

Risk Premia: Asymmetric Tail Risks and Excess Returns

Y. Lempérière, C. Deremble, T. T. Nguyen
P. Seager, M. Potters, J. P. Bouchaud
Capital Fund Management,
23 rue de l'Université, 75007 Paris, France

September 30, 2014

Abstract

We present extensive evidence that “risk premium” is strongly correlated with tail-risk skewness but very little with volatility. We introduce a new, intuitive definition of skewness and elicit a linear relation between the Sharpe ratio of various risk premium strategies (Equity, Fama-French, FX Carry, Short Vol, Bonds, Credit) and their negative skewness. We find a clear exception to this rule: trend following (and perhaps the Fama-French “High minus Low”), that has positive skewness and positive excess returns, suggesting that some strategies are not risk premia but genuine market anomalies. Based on our results, we propose an objective criterion to assess the quality of a risk-premium portfolio.

1 Introduction

One of the pillars of modern finance theory is the concept of “risk premium”, i.e. that a more risky investment should, on the long run, also be more profitable. The argument leading to such a conclusion looks, at least at first sight, rock solid: given the choice between two assets with the same expected return, any rational investor should prefer the less risky one. The resulting smaller demand for riskier assets should make their prices drop, and therefore their yields go up, until they become attractive enough to lure in investors in spite of their higher risk. This simple idea appears to explain why stock indices on the long run yield higher returns than monetary funds, for example. It is also at the heart of the well-known Capital Asset Pricing Model (CAPM), which predicts that the excess return of a given stock (over the risk free rate) is proportional to its so-called “ β ”, i.e. to its covariance with the market risk [1].

The idea has become so pervasive that its logic has actually inverted: efficient market theorists assume that any return in excess of the risk free rate should in fact be associated with a non-diversifiable risk factor of some sort – sometimes argued to be unobservable or “latent”. It is often invoked to do away with any market anomaly or discrepancies between market prices and theoretical prices, as it allows one to introduce an extra fitting parameter conveniently called the

“market price of risk”, see e.g. [2].¹

Although the argument underlying the existence of risk premia is cogent, estimating and rationalising their order of magnitude in financial markets is still very much a matter of debate. This is quite surprising in view of the importance of the issue for the asset management industry as a whole, and reveals how primitive our understanding of financial prices and returns is. Empirically, the equity risk premium is found to be in the range 3–9% depending on the period and country (see [3, 4] and our own data in Sect. II below). Mehra & Prescott have pointed out that these values are too high to be explained within a general equilibrium model of the economy and have coined this discrepancy the “equity premium puzzle” [5].² More recently, several authors have reported an *inverted* relation between the volatility of a stock (or its β) and its excess return. In other words, contrarily to intuition, less volatile stocks appear to be more profitable [6, 7].

As pointed out in numerous papers – many of them quite recent – the problem may reside in the very definition of risk. Classical theories identify risk with volatility σ . This (partly) comes from the standard assumption of Gaussian fluctuations of asset returns, which is entirely characterized by its first two moments: mean μ and variance σ^2 . But in fact fluctuations are known to be strongly non Gaussian, and investors are arguably not concerned by small fluctuations around the mean. Rather, they fear large negative drops of their wealth, induced by rare, but plausible crashes. These negative events are not captured by the r.m.s. σ but rather contribute to the negative skewness of the distribution. Therefore, an alternative idea that has progressively emerged in the literature is that a large contribution to the “risk premium” is in fact a compensation for holding an asset that provides positive average returns but may occasionally erase a large fraction of the accumulated gains. This was in fact proposed as a resolution of the Mehra & Prescott “equity premium puzzle” first by Rietz [8] and then in many follow up papers with a similar message [9, 10, 11, 12], see also [13, 14]. Risk premium might be more a question of skewness and negative tail events than related to volatility *per se*.

A more formal way to frame this story is in terms of the utility function supposed to describe investors’ preferences. Assuming that the “happiness” of the investor depends on his/her wealth W as a certain function $U(W)$, the preference for higher profits imposes that $U(W)$ is a monotonously increasing function of W , whereas the condition $U''(W) \leq 0, \forall W$ ensures that a less risky investment with the same expected gain is always preferred. Assuming that the proposed investment yields an uncertain final wealth W distributed around a mean $\mathbb{E}[W]$ with small fluctuations of variance σ^2 , and expanding in powers of

¹Note, however, that since the market price of risk itself is usually found to be strongly fluctuating in time, the predictive strength of these theories becomes frustratingly close to zero.

²However, one should note that the Mehra & Prescott framework also fails to explain the high level of the volatility of the stock market, i.e. it does not explain the “excess volatility puzzle” either.

the difference $W - \mathbb{E}[W]$, one finds that the expected utility is given by:³

$$\mathbb{E}[U(W)] = U(\mathbb{E}[W]) - \frac{1}{2}\lambda^2\sigma^2 + \dots \quad \lambda^2 := -U''(\mathbb{E}[W]) \geq 0. \quad (1)$$

Hence, within a small risk expansion, utility reduction is entirely described by the variance σ^2 . The above approximate formula is the starting point to most risk-return portfolio optimisation schemes in the classical literature.

But as intuition suggests, skewness matters. Lottery experiments [15, 16] demonstrate that agents prefer lotteries with very large potential gains, even when the expected gain is negative, to lotteries with small positive expected gain but no possibility of large payoffs. Symmetrically, one expects agents to shy away from investments with a high degree of negative skewness. Clearly, this cannot be captured by Eq. (1) above, where skewness is absent. A large amount of papers have therefore been devoted to generalising Markowitz’s portfolio optimisation and the CAPM to account for skewness preferences [9, 17, 18]. In terms of utility functions, one can show that the additional condition $U'''(W) \geq 0, \forall W$ ensures that more positively skewed investments are preferred [19].⁴

Our work is clearly in the wake of the above mentioned literature on skewness preferences and tail-risk aversion. We will present extensive evidence that “risk premium” is indeed strongly correlated with the skewness of a strategy but very little with its volatility, not only in the equity world – as was emphasized by previous authors – but in other sectors as well. We will investigate in detail many classical risk premium strategies (in equities, bonds, currencies, options and credit) and elicit a possibly universal linear relation between the Sharpe ratio of these strategies and their negative skewness. We will find however that some well-known strategies, such as trend following and to a lesser extent the Fama-French “High minus Low” factor, are clearly not following this rule, suggesting that these strategies are not risk premia but genuine market anomalies.

We will first start in Sect. 2 with the equity market as a whole and revisit the equity risk premium world-wide, and its correlation with the volatility. We introduce a simple way to plot the returns of a portfolio that reveals its skewness to the “naked eye” and suggests a new, intuitive definition of skewness that we use throughout the paper. We then focus on the Fama-French factors in Sect. 3 and study market neutral portfolios. We move on to the fixed income world (Sect. 4), where we again build neutral portfolios. Sect. 5 is devoted to an account of risk premia on currencies, and finally, in Sect. 6, to the paradigmatic case of options. We summarize our findings in Sect. 7 with a universal linear relation between the Sharpe ratio and the skewness of all the Risk Premium strategies investigated in the paper, and discuss some exceptions to the rule – i.e. positive Sharpe strategies with zero or positive skewness – that we define as “pure α ” strategies. We then end the paper by proposing a simple argument

³ A way to avoid this expansion is to assume an exponential (CARA) utility function $U(W) \equiv -\exp(-\lambda W)$ and Gaussian fluctuations for W . In this case,

$$\mathbb{E}[U(W)] \equiv -\exp[-\lambda\mu + \frac{1}{2}\lambda^2\sigma^2].$$

⁴ More precisely, this preference ordering requires in general the distributions to be “skewness comparable”, which means that their cumulative functions must cross twice and only twice [20, 19]. This turns out to be often the case in financial applications, see below and Appendix.

to rationalize for the Sharpe ratio/skewness tradeoff and an objective criterion to assess the quality of a risk premium portfolio in Sect. 8.

2 Back to Basics: Risk Premium in Global Equity Markets

2.1 No “volatility premium” in stock indices

According to the standard framework described in the previous section, investing in the stock market rather than in risk-free instruments only makes sense if stocks provide, in the long run, superior returns than the risk-free rate. The existence and the strength of this equity risk premium (ERP) has been much debated amongst both the academic community and the asset management industry. The general consensus (confirmed below) is that this premium exists and is in the range 3–9%, with large fluctuations, both across epochs and countries. Within a classical general equilibrium framework, this premium appears to be large: investors should put money more massively in equities, and the demand for bonds should be much weaker than what we observe, which should lead to a smaller risk premium. This is the “equity premium puzzle” brought to the fore by Mehra & Prescott in 1985 [5, 3].

In order to estimate the ERP, it is customary to compute the total return of indices (dividend included), and to subtract the risk-free rate. This is virtually equivalent to buying a future contract on indices. We have computed this quantity for various countries, using the 10Y government rate as a risk-free asset (using 3-month bills would only give a more flattering result, since short rates are usually lower than 10Y rates). The result is, as expected, overwhelmingly in favor of the existence of a risk premium, as can be seen from Table 1 and Fig. 1. All countries but one display positive excess return on the long run (some time series go back as far as 1870). The global t-stat of a strategy equi-invested in all available indices at any moment in time is above 4. The average value (weighted by the number of available years for each country) is found to be $\approx 5\%$ with a cross-sectional dispersion of $\pm 3\%$, compatible with numbers quoted above [3, 4].

In the last column of Table 1, we give the average annual volatility of the corresponding index. This allows us to compute the Sharpe ratio $\mathcal{S} = \mu/\sigma$ of each risk premium. Note indeed that the value of the risk premium itself is not meaningful since it can be leveraged. Therefore here and throughout the paper we use the a-dimensional Sharpe ratio to meaningfully compare different risk premia. Fig. 2 shows a scatter plot of \mathcal{S} as a function of the volatility of the index, together with a regression line. We see that the two quantities are in fact *negatively* correlated ($\rho \approx -0.45$). Even the risk premium itself (i.e. without dividing by the volatility σ) is found to be weakly anti-correlated with σ ($\rho \approx -0.1$) but this is not statistically significant. This is at odds with the standard interpretation of the risk premium.

Country	Start date	Risk Premium (%/y)	Vol. (%/y)	Skewness ζ^*
Australia	1983	4.77	16.7	-2.35
Canada	1934	4.44	14.8	-1.68
Germany	1870	4.23	19.6	-0.31
France	1898	6.65	19.7	1.05
Finland	1970	8.09	23.1	0.8
Italy	1925	6.24	26.6	1.30
Hong Kong	1996	5.24	26.1	-1.87
Hungary	1999	2.51	24.7	-1.79
India	1988	13.6	30.3	0.63
Indonesia	2009	13.8	20.4	-3.26
Japan	1980	2.04	18.7	-0.85
South Korea	2000	9.33	24.3	-0.08
Malaysia	1974	6.32	28.2	-0.14
Mexico	2001	10.1	18.2	-2.26
The Netherlands	1957	4.57	16.9	-1.7
New Zealand	1984	3.02	17.8	-1.67
Norway	1969	5.21	24.3	0.17
Philippines	1996	-1.96	26.6	-0.42
Poland	1999	2.29	23.3	-0.62
Singapore	1973	7.51	24.7	-0.86
South Africa	1983	5.26	20.4	-0.14
Spain	1948	5.25	18.3	-0.12
Sweden	1960	6.69	17.1	-0.42
Switzerland	1966	4.58	15.6	-2.05
Thailand	1997	5.25	35.1	-0.94
United Kingdom	1958	5.71	18.8	-2.31
United States	1871	5.33	15.2	-0.32

Table 1: Average excess return (“risk premium”), volatility and skewness ζ^* of various indices over the 10Y bond rates, based on monthly returns. Start date as indicated in the second column; end date Jan 2014. Data from Global Financial Data, (www.globalfinancialdata.com).




Figure 1: Cumulated ERP since 1870, corresponding to a strategy where one holds a long portfolio of indices, with equal weight on all available contracts at any instant of time, i.e. 1 in 1870 and 27 in 2014. The t-stat is 4.2.




Figure 2: Sharpe ratio of the excess return (over the 10Y bond) for each index listed in Table 1 as a function of the volatility of that index. We also plot the regression line, with a *negative* slope and a correlation coefficient $\rho \approx -0.45$. Note that the (flat) average Sharpe ratio over all contracts is $\bar{S} \approx 0.3$.

2.2 Downside tail risk and skewness premium

However, something interesting occurs when we suggestively re-plot the cumulated ERP (say for the risk managed SPX since 1928⁵) in the following way. Instead of considering the returns of the strategy in chronological order, we sort these returns in terms of their absolute value and plot the cumulated ERP $F(p)$ as a function of the normalized rank $p = k/N$, starting from the return with the smallest amplitude ($k = 1$) and ending with the largest ($k = N$). The result in the case of the US since 1928 is shown in Fig. 3. We immediately see that while the small returns contribute positively to the average, the largest returns, contrarily, lead to a violent drop of the P&L. Strikingly, the 5% largest returns wipe out half of the ERP of the 95% small-to-moderate returns!

In the same graph, we also show the P&L $F_s(p)$ that would have been observed if the distribution of returns was exactly symmetrical around the same mean – i.e. such that the final point $F_s(p = 1)$ coincides with $F(p = 1)$.⁶ In this case, it is easy to show that the P&L is (for large N) a monotonously increasing function of the normalized rank p . One finds that generically $F(p)$ behaves as p^3 at small p – explaining the strong curvature seen in Fig. 3. The comparison between real returns and symmetrized returns therefore reveals the *strongly skewed* nature of the ERP. This will be the main theme of our paper.

The ranked amplitude P&L representation in fact suggests a new general definition of skewness, as follows: one first normalizes the returns such that their average is zero and their rms is unity. Then the ranked amplitude P&L is plotted, and defines a certain function $F_0(p)$, such that $F_0(0) = F_0(1) = 0$. By construction, the symmetrized version $F_{0s}(p)$ is now zero (up to small fluctuations). Our definition of skewness (that we call ζ^*) is simply (minus) the area underneath the function $F_0(p)$:⁷

$$\zeta^* := -100 \int_0^1 dp F_0(p), \quad (2)$$

where the arbitrary factor 100 is introduced such that the skewness is of order unity. We give several properties of $F_0(p)$ in the Appendix and show how ζ^* can be related to other, more standard definitions of skewness when returns are independent, identically distributed variables. [Note that all the results reported in this paper are actually robust to the chosen definition of skewness.] We also show in the Appendix that whenever $F_0(p)$ has a *single hump*, as is usually the case in financial applications, then the cumulative distribution of returns crosses its symmetrized version *twice and only twice*. This means, as announced in the introduction, that investors with a utility U such that $U''' > 0$ will shy away from investments with negative ζ^* , i.e. will demand a higher risk premium. In the case of the US market factor, we find $\zeta^* \approx -1.47$ for daily returns (for a Sharpe ratio of 0.57), and $\zeta_{month}^* \approx -0.32$ for monthly returns.

⁵Here and below, we will consider “risk managed” strategies that run at approximately constant risk over time. This is achieved by normalising the position by a twenty-day exponential moving average of the volatility of the bare strategy itself. In the case of the SPX, this simply amounts to the recent volatility of the SPX.

⁶Technically, this amounts to transforming the returns r_t as: $m + \epsilon_t(r_t - m)$ where m is the average return and $\epsilon_t = \pm 1$ an independent random sign for each t .

⁷One could actually directly take the area between $F(p)$ and $F_s(p)$ as our definition of skewness, i.e. without first removing the average drift μ . Up to relative corrections of order $(\mu/\sigma)^2$, the result would be identical. For all purposes, this correction is extremely small, at least for daily returns (i.e. less than 10^{-3}).




Figure 3: The ranked amplitude P&L representation: plot of the cumulated daily P&L $F(p)$ for the SPX (at constant risk) since 1928, as a function of the normalized rank of the *amplitude* of the returns, $p = k/N$ (in red). The standard, chronological P&L corresponding to holding the US equity market index at constant risk is shown in black, that by construction ends at the exact same point. We also show for comparison (in green) $F_s(p)$, corresponding to returns symmetrized around the same global mean (i.e. $F_s(p = 1) = F(1)$, again by construction). Our definition of the skewness is, up to a negligible correction, minus the area enclosed between $F(p)$ and $F_s(p)$ (shown in yellow).

When analyzing the data given in Table 1, we do find a positive correlation between $-\zeta^*$ and the ERP across different countries, $\rho \approx 0.4$. This is encouraging but not particularly convincing. We have therefore extended our analysis to different contracts and different risk premia, and indeed find that in most cases, excess returns lead to a humped shape function $F_0(p)$, meaning that returns of large amplitude have a negative mean. This will lead us to propose a possibly universal definition of “risk premium” in terms of the skewness of the returns – in a precise sense, the premium compensates for a tail risk that is systematically biased downwards.

3 Zooming into stock markets: the Fama-French factors

We now consider the celebrated Fama-French model, which attributes excess stock returns to 4 components, corresponding to 4 orthogonal portfolios: the Equity Risk Premium discussed above (i.e. market minus the risk-free rate, MKT), SMB (“Small” caps minus “Big” caps), UMD (“Up” – previous winners

minus “Down” – previous losers) and HML (High book to price – “value” stocks minus Low book to price – “growth” stocks), each of which is interpreted as a reward for a specific risk, even if the origin of said risk is not always clear.

The first, long market portfolio is clearly not market neutral, while the other three are explicitly constructed in a market neutral way, in order to be orthogonal to the market risk factor MKT. SMB consists of being long the small cap stocks and short the large caps. This sounds like an intuitive risk premium strategy, since it seems plausible that small caps have a greater expected future growth but are more “risky” than large caps. However, if we look at Table 2, we see that the volatility of large or small cap indices is practically the same in the US and that, if anything, volatility is actually *smaller* for small caps than for large caps in all other countries! Therefore, volatility alone is unable to explain the existence of an SMB risk premium. This is known in the literature as the “low volatility puzzle”; many authors have reported that low-volatility stocks actually yield higher returns than high-volatility stocks.

But if we now observe the above defined skewness ζ^* of large, mid and small cap indices reported in Table 2, we clearly conclude that skewness does indeed become more negative as we move to smaller cap indices. This behaviour is furthermore quite universal: across the 10 cases we have investigated, only the Netherlands and New Zealand show a slight inversion of the skewness of mid and small cap stocks. Let us insist that we have also computed three other (more standard) estimates of the skewness with exactly the same conclusion. To add to the significance of this observation, we have built market neutral portfolios that are long 1\$ of the small-cap index, and short 1\$ of the large-cap index for each of the 10 countries we consider. We then globally risk manage these portfolios by rescaling the position with a 20-day moving average of their own volatility. In this way, we achieve constant volatility throughout the simulation. For each country in our pool, the skewness of these portfolios is found to be negative, and they all yield positive performance, except one of them (Australia) – see Table 2, last column. This is quite significant, since these portfolios are completely decorrelated from one another.

We can actually dwell further on this issue using the “decile” portfolios for the three Fama-French factors in the US equity market.⁸ We report in Table 3 the volatility, skewness and Sharpe ratio for being long each decile portfolio since 1950, for the three factors SMB, HML and UMD. Again, there is little evidence that the volatility plays any role at all for any of the 3 factors. The skewness, on the other hand, becomes more negative for the portfolios with the highest returns for the SMB and UMD (momentum) factors. The smallest caps have a skewness $\zeta^* = -1.83$, whereas the largest caps have much weaker skewness $\zeta^* = -0.39$, confirming the conclusion above. Winners have a skewness $\zeta^* = -0.90$ and are prone to large reversal events, whereas losers are essentially not skewed. For HML, surprisingly, there is no such effect; if anything, the skewness appears to be slightly less negative for high deciles, i.e. for value stocks with large book to price ratios. Notice also that the correlation between skewness and Sharpe ratios is clearly inverted for HML, see Table 3. As a visual confirmation, we plot in Fig. 4 the ranked amplitude P&L of the three Fama-French factors in the US since 1950. Clearly, SMB and UMD show the typical humped shape observed for the market ERP in the previous section, therefore corresponding

⁸Available at <http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/>

Index	Start date	Cap	σ (%/day)	ζ^*	SR of SMB
S&P	1990	B	1.1	-0.66	0.32
		M	1.2	-1.12	
		S	1.2	-1.16	
Russell Russell	1978	B	1.1	-0.58	0.07
		S	1.2	-1.43	
Japan	1980	B	1.3	0.08	0.02
		M	1.1	-0.63	
		S	1.0	-0.94	
France	1999	B	1.4	-0.78	0.40
		M	1.1	-1.66	
		S	0.9	-2.40	
The Netherlands	1995	B	1.4	-0.92	0.33
		M	1.3	-1.71	
		S	1.0	-1.68	
Italy	2001	B	1.2	-1.68	0.70
		M	1.2	-1.86	
		S	1.0	-2.98	
Australia	1994	B	1.0	-0.78	-0.27
		M	1.0	-0.88	
		S	1.0	-1.41	
New Zealand	1998	B	0.8	-0.90	0.39
		M	0.7	-1.23	
		S	0.5	-1.06	
Poland	1998	B	1.7	0.15	0.34
		M	1.2	-0.92	
		S	1.2	-1.17	
South Africa	1996	L	1.3	-0.68	0.26
		M	0.9	-1.36	
		S	0.7	-1.59	

Table 2: Volatility (in percent per day), and skewness of the daily returns of large(B)/mid(M)/small(S) cap indices in various countries. In the last column, we give the Sharpe ratio of portfolios that are long small caps (S) and short large caps (B), and neutral on medium caps (M).

	1	2	3	4	5	6	7	8	9	10
SMB σ (%/d)	0.83	1.00	1.00	0.98	0.97	0.92	0.93	0.94	0.93	0.96
SMB ζ^*	-1.83	-1.53	-1.49	-1.52	-1.42	-1.45	-1.26	-1.28	-0.93	-0.39
SMB SR	0.56	0.48	0.53	0.50	0.53	0.53	0.53	0.50	0.48	0.39
UMD σ (%/d)	1.48	1.21	1.05	0.99	0.95	0.93	0.92	0.94	1.01	1.23
UMD ζ^*	0.00	-0.14	0.00	-0.24	-0.24	-0.30	-0.58	-0.71	-0.79	-0.90
UMD SR	-0.07	0.17	0.33	0.37	0.39	0.46	0.45	0.62	0.53	0.67
HML σ (%/d)	1.05	0.97	0.93	0.95	0.94	0.92	0.91	0.97	0.98	1.10
HML ζ^*	-0.70	-0.69	-0.52	-0.68	-0.52	-0.66	-0.59	-0.65	-0.55	-0.36
HML SR	0.33	0.41	0.43	0.44	0.50	0.52	0.52	0.59	0.61	0.60

Table 3: Volatility, skewness and Sharpe ratio of the daily returns of the decile portfolios corresponding to SMB, UMD and HML in the US equity market since 1950. See also Fig. 7 for a plot of Sharpe ratios vs. skewness. Decile 1 corresponds, respectively, to Small caps, Down stocks and Low earning to price stocks, whereas decile 10 corresponds to Big caps, Up stocks and High book to price.

to a negative overall skewness equal, respectively, to $\zeta^* = -1.39$ and -1.38 . For UMD, the 5% largest events contribute to losses that amount to 20% of the gains accumulated by the 95% small to moderate returns, whereas for SMB nearly all the gains are erased by the 5% largest events! For HML however the situation is inverted: large returns contribute positively to the gains, therefore leading to a positive skewness $\zeta^* = +0.25$. We will discuss this feature of HML in Sect. 7-2.

Apart from the special HML case, the above findings give further credence to the idea that risk premium is actually *not* compensating investors for taking a symmetric risk (as measured by the volatility), but, as suggested by Rietz, Barro and others [8, 10, 9, 12], for carrying the risk of large losses. However, at variance with these authors, we do not attempt to define catastrophic macro-economic events and estimate their probability, but rather directly measure the skewness of these risk premia portfolios through our single humped function $F(p)$ and its integral $-\zeta^*$. These quantities immediately reveal the “tail risk” embedded in these strategies, i.e. that large returns systematically have a negative mean.

4 Risk Premia in the fixed-income world

In the world of fixed income, credit risk is gauged by the investment grade of the corresponding bond, and one should in principle expect higher yields for lower investment grades (i.e. riskier bonds). We have studied corporate bond portfolios, using Global Financial Data time series for US bond total return data.⁹ We have used the BofA Corp. total return index suite, that provides daily bond returns since 1997 for grades going from AAA (safe) to CCC and below (risky). We have summarized the average yield, volatility and skewness for each of these indices in Table 4. The first striking observation is that although

⁹In this paper, we only study US corporate debt, but we believe our results have a more universal scope. For example, we have checked that a portfolio of international 10Y bonds also exhibits the same features.




Figure 4: The ranked amplitude P&L $F(p)$ of the three Fama-French factors, SMB, UMD & HML. Note the familiar humped shape for SMB and UMD, revealing strong negative skewness (respectively $\zeta^* = -1.39$ and -1.38), but an inverted behaviour for HML ($\zeta^* = +0.43$). The corresponding Sharpe ratios are, respectively, 0.14, 0.67 and 0.64. We also show for reference the p^3 behaviour expected at small p (for clarity, only shown for UMD).

Grade	Yield (%/y)	σ (%/y)	ζ^*
AAA	4.7	5.1	-0.87
AA	4.8	4.6	-0.74
A	5.3	4.9	-0.83
BBB	6	4.8	-0.81
BB	7.7	4.0	-1.78
B	9.5	4.8	-1.99
CCC and below	15.5	7.3	-1.94

Table 4: Return and volatility (in percent), and skew of BofA US Corp. indices, using Global Financial Data (www.globalfinancialdata.com). All data since 1997




Figure 5: Left: Performance of a risk-managed portfolio long a high yield index (CCC or BB), and short its AAA counterpart. SR: 0.54/0.45, t-stat: 2.2/1.9, ζ^* :-0.84/-0.43. Right: Ranked amplitude P&L of the CCC/AAA portfolio, showing the characteristic humped shape of portfolios with downside tail risk.

we cover a variety of grades and yields (ranging from 4.7% annual for AAA grade to 15.5% annual for CCC and below), the volatility of these indices is found to be remarkably constant at around 4.8% annualised, except for the lowest grade which has a 7.3% volatility. Similar to what we found above for stock indices corresponding to different market caps, the skewness is much more revealing than the volatility as it clearly drops between BBB and BB grades (see Table 4).

Another way to elicit this skewness premium is to build a portfolio where we invest 1\$ every day in a high yield index (CCC or BB), and sell 1\$ of the AAA index. To achieve constant volatility throughout the simulation, the portfolio is (as above) risk-managed based on its past volatility, as measured by a 20-day moving average of the absolute returns of the strategy itself. The performance of the BB/AAA or CCC/AAA strategies is shown in Fig. 5, Left. The two strategies are similar and yield positive (but marginally significant) performance. While the explanation in terms of volatility difference is not very convincing, we again find that the ranked amplitude P&L plot has a clear hump shape (see Fig. 5, Right), indicating again that large returns are systematically biased downwards.

5 Risk Premia in the FX world: the “Carry” trade

Similarly to the previous section where we traded bonds of differing quality, one may want to buy short term debt in a country whose short term rate is high using borrowed money from a country whose short term rate is low. This idea is a textbook strategy in the world of FX finance called the carry trade. In the absence of adverse exchange rate moves, this allows the investor to pocket the difference between the high interest rate and the low interest rate. This excess return is again usually thought of as a risk premium: one invests in the economy of a “risky” country, while getting financed in a safe country. The questions we want to ask are, as above: a) is the excess return compensating for volatility? and b) is the strategy negatively skewed?

We have simulated this currency pair strategy on a pool of 20 developed countries using data from Global Financial Data (www.globalfinancialdata.com), both for spot currencies and for interest rates, from Jan 1974 to Jan 2014.¹⁰ We use intra-bank rates when they are available, and central bank discount rates otherwise. For each day, we consider all $20 \times 19/2$ currency ordered pairs, using the difference of interest rate in the two countries to order the pairs in a systematic way. We then make 10 deciles (following the Fama-French construction) corresponding to pairs with the smallest interest rate difference (i.e. the smallest “carry”) to the largest. We then compute, as above, the volatility, skewness and Sharpe ratio corresponding to these 10 portfolios. The data is shown in Table 5 and as small plusses in Fig. 7 below. We find again that skewness and Sharpe ratio are quite strongly correlated $\rho = -0.76$, with a correlation that is significantly larger in the period 1994-2014 than in the previous period (1974-1993). The Sharpe ratio of the fully diversified Carry strategy (i.e. with an equal weight over all ordered FX pairs) is $\mathcal{S} \approx 0.85$ for a skewness $\zeta^* \approx -0.94$. These results somewhat contrast with the recent conclusion of Jurek [21] who finds that the profitability of the FX carry trade is *not* related to the implied tail-risk, as provided by out-of-the-money currency options. However, hedging the carry trade with at-the-money options instead completely removes the carry P&L, suggesting a more complex picture where the skewness of the carry trade is related to intermediate (rather than extreme) adverse currency moves.

6 A paradigmatic example: Risk Premia in option markets

The textbook example of a risk premium strategy is selling options, i.e. insurance contracts against large movements of financial assets – for example stock indices. In this case, there is a built-in asymmetry between the risk taken by the option seller (the insurer) and the option buyer. The option seller bears the risk of arbitrarily large losses, and hence the skewness of the P&L of the insurer is (almost by definition) negatively skewed – even after hedging in a world where perfect replication is not possible, see the discussion in [22]. Conversely,

¹⁰The pool of countries is made up of: Australia, Canada, Switzerland, Czech Republic, Germany, Hungary, Israel, Iceland, Japan, Korea, Mexico, Norway, New Zealand, Poland, Sweden, Singapore, Thailand, UK, USA, South Africa.

	1	2	3	4	5	6	7	8	9	10
FX Carry σ (%/d)	0.45	0.47	0.45	0.50	0.55	0.58	0.64	0.68	0.74	0.93
FX Carry ζ^*	-0.41	-0.35	-0.60	-0.50	-0.87	-0.97	-0.86	-1.04	-0.81	-1.05
FX Carry SR	-0.28	-0.05	0.56	0.44	0.51	0.40	0.78	0.62	0.91	1.04

Table 5: Volatility, skewness and Sharpe ratio of the daily returns of the decile portfolios corresponding to FX Carry on G20 currencies since 1974. See also Fig. 7 for a plot of Sharpe ratios vs. skewness. Decile 1 corresponds to pairs with small interest rate difference whereas decile 10 corresponds to pairs with small interest rate difference. Note that in this case, there is a positive correlation between the Sharpe ratio and the volatility, on top of the now familiar negative correlation between Sharpe ratio and skewness.

the skewness of the option buyer is positive. Therefore, one should expect that option sellers require a “skew premium” as a compensation for the asymmetric risk they agree to cover. In order to confirm this rather basic intuition, we have computed the performance of selling variance swaps, which are a basket of options whose return is directly proportional to the difference between realized and implied variance. This difference is exactly the risk premium we seek to elicit, so these baskets are perfectly suited to our purpose.

We test this risk premium strategy on a variety of products, since the above skew premium argument should be a universal feature valid across the board, from commodity (Crude oil, Gold) and currency options (EURXUSD, GBPXUSD) to stock index variance swaps (S&P 500, DAX). Selling each of these variance swap contracts, and risk-managing using the last 20-day’s volatility of the P&L itself (as we did in the previous sections), leads to the global performance shown in Fig. 6. As we can see, the return is clearly positive, with a t-stat over 4, while the strategy is strikingly skewed, as expected. We have summarized the Sharpe ratio and the skewness ζ^* for all six contracts in Table 6. Clearly the skewness is indeed negative in all cases. In fact, the skewness ζ^* is substantially larger than all the examples considered so far in this paper and the performance of the individual short varswap strategies is also somewhat better than all the risk premia analyzed above. Interestingly, shorting the VIX future leads to a P&L that is both less skewed and with a lower Sharpe ratio (see Fig. 7 below). This again suggests that there might exist a systematic relation between risk premia and skewness. This is the point we now discuss.

7 Risk premium is Skewness premium

7.1 A remarkable plot

The consistent picture that seems to emerge from all the above empirical results is that “risk premia” are in fact very weakly (if at all) related to the volatility of an investment, but rather to its skewness, or more precisely to the fact that the largest returns of that investment are strongly biased downwards. This has already been discussed several times in the context of stock markets, from different standpoints (economic, econometric and financial). For example, the so called “low volatility” puzzle (also known as “Betting against Beta”) goes




Figure 6: Performance of a risk managed short-option strategy, equi-weighted on S&P 500, DAX, Crude oil, Gold, EURXUSD and GBPXUSD variance swaps from Jan 1996 until Jan 2014. The Sharpe ratio is 1.26. We also show the ranked P&L which reveals the strongly skewed nature of selling volatility.

Underlying	Sharpe Ratio	ζ^*
S&P 500	1.47	-4.64
DAX	0.79	-5.53
Crude oil	0.74	-5.28
Gold	0.71	-4.79
EURXUSD	0.88	-5.32
GBPXUSD	0.38	-5.01
Total	1.26	-4.61

Table 6: Sharpe ratio and skewness ζ^* of short varswap, risk managed strategies on different underlyings since Jan 1996.

Underlying	Vol/SR corr.	Skewness/SR corr
Bonds	-0.69	-0.36
Intl. IDX	-0.45	-0.38
SMB	-0.42	-0.89
UMD	-0.63	-0.85
FX Carry	+0.78	-0.76
HML	+0.03	+0.64
TREND	+0.23	+0.58

Table 7: Correlation coefficient ρ between volatility and Sharpe ratio, and between skewness and Sharpe ratio for all strategies investigated in this paper, and for a 50-day trend following strategy across various futures markets. In most cases, and counter-intuitively, correlation with volatility is found to be *negative* or zero. Correlation with skew, on the other hand, is systematically negative, showing that the main determinant of risk premium must be skewness and not volatility. Trend and HML are clear outliers.

against the intuitive argument that more (volatility) risk must be accompanied by higher average returns. Similarly, we find no evidence of a higher volatility in the small caps, large momentum, or high value stocks that would explain the excess returns associated to these factors, while skewness seems to play an important role. Perhaps obviously, the risk premium collected by shorting options appears to be compensating for the substantial skewness of an insurance P&L. Table 7 summarises the correlation between Sharpe ratios, volatility and skewness for all the strategies discussed above.

In order to reveal a possibly universal relation between excess returns and skewness, we have summarized all our results above in a single scatter plot, Fig. 7, where we show the Sharpe Ratios of different portfolios/strategies as a function of their (negative) skewness $-\zeta^*$. Quite remarkably, *all but one* fall roughly on the regression line $\mathcal{S} \approx 1/3 - \zeta^*/4$. This is the central result of our paper. The two parallel dashed lines correspond to a $2\text{-}\sigma$ channel, computed with the errors on the SR and the skewness of the Fama-French strategies (other strategies have even larger error bars), see Table 8. The glaring exception, that we discuss in more detail below, is the 50 day trend following strategy on a diversified set of futures contract since 1960 (see [23] for details). The other outlier, although less clear-cut, is the Fama-French HML factor, which also has positive skewness and, as illustrated in Table 3, a positive correlation between skewness and Sharpe ratio for the performance of the decile portfolios.

Interestingly, we have also considered the returns of a portfolio of four credit indices since 2004, and of the HFRX global hedge fund index, which provides us with daily data since 2003. Although this history is relatively short, the Sharpe ratio/skewness of both credit and of the hedge fund index fall in line with the global behaviour (once fees are taken into account in the case of the HFRX ¹¹). Let us insist once more that Fig. 7 would look qualitatively similar for other, more conventional definitions of skewness.

¹¹The Sharpe ratio of HFRX with fees is 0.48, and we have corrected this to yield a Sharpe ratio without fees of 0.85. This assumes that funds run at a volatility of around $\sim 10\%$ and have a classic 2 and 20 fee structure.

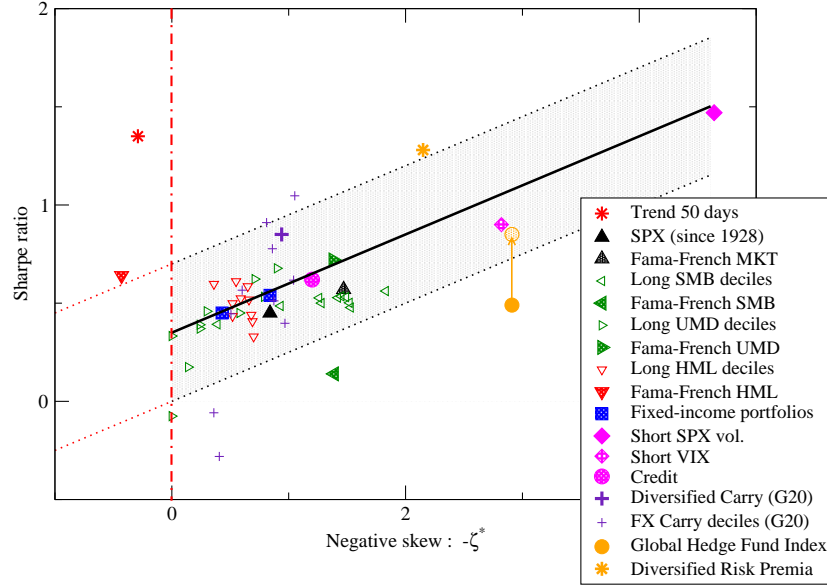


Figure 7: Sharpe ratio vs Skewness $-\zeta^*$ of all assets and/or strategies considered in this paper (see legend). The relatively low Sharpe of the HFRX index can be explained by the hedge fund fees (the vertical arrow corresponds to a tentative estimate of these fees). The plain line $\mathcal{S} \approx 1/3 - \zeta^*/4$ corresponds to the regression line through all risk premia marked with filled symbols (excluding trend following). The two dotted lines correspond to a $2\text{-}\sigma$ channel, computed with the errors on the SR and on the skewness of the Fama-French strategies. The small empty triangles correspond to the Fama-French decile portfolios. Trend following is a clear outlier marked as a red point, characterized by a positive skewness *and* a positive SR. The same is true, perhaps to a lesser extent, of HML. Note also the interesting difference between short varswap and short VIX strategies.

Strategy	Error on the skewness	Error on the SR
FXR pairs	0.19	0.18
Fixed income	0.27	0.24
Agg. varswaps	0.22	0.24
Global hedge fund index	0.41	0.30
Trend 50d fut.	0.16	0.14
Fama-French MKT	0.12	0.13
Fama-French SMB	0.13	0.13
Fama-French HML	0.12	0.13
Fama-French UMD	0.12	0.13

Table 8: Error bars on the SR and skewness of the strategies appearing in Fig. 7. The errors on the skewness are estimated using a standard bootstrap method.

Fig. 7 not only efficiently summarizes all our results, but also suggests both a classification and a grading of different investment strategies. It is tempting to *define* a risk premium strategy as one that compensates for the skewness of the returns, in the sense that its Sharpe ratio lies close to the “skew-rewarding” line $\mathcal{S} \approx 1/3 - \zeta^*/4$. Interestingly, the aggregate returns of hedge fund strategies fall on this line, suggesting that a substantial fraction of hedge fund strategies are indeed risk premia.

Strategies that lie significantly below this line take too much tail risk for the amount of excess return. Contrarily, strategies that lie significantly above this line, in particular those with *positive* skewness such as trend following, seem to get the best of both worlds. But by the same token, these strategies cannot meaningfully be classified as risk premia;¹² rather, as argued in [23], these excess returns must represent genuine market anomalies, or “pure α ”.

7.2 Two interesting exceptions: HML & Trend following

Let us dwell a little longer on our two outliers: trend following and HML. The daily skewness of these portfolios is significantly positive, respectively $\zeta^* = +0.43$ and $\zeta^* = +0.27$, and yet the performance is also highly significant (at the 2σ level for HML, and at almost 10σ for the trend).

The concept of risk premium for HML is indeed problematic. Intuitively, high book to price corresponds to “value” stocks that are usually considered safe and defensive, while low book to price corresponds to “growth” stocks, i.e. a risky bet on future earnings. This is confirmed by the results of Table 3 above, which shows that growth stocks are indeed more negatively skewed than value stocks, in line with our definition of risk, but making HML an outlier in Fig 7. However, to add to the confusion, the skewness of the monthly returns of the HML strategy, although more noisy, appears to change sign ($\zeta_{month}^* = -0.56$). This suggests that more work is necessary to fully understand the nature of the HML factor, which is actually known to have long periods of drawdowns (like during the Internet bubble) where the effect appears to be more in line with a standard risk premium interpretation.

Turning now to the nature of trend following, it could in fact have been anticipated that its skewness should be positive. This is because trend following is profitable, by definition, when the long term realized volatility is larger than the short term volatility [24]. In other words, trend following is akin to a “long gamma” strategy, and is thus expected to have an oppositely signed skewness to that of options. Correspondingly, the skewness of monthly returns is found to be even more positive than that of daily returns ($\zeta_{month}^* = +1.72$). Hence, the highly significant excess returns of trend following strategies in the last two centuries (see [23] for a recent discussion) seem difficult to ascribe to any reasonably defined risk premium. The profitability of trend following appears to be a genuine market anomaly, plausibly of behavioral origin.¹³ We find our plot Fig. 7 interesting in the sense that trend following and (albeit to a lesser

¹²Barring, of course, the existence of an unidentified risk factor that has never materialized, but that would contribute to a negative skewness if it did.

¹³However, the attentive reader might wonder why UMD has a negative skewness, since it should also be a long gamma strategy. Although we do not have a complete answer to this question, we believe that the crowdedness argument given in the next section might explain why outperforming stocks tend to have a negative skewness.

extent) HML can be clearly identified as outliers.

8 Discussion and Conclusions

Our central result, conveyed by Fig. 7 above, is that skewness is indeed the main determinant of risk premia, with an approximately linear relation between the Sharpe ratio \mathcal{S} of a risk premium strategy and its skewness ζ^* .

From a formal point of view, this finding can be qualitatively interpreted within the classical framework of utility theory – see e.g. [9]. Provided the third derivative of the utility function is positive, skewness-comparable P&Ls can be ordered, and negatively skewed strategies must be compensated by higher returns to remain attractive – as amply confirmed by lottery experiments [15, 16]. Although interesting, we feel that this argument falls short of giving a fully convincing picture of the mechanism enforcing the trade-off between skewness and excess returns.

Our view is that there may in fact be *two fundamentally different* ways that risk premia evolve with respect to the “skew-rewarding” line in Fig. 7:

- In many situations, one buys an asset because of an expected stream of payments, like dividends for stocks, coupons for bonds, interest rate differences for currency pairs, etc. This intrinsic source of returns attracts a crowd of investors that generate a price increase but simultaneously create the risk of a crash, induced by a self-fulfilling panic or “bank run” mechanism, due to the crowdedness of the trade. Schematically:

dividends \rightarrow crowdedness \rightarrow price increase \rightarrow downside tail risk.

As illustrated in Fig. 8, crowdedness decreases returns and increases downside tail risk. Both effects limit the number of additional investors and stabilize the market around an acceptable skewness/excess return trade-off.

- In other situations, downside tail risks pre-exist and it is their very existence that leads to excess returns. A perfect example is provided by option markets. Since options are insurance contracts, their payoff profile is *by construction* skewed - negatively for option sellers and positively for option buyers. In an efficient market, the fair price of options is such that their average payoff is zero, no risk premium exists and the Sharpe ratio of being long or short options is zero. However, the presence of fat tails and a natural investor preference for positive skewness leads to a rise in the market price of options as investors seek insurance against sudden adverse price movements. The negative skewness from the short option position, arising from the rare large payouts remains constant and ever present. The Sharpe ratio of the short option position in an inefficient market rises to reflect the extra premium required to act as insurer. We would therefore summarise this scenario as:

downside tail risk \rightarrow dearth of option sellers \rightarrow increased premium and Sharpe,

until the risk premium collected by selling options becomes large enough to lure in enough supply to cap the price of these options.




Figure 8: The “equilibrium” skew-rewarding line $\mathcal{S} \approx 1/3 - \zeta^*/4$ and different risk premium paths. A first path (I) starts with a symmetric asset that pays dividends, and therefore attracts investors. The trade becomes crowded and this creates potential de-leveraging spirals [25] and downside tail risk. A second path (II) corresponds to an initially skewed payoff (like options) that moves up in Sharpe ratio as a result of skewness aversion and a dearth of option sellers. The path indicated with a question mark (?) might describe the fate of pure alpha strategies as they become more crowded and prone to deleveraging spirals of the type discussed in [25].

A crucial assumption for the first mechanism above to work is that negative skewness can be induced by crowdedness and de-leveraging spirals, as was convincingly argued in [25]. While it seems to us highly plausible that this is the case, it would be interesting to have some direct empirical evidence. Interestingly, we have checked that the 50-day trend following strategy, that has attracted a large amount of assets in the last 20 years, is characterized by a positive skewness, as shown in Fig. 7, that has nevertheless been slowly drifting towards zero in recent years. Similarly, the skewness of the carry trade has significantly increased (i.e ζ^* has become more negative) in the period 1994-2014 as compared to the period 1974-1994, perhaps indicating that the carry trade is more popular than it used to be.

From a practical point of view, our results provide an objective definition of risk premia strategies and a criterion to assess their quality based on the comparison between Sharpe ratio and skewness. Clearly, not all excess returns can be classified as risk premia – trend following being one glaring counter-example. On the other hand, the fact that all well known risk premia strategies seem to lie around a unique line is striking. One tentative interpretation, following e.g. [8, 10, 12] is that all these risk premia in fact represent the same “catastrophic”

risk, that would – if realized – spread out over many different asset classes. Conversely, if different risk premia are somewhat decorrelated, then skewness can be diversified away, allowing risk premia portfolios to move *above* the regression line shown in Fig. 7. That this may be possible is illustrated by the orange “star” shown in Fig. 7, corresponding to a synthetic Diversified Risk Premia portfolio, with an equal weight on Long stock indices, Short Vol, FX Carry and CDS indices. We believe that this is what good “alternative beta” managers should strive to achieve.

Acknowledgements

This work is the result of many years of research at CFM. Many colleagues must be thanked for their insights, in particular: P. Aliferis, A. Beveratos, L. Dao, B. Durin, Z. Eisler, A. Egloff, Ph. Jordan, L. Laloux, A. Landier, P. A. Reigneron, G. Simon, D. Thesmar and S. Vial.

Appendix: Ranked Amplitude P&Ls and an alternative definition of Skewness

Let us consider a random variable r (the returns) with a certain probability density $P(r)$. We assume that r has been standardized (i.e. r has zero mean and unit variance). We will denote as $x = |r| \geq 0$ the amplitude of r . From the definition of the ranked P&L function $F_0(p)$, one has:

$$F_0(p) = \int_0^{x(p)} dy y [P(y) - P(-y)], \quad (3)$$

where $x(p)$ is the p -quantile of $|r|$, defined as $p = \int_0^{x(p)} dr (P(r) + P(-r))$. We will introduce the symmetric and antisymmetric contributions to P as:

$$P_s(r) = P(r) + P(-r), \quad P_a(r) = P(r) - P(-r). \quad (4)$$

Note that for generic distributions, $P(r) \approx_{r \rightarrow 0} P(0) - P'(0)r + \dots$, leading to lowest order (when $p \rightarrow 0$) to $x(p) \approx p/2P(0)$ and therefore

$$F_0(p) \sim -P'(0)/12 \times (p/P(0))^3, \quad (5)$$

i.e. a generic $\propto p^3$ behaviour for small p .

Now, our definition of skewness is:

$$\zeta^* := -100 \int_0^1 dp F_0(p) = -100 \int_0^\infty dx P_s(x) \int_0^x dy y P_a(y). \quad (6)$$

It is interesting to give an alternative, intuitive interpretation of this definition. After simple manipulations, one finds:

$$F_0(p) = E[|r||r < -x(p)] - E[|r||r > -x(p)]. \quad (7)$$

$F_0(p)$ therefore compares the average amplitude of large negative and large positive returns. ζ^* is an average of this difference over all possible quantile

choices. It might also be useful to relate ζ^* to the standard definition of skewness (defined through the third cumulant of a distribution) in the limit of weakly non-Gaussian distributions:

$$P(r) = \left[1 - \frac{\zeta}{3!} \frac{d^3}{dr^3} + \frac{\kappa}{4!} \frac{d^4}{dr^4} + \dots \right] \frac{1}{\sqrt{2\pi}} e^{-r^2/2}, \quad (8)$$

finally leading to

$$\zeta^* = \frac{25}{6\pi} \zeta \left(1 - \frac{\kappa}{24} + \dots \right) \approx 1.273 \zeta \left(1 - \frac{\kappa}{24} + \dots \right) \quad (9)$$

Note that ζ^* does not require the existence of the third moment of the distribution, and remains well defined for all distributions with a finite first moment. (In fact, an even weaker condition is sufficient).

Finally, let us assume that $F_0(p)$ is a humped function of p with a single maximum, corresponding to a negatively skewed distribution. This means that necessarily $P_a(y > 0)$ is negative for large enough $y > y^*$ and positive for smaller y s, and vice-versa for $y < 0$. Now the cumulative function of P compared to its symmetrized version is precisely the cumulative of $P_a(y)$:

$$G(y) = \int_{-\infty}^y dr [P(r) - P_s(r)] = \int_{-\infty}^y dr P_a(r) \quad (10)$$

Since $P_a(r)$ vanishes three times and is positive for large negative r , it is clear that $G(y)$ has two symmetric maxima located at $\pm y^*$ and a minimum for $y = 0$. Now using $0 = F_0(p = 1) < y^* \int_0^\infty dr P_a(r)$ and $G(0) = - \int_0^\infty dr P_a(r)$, one immediately finds that $G(0) < 0$ for non degenerate distributions. This proves that $G(y)$ crosses zero twice and only twice, and therefore that $P(r)$ and $P_s(r)$ are skewness-comparable in the sense of Oja [20] whenever $F_0(p)$ has a unique maximum (or minimum).

References

- [1] W. F. Sharpe, *Capital asset prices: A theory of market equilibrium under conditions of risk*, Journal of Finance, 19, 425-442 (1964)
- [2] In the context of option pricing, see e.g. J.-P. Fouque, G. Papanicolaou, R. Sircar, K. Sølna, *Multiscale Stochastic Volatility for Equity, Interest Rate, and Credit Derivatives*, Cambridge University Press, 2011.
- [3] R. Mehra, E. C. Prescott, *The Equity Premium in Retrospect*, in G. Constantinides, M. Harris, and R. M. Stulz (Eds.), Handbook of the Economics of Finance (Amsterdam: Elsevier, 2003).
- [4] A. Damodaran, *Equity Risk Premiums (ERP): Determinants, Estimation and Implications*, working paper, 2012
- [5] R. Mehra, E. C. Prescott, *The Equity Premium: A Puzzle*, Journal of Monetary Economics 15 (1985), 145-161.
- [6] M. Baker, B. Bradley, J. Wurgler, *Benchmarks as limits to arbitrage: understanding the low volatility anomaly*, Financial Analysts Journal 67, 40-54 (2011).

- [7] A. Frazzini, L. H. Pedersen, *Betting against beta*, Journal of Financial Economics 111, 1-25 (2014)
- [8] T. A. Rietz, *The Equity Premium Puzzle: A Solution?*, Journal of Monetary Economics 22, 117-131 (1988).
- [9] C. Harvey, A. Siddique, *Conditional skewness in asset pricing tests*, Journal of Finance 55, 1263-1295 (2000)
- [10] R. J. Barro, *Rare Disasters and Asset Markets in the Twentieth Century*, Quarterly Journal of Economics 121, 823-866 (2006).
- [11] T. Bollerslev, V. Todorov, *Tails, fears, and risk premia*. The Journal of Finance 66, 2165-2211 (2011)
- [12] P. Santa-Clara, S. Yan, *Crashes, volatility, and the equity premium: Lessons from S&P 500 options*, The Review of Economics and Statistics 92, 435-451 (2010).
- [13] M. R.C. van Oordt, C. Zhou, *Systematic tail risk*, working paper (2013)
- [14] R. Kozhan, A. Neuberger, P. Schneider, *The Skew Risk Premium in the Equity Index Market*, working paper (2013)
- [15] T. Garrett, S. Russell, *Gamblers Favor Skewness not Risk: Further Evidence from United States Lottery Games*, Economics Letters, 63, 85-90 (1999).
- [16] T. Astebro, J. Mata, L. Santos-Pinto, *Preference for skew in gambling, lotteries and entrepreneurship*, working paper (2008)
- [17] C. R. Harvey, J. C. Liechty, M. W. Liechty, P. Muller, *Portfolio selection with higher moments*, Quantitative Finance 10, 469-485 (2010).
- [18] G. Athayde, R. Flores, *Finding a Maximum Skewness Portfolio: A General Solution to Three-Moments Portfolio Choice*, Journal of Economic Dynamics & Control, 28, 1335-1352 (2004).
- [19] W. H. Chiu, *Skewness preference, risk taking and expected utility maximization*, The Geneva Risk and Insurance Review 2010, 1-22.
- [20] H. Oja, *On Location, Scale, Skewness and Kurtosis of Univariate Distributions*, Scandinavian Journal of Statistics, 8, 154-168 (1981)
- [21] J. W. Jurek, *Crash-Neutral Currency Carry Trades*, [ssrn-id1262934](https://ssrn.com/abstract=1262934), August 2013.
- [22] J.-Ph. Bouchaud, M. Potters, *Theory of Financial Risks and Derivative Pricing*, 2nd Edition, Cambridge University Press (2004), see in particular Ch. 14.
- [23] Y. Lempérière, C. Deremble, P. Seager, M. Potters and J.-Ph. Bouchaud *Two centuries of trend following* Journal of Investment Strategies, 3, 41-61 (2014), and refs. therein.

- [24] W. Fung, D. Hsieh, *Survivorship Bias and Investment Style in the Returns of CTAs*, Journal of Portfolio Management, 23, 30-41 (1997).
- [25] M. K. Brunnermeier, L. H. Pedersen, *Market Liquidity and Funding Liquidity* Rev. Financ. Stud. 22, 2201-2238 (2009).