thiagarajan, and B. V. Balachandran. An investigation of the informational role of short interest in the nasdaq market. *The Journal of Finance*, 57(5):2263–2287, 2002.

J. Driessen, P. J. Maenhout, and G. Vilkov. The price of correlation risk: Evidence from equity options. *The Journal of Finance*, 64(3):1377–1406, 2009.

L. Eeckhoudt and H. Schlesinger. Putting risk in its proper place. *The American Economic Review*, 96 (1):280–289, 2006.

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

E. F. Fama and K. R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116 (1):1–22, 2015.

N. Garleanu, L. H. Pedersen, and A. M. Poteshman. Demand-based option pricing. *Review of Financial Studies*, 22(10):4259–4299, 2009.

E. Gourier. Pricing of idiosyncratic equity and variance risks. *Working paper*, 2016.

P. S. Hagan, D. Kumar, A. S. Lesniewski, and Woodward D. E. Managing smile risk. *The Best of Wilmott*, page 249, 2002.

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

C. M. Jones and O. A. Lamont. Short-sale constraints and stock returns. *Journal of Financial Economics*, 66(2):207–239, 2002.

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

B. Kelly, H. Lustig, and S. Van Nieuwerburgh. Too-systemic-to-fail: what option market imply about sector-wide government guarantees. *The American Economic Review*, 106(6):1278–1319, 2016.

R. Kozhan, A. Neuberger, and P. Schneider. The skew risk premium in the equity index market. *Review of Financial Studies*, 26(9):2174–2203, 2013.

F. W. Li and Q. Zhu. Short selling etfs. *Working paper*, 2017.

T. Lin and X. Lu. How do short-sale costs affect put options trading? evidence from separating hedging and speculative shorting demands. *Review of Finance*, 20(5):1911–1943, 2015.

I. Martin. What is the expected return on the market? *The Quarterly Journal of Economics*, 132(1): 367–433, 2017.

R. C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2):125–144, 1976.

E. M. Miller. Risk, uncertainty, and divergence of opinion. *The Journal of Finance*, 32(4):1151–1168, 1977.

D. Muravyev, N. D. Pearson, and J. M. Pollet. Is there a risk premium in the stock lending market? evidence from equity options. *Working paper*, 2016.

S. Nagel. Short sales, institutional investors and the cross-section of stock returns. *Journal of Financial Economics*, 78(2):277–309, 2005.

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

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

I. Piatti. Heterogeneous beliefs about rare event risk in the lucas orchard. *Working paper*, 2015.

V. K. Pool, H. R. Stoll, and R. E. Whaley. Failure to exercise call options: An anomaly and a trading game. *Journal of Financial Markets*, 11(1):1–35, 2008.

M. Porras Prado, P. Saffi, and J. Sturgess. Ownership structure, limits to arbitrage, and stock returns: Evidence from equity lending markets. *The Review of Financial Studies*, 29(12):3211–3244, 2016.

P. Schneider and F. Trojani. Fear trading. *Swiss Finance Institute Research Paper*, 15(3), 2015.

P. Schneider, C. Wagner, and J. Zechner. Low risk anomalies? *Working paper*, 2016.

P. S. Stilger, A. Kostakis, and S. Poon. What does risk-neutral skewness tell us about future stock returns? *Management Science*, 2016.

G. Venter. Short-sale constraints and credit runs. *Working paper*, 2016.

A Tables

Table 1. Average return of the skewness swap on individual stocks.

Panel A of the table shows the median return of the skewness swap across the stocks in the pre-/post-crisis subsamples, as well as the median fixed leg and floating leg of the swap. For each stock I consider the median of the skewness swap returns across time; the table displays the average of these medians across stocks. The same averaging procedure is computed for the fixed leg and floating leg of the swap. The numbers in parentheses indicate the number of stocks that have a statistically significant positive swap return, where the confidence interval for the median swap return is calculated with a bootstrap technique with 2000 bootstrap samples. The table shows the skewness swap returns computed using American options, European options, and American options with bid/ask prices. The $\mathbb{Q}$ skewness ($\mathbb{P}$ skewness) is the fixed leg (floating leg) of the swap standardized by variance, that is, $\mathbb{Q}$ skewness $= \frac{\text{fixed leg}}{(\mathbb{Q}\text{-variance})^{3/2}}$ and $\mathbb{P}$ skewness $= \frac{\text{floating leg}}{(\mathbb{Q}\text{-variance})^{3/2}}$. In Panel A, I consider only the 209 stocks that have good option data coverage since the beginning of the sample. In Panel B (Panel C), I consider the median return of an equally weighted (value-weighted) portfolio of skewness swaps. In Panels B and C, I consider all the stocks of my sample in the portfolio construction.

Panel A: Summary statistics for individual skewness swaps

	Before crisis (2003–2007)	During crisis (2007–2009)	After the crisis (2009–2014)
Average return of the skewness swap	43% (154)	43% (91)	52% (202)
Average Sharpe ratio	1.54	0.81	1.55
Average return of the skewness swap with European options	43% (155)	42% (82)	52% (201)
Average return of the skewness swap with transaction costs	26% (94)	29% (49)	39% (181)
Average fixed leg of the swap ($\mathbb{Q}-$skewness)	-0.65	-0.82	-0.99
Average floating leg of the swap ($\mathbb{P}-$skewness)	0.13	-0.03	0.07

Panel B: Portfolio of equally weighted skewness swaps

	Before crisis (2003–2007)	During crisis (2007–2009)	After the crisis (2009–2014)
Median return of the portfolio	33.50%	43.49%	47.83%

Panel C: Portfolio of value-weighted skewness swaps

	Before crisis (2003–2007)	During crisis (2007–2009)	After the crisis (2009–2014)
Median return of the portfolio	32.40%	43.32%	45.80%

Table 2. Univariate portfolio analysis: skewness swap return vs variance swap return.

This table presents the results of two univariate portfolio sort analyses. In Panel A, the stock portfolios are formed each month according to the stocks skewness swap return with ptf1 (ptf3) denoting the portfolio with the lowest (highest) skewness swap return. Then, for each portfolio the table reports the average variance swap return. The variance and skewness swap returns are the monthly percentage gain of the variance and skewness swaps, respectively, and they are contemporaneous. The column labeled '3-1 t-stat' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that ptf3 and ptf1 have the same average variance swap return. In Panel B, the sorting variable is the variance swap return and the table reports the statistics for the skewness swap return of each variance-sorted portfolio. The t-statistics are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

Panel A: Skewness swap return				
	ptf1	ptf2	ptf3	3-1 t-stat
Variance swap return	-10.92%	-28.91%	-11.37%	-0.22
Panel B: Variance swap return				
	ptf1	ptf2	ptf3	3-1 t-stat
Skewness swap return	38.65%	38.25%	29.03%	-3.77

Table 3. The return of the index skewness swap, the basket of individual skewness swaps, and the coskewness swap.

This table presents the median monthly returns of the three following swap trading strategies in the pre-/post-crisis subsamples:

1. Index skewness swap return: $\left(\sum_i w_i r_i\right)^3$. This strategy trades the third moment of the returns of the index. It is computed with the skewness swap of Section 1 applied to the options of the index (S&P500 or SPDR sector ETF).
2. Basket skewness swap return: $\sum_i w_i^3 r_i^3$. This strategy trades the sum of the third moments of the index constituents. It is computed with a portfolio of skewness swaps applied independently to the options of each stock, each one with weight w_i^3.
3. Coskewness swap return: $\left(\sum_i w_i r_i\right)^3 - \sum_i w_i^3 r_i^3 = 3\sum_{i,j} w_i w_j^2 r_i r_j^2 + 6\sum_{i,j,k} w_i w_j w_k r_i r_j r_k$. This strategy is implemented by going long strategy 1 and short strategy 2.

The statistical significance of the median returns is assessed by computing the bootstrap confidence intervals with 2000 bootstrap samples at the levels of 5%, 1%, and 0.1%.

Panel A: Before the crisis 2003–2007			
	Index skewness swap return	Basket skewness swap return	Coskewness swap return
S&P500	78.91%***	32.81%***	78.88%***
Financial sector	60.49%***	45.15%***	59.86%***
Energy sector	44.12%**	40.28%***	38.97%*
Material sector	75.53%***	41.75%**	73.85%***
Industrial sector	68.78%***	13.07%*	56.59%***
Consumer discretionary sector	65.44%***	30.18%***	65.08%***
Consumer staples sector	67.52%***	39.84%***	59.81%**
Health care sector	54.00%**	30.22%***	52.98%**
Technology sector	45.64%**	28.18%***	45.46%***
Utility sector	79.34%***	59.68%***	79.15%***

Panel B: After the crisis 2009–2014			
	Index skewness swap return	Basket skewness swap return	Coskewness swap return
S&P500	84.17%***	49.13%***	84.13%***
Financial sector	60.39%***	61.62%***	59.77%***
Energy sector	74.98%***	67.27%***	72.20%***
Material sector	75.61%***	59.46%***	74.46%***
Industrial sector	76.01%***	64.43%***	75.38%***
Consumer discretionary sector	84.16%***	65.09%***	83.04%***
Consumer staples sector	80.97%***	66.89%***	75.50%***
Health care sector	79.77%***	62.57%***	78.59%***
Technology sector	72.57%***	52.50%***	70.87%***
Utility sector	85.68%***	66.91%***	83.36%***

Table 4. The statistical significance of the post-crisis increase in the skewness risk premiums.

This table shows the confidence level α at which I can reject the null hypothesis that the median skewness risk premium in the pre-crisis period (2003–2007) is equal to the median risk premium in the post-crisis period (2009–2014). The α is calculated as the confidence level under which the bootstrap confidence interval of the median skewness risk premium in the pre-crisis period does not overlap with the bootstrap confidence interval of the median skewness risk premium in the post-crisis period. The exercise is performed for each time series of index skewness swap returns, basket skewness swap returns, and coskewness swap returns of Table 3.

	α_{index}	α_{basket}	α_{cosk}
S&P 500	0.20	0.14	0.15
Financial sector	1.00	0.35	1.00
Energy sector	0.08	0.06	0.08
Material sector	0.97	0.04	0.90
Industrial sector	0.42	0.01	0.25
Consumer discretionary sector	0.28	0.01	0.28
Consumer staples sector	0.45	0.01	0.16
Health care sector	0.02	0.01	0.02
Technology sector	0.11	0.07	0.09
Utility sector	0.68	0.26	0.66

Table 5. Robustness check: Bakshi-Kapadia-Madan (2003) risk-neutral skewness.

The table shows the average risk-neutral skewness of Bakshi and Kapadia (2003) in the pre-/post-crisis subsamples for the S&P500 and for the cross-section of stocks. The risk-neutral skewness is computed daily for each stock by using the options with 30 days to maturity. The results are first averaged for each stock in each of the sample periods, and then they are averaged across stocks. The column labeled 't-stat diff' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that the average risk-neutral skewness in the pre-crisis period (2003–2007) is equal to the average risk-neutral skewness in the post-crisis period (2009–2014). The standard errors are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992). The number in parentheses indicates the number of stocks for which the change in the risk neutral skewness is statistically significant (t-stat bigger than 2).

The Bakshi-Kapadia-Madan (2003) risk-neutral skewness			
	Before crisis (2003–2007)	After the crisis (2009–2014)	t-stat-diff
S&P500	-0.53	-0.74	-14.25
Average across stocks	-0.13	-0.21	-4.89
			(310)

Table 6. Robustness check: the model-based skewness risk premium.

This table shows the median return of the model-based skewness swaps for the S&P500 and for the cross-section of stocks in the pre-crisis and post-crisis subsamples. The model-based skewness swap is constructed at each start date of the swap by using the option prices calculated according to the Merton jump-diffusion model that best fits the data. The Merton jump-diffusion model is calibrated separately for each stock, and it is recalibrated at each start date of the swap. Then, I consider an equispaced grid of 40 synthetic option prices, 20 calls and 20 puts, calculated with the Merton jump-diffusion model just calibrated, which spans a constant moneyness range $[-4SD, +4SD]$, where $SD = (\log(K/F))/(\sigma\sqrt{T})$ is the moneyness of the options measured in standard deviations. I finally construct my skewness swap with these synthetic options according to Equation 1 (fixed leg of the swap) and Equation 2 (floating leg of the swap).

The model-based skewness risk premium		
	Before crisis (2003–2007)	After the crisis (2009–2014)
S&P500	69.98%	72.83%
Average across stocks	40.68%	49.20%

Table 7. Idiosyncratic skewness risk premium.

The table shows the median idiosyncratic skewness risk premium across the stocks in the pre-/post-crisis subsamples. The idiosyncratic skewness risk premium is calculated separately for each stock as

$$\text{Idiosyncratic risk premium}_{i,t,\%} = \frac{RP_{i,t} - \widehat{\beta}_i RP_{S\&P500,t}}{C_{i,t}(\overline{S_A})},$$

where $RP_{i,t}$ ($RP_{S\&P500,t}$) is the dollar gain of the skewness swap strategy of Section 1 implemented on stock i (on the S&P500), $\widehat{\beta}_i$ is estimated with the median regression of Equation 14, and $C_{i,t}(\overline{S_A})$ is the capital needed to purchase the fixed leg of the swap of stock i on month t defined in Equation 11. The numbers in parentheses indicate the number of stocks that have a statistically significant positive and negative idiosyncratic risk premium, where the confidence interval for the median risk premium is calculated with a bootstrap technique and 2000 bootstrap samples.

The idiosyncratic skewness risk premium			
	Before crisis (2003–2007)	During crisis (2007–2009)	After the crisis (2009–2015)
Average across stocks	-1.75%	9.15%	5.98%
N significant positive	(113)	(73)	(157)
N significant negative	(45)	(10)	(62)

Table 8. Short interest ratio and idiosyncratic skewness risk premium.

This table presents the results of two univariate portfolio sort analyses and a Fama MacBeth regression. In **Panel A**, the portfolios are sorted each month according to the stocks' idiosyncratic skewness risk premium (computed with Equation 16), with ptf1 (ptf3) denoting the portfolio with the lowest (highest) risk premium. For each portfolio, the table reports the average short interest ratio, which is calculated for each stock at the start date of the swaps as the ratio of the short interest (i.e., the number of shares held short) against the total shares outstanding. The column labeled '3-1 t-stat' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that ptf3 and ptf1 have the same short interest ratio. In **Panel B**, the sorting variable is the short interest ratio and the table reports the statistics for the idiosyncratic skewness risk premium of each portfolio.**Panel C** reports the results of the Fama MacBeth regression where the Idiosyncratic skewness risk premium is regressed on the short interest ratio:

$$\text{Idiosyncratic risk premium}_{i,t,t+1,\%} = \alpha_t + \beta_{1,t} SI_{i,t} + \beta_{2,t} BidAsk_{i,t} + \beta_{3,t} VarRP_{i,t,t+1} + \epsilon_{i,t},$$

where $SI_{i,t}$ is the short interest ratio of stock i on day t, $BidAsk_{i,t}$ is the average bid ask spread of the options on stock i on day t, and $VarRP_{i,t,t+1}$ is the variance risk premium of stock i from month t to month $t+1$, which is contemporaneous to the idiosyncratic skewness risk premium.
The Fama MacBeth procedure has two steps:

1. For each month t, I compute a cross-sectional regression of the idiosyncratic skewness risk premium on the independent variables, and I obtain the estimates of the coefficients $\beta_{j,t}$ for $j = 1, 2, 3$. The table reports the standardized betas of the regressions, which are computed as the average $\widehat{\beta}$ multiplied by the average standard deviation of the regressor;
2. I test whether the time series of the coefficients $\beta_{j,t}$ is statistically different than zero. The t-statistics are presented in parentheses.

All the t-statistics are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

Panel A: Sort according to idiosyncratic skewness risk premium

	ptf1	ptf2	ptf3	3-1 t-stat
Short interest ratio	0.0362	0.0394	0.0446	6.59

Panel B: Sort according to short interest ratio

	ptf1	ptf2	ptf3	3-1 t-stat
Idiosyncratic skewness risk premium	-18.28%	-12.62%	0.89%	3.58

Panel C: Fama MacBeth regression

	Idiosyncratic skewness risk premium	
	(1)	(2)
Short interest ratio	0.0652***	0.0571***
	(7.30)	(6.47)
Bid ask		0.0968*
		(2.21)
Variance risk premium		-0.0897***
		(-6.34)
Average R^2	0.0014	0.0595

Table 9. Short interest ratio and maximum daily return over the previous month (MAX).

Panel A presents the average value of MAX for three stock portfolios. The portfolios are sorted each month according to the short interest ratio of the stocks, with SI 1 (SI 3) denoting the portfolio with the lowest (highest) short interest ratio. The short interest ratio is calculated for each stock as the ratio of the short interest (i.e., the number of shares held short) against the total shares outstanding. MAX is the maximum daily return over the previous month and it is calculated on the same day as the short interest ratio. The column labeled '3-1 t-stat' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that the portfolios SI 3 and SI 1 have the same MAX value. **Panel B** shows the idiosyncratic skewness risk premium of bivariate stock portfolios constructed on the basis of short interest ratio (SI) and the maximum daily return over the previous month (MAX). At each start date of the swaps t, stocks are independently sorted in ascending order according to their SI estimate and MAX estimate and are assigned to one of the three tercile portfolios for each sorting variable. The intersection of these classifications yields nine portfolios. I then calculate the average idiosyncratic skewness risk premium of these nine portfolios at the end of the following month $t+1$ (i.e., ex-post risk premium). The t-statistics of the difference portfolios are provided in parentheses. All the t-statistics are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

Panel A: Univariate portfolio sort

	SI 1 (low)	SI 2	SI 3 (high)	3-1 t-stat
MAX	0.0329	0.0386	0.0478	9.62

Panel B: Bivariate portfolio sort

	MAX 1 (low)	MAX 2	MAX 3 (high)	Difference
SI 1 (low)	-18.34%	-24.09%	-11.76%	6.58%
				(1.02)
SI 2	-18.54%	-18.72%	-5.83%	12.71%*
				(2.25)
SI 3 (high)	-10.26%	-4.07%	6.95%	17.21%**
				(3.05)
Difference	8.08%	20.02%**	18.71%***	
	(1.12)	(2.67)	(5.42)	

Table 10. Short interest ratio and put option volumes.

Panel A presents the average option volume ratio for three stock portfolios. The portfolios are sorted each month according to the short interest ratio of the stocks, with SI 1 (SI 3) denoting the portfolio with the lowest (highest) short interest ratio. The short interest ratio is calculated for each stock as the ratio of the short interest (i.e., the number of shares held short) against the total shares outstanding. The volume ratio (VR) is defined as the ratio of the volume of the out-of-the-money calls against the out-of-the-money puts. The two measures are calculated on the same day. The column labeled '3-1 t-stat' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that the portfolios SI 3 and SI 1 have the same value of VR. **Panel B** shows the idiosyncratic skewness risk premium of bivariate stock portfolios constructed on the basis of short interest ratio (SI) and the option volume ratio (VR). At each start date of the swaps t, stocks are independently sorted in ascending order according to their SI estimate and VR estimate and are assigned to one of the three tercile portfolios for each sorting variable. The intersection of these classifications yields nine portfolios. I then calculate the average idiosyncratic skewness risk premium of these nine portfolios at the end of the following month $t+1$ (i.e., ex-post risk premium). The t-statistics of the difference portfolios are provided in parentheses. All the t-statistics are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

Panel A: Univariate portfolio sort				
	SI 1 (low)	SI 2	SI 3 (high)	3-1 t-stat
Volume ratio	2.3678	2.1735	2.0695	-3.6181

Panel B: Bivariate portfolio sort				
	VR 1 (low)	VR 2	VR 3 (high)	Difference
SI 1 (low)	-17.08%	-25.07%	-13.96%	3.12% (0.38)
SI 2	-12.56%	-15.42%	-10.37%	2.19% (0.49)
SI 3 (high)	4.72%	-3.05%	0.43%	-4.29% (-1.56)
Difference	21.80%** (3.12)	22.02%** (2.56)	14.38%* (2.01)	

Table 11. Cross-sectional regression: financial crisis.

This table shows the results for the cross-sectional regression:

$$\Delta IdioRP_i = \alpha + \beta Friction_i + \epsilon_i,$$

where $\Delta IdioRP_i$ is the average change in the idiosyncratic skewness risk premium after the crisis for stock i. It is computed as the difference between the median idiosyncratic skewness risk premium in the post-crisis period and the median idiosyncratic skewness risk premium in the pre-crisis period for stock i. As measures of frictions, I consider the maximum of the following three variables achieved by each stock during the financial crisis (2007–2009): i) option implied lending fee of Muravyev et al. (2016); ii) short interest ratio; and iii) the overvaluation proxy of Bali et al. (2011), which is computed as the maximum daily return of the stock over the previous month.

	ΔIdiosyncratic skewness risk premium			
	(1)	(2)	(3)	(4)
MaxFee	0.9256**			1.0016**
	(2.95)			(2.84)
MaxSI		0.1601		-0.4044
		(0.58)		(-1.31)
MaxRet			0.4972**	0.4617**
			(2.91)	(2.64)
R^2	0.0227	-0.0019	0.0220	0.0398

B Figures

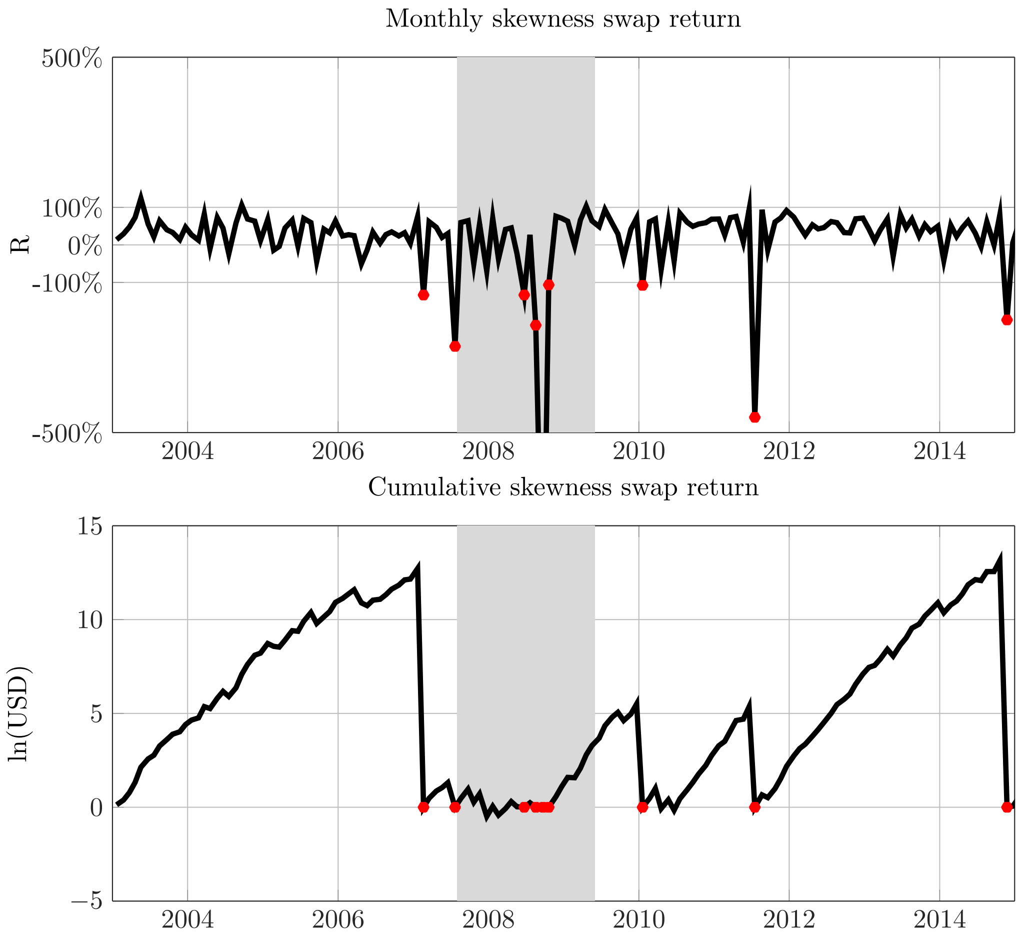

Figure 1. The monthly skewness swap returns and the cumulative skewness swap returns. The top panel shows the time series of the monthly return of the value-weighted portfolio of skewness swaps. The sample period starts on 1 January 2003 and ends on 31 December 2014. The points in red highlight the months where the return was less than -100%. The bottom figure plots the evolution of a one US dollar investment over time for the value-weighted portfolio of skewness swaps. Each month, all the money is reinvested in the swap strategy and when all the capital is lost (months highlighted in red dots), a new dollar is invested. The region in gray corresponds to the financial crisis period.

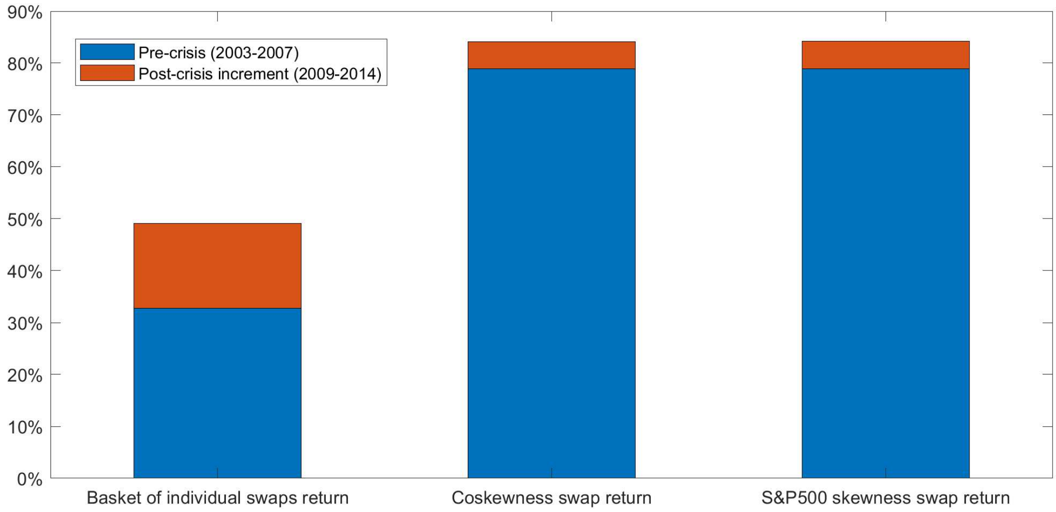

Figure 2. The return of the skewness swap on the S&P500 index, the basket of individual swaps, and the S&P500 coskewness swap. This bar graph shows the median monthly returns of the three following swap trading strategies applied to the S&P500 index and its costituents in the pre-/post-crisis subsamples:

1. S&P500 skewness swap return: $\left(\sum_i w_i r_i\right)^3$. This strategy trades the third moment of the returns of the S&P500 index. It is computed with the skewness swap of Section 1 applied to the options of the S&P500.

2. Basket of individual swaps return: $\sum_i w_i^3 r_i^3$. This strategy trades the sum of the third moments of the S&P500 constituents. It is computed with a portfolio of skewness swaps applied independently to the options of each stock, each one with weight w_i^3.

3. Coskewness swap return: $\left(\sum_i w_i r_i\right)^3 - \sum_i w_i^3 r_i^3 = 3\sum_{i,j} w_i w_j^2 r_i r_j^2 + 6\sum_{i,j,k} w_i w_j w_k r_i r_j r_k$. This strategy is implemented by going long strategy 1 and short strategy 2.

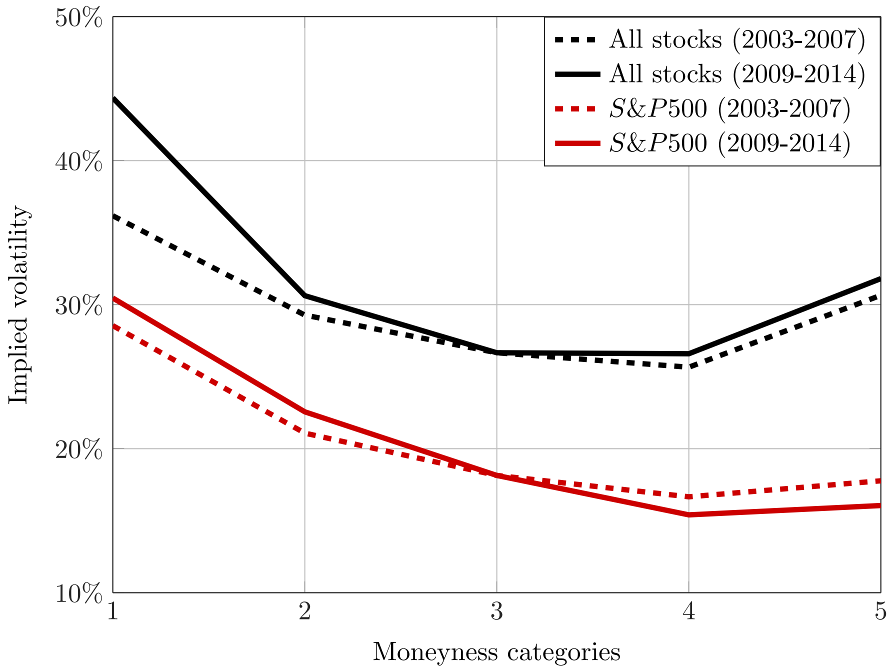

Figure 3. The implied volatility smile before and after the 2007–2009 financial crisis. The figure shows the average implied volatility smile before and after the financial crisis for the S&P500 and for the cross-section of stocks. For each stock, I compute the daily implied volatility smile and average the results in the pre-crisis sample and in the post-crisis sample. Then, I average the results across the stocks in each of the two samples. In order to build the daily average implied volatility smile for each stock, I follow Bollen and Whaley (2004) and divide all the options available on the day (both calls and puts) with maturity up to one year in five moneyness categories according to their delta values as follows:

- Moneyness category 1: $0.875 < \Delta_C \leq 0.980$ and $-0.125 < \Delta_P \leq -0.020$
- Moneyness category 2: $0.625 < \Delta_C \leq 0.875$ and $-0.375 < \Delta_P \leq -0.125$
- Moneyness category 3: $0.375 < \Delta_C \leq 0.625$ and $-0.625 < \Delta_P \leq -0.375$
- Moneyness category 4: $0.125 < \Delta_C \leq 0.375$ and $-0.875 < \Delta_P \leq -0.625$
- Moneyness category 5: $0.020 < \Delta_C \leq 0.125$ and $-0.980 < \Delta_P \leq -0.875$

where Δ_C indicates the delta of call options and Δ_P indicates the delta of put options. Finally, I average the implied volatility of all options in each category, where I use the implied volatility of the options provided by Optionmetrics. The average at-the-money implied volatility (implied volatility of the options of category 3) slightly increased in the post-crisis period for both the S&P500 and single stocks. The pre-crisis smiles are then shifted up in the graph in order to overlap the pre-/post-crisis smiles at their at-the-money volatility to better visualize the variation in their slopes.

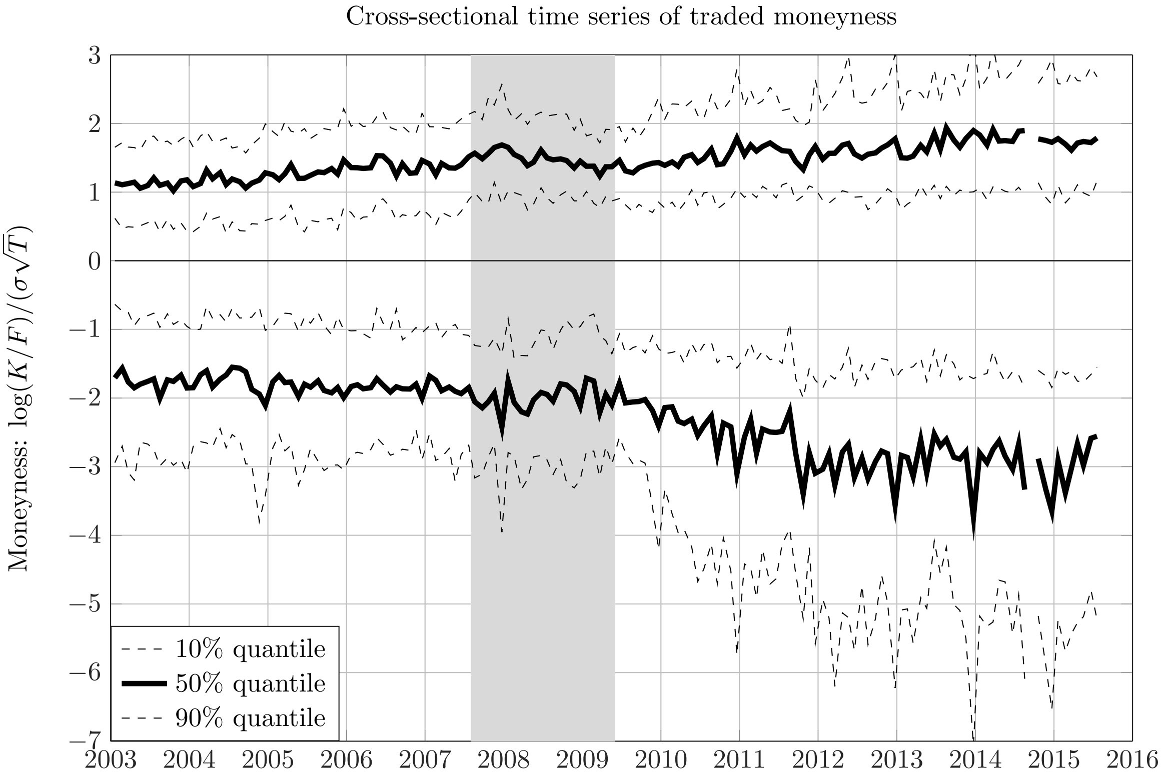

Figure 4. Cross-sectional moneyness range. The figure shows the cross-sectional times series of the minimum and maximum moneyness of the options traded for all the stocks of my sample. The moneyness of an option is defined in standard deviations as $(\log(K/F))/(\sigma\sqrt{T})$, where K is the strike price, F is the forward price of the stock, T is the time to maturity, and σ is the average implied volatility of the at-the-money options of the stock on that day. At each point in time, I compute for each stock the minimum and maximum moneyness traded, I sort the values across stocks, and I plot the 10% quantile, 50% quantile, and 90% quantile. The options used in this calculation are the same options used in the implementation of the skewness swaps, that is, out-of-the-money options with maturity of one month and standard filters as outlined in Section 2.2. The region in gray corresponds to the financial crisis period.

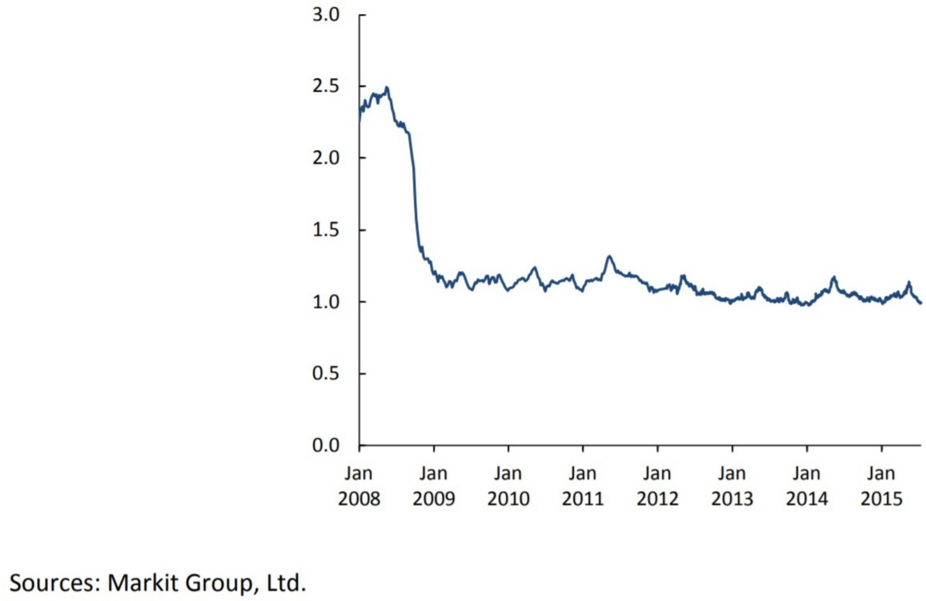

Figure 5. Volume of security on loan. This graph is taken from Figure 9 of the SEC report of Baklanova et al. (2015) and shows the total volume of the securities on loan from 2008 to 2015. The data source is Markit Group, Ltd. The numbers presented are aggregate numbers and include all the following securities: US equities, US Treasuries & Agencies, US Treasuries & MBS, Non-US equities, Non-US Sovereigns, Non-US Bonds.

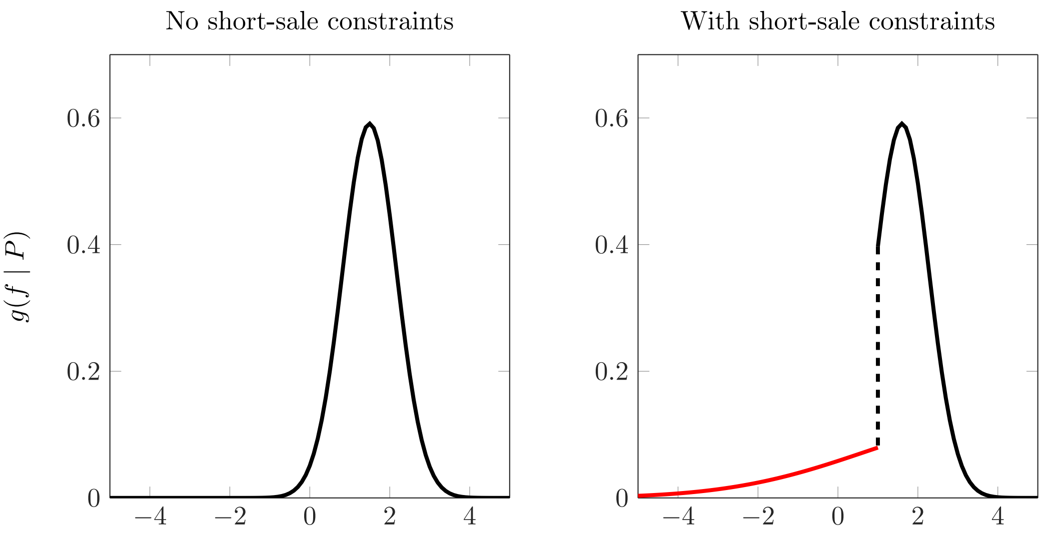

Figure 6. The predictions of the equilibrium model: distribution of final payoff. Distribution of the conditional final payoff of the stock predicted by the equilibrium model given the price, $f|P$, with short-sale constraints (right graph) and without short-sale constraints (left graph).

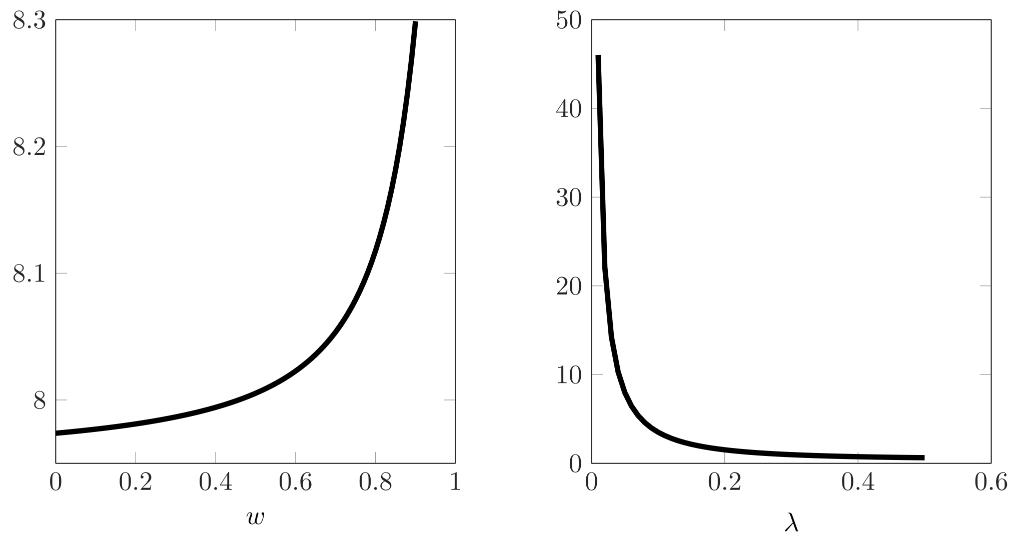

Figure 7. The predictions of the equilibrium model: the risk premium as function of short-selling constraints and asymmetry of information. These two panels show the value of the risk premium D as a function of short-selling constraints w (left graph) and asymmetry of information λ (right graph). In the model w, is the proportion of informed traders who are short-sale constrained and λ is the proportion of informed traders in the economy.

C Proofs

Proof of Proposition 1.

Proof.

$$
\begin{aligned}
\overline{S_A} &= E_0^Q\left[\left(\int_0^{F_{0,T}} \Phi^{''}(K)P_{A,T,T}dK + \int_{F_{0,T}}^{\infty} \Phi^{''}(K)C_{A,T,T}dK\right)\right] = \\
&= E_0^Q\left[\left(\int_0^{F_{0,T}} \Phi^{''}(K)(P_{E,T,T} + (P_{A,T,T} - P_{E,T,T}))dK\right)\right] + \\
&+ E_0^Q\left[\left(\int_{F_{0,T}}^{\infty} \Phi^{''}(K)(C_{E,T,T} + (C_{A,T,T} - C_{E,T,T}))dK\right)\right] = \\
&= E_0^Q\left[\left(\int_0^{F_{0,T}} \Phi^{''}(K)P_{E,T,T}dK + \int_{F_{0,T}}^{\infty} \Phi^{''}(K)C_{E,T,T}dK\right)\right] + \\
&+ E_0^Q\left[\left(\int_0^{F_{0,T}} \Phi^{''}(K)(P_{A,T,T} - P_{E,T,T})dK\right)\right] + \\
&+ E_0^Q\left[\left(\int_{F_{0,T}}^{\infty} \Phi^{''}(K)(C_{A,T,T} - C_{E,T,T})dK\right)\right] = \\
&= \overline{S_E} + \frac{1}{B_{0,T}}\left(\int_0^{F_{0,T}} \Phi^{''}(K)(P_{A,0,T} - P_{E,0,T})dK\right) + \\
&+ \frac{1}{B_{0,T}}\left(\int_{F_{0,T}}^{\infty} \Phi^{''}(K)(C_{A,0,T} - C_{E,0,T})dK\right).
\end{aligned}
$$

In the same way the American floating leg can be decomposed in the following way

$$
\begin{aligned}
\widetilde{S_A} &= \left(\int_0^{F_{0,T}} \Phi^{''}(K)P_{A,T,T}dK + \int_{F_{0,T}}^{\infty} \Phi^{''}(K)C_{A,T,T}dK\right) \\
&+ \sum_{i=1}^{n-1}\left(\Phi^{'}(F_{i-1,T}) - \Phi^{'}(F_{i,T})\right)(F_{T,T} - F_{i,T}) = \\
&= \widetilde{S_E} + \left(\int_0^{F_{0,T}} \Phi^{''}(K)(P_{A,T,T} - P_{E,T,T})dK\right) + \\
&+ \left(\int_{F_{0,T}}^{\infty} \Phi^{''}(K)(C_{A,T,T} - C_{E,T,T})dK\right)
\end{aligned}
$$

□

Proof of Proposition 2.

Proof. For every Φ

$$\overline{S_E} = E_0^Q[\widetilde{S_E}] = E_0^Q\left[\int_0^{F_{0,T}} \Phi^{''}(K)P_{E,T,T}dK + \int_{F_{0,T}}^{\infty} \Phi^{''}(K)C_{E,T,T}dK\right],$$

because $E_i^Q[F_{T,T} - F_{i,T}] = 0$ for every i, and hence the conditional expectation of the second term of the floating leg defined by Equation 2 is zero.

By applying the result of Carr and Madan (2001) to our case I obtain:

$$\int_0^{F_{0,T}} \Phi^{''}(K)P_{E,T,T}(K)dK + \int_{F_{0,T}}^{\infty} \Phi^{''}(K)C_{E,T,T}(K)dK = \Phi(F_{T,T}) - \Phi(F_{0,T}) - \Phi^{'}(F_{0,T})(F_{T,T} - F_{0,T}).$$

I substitute my definition of Φ of Equation 6 and by computing some calculations I obtain

$$\begin{aligned}
&\Phi(F_{T,T}) - \Phi(F_{0,T}) - \Phi^{'}(F_{0,T})(F_{T,T} - F_{0,T}) = \\
&= 72 - \frac{(72F_{T,T})}{F_{0.T}} - 24\left(\frac{F_{T,T}}{F_{0.T}}\right)^{0.5}\log\left[\frac{F_{T,T}}{F_{0.T}}\right] + 24\left(-8 + \left(\frac{F_{T,T}}{F_{0.T}}\right)^{0.5}\left(8 + \log\left[\frac{F_{T,T}}{F_{0.T}}\right]^2\right)\right) = \\
&= 72 - 72e^y - 24e^{y/2}y + 24(-8 + e^{y/2}(8 + y^2)) = \\
&= y^3 + O(y^5),
\end{aligned}$$

where $y = \log(F_{T,T}/F_{0,T})$ and the last equality is obtained by substituting e^y with its power series expansion $e^y = \sum_{n=0}^{\infty} \frac{y^n}{n!}$. □

Proof of Theorem 1.

Proof. The equilibrium price function is solved in three steps: i) find the optimal demands for all agents ii) conjecture a price function iii) verify that the price function solves the market clearing condition.

1. Obtain the optimal demands by all investors given the prices. The optimal demand for agent k is obtained by setting to zero the first order condition of the mean-variance utility function of the final wealth $W_k = W_0 + x^k(d-p)$:

$$x^k = \text{argmin} \quad U(W_k),$$

where

$$U(W_k) = E[W_k|I_k] - \frac{\rho}{2}Var[W_k|I_k] = W_0 + x^k(E[f|I_k] - p) - \frac{\rho}{2}\left((x^k)^2(Var[f|I_k] + \sigma_n^2)\right),$$

where I used the definition of the payoff $d = f + n$.
For the informed traders $E[f|I_k] = f$ and $Var[f|I_k] = 0$, hence $x^{uc} = \frac{f-P}{\rho\sigma_n^2}$ and $x^c = max\{\frac{f-P}{\rho\sigma_n^2}, 0\}$.
For the uninformed trader the optimal demand is

$$x^{ui} = \frac{E[f \mid P] - P}{\rho(Var[f \mid P] + \sigma_n^2)}$$

and it has to be made explicit by conjecturing a price function (next step).

2. Conjecture a price function in order to explicit the quantities $E[f \mid P]$ and $Var[f \mid P]$ and hence explicit the optimal demand of the uninformed traders. The conjectured price function is

$$P = f + \begin{cases} A(u-C) & u \leq C \\ B(u-C) & u > C \end{cases}$$

with $A, B > 0$. The idea is that the price is a function of the fundamental f plus a supply/demand clearing which is different if $P > f$ or $P < f$ due to short-sale constraints. Because the prior of f is uninformative, the uninformed investors learn about f only through the price, and the posterior of f is a piecewise gaussian:

$$f = P - \begin{cases} A(u-C) & u \leq C \\ B(u-C) & u > C \end{cases}$$

Thus:

$$\begin{aligned} E[f \mid P] &= E[\mathbb{I}_{u\leq C}(P - A(u-C))] + E[\mathbb{I}_{u>C}(P - B(u-C))] = \\ &= P + D, \end{aligned}$$

where $D = E[\mathbb{I}_{u\leq C}(-A(u-C))] + E[\mathbb{I}_{u>C}(-B(u-C))] = -A\int_{-\infty}^{C}(u-C)\phi(u)du - B\int_{C}^{\infty}(u-C)\phi(u)du$ and ϕ is the gaussian distribution with mean 0 and variance σ_u^2. D is always positive because it is the expectation of a positive random variable. It cannot be solved explicitly because it involves the solution of the gaussian integral. But we know that D is well defined and positive. Then we have to find $Var[f \mid P]$. I use the formula $Var[f \mid P] = E\left[(f \mid P)^2\right] - \left(E[f \mid P]^2\right)$. For the first term we know that

$$(f \mid P)^2 = P^2 + \begin{cases} A^2(u-C)^2 - 2AP(u-C) & u \leq C \\ B^2(u-C)^2 - 2BP(u-C) & u > C \end{cases}$$

So $E[(f \mid P)^2] = P^2 + \int_{-\infty}^{C} A^2(u-C)^2\phi(u)du - 2AP\int_{-\infty}^{C}(u-C)\phi(u)du + \int_{C}^{+\infty} B^2(u-C)^2\phi(u)du - 2BP\int_{C}^{+\infty}(u-C)\phi(u)du$. Hence

$$E[(f \mid P)^2] = P^2 + 2PD + E,$$

where $E = A^2\int_{-\infty}^{C}(u-C)^2\phi(u)du + B^2\int_{C}^{\infty}(u-C)^2\phi(u)du$. And finally

$$Var[f \mid P] = E\left[(f \mid P)^2\right] - \left(E[f \mid P]^2\right) = (P^2 + 2PD + E) - (P+D)^2 = E - D^2.$$

3. In the previous steps I computed the optimal demands of the three rational investors: the informed agents, informed constrained and uninformed. The last step consists in substituting these optimal demands in the market clearing condition and solve for the constants A, B and C which make the market clearing condition hold true $\forall u$.

For the uniqueness of the equilibrium see Venter (2016). □

D Numerical analysis

In this section I study the precision of my skewness swap in measuring the third moment of the asset returns in the Merton jump-diffusion model. The choice of the Merton model is by no means restrictive, because, as shown in Hagan et al. (2002), for time horizons less than one year the implied volatility smile generated by the Merton model can satisfactory reproduce the empirical one. The dynamics of the Merton model is the following:

$$ds_t = \left(r - \lambda\kappa - \frac{1}{2}\sigma^2\right) dt + \sigma dW_t + \log(\psi) dq_t, \tag{18}$$

where s_t is the logarithm of the stock price, r is the risk-free interest rate, and σ is the instantaneous variance of the return conditional on no jump arrivals. The Poisson process, q_t, is independent of W_t, and such that there is a probability λdt that a jump occurs in dt, and $(1-\lambda)dt$ probability that no jump occurs:

$$q_t = \begin{cases} 1, & \text{with probability } \lambda dt \\ 0, & \text{with probability } (1-\lambda)dt \end{cases} \tag{19}$$

The parameter λ represents the mean number of jumps per unit of time. The random variable ψ is such that $\psi - 1$ describes the percentage change in the stock price if the Poisson event occurs, and $\kappa = E[\psi - 1]$ is the mean jump size. I further make the standard assumption (for instance, see Bakshi et al. (1997)) that $\log(\psi) \sim N(\mu, \delta^2)$.

The characteristic function is given by:

$$\phi(u) = E[e^{iu(s_t - s_0)}] = e^{t(i(r-0.5\sigma^2-\lambda\kappa)u-0.5\sigma^2u^2+\lambda(e^{i\mu u-0.5\delta^2u^2}-1))}$$

and the moments can be recovered from the property of the characteristic function:

$$E[(s_t - s_0)^k] = (-i)^k \frac{d^k}{du^k}\phi(u)_{|u=0}.$$

I can then recover the third moment in closed form:

$$E[(s_t-s_0)^3] = 3\delta^2\lambda\mu t + \lambda\mu^3 t - 3(-\delta^2\lambda - \lambda\mu^2 - \sigma^2)(r + (-\kappa\lambda + \lambda\mu - 0.5\sigma^2))t^2 + (r - \kappa\lambda + \lambda\mu - 0.5\sigma^2)^3 t^3.$$

I chose the following standard parameter values for my simulation study: $r = 0$, $\mu = -0.05$, $\delta = 0.08$, $\sigma = 0.2$, $\lambda = 3$, $t = 30/365$. With these parameters the Merton implied volatility smile is left skewed as shown in Figure 8. I compute the fixed leg of my swap defined by Equation 1 with the discrete approximation of Equation 7 and Φ defined by Equation 6. I use the option prices implied by the Merton jump-diffusion model.

Table 12 shows the convergence of the fixed leg of my swap to the true model-based third moment when the number of options used increases from 8 to 100. I compare the performance of my swap defined by Φ_S with the swap of Schneider and Trojani (2015) defined by Φ_3. The precision of my swap is superior

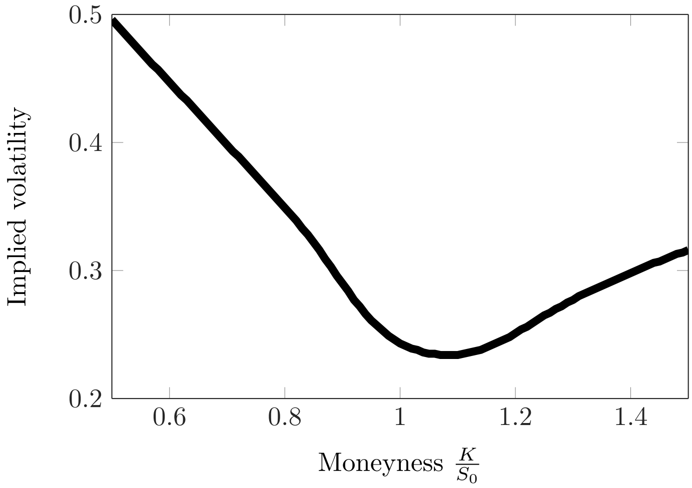

Figure 8. Merton implied volatility smile. The figure displays the one month implied volatility smile generated by the Merton model with the following set of representative parameters: $r = 0, \mu = -0.05, \delta = 0.08, \sigma = 0.2, \lambda = 3$.

because I can reach an error of less than 1%, while the swap defined by Φ_3 has always an error around 20%. This shows that the isolation of the third moment from the fourth is important in order to have a precise measure. It is also interesting to notice that in my swap even with only 8 strikes available (which is the average number of options that I have empirically) the error is only of the order of 5%.

Second, I check how the precision of the methodology depends on the range of moneyness available. In this contest I define the moneyness of an option in standard deviations (SD) as $\log(K/F_{0,T})/(\sigma\sqrt{T})$, where K is the strike price, $F_{0,T}$ is the forward price, σ is the at-the-money implied volatility and T is the time to maturity. I fix a number of options equal to 10 and I calculate the error that I make if my ten options span the moneyness range $\pm 1SD$, $\pm 2SD$, $\pm 3SD$ and $\pm 4SD$. Figure 9 shows that the precision changes a lot depending on the moneyness range available. I need at least three standard deviations in order to have an error around 10% and I need four standard deviations to have an error around 1%.

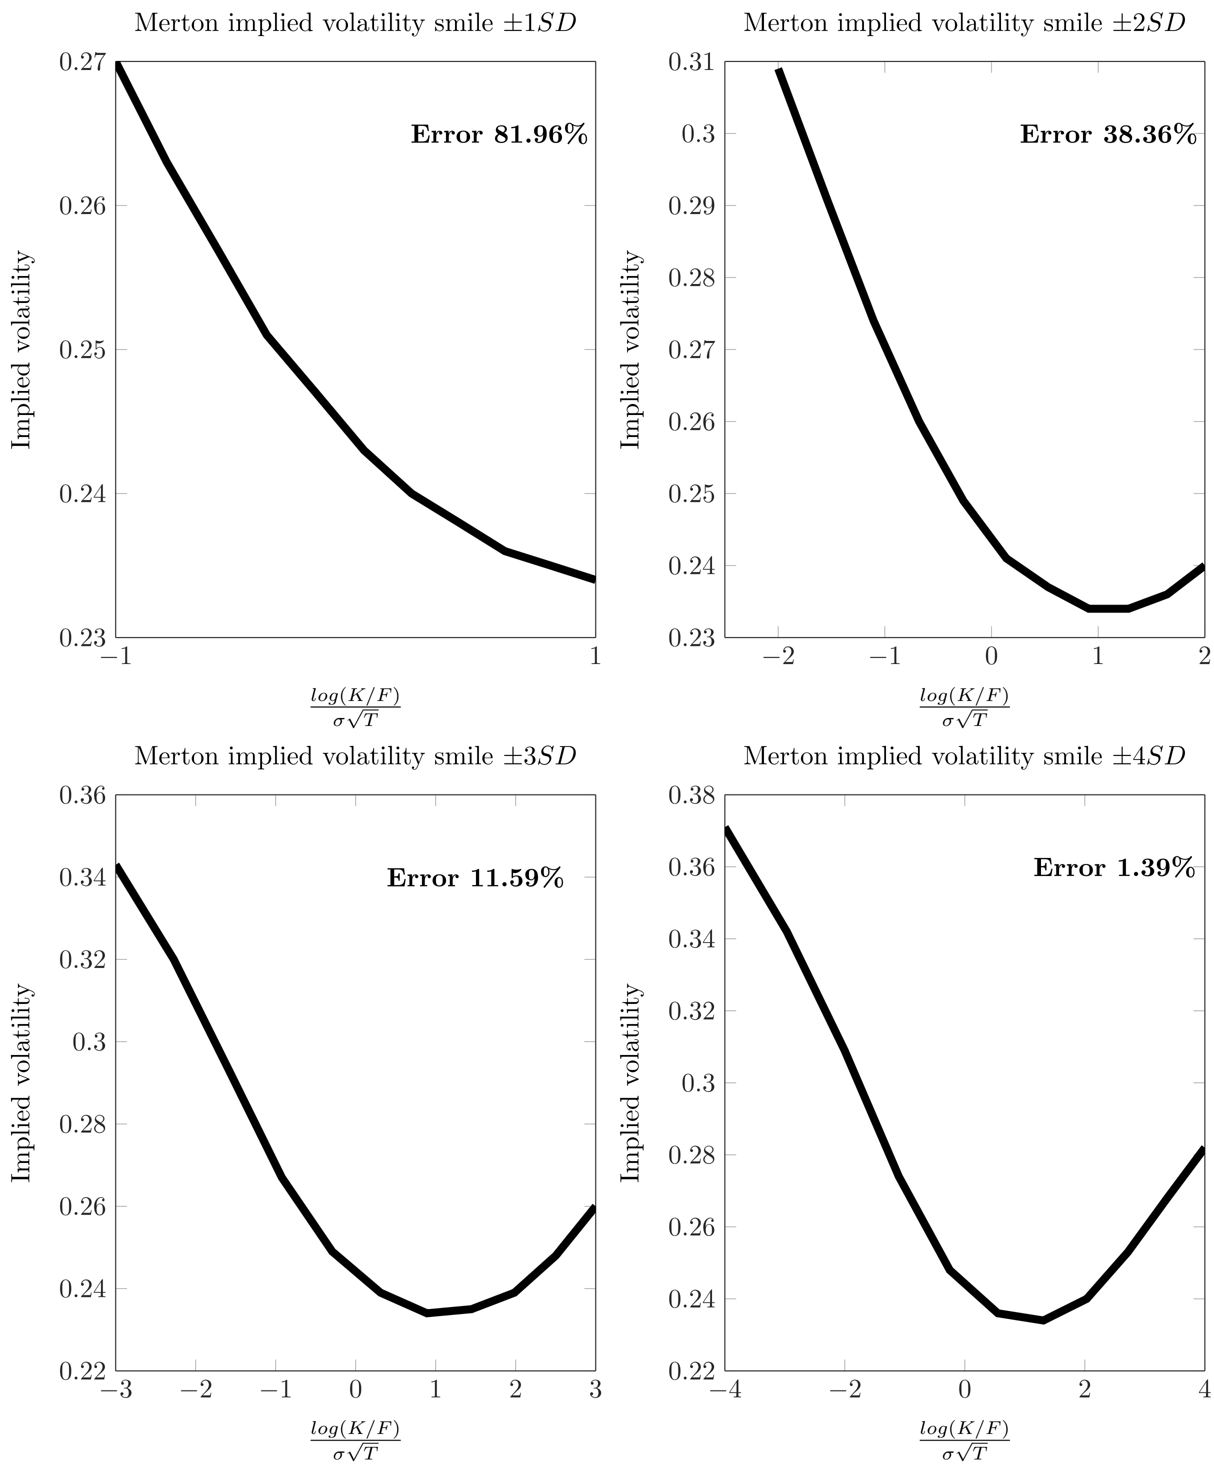

Figure 9. Convergence in the moneyness range. This figure shows the implied volatility smile implied by the Merton jump-diffusion process with parameters $\mu = -0.05$, $\delta = 0.08$, $\sigma = 0.2$, $\lambda = 3$, $r = 0$, $t = 30/365$, for increasing moneyness ranges. In the first plot I consider a moneyness range of $[-1SD, 1SD]$, in the second one $[-2SD, 2SD]$, the third one $[-3SD, 3SD]$ and the fourth one $[-4SD, 4SD]$, where SD is defined as $SD = \log(K/F_{0,T})/(\sigma\sqrt{T})$, in which K is the strike price, $F_{0,T}$ is the forward price, σ is the at-the-money implied volatility and T is the time to maturity. Each plot is a zoom out of the previous one by $1SD$. For each moneyness range I compute the risk-neutral third moment of the returns with my skewness swap according to Equation 7 and I compare it with the benchmark model based risk-neutral third moment which is recovered in closed form. At the top right of each graph is displayed the error in percentage.

Table 12. Convergence in the number of options.

This table shows the convergence of the fixed leg of the trading strategy to the third moment of the asset returns when the number of options increases. The error is computed as $\frac{|\text{True moment}-\text{Strategy fixed leg}|}{\text{True moment}}$ and it is displayed in percentage. The returns are assumed to follow a Merton jump-diffusion process with standard parameters, i.e. $\mu = -0.05$, $\delta = 0.08$, $\sigma = 0.2$, $\lambda = 3$, $r = 0$, $t = 30/365$. The true moment is computed in closed form and is equal to $-3.12 \cdot 10^{-4}$. In the second column I consider the skewness swap of Schneider and Trojani (2015) with $\Phi = \Phi_3$ while in the third column I consider my skewness swap with $\Phi = \Phi_S$.

Number of options	Error (%)	
	Φ_3	Φ_S
8	-18.81%	-5.16%
10	-21.10%	-1.70%
20	-22.70%	-0.72%
50	-22.80%	-0.86%
100	-22.80%	-0.87%

E Bakshi-Kapadia-Madan (2003) risk-neutral skewness

The risk-neutral skewness proposed by Bakshi et al. (2003) has the following expression:

$$SKEW(t,\tau) = \frac{e^{r\tau}W(t,\tau) - 3\mu(t,\tau)e^{r\tau}V(t,\tau) + 2\mu(t,\tau)^3}{[e^{r\tau}V(t,\tau) - \mu(t,\tau)^2]^{3/2}}$$

where $\mu(t,\tau) = e^{r\tau} - 1 - \frac{e^{r\tau}}{2}V(t,\tau) - \frac{e^{r\tau}}{6}W(t,\tau) - \frac{e^{r\tau}}{24}X(t,\tau)$, and $V(t,\tau)$, $W(t,\tau)$, $X(t,\tau)$ are respectively the prices of a volatility, cubic and quartic contracts which can be computed with option prices. The price of a volatility contract is given by

$$\begin{aligned} V(t,\tau) = &\int_{S(t)}^{\infty} \frac{2\left(1 - \log\left[\frac{K}{S(t)}\right]\right)}{K^2} C(t,\tau,K)dK \\ &+ \int_0^{S(t)} \frac{2\left(1 + \log\left[\frac{S(t)}{K}\right]\right)}{K^2} P(t,\tau,K)dK, \end{aligned}$$

where $C(t,\tau,K)$ ($P(t,\tau,K)$) is the price of an European call (put) option quoted on day t with time to maturity τ and strike K. $S(t)$ is the stock price on day t. Similarly, the prices of a cubic and quartic contract are given by the following portfolios of options:

$$\begin{aligned} W(t,\tau) = &\int_{S(t)}^{\infty} \frac{6\log\left[\frac{K}{S(t)}\right] - 3\left(\log\left[\frac{K}{S(t)}\right]\right)^2}{K^2} C(t,\tau,K)dK \\ &- \int_0^{S(t)} \frac{6\log\left[\frac{S(t)}{K}\right] + 3\left(\log\left[\frac{S(t)}{K}\right]\right)^2}{K^2} P(t,\tau,K)dK, \end{aligned}$$

$$X(t,\tau) = \int_{S(t)}^{\infty} \frac{12\left(\log\left[\frac{K}{S(t)}\right]\right)^2 - 4\left(\log\left[\frac{K}{S(t)}\right]\right)^3}{K^2} C(t,\tau,K)dK$$
$$+ \int_{0}^{S(t)} \frac{12\left(\log\left[\frac{S(t)}{K}\right]\right)^2 + 4\left(\log\left[\frac{S(t)}{K}\right]\right)^3}{K^2} P(t,\tau,K)dK.$$

I calculate the risk-neutral skewness of Bakshi et al. (2003) for each day of my sample period (01/01/2003 - 31/12/2014) and for each stock of my sample. I use the daily options quotes given by the optionmetrics interpolated implied volatility surface with 30 days to maturity. I convert the implied volatilities given by optionmetrics in European Black-Scholes prices and I use these prices in the algorithm ($C(t,\tau,K)$ and $P(t,\tau,K)$). The above expressions are written for a continuum of options, but in practice only a finite number of options is available. I therefore apply the following discrete approximation. Suppose that on day t, I have N_c out-of-the-money call options available with strikes $S(t) < K_{c,1} < \cdots < K_{c,N_c}$ and N_p out-of-the-money put options available with strikes $K_{p,1} < \cdots < K_{p,N_p} < S(t)$. The price of a volatility contract is discretized as follows

$$V(t,\tau) = \sum_{i=1}^{N_c} \frac{2\left(1 - \log\left[\frac{K_{c,i}}{S(t)}\right]\right)}{K_{c,i}^2} C(t,\tau,K_{c,i})\Delta K_i$$
$$\sum_{j=1}^{N_p} \frac{2\left(1 + \log\left[\frac{S(t)}{K_{p,j}}\right]\right)}{K_{p,j}^2} P(t,\tau,K_{p,j})\Delta K_j,$$

where

$$\Delta K_i = \begin{cases} K_1 - S(t) & \text{if } i = 1, \\ (K_i - K_{i-1}) & \text{if } i > 1. \end{cases}$$

and

$$\Delta K_j = \begin{cases} (K_{j+1} - K_j) & \text{if } j < N_p, \\ S(t) - K_{N_p} & \text{if } j = N_p. \end{cases}$$

The cubic and quartic contracts are discretized in the same way.

F Model-based skewness swap

As a robustness check, I implement a model-based skewness swap, where instead of using the actual option prices, I use the option prices calculated from a fitted model. This skewness swap is not tradable, because the fitted option prices are not real quotations, but it is nevertheless a useful econometric check to compare the model-based skewness swap returns with the real tradable skewness swap returns.

There are two main steps in the implementation of the model-based skewness swap: i) calibration of a model ii) implementation of the swap according to the calibrated model.

Calibration.

The model-based skewness swaps are implemented monthly, as the tradable skewness swaps, and they start and end on the third Friday of each month. I recalibrate the model at each start date of the swap and for each stock separately. In details, at each start date of the swap t and for each stock S_t, I consider all the out-of-the-money options with maturity 30 days provided by the Optionmetrics implied volatility surface file. This sample constitutes my calibration sample. I choose as benchmark model the Merton jump-diffusion model with gaussian jump-size distribution, whose dynamics is given by Equation 18. The variable $\log(\psi)$ (the size of the jump) is normally distributed with parameters $N(\mu, \delta^2)$. This model is simple and tractable, and many empirical studies (see e.g. Hagan et al. (2002)) show that it provides a good fit for short-term options data. I then calibrate the parameters of the Merton jump-diffusion model by minimizing the implied volatility mean squared error (IVMSE) as

$$IVMSE(\chi) = \sum_{i=1}^{n} (\sigma_i - \sigma_i(\chi))^2,$$

where $\chi = \{\lambda, \mu, \delta, \sigma\}$ is the set of parameters to estimate, $\sigma_i = BS^{-1}(O_i, T_i, K_i, S, r)$ is the market implied volatility provided by Optionmetrics and $\sigma_i(\chi) = BS^{-1}(O_i(\chi), T_i, K_i, S, r)$ is the model implied volatility, where $O_i(\chi)$ is the Merton model price of the option i. The model implied volatility is obtained by inverting the Black-Scholes formula where the option price is given by the Merton model price. In the Merton jump-diffusion model the option prices are available in closed form (see Merton (1976)). The price of a call option is given by:

$$C_{MRT}(t, \tau, K) = \sum_{n=0}^{\infty} e^{-\lambda'\tau + n\log(\lambda'\tau) - \sum_{i=1}^{n} \log n} C(S_t, K, r_n, \sigma_n),$$

where $\lambda' = \lambda(1+k)$, $k = e^{\mu + \frac{1}{2}\delta^2} - 1$, $C(S_t, K, r_n, \sigma_n)$ is the Black-Scholes price of an European call with volatility $\sigma_n = \sqrt{\sigma^2 + \frac{n\delta^2}{\tau}}$ and risk-free rate $r_n = r - \lambda k + \frac{n\log(1+k)}{\tau}$. The price of a put option is defined analogously. The choice of the implied volatility mean squared error (IVMSE) loss function follows the argumentation of Christoffersen and Jacobs (2004), where they show that the calibration made on implied volatilities is more stable out of sample. Table 13 displays the average calibrated parameters for the S&P500 index and for the cross-section of stocks.

Table 13. Calibrated parameters of the Merton jump-diffusion model.

This table displays the average calibrated parameters of the Merton jump-diffusion model for the S&P500 and for the cross-section of stocks. The model is calibrated separately for each stock and index and it is recalibrated monthly at each start date of the swap. The calibration sample includes all the out-of-the-money options with maturity 30 days quoted by the Optionmetrics interpolated volatility surface file on the calibration day. The numbers displayed are the average calibrated parameters across time and across stocks.

	λ	μ	δ	σ
S&P500	2.00	-0.10	0.02	0.10
All stocks	1.70	-0.11	0.13	0.20

Implementation of the model-based swap.

With the calibration of the Merton jump-diffusion model in the previous step, I estimate the parameters $\{\widehat{\lambda}, \widehat{\mu}, \widehat{\delta}, \widehat{\sigma}\}$ on day t for stock S_t. I then compute the Merton option prices of an equispaced grid of strikes covering the moneyness range $[-4SD, +4SD]$, where $SD = \frac{\log(K/S_t)}{\sigma\sqrt{T}}$ is the moneyness of the options measured in standard deviations. In this definition σ is calculated as the implied volatility of an at-the-money option, i.e. $\sigma = BS^{-1}(C_{MRT}, T, S_t, S_t, r)$, where C_{MRT} is the Merton price of a call option with strike equal to S_t. The equispaced grid is constructed as follows. First, I recover $K_{min} = S_t e^{-4\sigma\sqrt{T}}$ and $K_{max} = S_t e^{4\sigma\sqrt{T}}$, then I divide the range $[K_{min}, S_t]$ and $[S_t, K_{max}]$ in 20 intervals each, and finally I compute the Merton option prices of these out-of-the-money puts and calls. In this way I keep constant the number of options and the moneyness range, because I always have 40 option prices, 20 calls and 20 puts, which span the range $[-4SD, +4SD]$.

The fixed leg of the swap (Equation 1) is computed using these model-based option prices, where the discreteness of the options is addressed with the same quadrature-based approximation of the integral explained in Equation 7. The floating leg of the swap (Equation 2) is given by the sum of the payoff of the same option portfolio and a continuous delta-hedge in the forward market. The return of the swap is computed with Equation 11.

Table 6 reports the average return of the model-based skewness swap for the S&P500 index and for the cross-section of stocks in the pre-crisis and post-crisis subsample. While the risk premium for the S&P500 does not change much after the crisis, the change is substantial for single stocks for which the model-based skewness risk premium goes from 41% to 49%. To assess the statistical significance of the change, I compute the same econometric exercise computed in Section 3.1.2. In details, for each stock I calculate the bootstrap confidence interval for the median skewness risk premium in the pre-crisis and post-crisis subsamples separately, and I look for the minimum confidence level α under which the two intervals do not overlap. With this procedure I find that the post-crisis change in the model-based skewness risk premium is statistically significant (i.e. $\alpha < 0.05$) for only 31 stocks. This is not surprising because the standard errors of a crash measure like the skewness risk premium are very high, especially at the single stock level. In order to mitigate the effect of outliers, instead of considering the skewness swap returns separately, I consider the portfolio of all skewness swaps, where the weight of each skewness swap is proportional to the market capitalization of the stock: $w_{i,t} = \frac{S_{i,t}Sh_{i,t}}{\sum_j S_{j,t}Sh_{j,t}}$, where $S_{i,t}$ is the price of stock i on day t and $Sh_{i,t}$ is the number of shares outstanding for stock i on day t. The average return of this basket of skewness swaps is 25.79% before the financial crisis and 37.40% after the crisis, and the change is statistically significant at the $\alpha = 0.025$ level.

G The option implied lending fee of Muravyev et al. (2016)

The option implied lending fee h_{imp} defined by Muravyev et al. (2016) is based on the idea that an investor can alternatively short the stock in the option market by purchasing a put option and selling a call option with the same strike and same maturity. The fee for shorting the stock in the option market is given by the option implied lending fee h_{imp} which can be computed from a put-call pair with the following formula:

$$P_t - C_t = Ke^{-r(T-t)} - S_t e^{-h_{imp}(T-t)} + PVD, \tag{20}$$

where S_t is the current stock price, PVD is the present value of future dividends, C_t and P_t are the prices of the European call and put option respectively and K is the option strike price. Intuitively, the payoff of buying a put and selling a call is $K - S_T$. The equivalent strategy in the equity market is to borrow $Ke^{-r(T-t)}$ and short the stock by paying the lending fee to the borrower, which is paid as a continuous yield. With a reverse engineering of Equation 20 one can compute the lending fee implied in option prices:

$$h_{imp} = \frac{1}{T-t} \log\left(\frac{S_t}{C_t - P_t + Ke^{-r(T-t)} + PVD}\right). \tag{21}$$

Muravyev et al. (2016) show that the average option implied lending fee equals the average realized lending fee. Hence I use h_{imp} as proxy for short-selling costs.

I calculate the option implied lending fee $h_{imp,i,t}$ for each stock of my sample at each start date of the swaps t. For each stock-day combination I average the h_{imp} extrapolated with Equation 21 for all the put-call pairs quoted by Optionmetrics on day t which satisfy the following requirements: I consider the put-call pairs with maturity between 15 and 90 days and moneyness $0.9 \leq K/S \leq 1.1$. I consider only the pairs where both options have a positive open interest, positive implied volatility, positive bid-ask spread and I exclude the options with bid price equal to zero. The prices of the European options P_t and C_t are calculated with the Black-Scholes formula with the implied volatility provided by Optionmetrics. I also exclude the options which are used in the calculation of the fixed leg of the swap in order to avoid mechanical relations. With this methodology, each stock of my sample have a time-series of option implied lending fee $h_{imp,i,t}$.

The next tables are the equivalent of the tables presented in Section 3.2.1 on the short-interest ratio measure. The results are qualitatively similar, and they show that the stocks that are hard to short-sell have a higher idiosyncratic crash risk premium, they are more likely to be overpriced and have higher volumes in put options. However, econometrically the option implied lending fee is a very volatile and noisy variable and the results are statistically weaker. In addition, even if I excluded the options which are used in the calculation of the swap, there might be some endogeneity between the option implied lending fee measure, the skewness risk premium and the volume ratio, given that they are all computed using the Optionmetrics database. This is why in the main text I present the results for the short interest ratio, which is a measure that does not involve options and has better econometric properties.

Table 14. Option implied lending fee and idiosyncratic skewness risk premium.

This table presents the results of two univariate portfolio sort analyses. In **Panel A**, the portfolios are sorted each month according to the stocks' idiosyncratic skewness risk premium (computed with Equation 16), with ptf1 (ptf3) denoting the portfolio with the lowest (highest) risk premium. For each portfolio, the table reports the average option implied lending fee which is calculated for each stock at each start date of the swaps. In details, for a given day t and a given stock i, I extrapolate the lending fees h_{imp} from each put-call pair with the same strike and maturity with the formula:

$$h_{imp} = \frac{1}{T-t} \log \left(\frac{S_t}{C_t - P_t + Ke^{-r(T-t)} + PVD} \right).$$

I then average all the fees computed on day t for stock i in order to obtain $h_{imp,i,t}$. The column labeled '3-1 t-stat' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that ptf3 and ptf1 have the same implied fee. In **Panel B**, the sorting variable is the option implied lending fee, and the table reports the statistics for the idiosyncratic skewness risk premium of each portfolio. All the t-statistics are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

Panel A: Sort according to idiosyncratic skewness risk premium				
	ptf1	ptf2	ptf3	3-1 t-stat
Implied lending fee	0.31%	0.31%	0.43%	3.39
Panel B: Sort according to the implied lending fee				
	ptf1	ptf2	ptf3	3-1 t-stat
Idiosyncratic skewness risk premium	-10.71%	-12.07%	-6.97%	1.96

Table 15. Option implied lending fee and idiosyncratic skewness risk premium.

This table presents the results of the following Fama MacBeth regression:

$$\text{Idiosyncratic risk premium}_{i,t,t+1,\%} = \alpha_t + \beta_{1,t} h_{imp,i,t} + \beta_{2,t} BidAsk_{i,t} + \beta_{3,t} VarRP_{i,t,t+1} + \epsilon_{i,t},$$

where the Idiosyncratic risk premium$_{i,t,t+1\%}$ is calculated according to Equation 16 from month t to month $t+1$ for stock i. The implied lending fee is calculated for each stock at each start date of the swaps as average of the h_{imp} extrapolated from all the put call pairs with the formula

$$h_{imp} = \frac{1}{T-t} \log\left(\frac{S_t}{C_t - P_t + Ke^{-r(T-t)} + PVD}\right).$$

$BidAsk_{i,t}$ is the average bid ask spread of the options on stock i on day t, and $VarRP_{i,t,t+1}$ is the variance risk premium of stock i from month t to month $t+1$, which is contemporaneous to the idiosyncratic skewness risk premium.

The Fama MacBeth procedure has two steps:

1. For each month t, I compute a cross-sectional regression of the idiosyncratic skewness risk premium on the independent variables, and I obtain the estimates of the coefficients $\beta_{j,t}$ for $j = 1, 3$. The table reports the standardized betas of the regressions, which are computed as the average $\widehat{\beta}$ multiplied by the average standard deviation of the regressor;
2. I test whether the time series of the coefficients $\beta_{j,t}$ is statistically different than zero. The t-statistics, which are presented in parentheses, are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

	Idiosyncratic skewness risk premium	
	(1)	(2)
Implied lending fee	0.0336**	0.0231*
	(3.24)	(2.26)
Bid ask		0.1279***
		(4.05)
Variance risk premium		-0.1342***
		(-4.80)
Average R^2	0.0013	0.0920

Table 16. Univariate portfolio sort: maximum daily return over the previous month (MAX).

This table presents the average value of MAX for three stock portfolios. The portfolios are sorted each month according to the option implied lending fee of the stocks, with ptf1 (ptf3) denoting the portfolio with the lowest (highest) lending fee. The implied lending fee is calculated for each stock at each start date of the swaps. In details, for a given day t and a given stock i I extrapolate the lending fees h_{imp} from each put-call pair with the same strike and maturity with the formula:

$$h_{imp} = \frac{1}{T-t} \log \left(\frac{S_t}{C_t - P_t + Ke^{-r(T-t)} + PVD} \right).$$

I then average all the fees computed on day t for stock i in order to obtain $h_{imp,i,t}$. MAX is the maximum daily return over the previous month and it is calculated on the same day of the implied lending fee. The column labeled '3-1 t-stat' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that ptf3 and ptf1 have the same MAX value. The t-statistics are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

	Sort according to the implied lending fee			
	ptf1	ptf2	ptf3	3-1 t-stat
MAX	0.0425	0.0403	0.0458	4.77

Table 17. Bivariate portfolio sort: implied lending fee and maximum daily return over the previous month (MAX).

This table shows the idiosyncratic skewness risk premium of bivariate stock portfolios constructed on the basis of the option implied lending fee and the maximum daily return over the previous month (MAX). The implied lending fee is calculated for each stock at each start date of the swaps. In details, for a given day t and a given stock i I extrapolate the lending fees h_{imp} from each put-call pair with the same strike and maturity with the formula:

$$h_{imp} = \frac{1}{T-t} \log \left(\frac{S_t}{C_t - P_t + Ke^{-r(T-t)} + PVD} \right).$$

I then average all the fees computed on day t for stock i in order to obtain $h_{imp,i,t}$. MAX is the maximum daily return over the previous month and it is calculated on the same day of the lending fee. At each start date of the swaps t, stocks are independently sorted in ascending order according to their h_{imp} estimate and MAX estimate and are assigned to one of the three tercile portfolios for each sorting variable. The intersection of these classifications yields nine portfolios. I then calculate the average idiosyncratic skewness risk premium of these nine portfolios at the end of the following month $t+1$ (i.e., ex-post risk premium). The t-statistics of the difference portfolios are provided in parentheses and are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

	MAX 1 (low)	MAX 2	MAX 3 (high)	Difference
h_{imp} 1 (low)	-22.28%	-16.49%	-2.15%	20.13%**
				(2.67)
h_{imp} 2	-11.90%	-15.95%	-5.12%	6.78%
				(0.93)
h_{imp} 3 (high)	-25.37%	-11.77%	2.09%	26.94%**
				(2.77)
Difference	-4.44%	4.72%	4.24%	
	(-0.51)	(1.18)	(1.52)	

Table 18. Univariate portfolio sort: implied lending fee and volume ratio (VR)

This table presents the average volume ratio (defined as the ratio of the volumes of out-of-the-money calls against out-of-the-money puts) for three stock portfolios. The portfolios are sorted each month according to the option implied lending fee of the stocks, with ptf1 (ptf3) denoting the portfolio with the lowest (highest) lending fee. The implied lending fee is calculated for each stock at each start date of the swaps. In details, for a given day t and a given stock i, I extrapolate the lending fees h_{imp} from each put-call pair with the same strike and maturity with the formula:

$$h_{imp} = \frac{1}{T-t} \log\left(\frac{S_t}{C_t - P_t + Ke^{-r(T-t)} + PVD}\right).$$

I then average all the fees computed on day t for stock i in order to obtain $h_{imp,i,t}$. The column labeled '3-1 t-stat' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that ptf3 and ptf1 have the same volume ratio. The t-statistics are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

Sort according to the implied lending fee				
	ptf1	ptf2	ptf3	3-1 t-stat
Volume ratio	2.3746	2.1650	2.1394	-4.94

Table 19. Bivariate portfolio sort: implied lending fee and volume ratio (VR).

This table shows the idiosyncratic skewness risk premium of bivariate stock portfolios constructed on the basis of the option implied lending fee and the volume ratio (VR). The implied lending fee is calculated for each stock at each start date of the swaps. In details, for a given day t and a given stock i, I extrapolate the lending fees h_{imp} from each put-call pair with the same strike and maturity with the formula:

$$h_{imp} = \frac{1}{T-t} \log\left(\frac{S_t}{C_t - P_t + Ke^{-r(T-t)} + PVD}\right).$$

I then average all the fees computed on day t for stock i in order to obtain $h_{imp,i,t}$. VR is defined as the ratio of the volumes of out-of-the-money calls against out-of-the-money puts, and it is calculated on the same day of the lending fee. At each start date of the swaps t, stocks are independently sorted in ascending order according to their h_{imp} estimate and VR estimate and are assigned to one of the three tercile portfolios for each sorting variable. The intersection of these classifications yields nine portfolios. I then calculate the average idiosyncratic skewness risk premium of these nine portfolios at the end of the following month $t+1$ (i.e., ex-post risk premium). The t-statistics of the difference portfolios are provided in parentheses and are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

	VR 1 (low)	VR 2	VR 3 (high)	Difference
h_{imp} 1 (low)	-4.48%	-18.26%	-9.07%	-4.58%
				(-1.46)
h_{imp} 2	-10.20%	-13.20%	-12.67%	-2.47%
				(-0.43)
h_{imp} 3 (high)	-4.63%	-8.32%	-5.87%	-1.24%
				(-0.33)
Difference	-0.14%	9.95%*	3.20%	
	(-0.04)	(2.40)	(1.23)	

H Additional tables

Table 20. Idiosyncratic and systematic component of the skewness risk premium: quantile regressions.

The table shows the results of the time series regressions:

$$RP_{i,t} = \alpha_i + \beta_i RP_{SP500,t} + \epsilon_{i,t}$$

where RP_i ($RP_{SP500,t}$) is the monthly gain of the swap trading strategy defined in Section 1 for stock i (for the SP500). The regressions are run separately for each stock as quantile regressions, where I consider three quantiles: q=0.5 (median regression), q=0.25, q=0.75. Panel A presents the summary statistics for the parameter α, while Panel B presents the summary statistics for the parameter β. The field 'N significant positive' ('N significant negative') indicates the number of stocks for which the estimate is statistically significant positive (negative) at the 5% level.

Panel A: parameter α			
	q=0.25	q=0.5	q=0.75
Average estimate	-0.0003	0.0001	0.0006
N significant positive	6	115	226
N significant negative	74	27	0
Panel B: parameter β			
	q=0.25	q=0.5	q=0.75
Average estimate	1.0723	1.3199	1.8106
N significant positive	378	430	380
N significant negative	4	4	0

Table 21. Univariate portfolio sort: book leverage.

This table presents the average book leverage for three stock portfolios. The portfolios are sorted each month according to the stocks idiosyncratic skewness risk premium (computed with equation 16), with ptf1 (ptf3) denoting the portfolio with the lowest (highest) risk premium. The book leverage is calculated for each stock at each start date of the swaps as the logarithm of the ratio of the book value of the assets on the book value of equity. The column labeled '3-1 t-stat' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that ptf3 and ptf1 have the same book leverage. The t-statistics are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

Sort according to idiosyncratic skewness risk premium				
	ptf1	ptf2	ptf3	3-1 t-stat
Book leverage	1.0950	0.9216	0.9021	-6.29

Table 22. Univariate portfolio sort: CAPM beta.

This table presents the average CAPM-beta for three stock portfolios. The portfolios are sorted each month according to the stocks idiosyncratic skewness risk premium (computed with equation 16), with ptf1 (ptf3) denoting the portfolio with the lowest (highest) risk premium. The CAPM-beta is calculated for each stock at each start date of the swaps by regressing the daily excess returns of the stock in the previous month on the daily excess return of the S&P500. The column labeled '3-1 t-stat' reports the t-statistic of the two-sample test for equal means, which tests the null hypothesis that ptf3 and ptf1 have the same CAPM-beta. The t-statistics are adjusted using Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

Sort according to idiosyncratic skewness risk premium				
	ptf1	ptf2	ptf3	3-1 t-stat
CAPM-beta	1.1503	1.1719	1.1815	1.19

I Additional figures

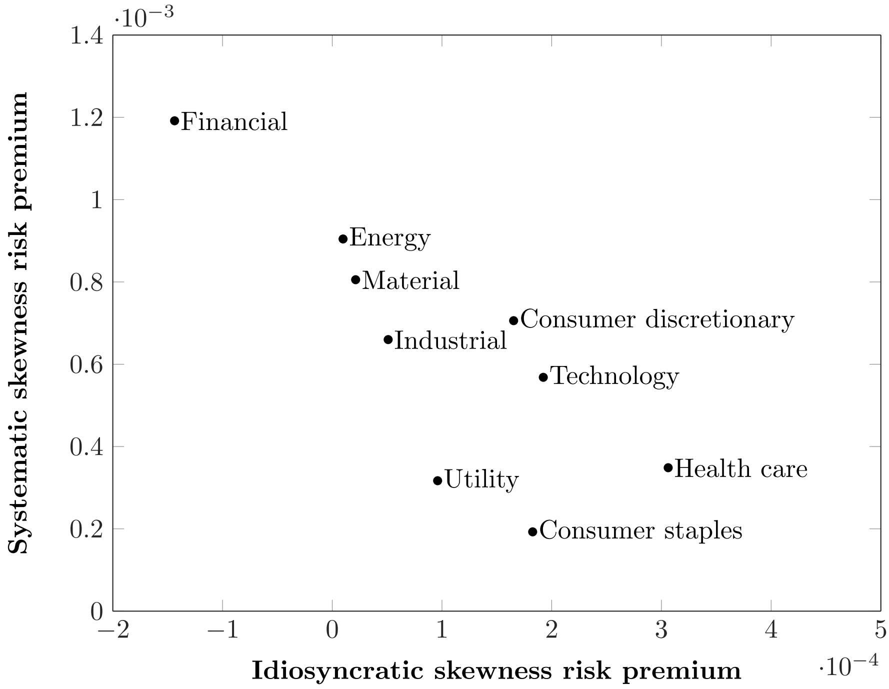

Figure 10. Idiosyncratic and systematic skewness risk premium - breakdown by sector. This figure shows for each sector the average estimate of the idiosyncratic and systematic skewness risk premium. For each stock separately, I run the median regression $RP_{i,t} = \alpha_i + \beta_i RP_{SP500,t} + \epsilon_{i,t}$ where $RP_{i,t}$ is the gain of my swap trading strategy for stock i (see Table 20 for details). I then divide the stocks of my sample in nine portfolios according to their sector and I compute the average sector idiosyncratic skewness risk premium as $\overline{\widehat{\alpha_i}}$ and the average sector systematic skewness risk premium as $\overline{\widehat{\beta_i}} E[RP_{SP500}]$.