The Cumulant Risk Premium*

Albert S. Kyle[a], Karamfil Todorov[b]

[a] *University of Maryland, Robert H. Smith School of Business, askyle4@umd.edu*
[b] *Bank for International Settlements, karamfil.todorov@bis.org*

Abstract

We develop a novel methodology to measure the risk premium of higher-order cumulants (closely related to the moments of a distribution) based on assets satisfying a single-factor setting. We show that single-factor linear pricing works only if the difference between physical and risk-neutral cumulants, which we call the cumulant risk premium (CRP), is zero. To illustrate our approach empirically, we study leveraged ETFs, which are assets with constant betas and exposure to a single factor only. We show that the CRP is different from zero across all assets studied: equities, bonds, commodities, currencies, and volatility. We quantify the even-order CRP by developing a simple strategy of shorting ETFs with opposite betas. The strategy mimics liquidity provision, earns Sharpe ratios above one, and can be used as a simple gauge of global market stress in real time. Our results have implications not only for factor models but also for portfolio theory, momentum strategies, option pricing, hedge funds, and leverage in general.

Keywords: Cumulants, leverage, ETF, CAPM, factor models, VIX

*The views expressed herein are those of the authors and do not necessarily reflect the views of the Bank for International Settlements.

1. Introduction

Many episodes of market turbulence, including the March 2020 COVID-19 crisis, show that asset returns are not normally distributed and that higher-order moments play an important role in financial markets. These facts raise two crucial questions: (1) What are the implications of higher-order moments for standard finance models? (2) How can the risk premium of higher-order moments across asset classes be measured in a tractable way? In this paper, we study these questions by revisiting single-factor models, a basic finance setting in which asset returns are a linear function of one factor. We develop a novel methodology to measure the risk premium of higher-order moments based on assets satisfying the single-factor setting, and quantify this risk premium empirically across asset classes. Our results have implications not only for factor models but also for portfolio theory, momentum strategies, option pricing, and leverage in general.

To fix ideas, assume a world where asset returns load on one factor, which is perfectly observable (addressing Roll (1997) critique). In addition, asset betas are known and constant over time (addressing Hansen and Richard (1987) critique). What are the returns of such constant-beta assets in this idealized single-factor setting, especially in a world with high leverage and higher-order moments? We show that the standard single-factor logic, which says that an asset's risk premium is linear in the risk premium of the factor, applies only if the factor returns are lognormal as in a Black and Scholes (1973) world. Single-factor linear pricing generally fails in any other setting with non-zero higher-order moments. In addition, we show that leverage exposes assets to higher-order moments and mimics momentum trading in the direction of factor movements.

To understand the failure of single-factor models in a setting with higher-order moment risk and the exposure of assets to this risk, consider the definition of beta as leverage with respect to the single factor. There are two methods to create a leveraged asset. The first one, a "static strategy", is to invest a fraction β of the portfolio in the factor with price P_t and a fraction $1-\beta$ in the safe asset and then do nothing as P_t changes. An example of such strategy is futures trading, which involves a leveraged exposure to a given asset by posting margin that is smaller than the value of the asset. The drawback of the static strategy is that it becomes more risky when the factor moves against the investor since leverage rises. The static strategy also exposes investors to bankruptcy risk: if P_t drops by more than $1/\beta$ (for $\beta > 0$), the strategy is bankrupt. The second method, "dynamic rebalancing", is to trade dynamically in order to keep the leverage constant at β. This method reduces bankruptcy risk since it maintains a constant leverage irrespective of factor changes. Thus, leverage does not increase when P_t moves against the investor. An example of dynamic rebalancing in practice is the trading by leveraged ETFs, which we explain below.

The reduction in bankruptcy risk comes at a cost since dynamic rebalancing requires trading continuously. This exposes the strategy to higher-order moments of the factor. For example, we show that assets with $\beta > 1$ or $\beta < 0$ conduct "destabilising", "momentum" trades which demand liquidity. Such strategies mechanically buy after price increases and sell after price decreases.

We show that higher (absolute) β exposes these assets disproportionately more to higher-order moments. To the contrary, assets with $0 < \beta < 1$ conduct "stabilizing", "rebalancing" trades which provide liquidity and are less exposed to higher-order moments. Standard single-factor models (e.g., the CAPM) fail to account for these exposures and assume that dynamic rebalancing is cost-free and that β is unaffected and constant over any period of time.

To quantify the exposure to higher-order moments and illustrate the failure of single-factor pricing, we use the concept of cumulants, which allows us to deal with non-normality in a tractable way. Cornish and Fisher (1938) describe cumulants in a general setting. Martin (2013) is one of the first researchers to apply cumulants in finance. Cumulants are simple and convenient way to summarise the main characteristics of a given distribution function. They are similar to the moments of a distribution, but are more intuitive to work with compared to non-central moments. Cumulants are also more convenient to use in the case of log-returns that appear over multiple periods in our setting and to model linear combinations of random variables. The first cumulant is the mean of the distribution; the second is variance σ^2; the third and fourth cumulants are skewness times σ^3 and excess kurtosis times σ^4. Higher-order cumulants are more complicated polynomial functions of the moments. For a lognormal distribution, there are only two cumulants, mean and variance, whereas for any other distribution, there are also higher-order cumulants beyond variance.

We show that in a non-lognormal setting, the risk premium on an asset is the sum of a linear term in asset's β and a non-linear one, which depends on higher-order powers of β, weighted by differences in physical and risk-neutral cumulants of the factor. We call the sum of these differences the **cumulant risk premium (CRP)**. Intuitively, since assets are leveraged loadings on the factor and need to rebalance when the factor moves, this exposes them to higher order cumulant risk. For example, we show that momentum strategies are positively exposed to implied variance (vega) and negatively exposed to realized variance (gamma). In a lognormal world, the net exposure is zero since variance is the same in the physical and risk-neutral worlds. As a result, the non-linear term in an asset's risk premium is zero, and simple linear beta pricing works for lognormal returns. In a non-lognormal setting however, assets are exposed not only to the variance premium (difference between physical and risk-neutral variance) but also to higher-order cumulant differences. This exposure makes the non-linear term different from zero and leads to the failure of simple linear beta pricing. In addition, we show that for larger β, the non-linear term dominates the linear one since the weight of higher-order terms generally increases polynomially in β. Thus, highly-levered strategies are extremely exposed to high-order cumulant differences. In contrast, even if higher-order cumulants are non-zero, strategies with small leverage ($0 < \beta < 1$) are less exposed to them since their loadings converge to zero much faster. This makes these strategies much less risky, even in a non-lognormal setting.

Our results explain why standard single-factor models like the Merton (1992) version of the

CAPM work well over any period of time and are not affected by higher-order moments. The crucial assumption of lognormality significantly simplifies such models and eliminates the need to deal with the complex effects arising from higher-order moments. However, empirical evidence shows that higher-order moments matter in practice and that financial assets have returns far different from lognormality, which makes the non-linear term in asset's risk premium especially important for highly-levered strategies. The many blow-ups of levered strategies in practice (e.g., the portfolio-insurance-induced stock market crash of 1987, the LTCM collapse, the financial crisis in 2008, and the Covid crash of March 2020) are a clear evidence of these effects.

We illustrate the failure of linear beta pricing and the effects of higher-order cumulants by relaxing the two crucial assumptions of lognormal models: constant volatility and no jumps. We then analyse two non-lognormal settings: stochastic volatility (the Heston (1993) model), and Poisson jumps. Intuitively, if state variables have their own risk premiums like variance in the Heston (1993) model, the risk premium on the asset is non-linear in the premium of the factor. An additional source of non-linearity which arises with jumps is due to the inability of investors to hedge returns on leveraged assets in a continuous way. This creates unhedgeable bankruptcy risks. When the factor has jumps, dynamic rebalancing cannot generally hedge jump payoffs because investors are unable to rebalance during the jump. Since the jump return can be worse than -100% for some large $|\beta|$, this creates bankruptcy risk which is reflected in the CRP.

In the last part of the theoretical section in the paper, we illustrate that the exposure of constant-β assets to higher-order cumulants can be used to construct a payoff that measures the risk of higher-order even cumulants (variance, scaled kurtosis, etc.). A simple "short-both" strategy of short-selling equal amounts of two constant-β assets with opposite β-s (e.g., -1 and 1) measures the **even-order CRP (CRPE)** (weighted by $\beta^n,\ n \geq 2$) since the exposure to odd-order cumulants is cancelled out. The strategy approximates liquidity provision since it takes the perspective of a market maker trading against assets with constant β-s. We show that this type of liquidity provision is then exposed to higher-order even cumulants. The short-both strategy can also be used as a bet on the difference between *all* even-order physical and risk-neutral cumulants, similar to variance-risk premium (VRP) trades which are bets on the second-order cumulant only.

Next, we turn to the empirical analysis. In practice, it is not trivial to find a setting that satisfies the main assumptions of the idealized single-factor environment and allows us to test our theoretical results. Usually, single-factor models like the CAPM are tested using stock returns. There are, however, several critiques of such an approach. First, it is unlikely that a firm's beta is constant over time and therefore, single-factor models would fail unconditionally (Jagannathan and Wang (1996)). Second, it is nontrivial to define all factors to which a particular stock's return is exposed. Third, even if the stock is exposed to one factor only (the market), this factor itself is non-observable and hence the single-factor model cannot be tested properly (Roll (1997)).

Instead of pursuing the usual factor-model approach, we take a different route. To test our analysis empirically, we make use of a setting which overcomes the three critiques outlined above: we use assets that have constant β over time and by construction are exposed to only one factor, which is perfectly observable. These assets are leveraged ETFs: securities that provide multiples of the daily return on their benchmark index. For example, a double-leveraged ETF ($\beta = 2$) should return 10% if the benchmark index goes up by 5%. Leveraged ETFs mechanically have constant β equal to their leverage and track the returns of one factor only (their benchmark) by doing dynamic rebalancing. An additional benefit of our empirical setting is that we go beyond the traditional equities-based factor-model analysis and apply our methodology across several asset classes: equities, bonds, commodities, currencies, and volatility (VIX).

Our empirical results show that constant-β assets are exposed to higher-order cumulants of the factor across all asset classes studied. The CRP is different from zero and linear beta pricing fails across all assets. The average CRP is -7.4% annualized across assets and is significant share (104%) of the factor risk premium (FRP) in each asset. In addition, we show non-parametrically that higher-order cumulant risk premiums beyond that on variance are needed to explain the empirical patterns in most assets.

For the short-both strategy, we find that the CRPE is negative across the majority of assets, which shows that liquidity providers earn positive expected returns. The average CRPE is -4.4% annualized, and spikes beyond 20% in many asset classes in times of market stress like the COVID-19 period. The premium is also significant relative to the FRP across assets (in absolute terms): e.g., it is 139% of the FRP for oil, 46% for the S&P 500 index, and 51% for long-term Treasuries. The short-both strategy returns are highly positively correlated with VIX. The Sharpe ratios of the strategy are above one in many asset classes, consistent with the idea that liquidity providers are compensated for their risk exposure to even-order cumulants.

We show that the first principal component (PC) of the short-both strategy returns across assets can be used as a simple index of global market stress. There are several advantages of this metric relative to other commonly-used measures of market turbulence like VIX or various spreads like the TED spread (the spread between 3-month LIBOR in USD and the interest rate of Treasury bills). First, in contrast to VIX and other single-asset-based indexes, our measure is based on several asset classes and takes the perspective of a liquidity provider who is exposed to higher-order cumulants globally. Our measure increases when even-order risk-neutral cumulants rise above physical ones in all markets studied. We show that our metric drives out VIX in explaining returns of non-equity assets and is particularly important in assets with non-linear payoffs like options and CDS. Second, our index is simple to calculate also in real-time from observed prices of leveraged ETFs. It does not involve more complex and less liquid option portfolios like VIX or option-based skewness and kurtosis indexes. Third, we do not make any assumptions about the driving distribution of asset returns and "let the data speak". Our index can be applied as measure of global market stress in further research. The short-both strategy returns can also be used as an indicator of market stress in a particular asset class in real time.

Several factors can explain the existence of the CRP. Market frictions that prevent investors from trading continuously and lead to jumps are likely to create a non-zero premium on higher-order cumulants. Risk limits, de-leveraging (e.g., Adrian and Shin (2010)), or crowded trades could create price spirals at times of large price movements and cause extreme values of the factor's return distribution, increasing the premium for exposure to higher-order cumulants. Limits to arbitrage (e.g., Shleifer and Vishny (1997)) can also explain the existence of the CRP in convergence trades (which are similar to trades involving assets with opposite β) since convergence traders' wealth effect can amplify price changes and prove destabilizing (Kyle and Xiong (2001) and Xiong (2001)). Cumulants can also arise due to trading patterns of momentum traders. For example, the daily rebalancing of leveraged ETFs to keep constant β can amplify price movements and increase cumulants if this rebalancing is large part of the market (Todorov (2019), Tuzun (2013)). Finally, the "natural" distribution of the factor's return could be one with non-zero higher-order cumulants: for example, it is reasonable to assume that volatility (VIX) has a positively-skewed and highly non-normal distribution with jumps.

Implications. Our main results have implications for factor models and portfolio theory. We show that multi-factor models could fit asset returns better than single-factor models purely because the additional factors capture the contribution of higher-order cumulants of the single factor. The fact that some standard factors like momentum are positively correlated with even-order cumulant differences, is consistent with this logic. This result has implications for a vast financial literature studying factor models to explain asset returns. Instead of adding more linear factors, our theory suggests that researchers need also to account for the higher-order cumulants of the single-factor (e.g., the market portfolio). In addition, a proper test of single-factor models should first compare the difference between cumulant-generating functions in the physical and risk-neutral distributions before testing linear beta pricing.

Our approach could also help explain the flatness of the securities market line (SML). Our results show that high-beta assets, which conduct momentum trades, generally have negative CRP, which makes their returns lower than what is predicted by the CAPM. In contrast, assets with small betas should have higher returns if their CRP is positive, which would make the overall SML flatter.

The findings in this paper have implications also for momentum strategies and leverage in general. Our findings show that trend-chasing "momentum" strategies are exposed to higher-order cumulants, which could explain why the returns on these strategies have sudden crashes and exhibit higher-order moments.

Our results have also implications about the risk of higher-order cumulants. A common misperception is that this risk declines as the number of higher-order terms increases, and thus higher-order moments beyond kurtosis are rarely researched in finance. This misperception is driven by the discounting of higher-order cumulant differences with $n!$, which makes the contribution of higher-order terms extremely small for large n. Our theory emphasizes that this argument is true for unleveraged strategies but is significantly flawed for leveraged strategies, for which the contribution of higher-order cumulants generally *increases* up to the β-th order cumulant. For example, the loadings of a $\beta = 10$ strategy peak at the 10-th order cumulant difference with a loading above 2700 on that difference. Thus, more leveraged strategies are more exposed to higher-order cumulants, and even tiny changes in these cumulants are magnified.

These results have implications for agents like hedge funds who use leverage to exploit mispricings between similar assets. These agents often use strategies that involve assets with opposite sensitivities to a given factor: for example, convergence trades or relative value strategies (e.g., spot-futures basis, see Aramonte et al. (2021)). Our results show that such trades are risky because they are exposed to the CRPE, even in the case of no limits to arbitrage or noise trader risk (e.g., Shleifer and Vishny (1997)).

Our findings have implications also for policy makers and practitioners. The first principal component of the short-both strategy can be a useful gauge for policy intervention since the indicator increases when the CRPE rises globally, which could be a proxy for times when capital constraints are binding. Our empirical results illustrate also that models without compensation for higher-order cumulant risk cannot explain the return dynamics across all asset classes studied. Our paper also draws attention to a relatively underexplored empirical setting for studying factor models and higher-order cumulant risk: leveraged ETFs. We show that these assets can be used to construct bets on higher-order cumulants in a cost-efficient way.

Related literature

Our paper contributes to the literature on CAPM and factor models. Many studies show that the CAPM fail in practice (e.g., Stambaugh (1982), Fama and French (1992), Lakonishok et al. (1994), Roll and Ross (1994), Fama and French (1995), Ang et al. (1997)). Roll (1997) argues that the CAPM can never be tested properly since the market portfolio is hard to estimate. Hansen and Richard (1987) and Jagannathan and Wang (1996) argue that the CAPM would not hold unconditionally if betas are time-varying. Several studies have extended the standard single-factor CAPM equation to accommodate preference for skewness (Kraus and Litzenberger (1976), Harvey and Siddique (2000)) and co-skewness (Schneider and Zechner (2020)) in the pricing kernel. While CAPM is perhaps the most prominent example of a one-factor model used in equities, single-factor models are also found to explain the returns in other asset classes, such as currencies (Lustig et al. (2011)). The existing studies on factor models, and the CAPM in particular, focus on modifying the pricing equation with more factors or question the exact empirical implementation of the CAPM. There is, however, a lack of research on the return properties of assets that satisfy the idealized single-factor setting in the sense that they have constant betas and are exposed to one factor only. This question is particularly interesting in a world with non-zero higher-order cumulants. In this paper, we fill this gap in the existing literature.

The research presented here is related to the macro-finance literature on rare disasters. Rietz (1988), Barro (2006), Barro et al. (2013) and Longstaff and Piazzesi (2004) show that large declines in aggregate consumption growth (macroeconomic disasters) can help explain the equity risk premium. From a macro-finance prospective, our results show the implications for asset returns in a setting where consumption growth is non-normally distributed. We contribute to this literature by showing in a model-free way that processes with non-zero higher-order cumulants like those with jumps, are needed to explain the empirical findings across most asset classes studied.

Our study is also related to the literature on asset pricing with stochastic volatility and higher-order moments. Martin (2013) applies cumulants to extend the Epstein-Zin lognormal consumption-based asset-pricing model and allow for general independent and identically distributed (i.i.d.) consumption growth. Backus et al. (2011) use cumulants to show that options imply smaller probabilities of extreme outcomes than the estimates from macroeconomic data. Han and Kyle (2017) develop a rational expectations equilibrium model to show that even modest differences in higher-order beliefs may have large price effects. Bakshi et al. (2003) propose a framework to recover higher-order risk-neutral moments from option prices and to connect them to physical moments. Liu et al. (2003) show that jump risk significantly changes the optimal portfolio choice. Several papers analyse models that combine jumps with stochastic volatility (Pan (2002), Duffie et al. (2000), Eraker (2004), Bakshi and Kapadia (2003)). Carr and Wu (2009) show that variance risk premium is significant in US equities, whereas Dew-Becker and Giglio (2022) find that the premium is close to zero after 2010. Bollerslev and Todorov (2011) and Bollerslev et al. (2015) show that the compensation for jump risk accounts for a large fraction of the VRP. We show that a non-zero CRP makes linear pricing inapplicable.

Our paper contributes also to the literature on leveraged ETFs. Cheng and Madhavan (2009) show that the returns on leveraged ETFs are path-dependent, whereas Todorov (2019) illustrates that leveraged ETFs transmit non-fundamental price pressure in VIX and commodities. Our paper contributes to the field by showing that leveraged ETFs can be used to study factor models and to measure the risk of higher-order cumulants.

The rest of the paper is organized as follows. Section 2 illustrates the basic concepts—including constant beta assets, cumulants, linear pricing—and introduces the CRP. Section 3 studies constant-β strategies and shows that they can be used to measure the CRP. Section 4 presents the empirical results. Section 5 studies the economic implications of our main results and presents some extensions. Section 6 concludes.

2. Constant Beta Assets, Linear Beta Pricing, and Cumulants

This section lays out the fundamental concepts used in our paper: constant beta assets, linear beta pricing, and cumulants.

2.1. *Two methods to construct leveraged assets*

There are two ways to define a leveraged asset. The first method, "static strategy", is to invest fraction β of a portfolio in the factor P_t and fraction $1-\beta$ in the safe asset, and then do nothing as P_t changes. The drawback of the static strategy is that it becomes more risky when the factor moves against the investor since leverage rises. The static strategy also exposes investors to bankruptcy risk: if P_t drops by more than $1/\beta$ (for $\beta > 0$), the strategy is bankrupt.[1] The second method, "dynamic rebalancing", is to trade dynamically in order to keep the leverage constant at β. This method reduces bankruptcy risk since it maintains a constant leverage irrespective of factor moves. We define assets implementing dynamic rebalancing as "constant-beta assets" since they aim to keep constant β.

Let R_T denote the gross return on the factor (with $\beta = 1$), and let $R_{f,T}$ denote the gross return on the risk-free asset (with $\beta = 0$). Since the constant beta assets rebalance continuously, they achieve a payoff which is an exponential function of the gross unlevered return. Let $P_T(\beta) := R_T^{\beta} R_{f,T}^{1-\beta} := \mathrm{e}^{r_{f,T}+\beta(r_T - r_{f,T})}$ denote this payoff over any period T. Let $P_0(\beta) := G_T^*(\beta)$ denote the cost of buying an asset which pays off $P_T(\beta)$. [2] The returns on constant-beta assets can be written in several equivalent ways:

$$R_T(\beta) = \mathrm{e}^{r_{\beta,T}} = \frac{P_T(\beta)}{P_0(\beta)} = \frac{\mathrm{e}^{r_{f,T}+\beta(r_T-r_{f,T})}}{G_T^*(\beta)}, \qquad \text{with} \qquad r_{\beta,T} := r_{f,T} + \beta(r_T - r_{f,T}) - \log G_T^*(\beta). \tag{1}$$

The function $G_T^*(\beta)$ defines the prices of the derivative securities as a function of their leverage β. Since $\mathrm{E}^*[R_T(\beta)/R_{f,T}] = 1$ for the risk-neutral distribution, we infer $\mathrm{E}^*[\mathrm{e}^{\beta(r_T - r_{f,T})}] = G_T^*(\beta)$. This means that the function $G_T^*(\beta)$ is the moment-generating function (MGF) for the risk-neutral distribution of the random excess log-return $r_T - r_{f,T}$. Since the moment-generating function of a random variable defines the random variable uniquely, the function $G_T^*(\beta)$ uniquely defines the risk-neutral distribution for the random log-return r_T and therefore the random return R_T. For this risk-neutral distribution to define arbitrage-free pricing, it is also necessary that it be an equivalent martingale measure. It is a martingale measure since $\mathrm{E}^*[\mathrm{e}^{-r_{f,T}} R_T(\beta)] = 1$ by construction. Intuitively, it is "equivalent" if the risk-neutral distribution implied by $G_T^*(\beta)$ agrees with the zero-probability events of the physical distribution, which has moment-generating function defined by $G_T(\beta)$.

[1] For $0 \le \beta \le 1$, the return on the portfolio never blows up as long as the risky and safe assets have limited liability ($R_t > 0$ and $R_{f,t} > 0$).

[2] We use log-returns in line with most papers in the literature (e.g., Barro (2006), Martin (2013) among many others). This allows writing the single-factor equation for an arbitrary horizon T. Note also that the standard extension of the CAPM over multiple periods is typically the Merton (1992) version of the model that uses log-returns.

2.2. Linear Beta Pricing

Define "linear beta pricing" as $\mathrm{E}[R_T(\beta)] = \mathrm{e}^{r_{f,T}+\beta\pi_T}$, where π_T is the risk premium on the factor R_T. Linear beta pricing says that assets are priced by discounting the expected return at a continuously compounded rate which is linear in the asset's risk, as measured by its beta: $P_0(\beta) = \mathrm{e}^{-r_{f,T}-\beta\pi_T}\mathrm{E}[P_T(\beta)]$. Using $P_0(\beta) = G_T^*(\beta)$, we can rewrite linear beta pricing in terms of MGFs as:

$$G_T(\beta) = \mathrm{e}^{\beta\pi_T} G_T^*(\beta). \tag{2}$$

Our emphasis on log-returns and the above condition for linear beta pricing suggests that important intuition is associated with the log of the MGF $G_T(\beta)$.

2.3. Cumulants

The cumulant-generating function (CGF) is defined as the logarithm of the MGF:

$$c(\beta) = \log G(\beta) = \log \mathrm{E}\left[\mathrm{e}^{\beta X}\right]. \tag{3}$$

Recall that the n-th order moment of a random variable X is simply $g_n = \mathrm{E}\left[X^n\right]$. Applying a Taylor expansion of the moment-generating function (MGF), $G(\beta) = \mathrm{E}[\mathrm{e}^{\beta X}]$ around zero, is a convenient way to combine all of the moments of X into a single expression:

$$G(\beta) = \mathrm{E}\left[\mathrm{e}^{\beta X}\right] = 1 + \sum_{n=1}^{\infty} \frac{g_n \beta^n}{n!}.$$

Similarly to the MGF $G(\beta)$, the CGF $c(\beta)$ can also be expanded as a power series in terms of its **cumulants**:

$$c(\beta) = \sum_{n=1}^{\infty} \frac{\kappa_n \beta^n}{n!}.$$

The n-th order cumulant κ_n is obtained by computing the n-th order derivative of the CGF $c(\beta)$ at zero: for example, we have $\kappa_1 = c'(\beta)|_{\beta=0} = \frac{\mathrm{E}[Xe^{0X}]}{\mathrm{E}[\mathrm{e}^{0X}]} = \mathrm{E}[X]$.

Cumulants are convenient to use for three reasons. First, higher-order cumulants are easier to work with compared to non-central moments. Second, since the log of the expected value appears in the condition for linear beta pricing derived in Theorem 1 below, it is more convenient to use the CGF than the MGF. Third, the CGF of the sum of independent and identically distributed (i.i.d.) random variables is the sum of the individual CGFs. This feature of the CGF makes it convenient to model combinations of random variables, e.g., Poisson jumps with a different distribution for the size of jumps.[3] To illustrate the CGF, we next derive it for some simple distributions.

[3] In fact, this feature can also be used to prove the central limit theorem with cumulants.

Example 1: CGF for the normal distribution. For a normally distributed random variable $X \sim \mathcal{N}(\mu, \sigma^2)$, it is straightforward to show that the cumulant-generating function is quadratic in β:

$$c(\beta) = \beta\mu + \tfrac{1}{2}\beta^2\sigma^2.$$

The expression shows that $\kappa_1 = c'(\beta)|_{\beta=0}$ is the mean μ, $\kappa_2 = c''(\beta)|_{\beta=0}$ is the variance σ^2, and $\kappa_n = 0$ for all $n > 2$.[4] The normal distribution is the only one with a finite number of non-zero cumulants (e.g., Marcinkiewicz (1935)). The latter fact is yet another reason why the CGF is more convenient to work with compared to the MGF, since a normal distribution with non-zero mean has generally non-zero higher-order non-central moments of all orders, which show up in the Taylor series expansion of the MGF.

Example 2: CGF for the Bernoulli distribution

Recall that a Bernoulli random variable takes the value of 1 with probability p and 0 with probability $1-p$. Then from Equation 3, we obtain:

$$c(\beta) = \log\left(pe^{\beta} + (1-p)\right) = \log\left(1 + p(e^{\beta}-1)\right).$$

Then, $\kappa_1 = p$, $\kappa_2 = p(1-p)$, and $\kappa_n = p(1-p)\frac{d\kappa_{n-1}}{dp}$ for all $n > 2$.

Example 3: CGF for the Poisson distribution

Poisson random variable has a probability distribution defined by $\text{Prob}[X=n] = \frac{e^{-\lambda}\lambda^n}{n!}$, where λ is the arrival rate (and the mean and variance of the distribution). Then:

$$c(\beta) = \log\left(e^{-\lambda}\sum_{n=0}^{\infty}\frac{\lambda^n e^{\beta n}}{n!}\right) = \log\left(e^{-\lambda}e^{\lambda e^{\beta}}\right) = \lambda(e^{\beta}-1).$$

All cumulants are equal to λ: $\kappa_n = \lambda$ for all $n \geq 1$.

2.3.1. Example 4: Compound Poisson Distribution (CPP)

Suppose $\log X_t$ follows a compound Poisson process which has jumps r_n^J arriving randomly at rate λ. Using the independence of the CGF of jumps and that of the Poisson process, it is easy (see Section A.1 in the Appendix for an example) to derive the CGF of the CPP over T periods:

$$c_T(\beta) = \lambda T(e^{c^J(\beta)} - 1),$$

[4]For the third-order cumulant, we have $\kappa_3 = g_3\sigma^3 = 0$, since $g_3 = 0$ for the normal distribution. The forth-order cumulant of the normal distribution is zero even though the corresponding central moment is not zero: $\kappa_4 = (g_4 - 3)\sigma^4 = 0$, since $g_4 = 3$ for the normal distribution. In comparison, the power series expansion of the MGF gives more complicated non-central moments $g_0 = g_1 = \mu$, $g_2 = \sigma^2 + \mu^2$, and higher-order non-central moments are more complicated functions of the central moments (see, e.g., Ouimet (2021)).

where $c^J(\beta)$ is the CGF of the individual jump process.

Cumulants provide a tractable framework to quantify the properties of a distribution. The examples in this section show how to obtain the CGF for some popular distributions but the CGF is not always easy to derive. Some processes require more intensive computations to calculate cumulants over multiple periods and might not have a CGF in a simple closed form.

2.4. Risk premium through the lens of the CGF

The CGF provides a convenient way to compute risk premiums and to define linear beta pricing for a general distribution. The risk premium on an asset with leverage β can be expressed as the difference between physical and risk neutral cumulants:

$$\text{Risk Premium} = c_T(\beta) - c_T^*(\beta) = \sum_{n=1}^{\infty} \frac{\beta^n(\kappa_{n,T} - \kappa_{n,T}^*)}{n!}. \tag{4}$$

This result is easy to obtain by taking logs in the definition of linear beta pricing in Equation 2.

Linear beta pricing in terms of CGF. Linear beta pricing is valid if and only if this risk premium, $c_T(\beta) - c_T^*(\beta) = \sum_{n=1}^{\infty} \frac{\beta^n(\kappa_{n,T}-\kappa_{n,T}^*)}{n!}$ is equal to $\beta\pi_T$ for all β. Since $\beta\pi_T = c_T(1) - c_T^*(1)$, linear beta pricing holds if and only if

$$c_T(\beta) - c_T^*(\beta) = \beta(c_T(1) - c_T^*(1)). \tag{5}$$

Let us illustrate linear beta pricing in a one-period ($T = 1$) example with lognormal returns. The risk-neutral CGF $c^*(.)$ and physical CGF $c(.)$ are given by:

$$c(\beta) = \log \mathrm{E}[\mathrm{e}^{\beta(r-r_f)}] = \beta(\pi - \tfrac{1}{2}\sigma^2) + \tfrac{1}{2}\beta^2\sigma^2, \qquad c^*(\beta) = \log \mathrm{E}^*[\mathrm{e}^{\beta(r-r_f)}] = -\tfrac{1}{2}\beta\sigma^2 + \tfrac{1}{2}\beta^2\sigma^2.$$

Then, the risk premium on the asset is $c(\beta) - c^*(\beta) = \beta\pi$ and is linear in π for all β.

The theorem below shows that linear beta pricing is valid in a general setting with any distribution and over any period if and only if the difference between the physical and risk-neutral cumulants of the leveraged asset is a linear function of leverage β. We define the **Cumulant Risk Premium (CRP$_T$)**, $\sum_{n=2}^{\infty} \frac{\kappa_{n,T}-\kappa_{n,T}^*}{n!}$, as the sum of all the terms in the risk premium except for the term which is linear in β. This definition allows us to split the factor risk premium π_T into a first-order log risk premium (LRP), and the higher-order CRP:

$$\pi_T = c_T(1) - c_T^*(1) = \underbrace{\mathrm{E}(\log R_T) - \mathrm{E}^*(\log R_T)}_{\text{Log risk premium (LRP)}} + \underbrace{\sum_{n=2}^{\infty} \frac{\kappa_{n,T} - \kappa_{n,T}^*}{n!}}_{\textbf{Cumulant risk premium (CRP)}}. \tag{6}$$

The LRP captures the first-order component of the risk premium (mean), whereas the CRP captures higher-order components. The first term of CRP_T, given by $\frac{1}{2}(\kappa_{2,T} - \kappa_{2,T}^*) = \frac{1}{2}(\mathrm{Var}[\log R_T] - \mathrm{Var}^*[\log R_T])$, is related to the familiar variance risk premium, $VRP_T = \mathrm{E}[RV_T] - \mathrm{E}^*[RV_T]$, where RV_T is realized variance of the factor. All terms with $n > 2$ capture the gap between higher-order moments of the distribution.[5]

If $CRP_T(\beta) = \sum_{n=2}^{\infty} \frac{(\beta^n - \beta)(\kappa_{n,T} - \kappa^*_{n,T})}{n!}$ is non-zero for some asset, the asset's risk premium is non-linear in the factor risk premium since it depends on higher-order CRP terms, and hence linear pricing does not hold. We formalize this statement in the theorem below.

Theorem 1. *Linear beta pricing is valid with any distribution and over any period T if and only if the CRP is identically zero for all β:*

$$CRP_T(\beta) = \sum_{n=2}^{\infty} \frac{(\beta^n - \beta)(\kappa_{n,T} - \kappa^*_{n,T})}{n!} = 0 \ \textit{for all}\ \beta. \tag{7}$$

Linear beta pricing therefore requires that all higher order cumulants in the physical and risk neutral distributions be identical: $\kappa_{n,T} = \kappa^*_{n,T}$ for all $n \geq 2$.

Proof. As we showed before, linear beta pricing can be stated in terms of CGFs as

$$c_T(\beta) - c^*_T(\beta) = \beta\left(c_T(1) - c^*_T(1)\right).$$

Using the definition of the CGF, we can rewrite the equation in terms of cumulant differences:

$$\begin{aligned} \sum_{n=1}^{\infty} \frac{\beta^n(\kappa_{n,T} - \kappa^*_{n,T})}{n!} &= \beta \sum_{n=1}^{\infty} \frac{\kappa_{n,T} - \kappa^*_{n,T}}{n!} \\ \Longleftrightarrow \sum_{n=2}^{\infty} \frac{(\beta^n - \beta)(\kappa_{n,T} - \kappa^*_{n,T})}{n!} &= 0. \end{aligned}$$

□

Equation 7 shows that we can rewrite the risk premium on the asset as a sum of a linear component and a non-linear $CRP_T(\beta)$:

$$c_T(\beta) - c^*_T(\beta) = \beta\left(c_T(1) - c^*_T(1)\right) + CRP_T(\beta). \tag{8}$$

Intuition for Theorem 1: The theorem proves that for linear beta pricing to be valid, the shape of the physical and risk-neutral distributions (as characterized by higher-order cumulants) should be the same for any leveraged asset. As we saw above, this condition is satisfied for the Black-Scholes model since with a lognormal return, there is only one higher-order cumulant (variance), which is the same under the physical and risk-neutral measures. Cumulants above the second are all zero as we showed above, and hence the $\beta^n - \beta$ terms do not make the polynomial in Theorem 1

[5] The CRP is closely related to entropy: it is equal to the difference between physical and risk-neutral entropy $c_T(1) - \kappa_{1,T} - (c^*_T(1) - \kappa^*_{1,T})$.

non-linear in β. This fact makes the assumption of lognormal distribution convenient for using linear beta pricing models like the CAPM. In practice, this assumption is empirically unrealistic; we show in section 4 that financial assets' returns exhibit occasional jumps or other deviations from lognormality, which makes the classic CAPM equation invalid.

Our theory shows that in order to test linear beta pricing models, researchers need to first verify that the shape of the physical and risk-neutral distributions is the same by computing the CGF. This result has implications for a vast financial literature that tests the CAPM, but implicitly ignores the impact of higher-order cumulants by using the classic CAPM equation, which assumes that the underlying asset's return process follows the lognormal distribution. The fact that the CAPM fails in those tests does not necessarily mean that more factors are needed to explain asset returns. Even in our setting, where asset returns load on one factor only, the CAPM still fails because of the impact of higher-order cumulant differences rather than additional linear factors.

Next, we briefly illustrate how to check the condition of Theorem 1 in several standard settings: Black–Scholes, stochastic volatility, and jumps.

2.5. *CRP in different settings*

2.5.1. *Black–Scholes*

Under the standard Black–Scholes assumptions, the market return follows a geometric Brownian motion (GBM) with physical mean μ and constant volatility σ. Then, by applying the diffusion invariance property that σ must be the same in the risk-neutral and the physical distributions for there to be no arbitrage, we obtain

$$\kappa_{2,T} = \kappa^*_{2,T} = \sigma^2 T;\ \kappa_{n,T} = \kappa^*_{n,T} = 0 \ \text{ for all } n > 2. \tag{9}$$

Hence, $CRP_T(\beta) = 0$ for all β and linear beta pricing holds as per Theorem 1. The same holds also when σ_T is a deterministic function of time (not simply constant).

2.5.2. *Stochastic volatility*

To account for non-normality of asset returns, one strand of the literature relaxes the assumption of constant volatility and models volatility as a stochastic process. We illustrate this using the standard stochastic volatility model of Heston (1993). Under the physical measure, the log-market price process $\log P_t$ and its variance v_t follow

$$\begin{aligned} \mathrm{d}\log P_t &= (\mu - \tfrac{1}{2} v_t)\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}B^1_t, \\ \mathrm{d}v_t &= \lambda(\bar{v} - v_t)\,\mathrm{d}t + \sigma\sqrt{v_t}\,\mathrm{d}B^2_t, \end{aligned}$$

where λ is the mean-reversion speed, $\bar{v}$ is the long-term mean of volatility, σ is now the volatility of volatility, and B^1_t, B^2_t are correlated Brownian motions $\mathrm{d}B^1_t\,\mathrm{d}B^2_t = \rho\,\mathrm{d}t$. Under the risk-neutral measure, the corresponding equations are

$$
\begin{aligned}
\mathrm{d}\log P_t &= (r_f - \tfrac{1}{2} v_t)\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}B_t^{1*},\\
\mathrm{d}v_t &= \lambda^*(\bar{v}^* - v_t)\,\mathrm{d}t + \sigma\sqrt{v_t}\,\mathrm{d}B_t^{2*}.
\end{aligned}
$$

In the Heston model, the CGFs are

$$
\begin{aligned}
c_T(\beta) &= \mu\beta T + a(\beta,T) + b(\beta,T)\,v_t,\\
c_T^*(\beta) &= r_f\beta T + a^*(\beta,T) + b^*(\beta,T)\,v_t,
\end{aligned}
\tag{10}
$$

where $a(\beta,T)$, $b(\beta,T)$, $a^*(\beta,T)$, and $b^*(\beta,T)$ are complicated functions of the model parameters. Closed-form derivation are in subsection A.4.

For an unlevered asset ($\beta = 1$), we obtain $a(1,T) = a^*(1,T) = b(1,T) = b^*(1,T) = 0$ and thus $c_T(1) - c_T^*(1) = (\mu - r_f)T$. In other words, even in this stochastic volatility setting, the factor risk premium π_T, given by $(\mu - r_f)T$, does not depend on variance (as in the continuous-time CAPM).

For an asset with arbitrary leverage ($\beta \neq 1$), linear beta pricing breaks down. Since generally $\lambda \neq \lambda^*$ and $\bar{v} \neq \bar{v}^*$, we have that $a(\beta,T) - a^*(\beta,T) + (b(\beta,T) - b^*(\beta,T))\,v_t \neq 0$. This implies that the risk premium $c_T(\beta) - c_T^*(\beta) = \beta(\mu - r_f)T + a(\beta,T) - a^*(\beta,T) + (b(\beta,T) - b^*(\beta,T))\,v_t$ is different from $\beta(c_T(1) - c_T^*(1))$. Fig. 1 shows that $c_T(\beta) - c_T^*(\beta)$ is not linear in β for typical values of the model parameters. Hence, the condition of Theorem 1 is not satisfied and linear beta pricing does not work since $\mathrm{CRP}_T(\beta) \neq 0$.[6]

Expanding the premium for a general leveraged asset shows that assets with $0 < \beta < 1$ load negatively on the VRP $\mathrm{E}[v_T] - \mathrm{E}^*[v_T]$, whereas those with $\beta < 0$ or $\beta > 1$ load positively on the premium. In addition, the loadings on higher-order cumulant premiums explode for assets with large $\beta < 0$ or $\beta > 1$, whereas the loadings for assets with $0 < \beta < 1$ converge to zero for larger n. The loadings are zero and hence there is no $\mathrm{CRP}_T(\beta)$ for the unlevered risky asset ($\beta = 1$) and the risk-free asset ($\beta = 0$). The premium on a leveraged asset is

$$
c_T(\beta) - c_T^*(\beta) = \beta[\mu - r_f]T + \tfrac{1}{2}\beta(\beta - 1)\left(\mathrm{E}[v_T] - \mathrm{E}^*[v_T]\right) + \sum_{n=3}^{\infty} \frac{(\beta^n - \beta)(\kappa_{n,T} - \kappa_{n,T}^*)}{n!}.\text{[7]}
\tag{11}
$$

Empirically, the VRP is typically negative (see, e.g., Carr and Wu (2009)) since investors hedge against variance risk and inflate $\mathrm{E}^*(v_T)$ relative to $\mathrm{E}[v_T]$. In the context of the Heston (1993) model, this typically means that variance has higher mean and a slower rate of mean reversion under the risk-neutral measure ($\bar{v} < \bar{v}^*$, $\lambda^* < \lambda$). The fact that the VRP is negative means that if

[6]The higher-order cumulants of the leveraged asset can be found by evaluating the n-th order derivatives $a^{(n)}(\beta,T)|_{\beta=0}$ and $b^{(n)}(\beta,T)|_{\beta=0}$, which shows that leveraged assets load on the variance risk premium and on higher-order cumulants.

[7]Note that $\kappa_{2,T} = \mathrm{E}[v_T]$ in this case since the logreturn on the factor follows a diffusion.

higher-order cumulant risk premiums are small (i.e., the last term in Equation 11 is small), assets with $0 < \beta < 1$ have larger returns than in a world with no VRP (e.g., a Black-Scholes setting), whereas those with $\beta < 0$ or $\beta > 1$ have smaller returns. We return to this point when discussing the flatness of the securities market line later in Section 5.

It is not surprising that $c_T(\beta) - c_T^*(\beta) \neq \beta(c_T(1) - c_T^*(1))$ in the Heston model because the VRP influences also higher-order cumulants. As we outlined in the introduction, κ_3 = skewness $\cdot\, v^{3/2}$, κ_4 = kurtosis$\cdot\, v^2$, and similarly for higher-order cumulants. This means that the difference between higher-order physical and risk-neutral cumulants loads on the difference between physical and risk-neutral v, or, in other words, loads on the VRP.

2.5.3. Jumps

Another way of modelling non-normal returns is to assume that the factor follows a non-smooth jump process. A typical example of such a setting with non-zero higher-order cumulants assumes the log-return on the factor is the sum of a normally distributed component and Poisson jumps with a normal distribution. Each period, J normally distributed jumps with mean $-b$ and variance s^2 arrive, where $J \sim \text{Poisson}(\lambda)$. This setting is commonly used in option pricing (e.g., Merton (1976)) and macro-finance (e.g., Martin (2013), Backus et al. (2011)), and it is a particular case of the rare disaster setup in Barro (2006). The CGF of the log-return on the asset over T periods in this case is (see Section A.2 in the Appendix)

$$c_T(\beta) = (\mu - \tfrac{1}{2}\sigma^2)\beta T + \tfrac{1}{2}\sigma^2\beta^2 T + \lambda T(\mathrm{e}^{-b\beta + \frac{1}{2}s^2\beta^2} - 1) \tag{12}$$

and

$$\begin{aligned} c_T(1) - c_T^*(1) &= (\mu - r_f)T + T\left(\lambda(\mathrm{e}^{-b+s^2/2} - 1) - \lambda^*(\mathrm{e}^{-b^*+s^2/2} - 1)\right) \\ \kappa_{1,T} - \kappa_{1,T}^* &= (\mu - r_f + \lambda^* b^* - \lambda b)T \\ \kappa_{n,T} - \kappa_{n,T}^* &= (-1)^n(\lambda b^n - \lambda^*(b^*)^n)\,T \text{ for all } n \geq 2, \end{aligned} \tag{13}$$

where λ^* and b^* are the mean arrival rate and the size of the jump under the risk-neutral measure, respectively. $CRP_T(1)$ is $T\left(\lambda(\mathrm{e}^{-b+s^2/2} - 1 + b) - \lambda^*(\mathrm{e}^{-b^*+s^2/2} - 1 + b^*)\right)$. In the general case when $\lambda^* \neq \lambda$ and $b \neq b^*$, $CRP_T(1)$ is different from zero. In other words, if agents disagree on the intensity and the size of jumps under the physical and risk-neutral worlds, the risk premium on the factor $CRP_T(1)$ is different from zero. This distinguishes the jump model from the Heston model since even the premium of the factor depends on higher-order cumulants, in contrast to the Heston model. Equation 7 is not satisfied if $\lambda^* \neq \lambda$ and $b \neq b^*$ for any β because

$$\sum_{n=2}^{\infty} \frac{(\beta - \beta^n)\left((-1)^n(\lambda b^n - \lambda^*(b^*)^n)\,T\right)}{n!} \neq 0 \tag{14}$$

and hence, linear beta pricing does not hold with jumps. Fig. 1 illustrates this point graphically by showing that the premium of a leveraged asset is non-linear in the premium of the factor.

To summarise, there are at least two reasons why linear beta pricing fails in a non-lognormal world. First, if higher-order cumulant differences are different from zero, any leveraged asset would load non-linearly on those differences as in the case of the Heston model. Simple linear model would then fail to capture these effects. In a Black-Scholes world, this non-linearity does not arise since those differences are zero for cumulants above the first. Second, in a setting with jumps, the discontinuity prevents investors from trading continuously and hedging perfectly constant-β assets. Investors then require an additional premium to bear the discontinuity risk, which is reflected in the CRP. In practice, the CRP might reflect also the inability of market makers to hedge non-linear CRP terms similar to Garleanu et al. (2009).

3. Constant-β strategies

In this section, we study constant-β strategies and show that their exposures to higher-order cumulants can be used to measure the $CRP_T(\beta)$. Securities with constant β-s are important to analyse not only because these assets are a crucial element of factor models, but also because they provide useful insights about other finance phenomena. For example, we show that assets with $\beta > 1$ or $\beta < 0$ conduct momentum, "destabilising" trades which demand liquidity, in contrast to assets with $0 < \beta < 1$, which conduct "rebalancing", "stabilizing" trades that provide liquidity.

To illustrate the latter point, note that to keep the same β in a multi-period setting, constant-β strategies require rebalancing every time the market (factor) moves. This makes these strategies different from static strategy that starts with a given leverage=β but does not rebalance to maintain the exposure constant as the factor moves. Let us illustrate the difference with a simple example. Suppose an investor with \$100 starts with a $\beta = 2$, and therefore borrows \$100 at the risk-free rate to invest \$200 in the market. Assume that the market return is -10% in the next period: in that case, the portfolio of the investor consists of \$180 in the market and -\$100 at the risk free rate for a $\beta = \frac{180}{180-100} = 2.25 > 2$. A static strategy then becomes more risky because the leverage increased: if the market keeps dropping in future periods, the leverage increases further and the investor risks being bankrupt. In contrast, a constant-β strategy maintains the same β by rebalancing as the market moves. In this example, the strategy requires the investor to sell \$20 of his market exposure in the next period and to use the cash to repay part of the debt so that β is maintained constant: $\beta = \frac{160}{80} = 2$.

This simple example illustrates the main logic of this section: constant-β strategies need to rebalance every period, which exposes them to momentum/reversal effects. The rebalancing can also amplify price movements and contribute to larger cumulants, increasing the CRP.

3.1. Constant-β assets in discrete time with two periods

Let us first study the exposures of constant-β strategies by analyzing their dynamics over two discrete periods. The simple return on the market (factor) is $(1 + r_{0\to 2}) = (1 + r_1)(1 + r_2)$. The difference between the return on the constant-β strategy and the static strategy that does not keep constant β, $1+\beta r_{0\to 2}$ is (assume zero risk-free rate for simplicity):

$$\Delta = (1+\beta r_1)(1+\beta r_2) - \Big(1+\beta\big((1+r_1)(1+r_2)-1\big)\Big) = \beta(\beta-1)r_1 r_2. \tag{15}$$

Equation 15 shows that strategies with $\beta > 1$ and $\beta < 0$ outperform the static strategy over two periods ($\Delta > 0$) in case of momentum ($r_1 r_2 > 0$) but underperform it in case of reversal ($r_2 r_2 < 0$) since $\beta(\beta-1) > 0$ for all $\beta > 1$ or $\beta < 0$. These strategies need to buy when the market goes up and sell when it goes down: they demand liquidity. If the market continues the trend from the previous period, the strategies benefit. To the contrary, strategies with $0 < \beta < 1$ underperform the static strategy over two periods in case of momentum but outperform it in case of reversal. The intuition is that these strategies buy when the market goes down and sell when the market goes up: they provide liquidity and benefit from trend reversals.

The two-periods example shows that the rebalancing of constant-β assets exposes them to momentum/reversal effects. This exposure differentiates them from static strategies, which have higher bankruptcy risk over multiple periods. For example, if the market goes down by 25% each period, a passive strategy with initial $\beta = 2$ would be bankrupt after the second period, whereas a constant-β strategy would not be.[8] With more than two periods, the momentum/reversal effect of constant-β strategies translates to exposure to higher-order cumulants.

3.2. Constant-β assets with more than two periods: continuous time

The logic from the two-periods example extends to more periods, but the algebra is more tedious in discrete time. To illustrate the effects over many periods parsimoniously, we derive the results in continuous time. Constant-β strategies invest wealth fraction β in the market P_t, and the rest $(1-\beta)$ at the constant risk-free rate r_f. For simplicity, suppose that the Black–Scholes assumptions hold and the market follows GBM. Then, the value of a constant-β strategy is (see Section A.3 in the Appendix):

$$P_T(\beta) = P_0(\beta)\left(\frac{P_T}{P_0}\right)^{\beta} \mathrm{e}^{\left((1-\beta)r_f - \frac{1}{2}\beta(\beta-1)\sigma^2\right)T}. \tag{16}$$

Equation 16 shows that strategies with $\beta > 1$ and $\beta < 0$ are negatively exposed to variance but positively exposed to squared realized returns: $\frac{\partial P_T(\beta)}{\partial \sigma^2} < 0$ and $\frac{\partial^2 P_T(\beta)}{\partial r^2_{T,\text{simple}}} > 0$, where $r_{T,\text{simple}} = \frac{P_T}{P_0} - 1$ is the simple realized return. The negative exposure to variance is reflected in the "slippage" term $-\frac{1}{2}\beta(\beta-1)\sigma^2 T$. In options terminology, the strategies are long-gamma, short-vega (long

[8] In our empirical tests, we find that a static strategy that does not rebalance daily but once in a quarter, for example, goes bankrupt for many assets: Nasdaq, financials, VIX, Russell 2000, real estate, natural gas, oil and silver. In contrast, the daily rebalancing strategy avoids bankruptcy for all assets studied.

realized variance but short implied variance). This exposure illustrates that the strategies benefit from higher physical compared to risk-neutral even-order cumulants (variance in this setting since higher-order cumulants are zero in a Black–Scholes economy) over multiple periods. If the realized path of the market return is one with higher variance than the one implied by σ^2, these strategies are profitable. This is illustrated by the difference between $r^2_{T,\text{simple}}$ and $\sigma^2 T$ even though in expectation, the difference between realized and risk-neutral variance of the log-return is zero in the Black–Scholes setting under diffusion invariance. However, if variance has non-zero risk premium as in the Heston (1993) model, assets with $\beta > 1$ and $\beta < 0$ load positively on that premium as illustrated in Equation 11.

As explained in the discrete two-period case, strategies with $\beta > 1$ and $\beta < 0$ are "momentum" strategies requiring buying when prices rise and selling when prices fall. With multiple periods of infinitely small length, the negative exposure to "back-and-forth", reversal moves translates to negative vega, whereas the positive exposure to momentum translates to positive gamma. By similar argument, rebalancing strategies with $0 < \beta < 1$ have positive vega and negative gamma. In practice, momentum strategies could incur additional transaction costs since they are destabilizing trades that demand liquidity.

Note that constant-β assets have the "right" expected return even though they appear to have a slippage: taking expectations in Equation 16 yields $\mathrm{E}[R_T(\beta)] = \mathrm{e}^{\beta\mu T + (1-\beta) r_f T}$ and there is no slippage term $-\frac{1}{2}\beta(\beta-1)\sigma^2 T$ that drags returns lower. However, in practice, these assets have a performance drag as we show in Section 4. Our theory shows that this under-performance is not due to the slippage term in Equation 16, since this term cancels out when taking the expectation in a lognormal world. Instead, the under-performance arises through cumulant differences since σ^2 is different from σ^{*2} in a setting with a VRP as in the Heston model, for example.

3.3. Constant-β assets and the CRP

Linear beta pricing in Equation 5 shows that assets with larger absolute β-s are exposed to a larger number of higher-order cumulants. For example, an asset with $\beta = 3$ has a loading above one on CRP terms up to the sixth order since $\frac{3^n}{n!} > 1$ for $n \leq 6$. Thus, even a small difference in physical and risk-neutral cumulants is magnified. Assets with $0 < \beta < 1$ are less dependent on higher-order cumulants since their loadings converge to zero much quicker.

Equation 5 also illustrates that the returns on constant-β assets are, effectively, bets on physical (realized) versus risk-neutral (implied) cumulants of the factor. Assets with $\beta > 0$ are long physical and short risk-neutral cumulants: they are long the CRP. Assets with $\beta < 0$, which are typically considered a hedge against market downturns, are long even-order physical cumulants but short even-order risk-neutral cumulants: they are long the **even-order CRP (CRPE)**. In contrast to assets with $\beta > 0$, securities with $\beta < 0$ are short the odd-order CRP (CRPO): short odd-order physical cumulants but long odd-order risk-neutral cumulants.

These exposures can be used to construct some important payoffs, e.g, one that gives the CRPE. Let us denote that payoff by:

$$\mathbf{CRPE}_T(\beta) = \sum_{\text{n} \geq 2, \text{ even}}^{\infty} \beta^n \frac{\kappa_{n,T} - \kappa^*_{n,T}}{n!}. \tag{17}$$

Studying the contribution of even cumulants by extracting the $CRPE_T(\beta)$ is worth for at least two reasons. First, the negative of this payoff proxies what market-makers earn by providing liquidity and trading against constant-β assets over multiple periods. Second, one can construct a bet on implied vs. realized even-order cumulants to harvest the CRPE. This bet is similar to the traditionally studied trade of implied vs. realized variance (e.g., Carr and Wu (2009)) to earn the VRP, but the CRPE is a bet on *all* even-order cumulants as opposed to *the second-order cumulant only* (variance). To illustrate these effects, we consider a simple trade: short-sell equal amounts of two constant-β assets with opposite $\beta \geq 1$ (e.g., -1 and 1). Such a "short-both" strategy approximates liquidity provision or trading against assets with constant β-s.

Since the two assets have exactly the opposite β-s, then selling both of them cancels the exposure to odd-order cumulants and the strategy returns are proportional to the negative of the CRPE.[9] Assume that the cash amount from the short position is invested at the risk-free rate for simplicity. The returns on the short-both strategy are then:

$$\begin{aligned} r_{\text{SB},T}(\beta) =& 2r_{f,T} - \left(r_{f,T} + \sum_{n=1}^{\infty} \frac{\beta^n(\kappa_{n,T} - \kappa^*_{n,T})}{n!} \right) - \left(r_{f,T} + \sum_{n=1}^{\infty} \frac{(-\beta)^n(\kappa_{n,T} - \kappa^*_{n,T})}{n!} \right) \\ =& -2 \sum_{\text{n} \geq 2, \text{ even}}^{\infty} \beta^n \underbrace{\frac{\kappa_{n,T} - \kappa^*_{n,T}}{n!}}_{CRPE_T} = -2CRPE_T(\beta) \end{aligned} \tag{18}$$

The strategy earns twice the negative of the $CRPE_T(\beta)$: it benefits from even-order risk-neutral cumulants, but is negatively exposed to even-order physical cumulants. This is intuitive, since the strategy mimics market-making trade providing liquidity to momentum-like assets, which are long physical and short risk-neutral cumulants. Fig. B.1 illustrates the intuition using a simple binomial tree example. If the realized path of the benchmark has low even-order physical cumulants (variance in this example), then the strategy earns positive return, which is illustrated by the green cells. The fact that the returns of a liquidity provision are long even-order risk-neutral cumulants and short physical ones echoes the result of Nagel (2012), who shows that market-making profits in US stocks are proportional to VIX. This result is consistent with our theory since VIX^2 is a measure of risk-neutral entropy (sum of higher-order risk-neutral cumulants)

[9]With simple returns, the strategy also cancels the effect of dividends in case of equity ETFs that are used in the empirical section.

of the S&P 500 index as Martin (2015) shows.

In a Black–Scholes world, the short-both strategy earns zero returns in expectation since all even cumulants of order four and above are zero, whereas the second-order physical and risk-neutral cumulants are identical: the CRPE is zero. In a general setting, the strategy is a bet that even-order risk-neutral cumulants are larger than physical ones. In a model with compensation for liquidity provision, the strategy also earns the corresponding market-making costs. Next, we apply our approach in practice and quantify the $CRP_T(\beta)$ and $CRPE_T(\beta)$ by studying assets with constant β-s.

4. Empirical evidence: leveraged ETFs

Empirically, it is not trivial to find a setting that satisfies the main assumptions of single-factor models. The most prominent such model, the CAPM, is typically tested by using stock returns. However, there are several critiques for such an approach. First, it is unlikely that a firm's beta is constant over time and therefore, the CAPM would fail unconditionally (Jagannathan and Wang (1996)). Second, it is nontrivial to define all factors to which a particular stock's return is exposed. Third, even if the stock is exposed to one factor only (the market), this factor itself is non-observable and hence the CAPM cannot be tested properly (Roll (1997)).

Instead of pursuing the usual single-factor models approach, we take a different route. To test our analysis empirically, we make use of a setting which overcomes the three critiques outlined above: we use assets that have constant β over time and by construction are exposed to only one factor, which is perfectly observable. These assets are leveraged ETFs.

Leveraged ETFs are securities that provide multiples of the daily return on their benchmark index. For example, a double-leveraged ETF ($\beta = 2$) should return 10% if the benchmark index goes up by 5%, whereas a double inverse-leveraged ETF ($\beta = -2$) should return -10%. Leveraged ETFs are a useful application of our methodology since these assets mechanically have constant β equal to their leverage, which is fixed in the prospectus of each ETF. Moreover, by construction they are exposed to one factor only: their benchmark index. Leveraged ETFs are present in variety of asset classes: not only equities but also bonds, commodities, currencies and volatility. This allows us to test our approach in more asset classes compared to the traditional equities-based analysis.

4.1. Data

We use data on leveraged ETFs that track indexes in the main asset classes: US equity (S&P 500, Nasdaq, Russell 2000, basic materials, consumer services, financials, industrials, real estate and utilities), emerging market equity, mid-term (7-10 years) and long-term (more than 20 years) US Treasuries, US high yield corporate bonds, commodities (gold, silver, oil and natural gas), currencies (euro and Japanese yen), and volatility (VIX). Prices of these ETFs and their benchmarks are from Bloomberg at a daily frequency and span the period from the first leveraged ETF introduction date in a given asset class (earliest is June 2006 for the S&P 500 Index-Table B.2) until April 2021 (or the latest available date).

4.2. The role of higher-order cumulants across assets

First, we identify potential episodes of higher-order cumulants by simply calculating the difference between simple returns and log-returns on the benchmark. This difference would give the contribution of realised higher-order moments (and as a result, realised higher-order cumulants) as can be seen from the Taylor approximation of r_{simple} around zero: $\log(1 + r_{\text{simple}}) - r_{\text{simple}} = -\sum_{n=2}^{\infty} \frac{r_{\text{simple}}^n (-1)^n}{n}$. The red lines in Fig. 2 show the result for several assets. The plots illustrate that the difference is volatile over time, and is particularly large in times of extreme price movements, e.g., during the 2008 financial crisis, the COVID-19 crisis in March 2020 and in some idiosyncratic crises like for oil in April 2020. A large part of the contribution of higher-order cumulants is due to the second cumulant (variance), but the role of cumulants of order three and above is also significant in times of market stress as illustrated by the red lines in Fig. 3. The role of higher-order cumulants for leveraged assets is particularly evident from the difference between simple returns and log-returns for ETFs with $|\beta| > 1$. The blue lines in Fig. 2 and Fig. 3 show that the contribution of higher-order cumulants is magnified for those assets since the CRP terms are multiplied by β^n.

We next study whether the Black–Scholes model or any other model with no higher-order cumulants above variance can explain the empirical findings. To do so, we first plot the difference between the return on the ETF and the return on the β-weighted factor (benchmark) less the risk-free rate: $r_T - \beta r_{\text{M},T} - (1-\beta) r_{f,T}$. This difference shows the correction to the linear-beta-pricing return at the daily frequency and is closely related to the $CRP_T(\beta)$, which we estimate below. Fig. 4 illustrates that the difference jumps in the episodes of larger cumulants seen in Fig. 2. At times when even-order physical cumulants are smaller than risk-neutral cumulants and the $CRPE_T(\beta)$ is negative, both long and inverse ETFs lose wealth.[10] Prominent examples are the 2008 crisis and the COVID-19 crisis in equities when both the red and the blue lines are below zero.

Visually, the plots in Fig. 4 show that the correction to the linear beta pricing formula in Equation 8 could be explained both by a model with no higher-order CRP terms beyond variance ("only-VRP model"), and by a model with a compensation for higher-order cumulant risk like the Heston (1993) model or a setting with jumps. We construct a simple test to see if any model without higher-order cumulants above the second can explain the empirical patterns. This test covers any

[10]With simple returns, the difference is sometimes positive for long ETFs if the $CRP_T(\beta)$ is negative (Fig. B.2). In contrast, the difference is often negative for inverse ETFs ($\beta < 0$) since they are negatively exposed to the CRPO. Part of the differences in Fig. B.2 can also be due to fund management fees, other expenses, and the fact that leveraged ETFs do not pay the multiple of the benchmark's dividend in practice. For example, the quarterly "zig-zag" pattern in equity ETFs is consistent with the effect of dividends.

time-varying volatility model that has only VRP but no higher-order CRP terms. The test is to use ETFs with opposite β-s to check a simple necessary condition that must be satisfied if an only-VRP model fits the empirical patterns. If the model is a good fit, then for two ETFs with opposite β-s (e.g., 2 and -2), the ratio of $CRP_T(\beta) = VRP_T(\beta)$ to $CRP_T(-\beta) = VRP_T(-\beta)$ should be $\frac{\beta-1}{\beta+1}$ (see Equation 7 and Equation 11). Note that this condition does not depend on the form of the VRP as the VRP cancels out.

We find that this condition is not satisfied for all assets except high yield bonds, the euro, and the Japanese yen. This observation squares well with the fact that these assets have very low realised higher-order moments beyond variance as seen from Fig. 3 (the pictures for high yield bonds and the euro are not reported for brevity). These facts illustrate that only-VRP model is unable to describe the empirical results across most assets and that higher-order CRP terms beyond the second (the VRP) have significant impact.

4.3. Quantifying $CRP_T(\beta)$

Next, we estimate the $CRP_T(\beta)$ by running regressions of ETF returns on their benchmark returns after controlling for the risk-free rate. The intercept in such a regression captures the $CRP_T(\beta)$ as seen from Equation 8 and should be zero if linear beta pricing holds. Table 1 shows that the $CRP_T(\beta)$ is different from zero across most assets and leverages.[11] The average $CRP_T(\beta)$ is -7.4% annualized across assets and β-s with significant $CRP_T(\beta)$ estimates. The size of the premium is generally larger in absolute value for assets with $\beta < 0$ and is of the order negative 10-13% annualized for many equity indices like small-cap stocks, financials and utilities. The $CRP_T(\beta)$ is the largest for oil ETFs, reaching a level of -54% annualized (significant at the 7% level). The $CRP_T(\beta)$ is significant share of the FRP in each asset: it is 104% of the FRP, on average (in absolute value among the significant estimates), and sometimes reaches levels above 200% of the FRP as shown in Table B.1. The plots of the $CRP_T(\beta)$ in Fig. 5 illustrate that the premium is significantly different from zero for most periods across equities, bonds, commodities, and volatility.

The empirical evidence shows that linear beta pricing fails in practice due to non-zero $CRP_T(\beta)$. Neither the Black and Scholes (1973) model nor any other model with only VRP but no higher-order CRP terms, can explain the patterns in most asset classes. These facts show that processes with non-zero higher-order CRP terms beyond variance like those with stochastic volatility (e.g., the Heston (1993) model) or jumps, are needed to account for the data findings across asset classes.

[11] In practice, the fact that $CRP_T(\beta) \neq 0$ means that ETFs have tracking error. Our theory explains that this tracking error is due to the risk of higher-order cumulants since ETFs are exposed to non-linearities. ETFs are incentivized to keep their tracking error low since the compensation of ETF managers and the performance evaluation of the fund are typically linked to that error. Therefore, it is unlikely that ETFs deliberately manipulate their tracking error.

4.4. Quantifying $CRPE_T(\beta)$

In Section 3.3 we showed that selling two ETFs with opposite leverages earns twice the $CRPE_T(\beta)$: the strategy benefits from even-order risk-neutral cumulants, but is negatively exposed to even-order physical cumulants. We now construct this shot-both strategy to measure the $CRPE_T(\beta)$.

4.4.1. $CRPE_T(\beta)$ across assets

Fig. 6 illustrates the performance of the short-both strategy for several assets. The figure shows that the strategy returns jump up in times of market stress, when even-order risk-neutral cumulants are larger than physical ones as illustrated by the COVID-19 shock, and some idiosyncratic shocks as for oil in April 2020. The returns on the strategy are significant and positive for each year in the sample for most equity indices, Treasuries, volatility, and commodities like oil and natural gas. Table 2 shows that the average return on the strategy is 8.9% annualized across assets and β-s, and the average $CRPE_T(\beta)$ is -4.4%. Implementing the strategy with $\beta > 1$ delivers more negative $CRPE_T(\beta)$ as shown in the table. Assets with $\beta = 3$ are a good illustration: for example, financials have an annualized $CRPE_T(\beta)$ of -6.9%, whereas some commodities like oil and natural gas have $CRPE_T(\beta)$ of -11.9% and -12.6%, respectively. The fact that the $CRPE_T(\beta)$ is negative means that market-makers earn a premium for trading against assets with opposite β-s.

To make use of the higher frequency of our data and identify episodes of higher even-order cumulants on a daily basis, we also construct the short-both strategy using daily log-returns. The plots in Fig. 7 illustrate such episodes and can be used as a simple tool to identify stress periods in a given asset, even in real time. The daily returns will be useful also for the construction of our global stress index based on cumulants in Section 4.4.2 below.

For practical implementation of the strategy, one can construct the strategy also with daily simple returns. The last six columns in Table 2 show that the daily returns are positive on average, but volatile and positively-skewed, since the mean is larger than the median. The strategy earns Sharpe ratios above one in many markets: e.g., 2.42 for high yield bonds, 1.56 for Financials, 1.49 for Russell 2000, and 1.31 for natural gas.[12]

Since the short-both strategy returns are a bet on higher risk-neutral vs. physical cumulants, the returns should increase when risk-neutral cumulants rise and decrease when physical cumulants increase. Generally, measures of risk-neutral cumulants across asset classes are not easily available and to proxy for risk-neutral cumulants, we use VIX^2. As explained before, VIX^2 is a measure of risk-neutral cumulants above the second for the S&P 500 index. Since variance and illiquidity in other markets than the S&P 500 generally increase at times when VIX spikes (Bao

[12] We use close-to-close returns as opposed to open-to-close (intra-day) returns. The Sharpe ratios with intra-day returns are even larger. The intra-day strategy is also more profitable after accounting for transaction costs since the trader does not have to pay ETF borrowing fees. Since leveraged ETFs are highly liquid, bid-ask spreads are usually extremely low for most assets.

et al. (2011)), the premium for liquidity provision and the $CRPE_T(\beta)$ in other markets could also increase (in absolute terms) when VIX is higher. Table 3 shows that the returns on the strategy are positively-exposed to VIX^2 across several assets, in line with this intuition. The returns on the strategy are also generally negatively exposed to realised higher-order moments as proxied by $r_{\text{simple}} - \log(1 + r_{\text{simple}})$, where r_{simple} is the simple return on the factor.[13]

4.4.2. $CRPE_T(\beta)$ measures global market stress

The short-both strategy returns across assets can be used as a gauge of global market stress since they increase when even-order risk-neutral cumulants are above physical ones, and when the premium for providing liquidity rises across equities, bonds, commodities, currencies and volatility. To illustrate this fact, we do a principal component (PC) analysis of the daily short-both strategy returns across all assets to quantify the impact of higher cumulants across assets at a high frequency. The variance-covariance matrix of returns does not have a particularly strong factor structure: the first PC explains about 19% of the variation in returns, the first six PCs explain about 52%, and 16 PCs are needed to explain 90% of that variation. This result shows that there are common components to the $CRPE_T(\beta)$ across assets, but the role of asset-specific factors is also significant. The first PC captures mostly the variation in $CRPE_T(\beta)$ of equities, commodities and Treasuries, whereas higher-order PCs capture better the residual variation in high yield corporate bonds, currencies and volatility.

Fig. 8 shows that the first PC spikes in periods of market stress and is highly correlated with VIX with a correlation of 70%. The average return on the strategy across assets is also highly correlated with VIX with a correlation of 66%. These facts show that times when risk-neutral cumulants are above physical ones across assets, as captured by the returns on the short-both strategy, are positively related to periods of market stress when VIX is higher.

The first PC (and higher-order PCs) and the average return on the short-both strategy across assets can be used as a simple index of global market stress. There are several advantages of these metrics relative to other commonly used measures of market turbulence like VIX or various spreads like the TED spread. First, our measures are based on several asset classes and take the prospective of a liquidity provider who is exposed to higher-order cumulants globally. As we show in Section 4.4.3 below, our metric drives out VIX in explaining returns of non-equity assets and is particularly important in assets with non-linear payoffs like options and CDS. Second, our measures are simple to calculate also in real-time from observed prices of leveraged ETFs. The measures are easy to compute also for individual markets and can be used to capture market stress in particular asset class at a high frequency (Fig. 7). Third, we do not make any assumptions about the driving distribution of asset returns and "let the data speak".

[13] VIX^2 and $r_{,\text{simple}} - \log(1 + r_{,\text{simple}})$ capture *all* higher-order cumulants/moments of the factor as opposed to *all even* cumulants/moments, which would be the relevant factors for the short-both strategy returns.

4.4.3. Relation to standard risk factors and cross-sectional asset-pricing

Table 3 shows that some standard risk factors like size, value, and momentum are significantly correlated with the short-both strategy returns for equities. This fact illustrates that these factors might span some of the variation in higher-order cumulants as captured by the $CRPE_T(\beta)$. In turn, this fact could also help explain why the standard linear CAPM logic does not capture the full variation in asset returns.

We next test if higher-order cumulants are priced factors across asset classes. Table 4 shows the results from cross-sectional asset-pricing regressions using the average returns on the short-both strategy across all assets, and the market, as the two factors.[14] The table shows that the price of risk associated with the short-both strategy returns is positive and statistically significant for government bonds, and particularly for options and CDS. The inclusion of the short-both strategy return makes insignificant the return on VIX (proxied by the largest long VIX ETF since VIX is not directly tradable) in all asset classes except US stocks. This is perhaps not surprising since VIX measures risk-neutral entropy of the S&P 500 equity index, whereas our strategy is based on more asset classes beyond equities. The fact that the short-both strategy is particularly important in assets with highly non-linear payoffs like options and CDS shows that these asset classes are more exposed to higher-order cumulant risk.

4.4.4. Comparison to FRP and VRP

It is useful to compare the magnitude of the $CRPE(\beta)$ to that of the FRP. Column 5 of Table 2 shows that the $CRPE(\beta)$ is significant relative to the FRP (in absolute values): it is 46% of the FRP for the S&P 500 index (with $\beta = 3$), 47% for VIX, 51% for long-term Treasuries (with $\beta = 3$), and 139% for oil.

Another interesting benchmark for comparison is the VRP. Carr and Wu (2009), Bakshi and Kapadia (2003), Heston and Li (2020) and Heston and Todorov (2022) show that the VRP is negative whereas Bollerslev and Todorov (2011) show that compensation for jump risk accounts for a large fraction of this premium. Our results show that the $CRPE(\beta)$ is also negative across markets, on average, and that higher-order terms have a non-negligible contribution, particularly during crisis times. The magnitude of the $CRPE(\beta)$ is generally smaller than the VRP, since our measure is different as the $CRPE(\beta)$ depends on cumulants above variance, some of which could have positive risk premium. In addition, our empirical tests rely on assets with leverage between -3 and 3, whereas options involved in the calculation of the VRP have typically larger (absolute) leverages. With a higher leverage, the $CRPE(\beta)$ is also higher as seen from Table 2. The Sharpe ratios of the short-both strategy to extract the $CRPE(\beta)$ are above one in some asset classes, similar to Sharpe ratios of VRP strategies.

[14] One limitation of the analysis is that the data is monthly instead of daily, and includes only a subsample from June 2006 to December 2012.

5. Economic implications and possible extensions

5.1. Implications

The main results in this paper have implications for the standard CAPM and portfolio theory in general, and for factor models. In addition, our findings have implications for momentum, leverage, hedge funds, and option pricing.

5.1.1. Implications for factor models and portfolio theory

Our main results have implications for factor models and portfolio theory. We show that multi-factor models could fit asset returns better than single-factor models purely because the additional factors capture the contribution of higher-order cumulants of the single factor. The fact that some standard factors like momentum are positively correlated with even-order cumulant differences (Table 3), is consistent with this logic. This result has implications for a vast financial literature studying factor models to explain asset returns. Our theory suggests that instead of adding more linear factors, researchers also need to account for the higher-order cumulants of the single-factor (e.g., the market portfolio). In addition, a proper test of single-factor models should first compare the difference between cumulant-generating functions in the physical and risk-neutral worlds before testing linear beta pricing.

The results in this paper have important consequences also for standard portfolio theory. We show that many classic single-factor results hold only in a lognormal world. For example, the standard CAPM logic that asset returns are linear in market returns, holds only in a lognormal world. Another classic portfolio theory result states that by combining two assets with opposite betas, one can construct a risk-free return. Our analysis shows that this is no longer true in a general setting with non-zero higher-order cumulants: such a portfolio would be exposed to the $CRPE(\beta)$ and would not be risk-free.

5.1.2. The flatness of the securities market line (SML)

Our approach could help explain the flatness of the securities market line (SML). Equation 8 shows that an asset with $CRP(\beta) < 0$ has lower return than the one predicted by the CAPM, whereas an asset with $CRP(\beta) > 0$ has a larger return. If $CRP(\beta) > 0$ for assets that have low CAPM betas, whereas $CRP(\beta) < 0$ for assets with high CAPM betas, this fact could explain why the SML is flatter than predicted by the standard CAPM formula. As shown before, assets with $\beta > 1$ conduct momentum trades and would have lower returns than predicted by the CAPM if market makers charge a premium for providing liquidity. Equation 11 shows that such assets load positively on the VRP and if market-makers charge a premium for being short the VRP, that would make estimated $\hat{\beta} < \beta$ for these assets.

There is some evidence in Table 1 that is consistent with this conjecture as several assets with $\beta = 3$ have $CRP(\beta) < 0$ but these results are inconclusive since we do not observe ETFs with $0 < \beta <$

1, and since some ETFs with $\beta = 2$ have $CRP(\beta) > 0$. We leave the test of the SML's flatness through the lens of the CRP for future research.

5.1.3. Implications for momentum, leverage, hedge funds and policy makers

Our findings have implications also for momentum strategies. We show that trend-chasing "momentum" strategies are exposed to the VRP and higher-order cumulants, which could explain why the returns on these strategies have sudden crashes and exhibit higher-order moments.

Our results have also implications about the risk of higher-order cumulants. A common misperception is that this risk declines as the number of higher-order terms grows and thus higher-order moments (typically, beyond kurtosis) are rarely researched in finance. This misperception is driven by the discounting of higher-order cumulant differences with $n!$ (see Equation 5), which makes the contribution of higher-order terms extremely small for larger n. Our theory emphasises that this result is true for unleveraged strategies (and even more pronounced for strategies with $0 < \beta < 1$), but is not true for leveraged strategies, for which the contribution of higher-order cumulants generally *increases* up to the β-th order cumulant. For example, the loadings of strategy with a leverage of $\beta = 10$ are increasing up to the 10th order term as illustrated in the left panel of Fig. 9. In contrast, the loadings of an unleveraged strategy quickly die out as illustrated in the right panel. Thus, more leveraged strategies are more exposed to higher-order cumulants.

These results have implications for agents like hedge funds who use leverage to exploit mispricings between similar assets. These agents often use strategies that involve assets with opposite sensitivities to a given factor: for example, convergence trades or relative value strategies (e.g., spot-futures basis, see Aramonte et al. (2021)). Our results show that such trades are risky because they are exposed to the CRPE, even in the case of no limits to arbitrage or noise trader risk (e.g., Shleifer and Vishny (1997)). For example, Equation 18 shows that a leverage of two has loadings of $2(= 2^2/2!)$, 0.67 and 0.09 on the second, forth, and sixth order CRP terms, respectively. In contrast, a leverage of ten, which is often used by hedge funds in such trades, has loadings of 50, 417, and 1389 on these CRP terms. The loadings on higher-order terms are even larger and are above one up to the 24th CRP term, which illustrates that even tiny changes in cumulant differences are magnified due to the explosive contribution of 10^n-weighted CRP terms. This reflects the enormous exposure of such levered trades to higher-order cumulants.

Our results have implications also for policy makers and practitioners. The first PC of the short-both strategy can be a useful gauge for policy intervention since the indicator increases when the $CRPE(\beta)$ rises, which could be a proxy for times when capital constraints are binding as we explain in Section 5.2. One benefit of the CRP-inspired approach is that it is based on several asset classes and incorporates information on the difference of *all* higher-order even cumulants, in contrast to indicators for policy intervention based on variance only. This benefit is evident if an increase in volatility reflects other factors than capital constraints. For example, an increase in uncertainty would raise variance but might not affect capital constraints or the risk of price spirals. Such an increase may not affect the $CRPE(\beta)$, however, if variance rises by the same amount under the physical and the risk-neutral measure and if higher-order cumulant differences are unchanged. In that case, variance-based indicators for policy intervention would increase, whereas the $CRPE(\beta)$ would remain unchanged.

5.1.4. Implications for option pricing

Our results show that out-of-the-money (OTM) put options are more expensive than what linear beta pricing would predict. Since these options have $\Delta < 0$ and thus, leverage $\beta < 0$, they are similar to momentum assets, and would load positively on the VRP in a stochastic volatility setting, for example. As the VRP is negative in practice, the returns on OTM put options would be more negative (equivalently, the options will be more expensive) than predicted by linear beta pricing as in a standard CAPM model, especially for OTM puts with more negative β-s. However, two caveats are that the standard linear pricing logic is presumably not applicable for options given that these assets do not load linearly on the underlying asset, and that higher-order cumulants beyond variance can have a positive risk premium.

5.2. Sources of the $CRP(\beta)$

Equation 7 shows that the $CRP(\beta)$ arises if cumulant differences, weighted by $(\beta^n - \beta)$, are larger than zero. These differences could arise if liquidity providers charge a premium for trading in the opposite direction of constant beta assets. What factors can create such a premium? Trading restrictions or other forms of market incompleteness are likely to give rise to higher-order cumulants since market makers cannot perfectly hedge those and would require an additional premium, which would be reflected in the $CRP(\beta)$. For example, limited trading hours create discontinuities in trading and could lead to higher-order cumulants being relevant since the return distribution is no longer continuous. In addition, risk limits like value-at-risk constraints, de-leveraging (e.g., Adrian and Shin (2010)), or crowded trades could create price spirals at times of large price movements and cause extreme values of the factor's return distribution.

Limits to arbitrage and costly capital could also give rise to the $CRP(\beta)$. Kyle and Xiong (2001) and Xiong (2001) show that convergence traders' wealth effect can amplify price changes and volatility, and prove contagious. Convergence trades to extract the $CRP(\beta)$ risk being liquidated prematurely if limits to arbitrage make raising capital costly and force traders to close out these trades before prices converge. Such liquidation amplifies further price drops and raises the $CRP(\beta)$ by increasing higher-order cumulants of the return distribution. The risk of future price spirals could then prevent traders from arbitraging away the $CRP(\beta)$. When limits to arbitrage are binding, the $CRP(\beta)$ should become larger. The fact that the $CRP(\beta)$ increases in times of market stress is consistent with this explanation.

The $CRP(\beta)$ can also arise due to trading patterns of momentum traders. For example, the daily rebalancing of ETFs to keep constant β can amplify price movements and increase cumulants if this rebalancing is large part of the market. An important point is that the rebalancing of strategies with $\beta > 1$ and $\beta < 0$ is in the same direction, which means that ETFs can amplify price changes even if the size of long ETFs is equal to that of inverse ETFs. Todorov (2019) shows that this rebalancing is significant share of the market in VIX and commodity markets, and can lead to sharp price changes and larger cumulants, as in February 2018 for VIX and April 2020 for oil.

Another explanation for the existence of the $CRP(\beta)$ is that the "natural" distribution of the factor's return could be one with a complicated form of non-zero higher-order cumulants: for example, it is reasonable to assume that volatility (VIX) has a positively-skewed and highly non-normal distribution with jumps. Whatever the reason for cumulants, risk-averse market makers would require a compensation for providing liquidity and bearing the cumulant risk. We leave the derivation of the $CRP(\beta)$ in the settings discussed here for future research.

5.3. Robustness: incorporating ETF fees

In the main empirical analysis, we used observed market prices of ETFs to construct the short-both strategy as these prices would be used by a trader who implements the strategy in practice. To address the concern that the $CRPE(\beta)$ is purely driven by ETF fees, we also repeat the analysis using before-fees returns in Table B.3. The table shows that the $CRPE(\beta)$ is slightly smaller but the Sharpe ratios of the strategy are still above one for some assets like high yield bonds and financials. This analysis shows that the $CRPE(\beta)$ is not mechanically driven by ETF fees. We also repeat Table 1 with fees in Table B.4 in the Appendix. The results show that the $CRP(\beta)$ is still significantly different from zero.

6. Conclusion

Higher-order cumulants play an important role in practice. We show that single-factor models work only when the difference of higher-order physical and risk-neutral cumulants is zero. In any other setting, the standard linear factor pricing equation should be adjusted for this difference, which we call the cumulant risk premium (CRP). To illustrate our approach, we study assets with constant betas and exposure to a single factor: leveraged ETFs. We show that the CRP is different from zero and that linear pricing fails across all assets studied. The CRP is a large part of the factor risk premium. These empirical findings cannot be explained by the Black–Scholes model or by any model without compensation for higher-order cumulant risk, but might be explained by a model with jumps or stochastic volatility.

We develop a simple strategy of shorting ETFs with opposite betas to measure the even-order CRP across asset classes. This strategy mimics liquidity provision and can be used to construct a bet on risk-neutral vs. physical even-order cumulants (variance, scaled kurtosis, etc.). The strategy earns Sharpe ratios above one.

Our findings have important implications not only for factor models but also for portfolio theory, momentum strategies, option pricing, hedge funds, and leverage in general. We show that standard portfolio theory results do not hold in a general setting with non-zero higher-order cumulants and that highly leveraged strategies employ momentum strategies. These findings have implications for asset managers and hedge funds who use large leverage to exploit mispricings between similar assets. Leveraged trades involving assets with opposite betas are exposed to higher-order cumulants and even tiny changes in these cumulants can be magnified enormously. Our cumulant-based index can be used as a simple, real-time gauge of market stress across asset classes.

References

Adrian, T., Shin, H.S., 2010. Liquidity and leverage. Journal of Financial Intermediation 19, 418–437.

Andrews, D., 1991. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. Econometrica 59, 817–858.

Ang, A., Hodrick, R.J., Xing, Y., Zhang, X., 1997. The cross-section of volatility and expected returns. Journal of Finance 61, 259–299.

Aramonte, S., Schrimpf, A., Shin, H.S., 2021. Non-bank financial intermediaries and financial stability BIS working paper N 972.

Backus, D., Chernov, M., Martin, I., 2011. Disasters implied by equity index options. Journal of Finance 66, 1969–2012.

Bakshi, G., Kapadia, N., 2003. Delta-hedged gains and the negative market volatility risk premium. The Review of Financial Studies 16, 527–566.

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

Bao, J., Pan, J., Wang, J., 2011. The illiquidity of corporate bonds. Journal of Finance 66, 911–946.

Barro, R.J., 2006. Rare disasters and asset markets in the twentieth century. Quarterly Journal of Economics 121, 823–867.

Barro, R.J., Nakamura, E., Steinsson, J., Ursua, J.F., 2013. Crises and recoveries in an empirical model of consumption disasters. American Economic Journal: Macroeconomics 5, 35–74.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–654.

Bollerslev, T., Todorov, V., 2011. Tails,fears and risk premia. Journal of Financial Economics 66, 2165–2211.

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

Borri, N., Verdelhan, A., 2012. Sovereign risk premia. Working paper .

Carr, P., Wu, L., 2009. Variance risk premiums. Review of Financial Studies 22, 1311–1341.

Cheng, M., Madhavan, A., 2009. The dynamics of leveraged and inverse exchange-traded funds. Journal of Investment Management 7, 43–62.

Cornish, E.A., Fisher, R.A., 1938. Moments and cumulants in the specification of distributions. Review of the International Statistical Institute 5, 307–320.

Dew-Becker, I., Giglio, S., 2022. Risk preferences implied by synthetic options. Unpublished working paper, Northwestern University and Yale University.

Duffie, D., Pan, J., Singleton, K.J., 2000. Transform analysis and asset pricing for affine jump-diffusions. Econometrica 68, 1343–1376.

Eraker, B., 2004. Do stock prices and volatility jump? reconciling evidence from spot and option prices. Journal of Finance 59, 1367–1404.

Fama, E.F., French, K.R., 1992. The cross-section of expected stock returns. Journal of Finance 47, 427–465.

Fama, E.F., French, K.R., 1995. Size and book-to-market factors in earnings and returns. Journal of Finance 50, 131–155.

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

Han, J., Kyle, A.S., 2017. Speculative equilibrium with differences in higher-order beliefs. Management Science 64, 4317–4332.

Hansen, L.P., Richard, S.F., 1987. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. Econometrica 55, 587–613.

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

He, Z., Kelly, B., Manela, A., 2017. Intermediary asset pricing: new evidence from many asset classes. Journal of Financial Economics 126, 1–35.

Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies 6, 327–343.

Heston, S.L., Li, S., 2020. Option momentum. Unpublished working paper .

Heston, S.L., Todorov, K., 2022. Exploring the variance risk premium across assets. Unpublished working paper .

Jagannathan, R., Wang, Z., 1996. The conditional capm and the cross-section of expected returns. The Journal of Finance 51, 3–53.

Kraus, A., Litzenberger, R.H., 1976. Skewness preference and the valuation of risk assets,. Journal of Finance 31, 1085–1100.

Kyle, A.S., Xiong, W., 2001. Contagion as a wealth effect. Journal of Finance 56, 1401–1440.

Lakonishok, J., Shleifer, A., Vishny, R.W., 1994. Contrarian investment, extrapolation, and risk. Journal of Finance 49, 1541–1578.

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

Liu, J., Longstaff, F.A., Pan, J., 2003. Dynamic asset allocation with event risk. Journal of Finance 58, 231–259.

Longstaff, F.A., Piazzesi, M., 2004. Corporate earnings and the equity premium. Journal of Financial Economics 74, 401–421.

Lustig, H., Roussanov, N., Verdelhan, A., 2011. Common risk factors in currency markets. The Review of Financial Studies 24, 3731–3777.

Marcinkiewicz, J., 1935. On the convergence of fourier series. Journal of the London mthematical Society , 264–268.

Martin, I., 2013. Consumption-based asset pricing with higher cumulants. Review of Economic Studies 80, 745–773.

Martin, I., 2015. Simple variance swaps. Unpublished working paper, London School of Economics.

Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012. Carry trades and global foreign exchange volatility. Journal of Finance 67, 681–718.

Merton, R., 1992. Continuous time finance. Blackwell.

Merton, R.C., 1976. Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics 3, 125–144.

Nagel, S., 2012. Evaporating liquidity. The Review of Financial Studies 25, 2005–2039.

Newey, W.K., West, K.D., 1987. A simple, positive semi-definite heteroscedasticity and autocorrelation-consistent covariance matrix. Econometrica 55, 703–708.

Ouimet, F., 2021. General formulas for the central and non-central moments of the multinomial distribution. Unpublished working paper .

Pan, J., 2002. The jump-risk premia implicit in options: Evidence from an integrated time-series study. Journal of Financial Economics 63, 3–50.

Rietz, T.A., 1988. The equity risk premium: A solution. Journal of Monetary Economics 22, 117–131.

Roll, R., 1997. A critique of the asset pricing theory's tests' part i: On past and potential testability of the theory. Journal of Financial Economics 4, 129–176.

Roll, R., Ross, S., 1994. On the cross-sectional relation between expected returns and betas. Journal of Finance 49, 101–121.

Schneider, Paul, C.W., Zechner, J., 2020. Low risk anomalies? Journal of Finance 75, 2673–2718.

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

Shleifer, A., Vishny, R.W., 1997. The limits of arbitrage. Journal of Finance 52, 35–55.

Stambaugh, R.F., 1982. On the exclusion of assets from tests of the two-parameter model: A sensitivity analysis. Journal of Financial Economics 10, 237–268.

Todorov, K., 2019. Passive funds affect prices: Evidence from the most ETF-dominated markets. Working paper .

Tuzun, T., 2013. Are leveraged and inverse ETFs the new portfolio insurers?. Unpublished working paper .

Xiong, W., 2001. Convergence trading with wealth effects: an amplification mechanism in financial markets. Journal of Financial Economics 62, 247–292.

Tables and Figures

Table 1
Estimating $CRP(\beta)$. The table shows the annualized $CRP(\beta)$ in %, estimated as α from regression $r_{\text{ETF},t}(\beta) = \alpha + \beta r_{bmk,t} + (1-\beta) r_{f,t} + \epsilon_t$ for several markets and leverages, where $r_{\text{ETF},t}(\beta)$ is the return on an ETF with leverage β, $r_{bmk,t}$ is the return on the ETF benchmark, and β is the ETF leverage. We estimate $r_t = \log \text{E}[R_t]$ by first calculating $\text{E}[R_t]$ as the average daily return, and then running monthly regressions of $\log \text{E}[R_{\text{ETF},t}(\beta)]$ on $\log \text{E}[R_{bmk,t}]$. Here and in the subsequent analysis $r_{f,t}$ is the 1-month Treasury rate and standard errors (shown in brackets) are computed using the Newey and West (1987) estimator with lag selection based on the Bartlett kernel (e.g., Andrews (1991)). The **bold** coefficients are statistically different from zero at the 5% level. The sample is from the first leveraged ETF inception date in a given market to April 2021.

	Annualized $CRP_T(\beta)$ in %				
Leverage= β	-3	-2	-1	2	3
S&P 500	**-9.10**	**-6.76**	**-3.94**	**1.37**	1.41
	(0.63)	(1.02)	(0.88)	(0.40)	(0.97)
Nasdaq	**-5.55**				-0.03
	(0.46)				(0.57)
Russell 2000	**-10.38**				-0.54
	(0.83)				(1.29)
Financials	**-12.86**				-2.06
	(1.26)				(1.47)
Consumer services		**-6.68**		0.04	
		(2.28)		(0.9)	
Basic materials		**-8.86**		-0.87	
		(2.67)		(0.95)	
Technology	-2.34				**-5.84**
	(1.71)				(1.61)
Utilities		**-13.38**		**2.59**	
		(2.96)		(0.93)	
Industrials		**-8.45**		0.33	
		(3.11)		(0.95)	
Real Estate	1.09	**-13.10**	**-5.12**	**2.76**	**-14.62**
	(2.04)	(1.36)	(0.40)	(1.25)	(3.85)
Emerging Markets	-5.86	**-8.43**	**-7.17**	3.33	-3.43
	(3.46)	(2.3)	(1.65)	(2.21)	(3.27)
VIX			**-37.14**	**-31.26**	
			(12.97)	(6.87)	
Treasuries 7-10 yr	**-2.81**	**-1.32**	**-1.93**	**-1.89**	**-4.38**
	(0.66)	(0.28)	(0.3)	(0.38)	(2.08)
Treasuries more 20 yr	**-2.39**	**-2.17**	**-1.63**	**-1.82**	**-4.13**
	(0.96)	(0.63)	(0.32)	(0.68)	(1.08)
High Yield			**-2.54**	-1.33	
			(0.65)	(2.52)	
Gold	**-3.99**	**-3.80**		-0.96	-0.45
	(1.58)	(0.81)		(0.84)	(1.40)
Silver	**-8.90**	**-10.71**		-3.91	-3.86
	(2.04)	(4.96)		(4.83)	(2.02)
Nat gas	**-12.79**	**-6.93**		**-3.76**	**-9.96**
	(3.25)	(1.7)		(1.63)	(3.0)
Oil	30.03	6.43		**-24.09**	-54.09
	(28.3)	(6.85)		(10.83)	(29.22)
Euro/US Dollar		-1.09		-0.89	
		(0.81)		(0.69)	
Yen/US Dollar		**-0.69**		-0.99	
		(0.33)		(0.81)	

Table 2

Returns on the short-both strategy and the FRP. The second column shows the average annualized return on the short-both strategy $r_{\mathrm{SB},t}(\beta) = -(\log \mathrm{E}[R_{\mathrm{ETF},t}(\beta)] + \log \mathrm{E}[R_{\mathrm{ETF},t}(-\beta)])$, where $\log \mathrm{E}[R_{\mathrm{ETF},t}(\beta)]$ is the return on an ETF with leverage β. $\mathrm{E}[R_{\mathrm{ETF},t}(\beta)]$ is the average daily return in a given month as in Table 1. Column 3 shows the average annualized $CRPE_T(\beta)$ $(= -\frac{1}{2} r_{\mathrm{SB},t}(\beta))$. Column 4 shows the average annualized factor risk premium ($FRP_T = \log \mathrm{E}[R_T(1)] - \log \mathrm{E}[R_{f,T}]$). Column 5 shows the ratio of the $CRPE_T(\beta)$ to the FRP_T. The last six columns show summary statistics of the short-both strategy with daily simple returns. Columns 2–5 are in %, 6–11 in basis points. The sample is from the first inverse ETF inception date in a given market to April 2021 (February 2018 for VIX, June 2020 for gold and gas since some long and inverse ETFs were delisted on those dates).

Asset	β (1)	Mean SB annual (2)	Mean $CRPE_T(\beta)$ annual (3)	Mean FRP annual (4)	$CRPE_T(\beta)/FRP_T$ (5)	Mean SB daily (6)	S.d. SB daily (7)	Median SB daily (8)	Min SB daily (9)	Max SB daily (10)	Sharpe annual (11)
S&P 500	1	2.20	-1.10	7.11	-0.15	0.88	27.98	0.00	-370.46	1384.66	0.29
S&P 500	3	6.55	-3.28	7.11	-0.46	2.55	32.47	0.60	-476.84	1157.92	1.16
Nasdaq	3	4.68	-2.34	10.23	-0.23	1.44	17.11	0.27	-227.33	242.19	1.13
Russell 2000	3	9.65	-4.83	10.60	-0.46	3.58	36.65	1.32	-436.64	1133.65	1.49
Financials	3	13.80	-6.90	9.70	-0.71	4.71	46.42	1.42	-339.22	1166.94	1.56
Consumer services	2	5.25	-2.63	9.10	-0.29	1.91	152.07	0.00	-2147.13	2724.9	0.17
Basic materials	2	8.53	-4.26	8.80	-0.48	3.04	88.27	0.26	-518.3	3918.69	0.49
Technology	3	7.55	-3.78	16.37	-0.23	3.06	41.24	0.74	-347.57	1435.52	1.11
Utilities	2	9.60	-4.80	3.83	-1.25	3.74	111.27	0.00	-587.54	3967.19	0.49
Industrials	2	6.80	-3.40	7.90	-0.43	2.51	129.74	0.00	-844.05	4541.52	0.27
Real Estate	3	9.40	-4.70	9.94	-0.47	3.68	71.05	1.49	-1441.85	2921.62	0.78
Emerging Markets	1	3.65	-1.83	7.08	-0.26	1.38	34.01	0.00	-649.28	1044.11	0.54
Emerging Markets	3	8.78	-4.39	7.08	-0.62	3.26	45.41	0.74	-481.70	1450.78	1.08
VIX	1	33.05	-16.53	-35.26	0.47	8.17	231.65	0.95	-1991.37	9582.18	0.55
Treasuries 7-10 yr	1	2.25	-1.13	3.75	-0.30	0.88	21.39	0.00	-177.77	191.40	0.49
Treasuries 7-10 yr	3	6.45	-3.23	3.75	-0.86	2.73	83.65	1.30	-1619.78	1358.75	0.49
Treasuries more 20 yr	1	3.50	-1.75	5.68	-0.31	1.35	17.10	0.23	-476.69	421.95	1.07
Treasuries more 20 yr	3	5.80	-2.90	5.68	-0.51	2.71	43.18	0.88	-224	2074.34	0.93
High Yield	1	7.15	-3.58	5.28	-0.68	2.86	17.28	0.80	-96.95	107.24	2.42
Gold	2	4.18	-2.09	4.83	-0.43	1.47	18.90	0.57	-132.3	207.44	1.08
Silver	3	12.05	-6.03	-4.45	1.35	5.85	215.61	0.79	-3016.03	4154.14	0.25
Nat gas	3	25.20	-12.60	-15.57	0.81	7.80	91.15	3.02	-1379.3	1251.68	1.31
Oil	3	23.88	-11.94	8.60	-1.39	7.48	361.02	0.00	-5995.73	6830.22	0.32
Euro/US Dollar	2	1.38	-0.69	-0.80	0.86	0.43	32.97	0.00	-280.27	286.97	0.11
Yen/US Dollar	2	0.83	-0.41	0.52	-0.79	0.49	43.12	0.00	-339.89	289.44	0.11

Table 3

Regressions of the short-both strategy (simple returns) r_{SB} on VIX^2, $r_{\text{simple}} - \log(1 + r_{\text{simple}})$ (a measure of physical higher-order moments of the factor), the Fama-French 5 factors and momentum. Sample: daily, from the first date of an introduction of an inverse ETF until December 2018 (February 2018 for VIX). VIX^2 is scaled by 100 for comparison. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Dependent variable: r_{SB}									
	S&P 500	Nasdaq	Russell 2000	Financials	Consumer services	Basic materials	Technology	Utilities	Industrials	Real estate
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
VIX^2	1.13*	0.36***	0.12	0.09	1.97	2.54	−0.26	2.80***	4.41*	−0.54*
	(0.64)	(0.09)	(0.41)	(0.15)	(1.33)	(1.60)	(0.31)	(1.00)	(2.57)	(0.30)
$r_M - \log(1 + r_M)$	−2.82*	−0.20	0.40	3.35***	−4.72	−1.62	2.90	−5.34**	−10.08	1.19*
	(1.70)	(0.33)	(1.32)	(0.18)	(4.25)	(2.13)	(1.88)	(2.20)	(7.55)	(0.65)
$R_{Mkt} - R_f$	0.02	0.02*	0.01	0.04***	−0.03	0.07	0.05***	0.10	0.07	−0.05**
	(0.02)	(0.01)	(0.03)	(0.01)	(0.04)	(0.05)	(0.01)	(0.07)	(0.07)	(0.02)
SMB	0.05	−0.01	0.03	−0.04	0.28***	0.13	0.01	0.22*	0.04	0.03
	(0.04)	(0.01)	(0.03)	(0.03)	(0.10)	(0.12)	(0.01)	(0.12)	(0.11)	(0.03)
HML	0.16	−0.02*	0.08	0.02	0.28	0.38	−0.06*	0.46	0.66	−0.02
	(0.11)	(0.01)	(0.06)	(0.02)	(0.21)	(0.27)	(0.04)	(0.28)	(0.40)	(0.05)
RMW	−0.01	−0.01	−0.03	−0.07*	0.02	−0.01	−0.01	−0.08	0.01	0.02
	(0.03)	(0.01)	(0.04)	(0.04)	(0.10)	(0.06)	(0.02)	(0.15)	(0.09)	(0.04)
CMA	−0.19	0.02	−0.19	−0.03	−0.37	−0.54	0.07*	−0.74*	−1.02*	−0.01
	(0.16)	(0.02)	(0.14)	(0.03)	(0.28)	(0.39)	(0.04)	(0.43)	(0.55)	(0.06)
Momentum	0.05*	−0.01**	−0.04*	0.06**	0.01	0.11	0.002	0.15	0.17	−0.07***
	(0.03)	(0.01)	(0.02)	(0.03)	(0.09)	(0.09)	(0.01)	(0.13)	(0.10)	(0.02)
Observations	150	106	121	121	142	142	120	142	142	27
R^2	0.26	0.40	0.26	0.97	0.29	0.29	0.20	0.35	0.31	0.29

	Emerging markets	VIX	Treasuries 7-10 yr	High yield	Gold	Silver	Nat gas	Oil	Euro/US Dollar	Yen/US Dollar
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
VIX^2	0.39*	1.83	0.04	0.08	0.38***	0.96	−0.73	1.17	0.06	−0.35**
	(0.21)	(3.74)	(0.23)	(0.08)	(0.06)	(0.60)	(0.81)	(1.50)	(0.09)	(0.18)
$r_M - \log(1 + r_M)$	−0.50	1.87	−1.92	3.69	−0.29	0.71	0.45	2.74*	2.87	4.16
	(0.42)	(1.17)	(10.30)	(3.54)	(0.35)	(1.20)	(1.19)	(1.54)	(3.75)	(3.42)
$R_{Mkt} - R_f$	−0.01	−1.06	−0.01	0.03**	0.005	0.05	−0.20**	0.24**	−0.003	−0.02
	(0.01)	(1.02)	(0.02)	(0.01)	(0.01)	(0.05)	(0.09)	(0.11)	(0.02)	(0.02)
SMB	0.07	1.36	−0.01	0.01	0.01	0.03	0.15	0.08	−0.03*	−0.03
	(0.04)	(1.19)	(0.02)	(0.01)	(0.01)	(0.06)	(0.13)	(0.08)	(0.02)	(0.03)
HML	0.10	1.39	−0.01	0.03*	0.05**	0.14*	−0.18	0.06	−0.03	−0.01
	(0.07)	(1.18)	(0.01)	(0.02)	(0.03)	(0.07)	(0.13)	(0.09)	(0.04)	(0.03)
RMW	−0.05	0.81	0.004	0.04	−0.01	0.12	0.05	0.20	−0.01	−0.09*
	(0.04)	(0.75)	(0.02)	(0.03)	(0.02)	(0.09)	(0.16)	(0.18)	(0.02)	(0.05)
CMA	−0.18	−3.95	0.01	0.01	−0.07**	0.05	0.06	0.17	0.10	−0.04
	(0.15)	(3.28)	(0.03)	(0.03)	(0.03)	(0.12)	(0.14)	(0.12)	(0.06)	(0.06)
Momentum	0.03	0.25	0.003	0.02*	0.003	0.06	0.04	0.02	0.004	−0.03*
	(0.02)	(0.31)	(0.02)	(0.01)	(0.01)	(0.05)	(0.09)	(0.07)	(0.01)	(0.02)
Observations	133	86	92	93	120	86	82	23	121	121
R^2	0.16	0.23	0.04	0.14	0.46	0.20	0.09	0.32	0.13	0.26

Table 4
Cross-sectional asset pricing. The table reports risk price estimates for the equal-weighted average return on the short-both strategy, the excess return on the market, and the return on VIX ETF. Risk prices are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on risk exposures (betas), reported in percentage terms. Betas are estimated in a first-stage time-series regression. The portfolios of assets are from Asaf Manela's website and are based on He et al. (2017). Stocks are 25 portfolios sorted by size and book-to-market, US gov. bonds are 10 maturity-sorted US government bond portfolios with maturities from six months to five years. Sov. bonds are the six portfolios from Borri and Verdelhan (2012). Options are S&P 500 index options sorted on moneyness and maturity. CDS are 20 portfolios sorted by spreads. FX are 6 currency portfolios sorted on interest rate differential (Lettau et al. (2014)) and 6 currency portfolios sorted on momentum (Menkhoff et al. (2012)). Shanken (Shanken (1992)) standard errors are in parentheses. Monthly frequency, from June 2006 to December 2012.

Asset	Stocks (1)	US gov. bonds (2)	Sov. bonds. (3)	Options (4)	CDS (5)	FX (6)
Short-both	-0.44	1.29***	3.20**	1.40***	1.38***	0.51
	(0.43)	(0.24)	(1.59)	(0.27)	(0.20)	(0.81)
Market	-0.33	1.33*	2.69	-0.64	0.46	1.93***
	(0.61)	(0.81)	(0.93)	(0.73)	(0.82)	(0.52)
VIX	-2.13*	-0.94	0.50	-0.49	-3.16	0.33
	(1.05)	(1.14)	(0.30)	(0.90)	(1.96)	(0.20)
Intercept	0.94*	0.18***	-1.38	-1.13***	-0.07**	-0.15
	(0.50)	(0.05)	(0.50)	(0.36)	(0.03)	(0.23)
Number of portfolios	25	20	6	18	20	12
R^2	0.08	0.66	0.88	0.99	0.88	0.55

Fig. 1. Cumulant risk premium in the Heston model and with jumps. The left panel shows the cumulant-generating functions in the physical $(c(\beta))$ and risk-neutral $(c^*(\beta))$ worlds for the Heston model, the right for the compound Poisson process (CPP). The parameters for the Heston model are: $\mu = 0.05$, $r_f = 0.02$, $\lambda = 2$, $\lambda^* = 1$, $\bar{v} = 0.01$, $\bar{v} = 0.04$, $\rho = -0.7$, $\sigma = 0.1$, $T - t = 1$. The parameters for the CPP are: $\mu = 0.05$, $r_f = 0.02$, $\lambda = 0.1$, $\lambda^* = 0.2$, $b = -0.05$, $b^* = -0.2$, $T - t = 1$.

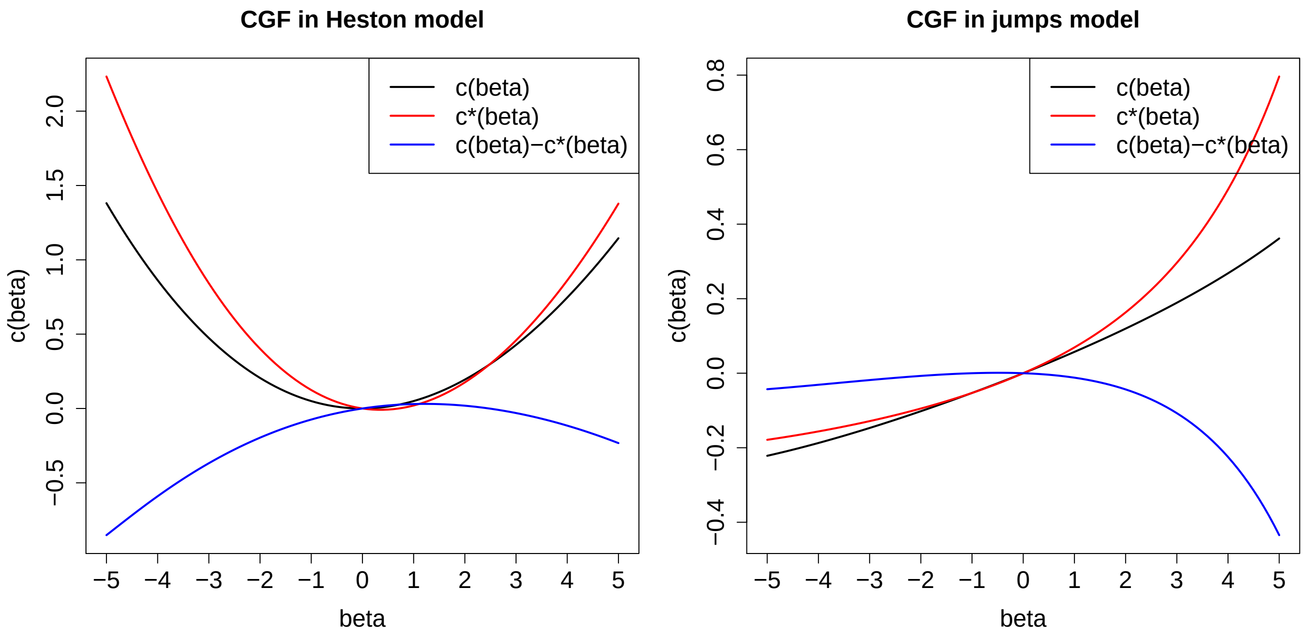

Fig. 2. Realised higher-order moments (second and above) for S&P 500 ($\beta = 3$), Nasdaq ($\beta = 3$), Small cap stocks ($\beta = 3$), Emerging market stocks ($\beta = 3$), Treasuries 20yr+ ($\beta = 3$), Japanese Yen/US Dollar ($\beta = 2$), Natural gas ($\beta = 3$), Oil ($\beta = 2$) and Gold ($\beta = 3$). The graphs show cumulative 1-month annualized differences between r_{simple} and $\log(1 + r_{\text{simple}})$ for the benchmark index (in red) and the ETF with the particular β (in blue).

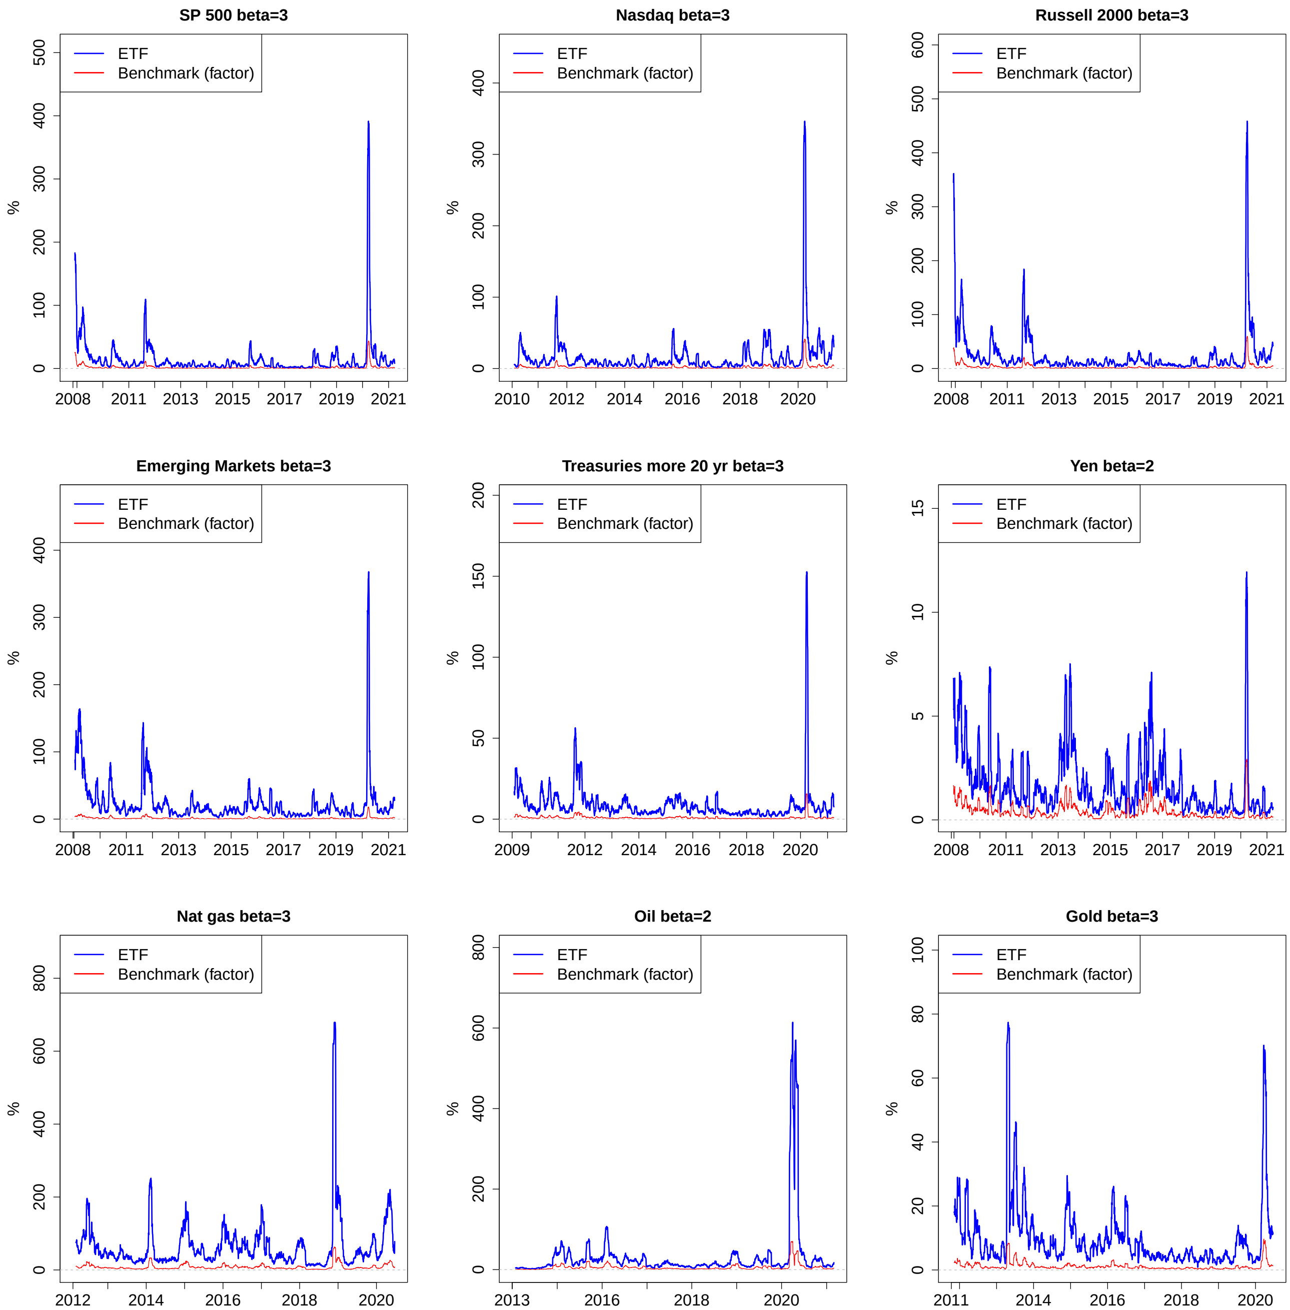

Fig. 3. Realised higher-order moments (third and above) for S&P 500 ($\beta = 3$), Nasdaq ($\beta = 3$), Small cap stocks ($\beta = 3$), Emerging market stocks ($\beta = 3$), Treasuries 20yr+ ($\beta = 3$), Japanese Yen/US Dollar ($\beta = 2$), Natural gas ($\beta = 3$), Oil ($\beta = 2$) and Gold ($\beta = 3$). The graphs show cumulative 1-month annualized differences between $\log(1 + r_{\text{simple}})$ and $r_{\text{simple}} - \frac{1}{2} r_{\text{simple}}^2$ for the benchmark index (in red) and the ETF with the particular β (in blue).

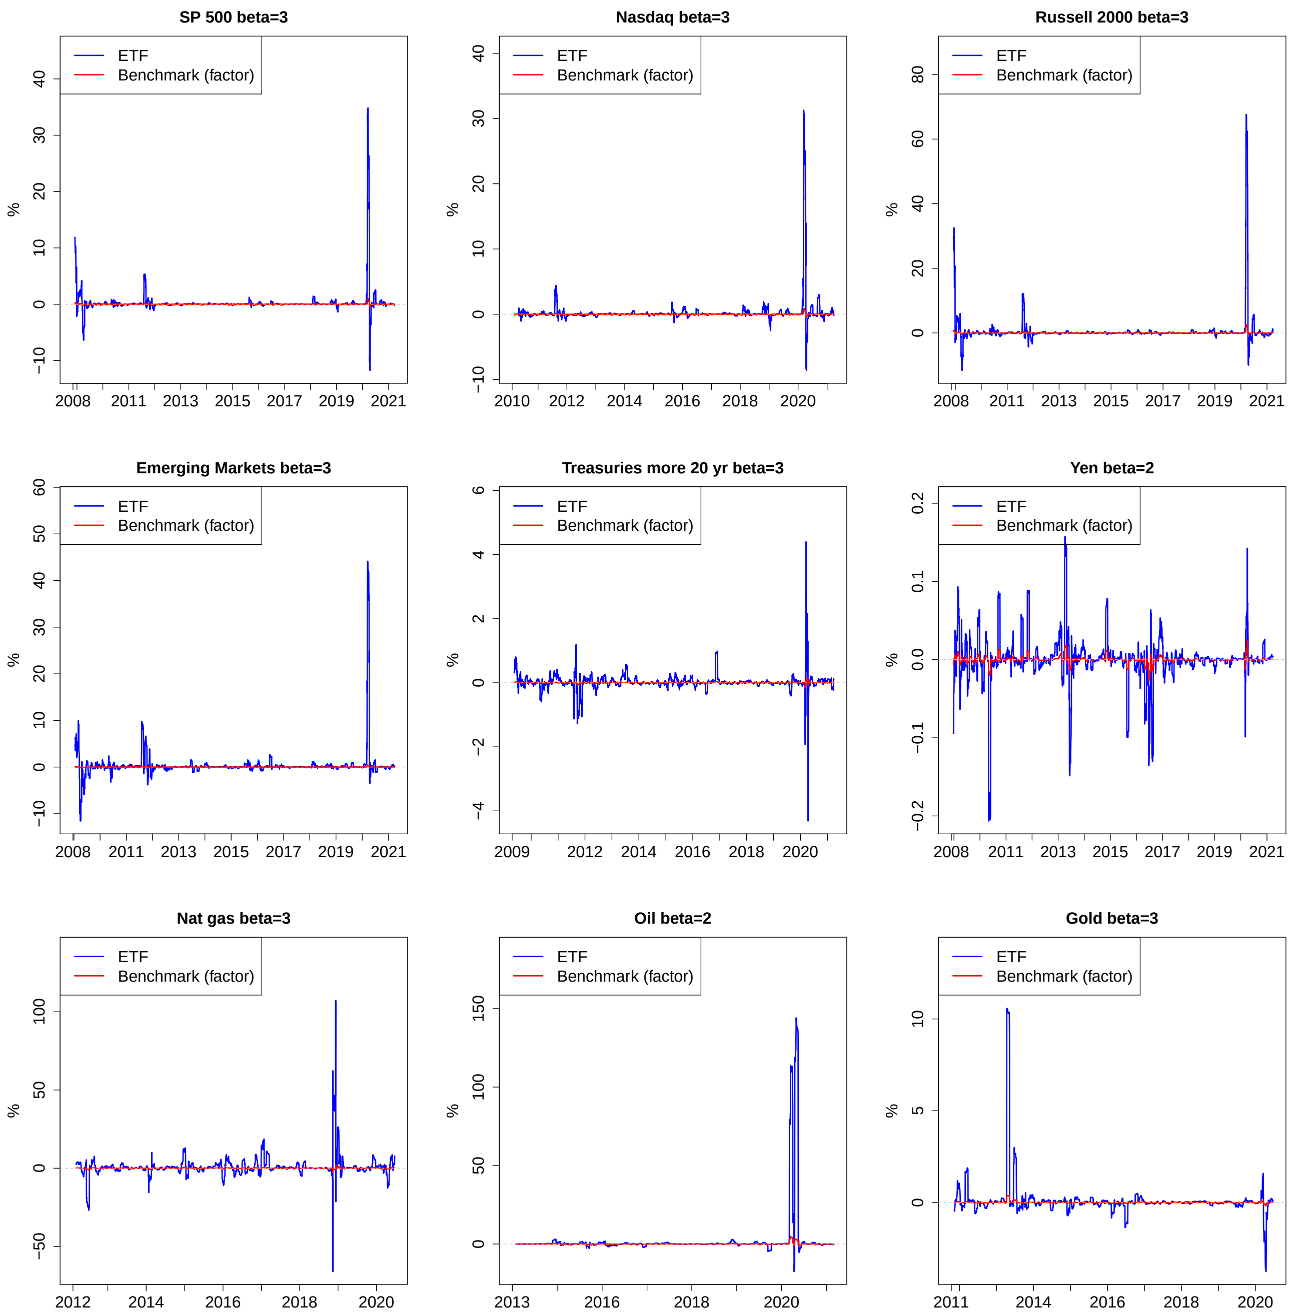

Fig. 4. Difference between the ETF return and the return implied from linear beta pricing for S&P 500 ($\beta = 3$), Nasdaq ($\beta = 3$), Small cap stocks ($\beta = 3$), Emerging market stocks ($\beta = 3$), Treasuries 20yr+ ($\beta = 3$), Japanese Yen/US Dollar ($\beta = 2$), Natural gas ($\beta = 3$), Oil ($\beta = 2$) and Gold ($\beta = 3$). The graphs show 1-month cumulative differences between the return on the ETF and the sum of the return on the leveraged benchmark and the risk-free rate: $r_{ETF} - (\beta r + (1-\beta) r_f)$. Blue lines are long ETFs ($\beta > 0$), red lines are inverse ETFs ($\beta < 0$).

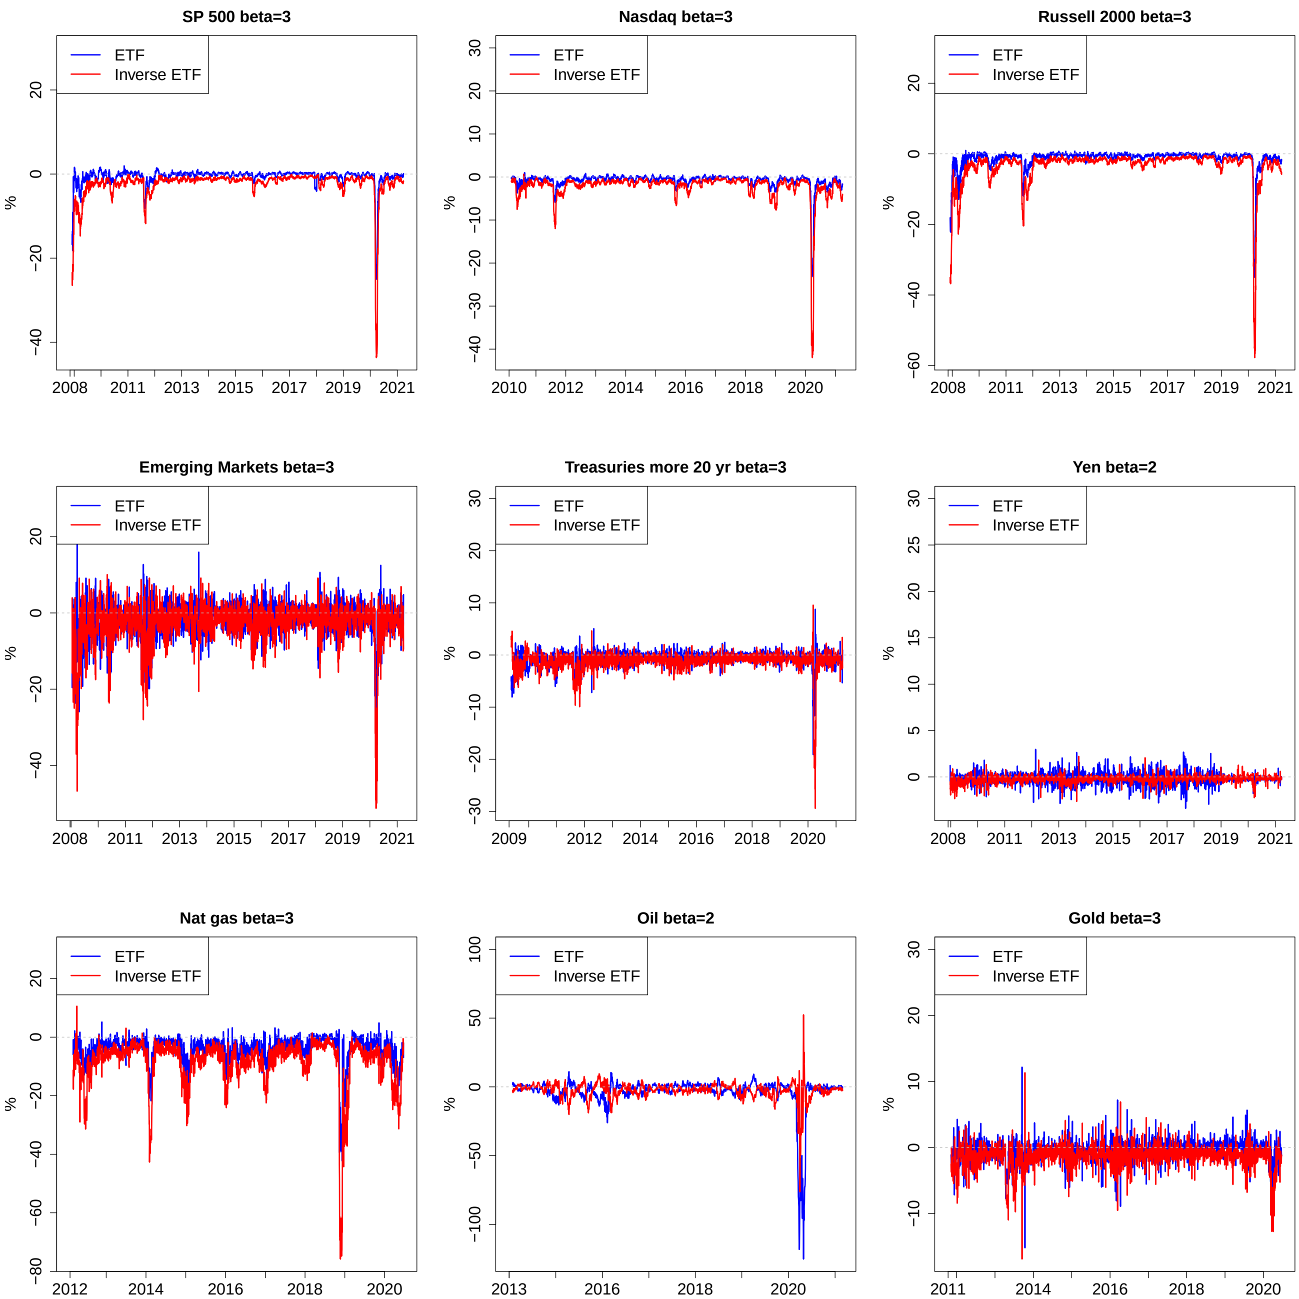

Fig. 5. $CRP(\beta)$ for S&P 500 ($\beta = 3$), Nasdaq ($\beta = 3$), Small cap stocks ($\beta = 3$), Emerging market stocks ($\beta = 3$), Treasuries 20yr+ ($\beta = 3$), Japanese Yen/US Dollar ($\beta = 2$), Natural gas ($\beta = 3$), Oil ($\beta = 2$) and Gold ($\beta = 3$). The figure shows the 12-months rolling annualized $CRP(\beta)$ in %, together with 95% confidence intervals. The $CRP(\beta)$ is estimated as α from regression $r_{\text{ETF},t}(\beta) = \alpha + \beta r_{bmk,t} + (1-\beta) r_{f,t} + \epsilon_t$ for several markets and leverages, where $r_{\text{ETF},t}(\beta)$ is the return on an ETF with leverage β and $r_{bmk,t}$ is the return on the ETF benchmark. We estimate $r_t = \log \mathrm{E}[R_t]$ by first calculating $\mathrm{E}[R_t]$ as the average daily return, and then running monthly regressions of $\log \mathrm{E}[R_{\text{ETF},t}(\beta)]$ on $\log \mathrm{E}[R_{bmk,t}]$.

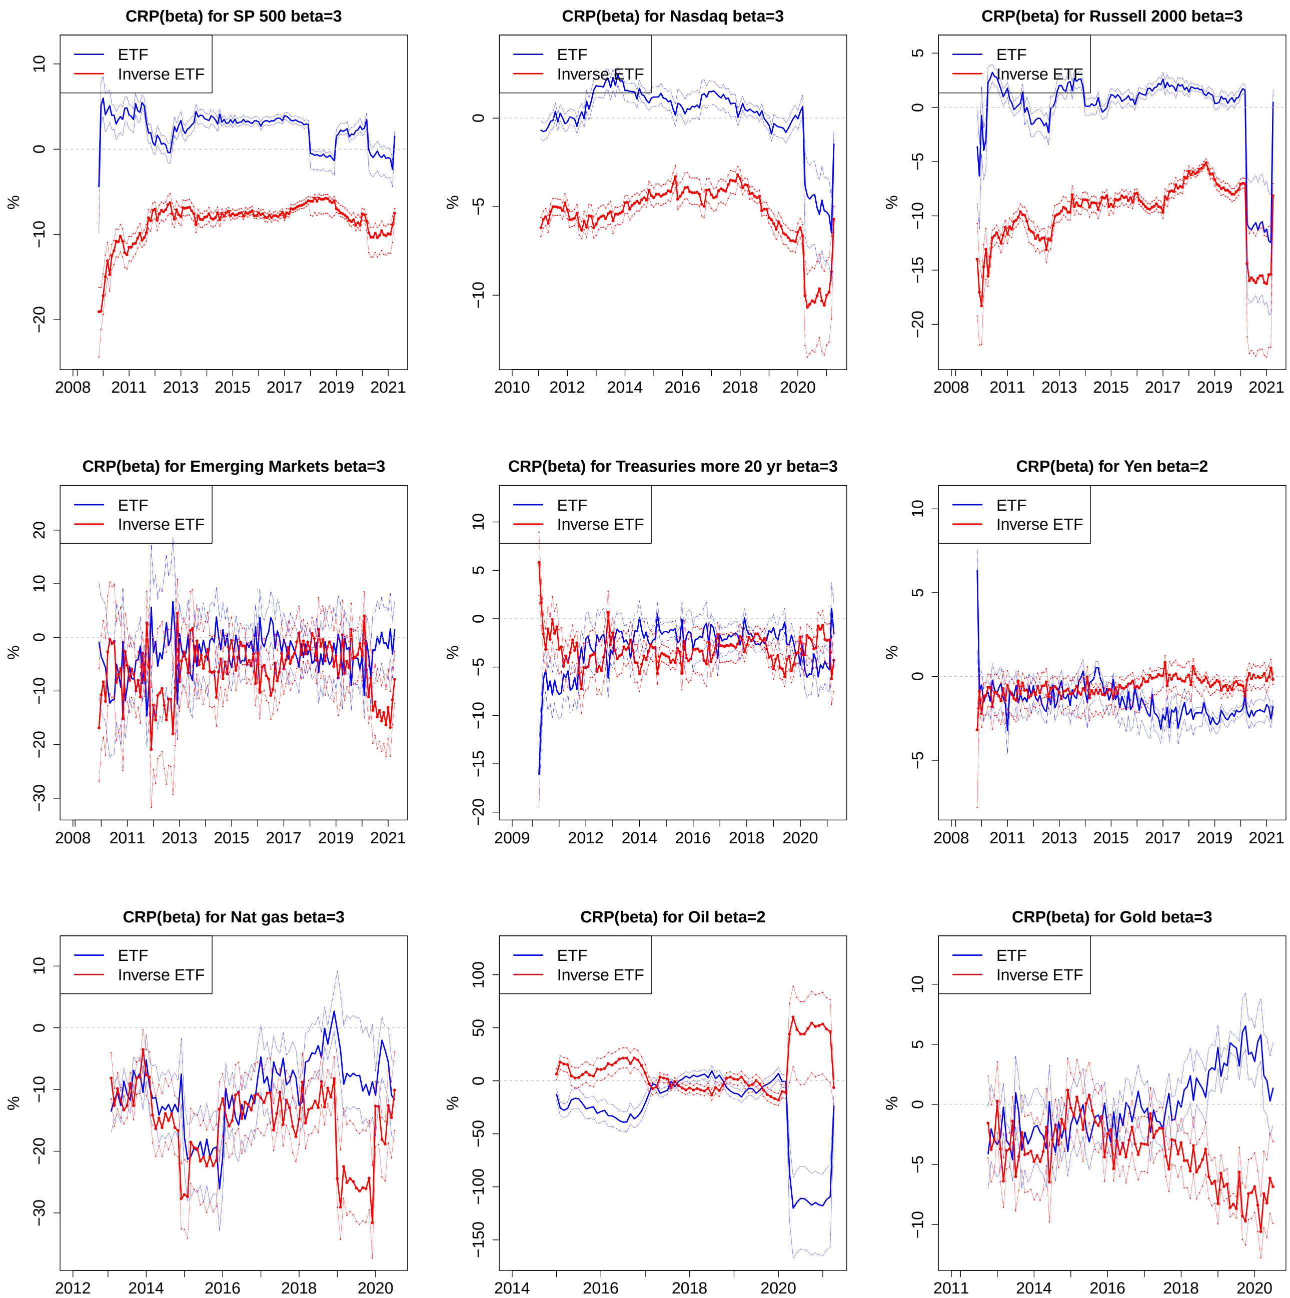

Fig. 6. Returns on the short-both strategy with log-expected-returns (log of the average daily return in a given month) for S&P 500 ($\beta = 3$), Nasdaq ($\beta = 3$), Small cap stocks ($\beta = 3$), Emerging market stocks ($\beta = 3$), Treasuries 20yr+ ($\beta = 3$), Japanese Yen/US Dollar ($\beta = 2$), Natural gas ($\beta = 3$), Oil ($\beta = 2$) and Gold ($\beta = 3$). Plots are 12-months rolling annualized returns, together with 95% confidence intervals.

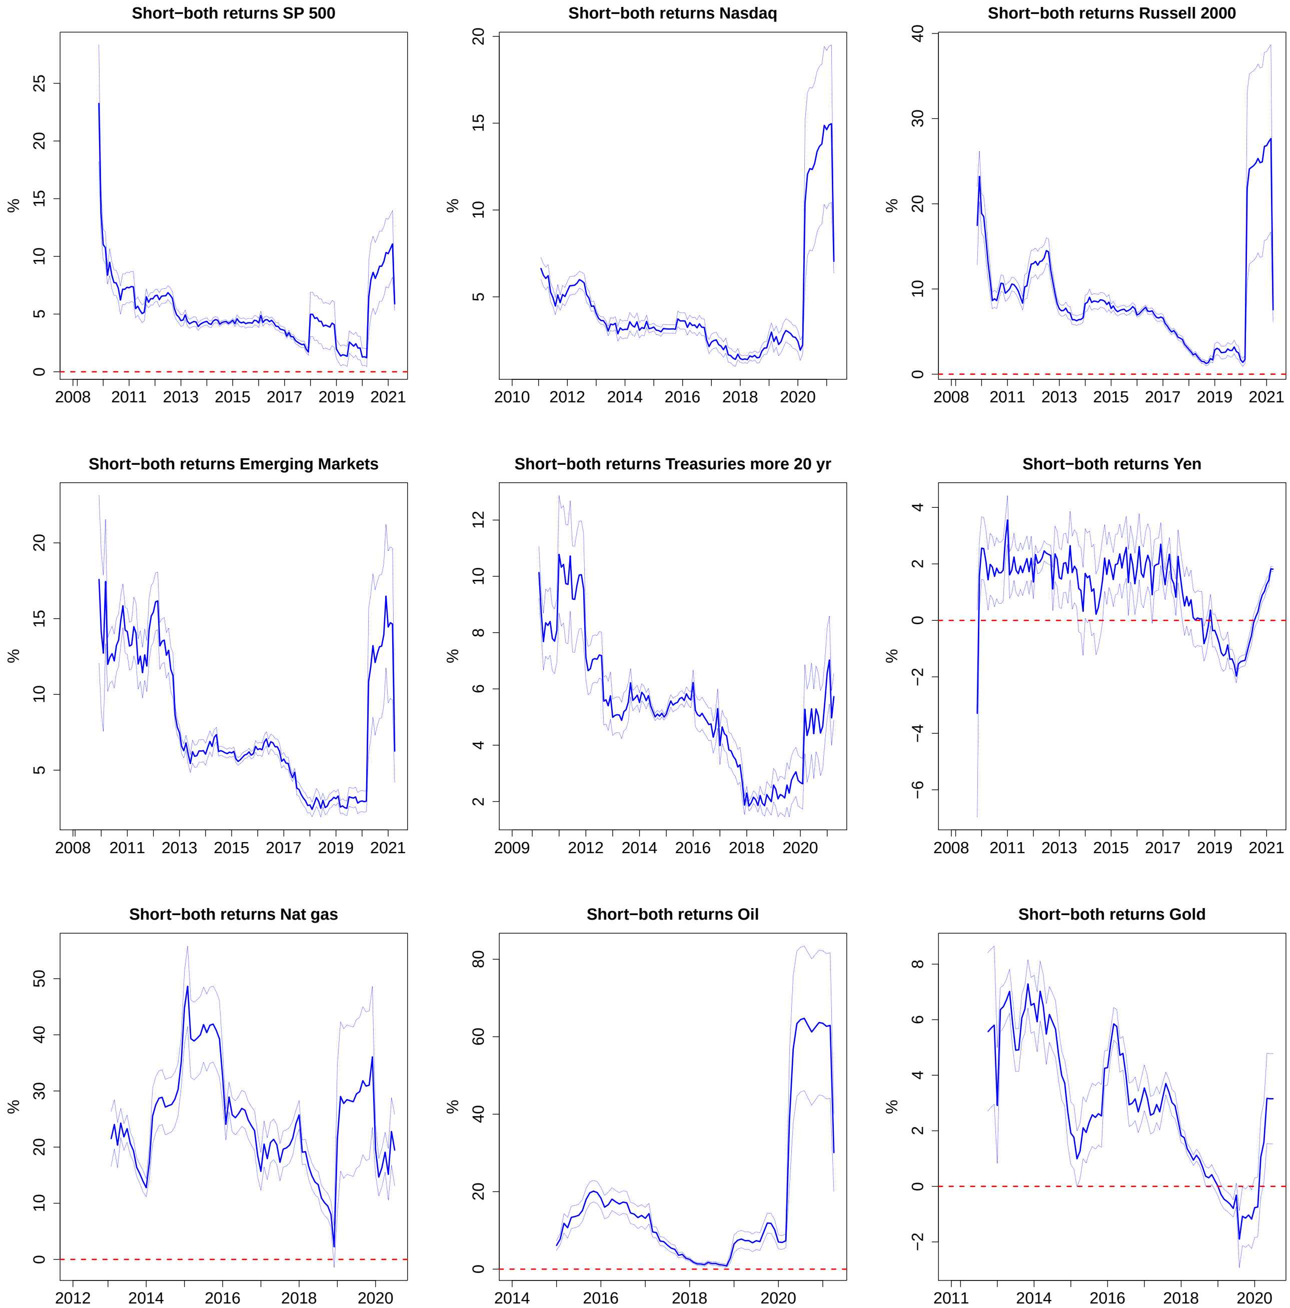

Fig. 7. Returns on the short-both strategy with daily log-returns for S&P 500 ($\beta = 3$), Nasdaq ($\beta = 3$), Small cap stocks ($\beta = 3$), Emerging market stocks ($\beta = 3$), Treasuries 20yr+ ($\beta = 3$), Japanese Yen/US Dollar ($\beta = 2$), Natural gas ($\beta = 3$), Oil ($\beta = 2$) and Gold ($\beta = 3$). Plots are 1-month cumulative returns.

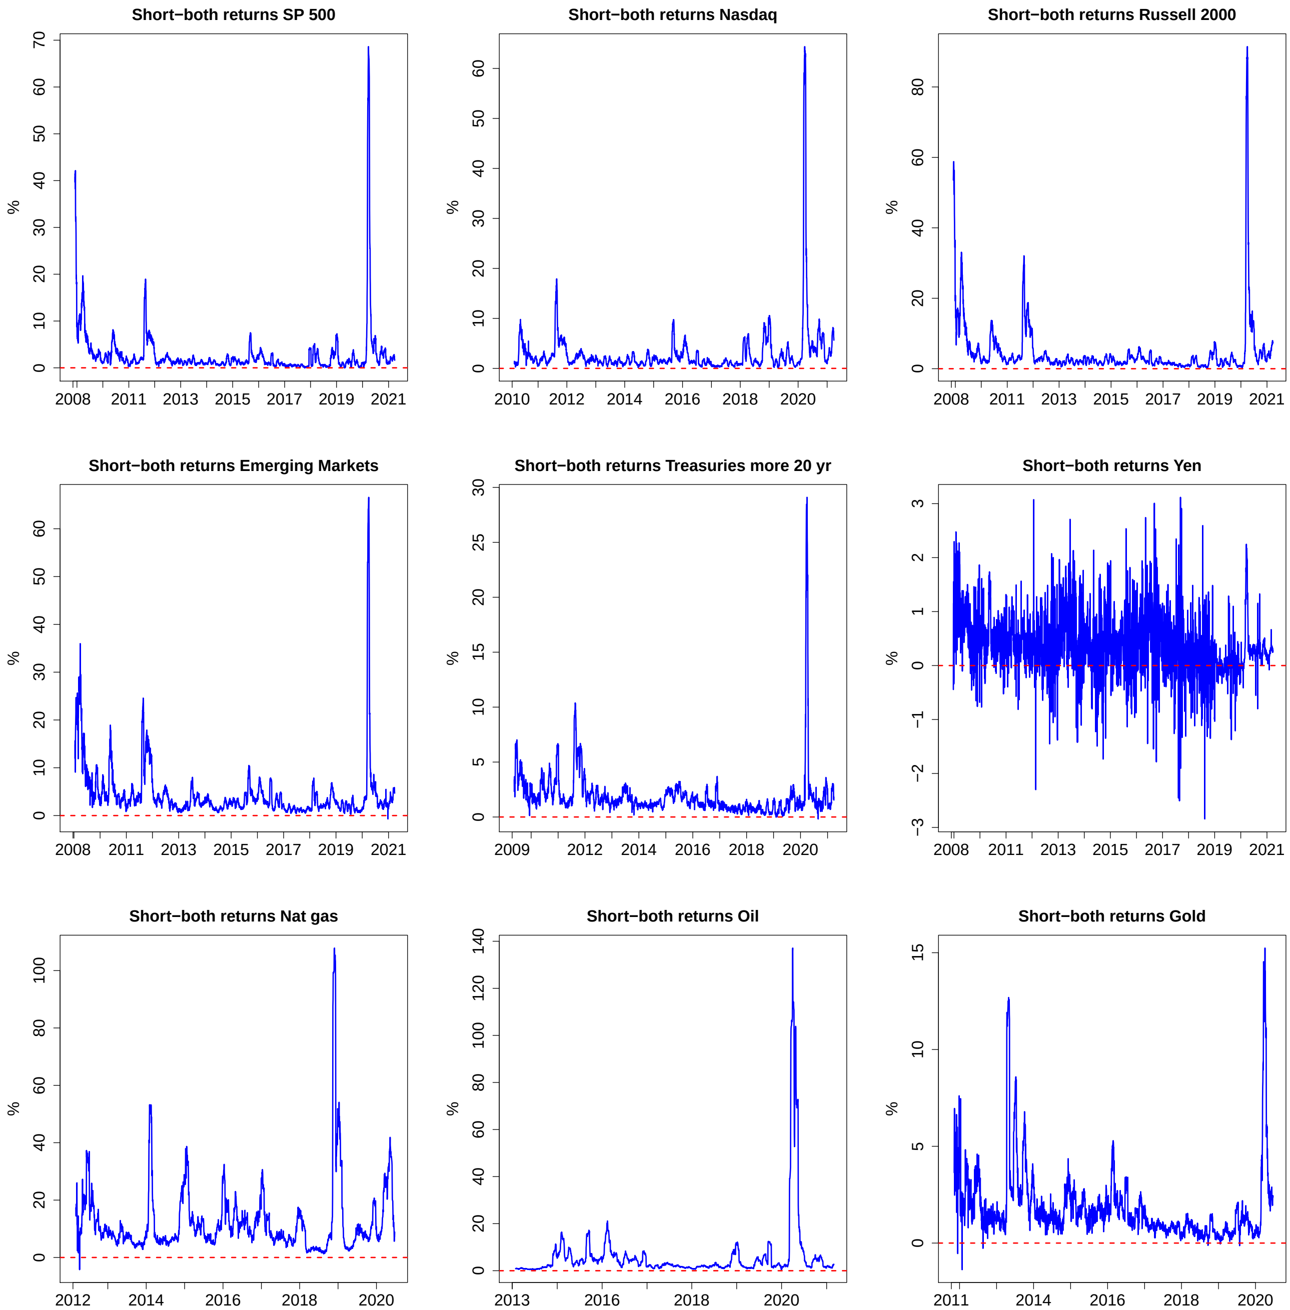

Fig. 8. First principal component and average return of the short-both strategy. The figure shows the first PC of the variance-covariance matrix of short-both strategy returns, VIX (left y-axis), and the average return of the strategy across assets (right y-axis). The plots of the PC1 and the average return are 1-month rolling sums and the maximum values are 308 and 35, respectively (the plots are truncated for better visibility). The assets we use are: S&P 500, Nasdaq, Russell 2000, Financial stocks, Consumer services, Basic materials, Technology, Utilities, Real estate, Emerging market stocks, VIX, Treasuries 7-10 yr, Treasuries more than 20yr, High yield corporate bonds, Japanese Yen/US Dollar, Euro/US Dollar, Natural gas, Oil, Silver and Gold.

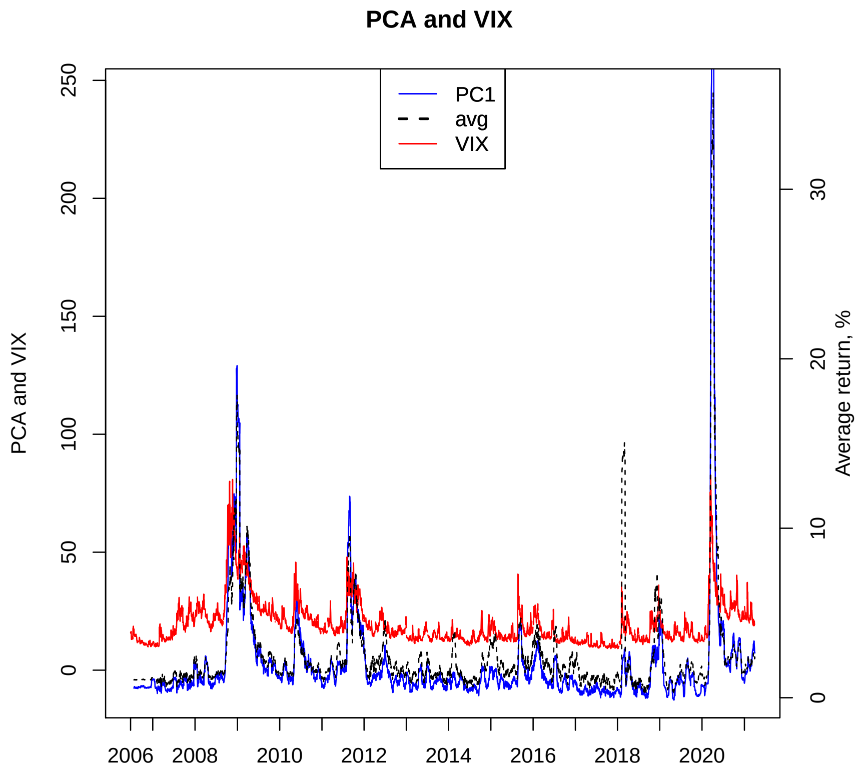

Fig. 9. Loadings of leveraged strategies on higher-order cumulant terms. The graphs show $\beta^n/n!$, which are the loadings on higher-order cumulant differences from Equation 5. The left panel shows these loadings for a leveraged strategy with $\beta = 10$, the right for an unleveraged one with $\beta = 1$.

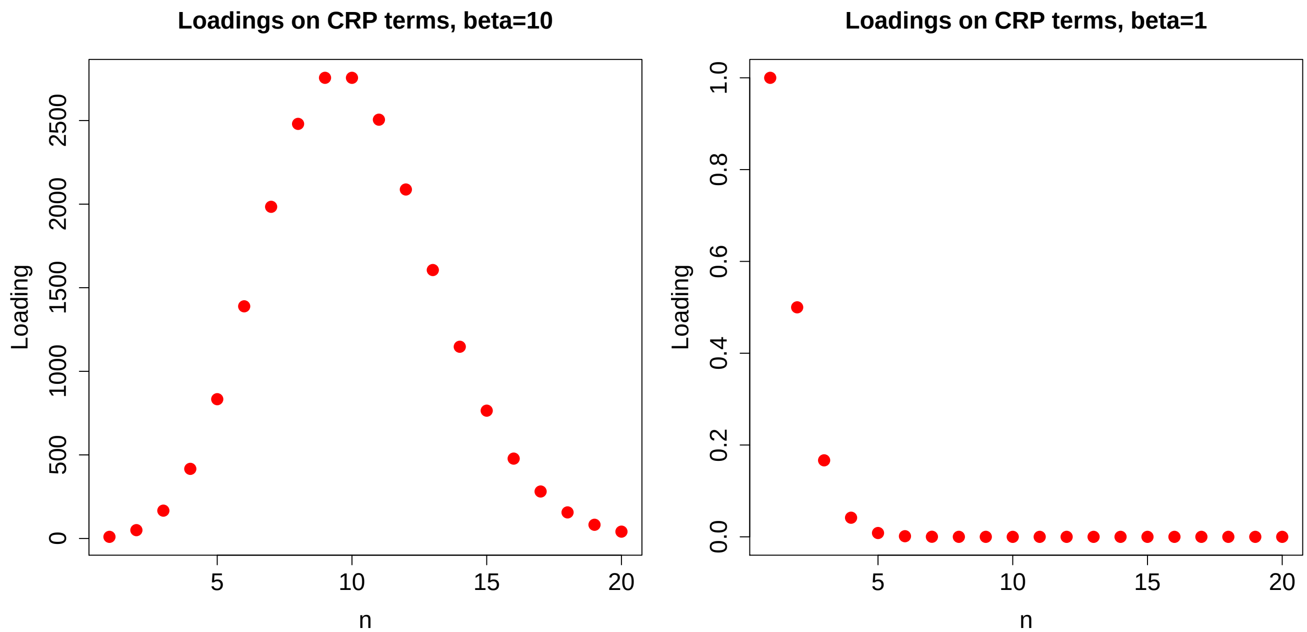

7. Appendix A – derivations and additional robustness checks

A.1. CGF of compound Poisson process with normal-sized jumps

Let $X = \sum_{j=1}^{J} Y_j$, where $J \sim \text{Poisson}(\lambda)$ and Y_j are i.i.d. normal conditional on the number of jumps j: $Y_j / j \sim \mathcal{N}(\mu, \sigma^2)$. Using the independence of Y_j given j, we can write the MGF $G_X(\beta)$ as a function of $G_{Y_j}(\beta) = G_{Y_1}(\beta)$:

$$G_X(\beta) = \sum_{n=0}^{\infty} e^{-\lambda} \frac{\lambda^n (G_{Y_1}(\beta))^n}{n!} = \mathrm{e}^{-\lambda} \mathrm{e}^{\lambda G_{Y_1}(\beta)} = \mathrm{e}^{\lambda(G_{Y_1}(\beta)-1)}.$$

Then, the CGF is

$$c_X(\beta) = \lambda(G_{Y_1}(\beta) - 1). \tag{19}$$

Using the normality of Y_1, we can then write

$$c_X(\beta) = \lambda(\mathrm{e}^{c_{Y_1}(\beta)} - 1) = \lambda(\mathrm{e}^{\mu\beta + \sigma^2\beta^2/2} - 1).$$

A.2. CGF in the setting with lognormal and Poisson component

Let X be a sum of a normal component and an independent Poisson jump: $X = Z + \sum_{j=1}^{J} Y_j$, where $Z \sim \mathcal{N}(\mu - \frac{1}{2}\sigma^2, \sigma^2)$, $J \sim \text{Poisson}(\lambda)$ and Y_j are i.i.d. normal conditional on the number of jumps j: $Y_j / j \sim \mathcal{N}(-b, s^2)$. Both J and Y_j are independent from Z, and hence (subscript in the CGF denotes the respective random variable):

$$c_X(\beta) = c_Z(\beta) + c_{\sum_{j=1}^{J} Y_j}(\beta) = (\mu\beta - \tfrac{1}{2}\sigma^2\beta) + \tfrac{1}{2}\sigma^2\beta^2 + c_{\sum_{j=1}^{J} Y_j}(\beta).$$

Equation 19 shows that the last term is $\lambda(G_{Y_1}(\beta) - 1)$. Using the normality of Y_1, we can then write $c_{\sum_{j=1}^{J} Y_j}(\beta) = \lambda(\mathrm{e}^{-b\beta + s^2\beta^2/2} - 1)$.

A.3. Derivations of constant-β strategies in a GBM setting

The value of a constant-β strategy that invests fraction β in the factor M and the rest in the risk-free rate evolves as:

$$\begin{aligned}
&\frac{\mathrm{d}P_t}{P_t} = \beta \frac{\mathrm{d}P_{t,\mathrm{M}}}{P_{t,\mathrm{M}}} + (1-\beta) r_f \,\mathrm{d}t \\
&\mathrm{d}\log P_t = (\beta\mu - \tfrac{1}{2}\beta^2\sigma^2 + (1-\beta) r_f)\,\mathrm{d}t + \beta\sigma\,\mathrm{d}B_t \\
&\iff P_T = P_0 \mathrm{e}^{(\beta\mu - \frac{1}{2}\beta^2\sigma^2 + (1-\beta) r_f) T + \beta\sigma B_T} \\
&\iff P_T = P_0 (\frac{P_{\mathrm{M},T}}{P_{\mathrm{M},0}})^{\beta} \mathrm{e}^{((1-\beta) r_f - \frac{1}{2}\beta(\beta-1)\sigma^2) T}
\end{aligned} \tag{20}$$

The last line is obtained from the previous one by adding and subtracting $\frac{1}{2}\beta\sigma^2 T$ in the power of e.

$$
\begin{aligned}
R_T &= R_{\mathrm{M},T}^{\beta} e^{(1-\beta)r_{f,T}-\frac{1}{2}\beta(\beta-1)\sigma_T^2} \\
\log R_T &= \beta \log R_{\mathrm{M},T} + (1-\beta)\log R_{f,T} - \underbrace{\tfrac{1}{2}\beta(\beta-1)\sigma_T^2}_{\text{deterministic}}.
\end{aligned} \tag{21}
$$

Note that the CAPM holds in this case with log-returns: taking expectations in the last line of Equation 20 yields $\mathrm{E}[R_T] = \mathrm{e}^{\beta\mu T + (1-\beta)r_f T}$, or $\log \mathrm{E}[R_T] = \beta \log \mathrm{E}[R_{\mathrm{M},T}] + (1-\beta)r_f T$.

A.4. CRP in a setting with stochastic volatility: the Heston model

Let us derive the CRP in a setting with stochastic volatility (e.g., Heston (1993)). The log-factor price process $x_t = \log P_{t,M}$ follows:

$$
\begin{aligned}
\mathrm{d}x_t &= (\mu - \tfrac{1}{2}v_t)\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}B_t^1 \\
\mathrm{d}v_t &= \lambda(\bar{v} - v_t)\,\mathrm{d}t + \sigma\sqrt{v_t}\,\mathrm{d}B_t^2
\end{aligned} \tag{22}
$$

under the physical measure, where v_t is the volatility of the log-factor price, λ is the mean-reversion speed, $\bar{v}$ is the long-term mean of volatility, σ is now the volatility of volatility, and B_t^1, B_t^2 are correlated Brownian motions $\mathrm{d}B_t^1\,\mathrm{d}B_t^2 = \rho\,\mathrm{d}t$.

By applying the Feynman-Kac theorem to the characteristic function $\psi(\beta) = \mathrm{E}[\mathrm{e}^{i\beta x_T}]$, we get a partial differential equation (PDE) for the MGF $G(\beta, x_t, v_t, t, T) = \psi(-i\beta) = \mathrm{E}[\mathrm{e}^{\beta x_T}]$:[15]

$$
G_t + (\mu - \tfrac{1}{2}v_t)G_x + \lambda(\bar{v} - v_t)G_v + \tfrac{1}{2}v_t G_{xx} + \tfrac{1}{2}\sigma^2 v_t G_{vv} + \rho\sigma v_t G_{xv} = 0 \tag{23}
$$

with a boundary condition $G_T = \mathrm{e}^{\beta x_T}$.

We guess the solution is exponentially affine of the form:

$$
G = \mathrm{e}^{\beta x_t + \mu\beta(T-t) + a(\beta,t,T) + b(\beta,t,T)v_t}. \tag{24}
$$

Substituting this form in Equation 23, simplifying and regrouping with respect to the state variable v_t, we obtain (we write b instead of $b(\beta,t,T)$, b_t for $\frac{\partial b}{\partial t}$, and similarly for a for ease of notation):

$$
\begin{aligned}
&-\mu\beta + a_t + b_t v_t + (\mu - \tfrac{1}{2}v_t)\beta + \lambda(\bar{v} - v_t)b + \tfrac{1}{2}v_t\beta^2 + \tfrac{1}{2}\sigma^2 v_t b^2 + \rho\sigma v_t\beta b = 0, \\
&v_t(b_t - \tfrac{1}{2}\beta - \lambda b + \tfrac{1}{2}\beta^2 + \tfrac{1}{2}\sigma^2 b^2 + \rho\sigma\beta b) + a_t + \lambda\bar{v}b = 0
\end{aligned} \tag{25}
$$

[15] It is easier to work with MGF than the CGF since MGF is a simpler function of the characteristic function, whereas CGF involves the log, and the derivations are more algebraically complex. It is easier to solve using the MGF and then apply the log to the solution to obtain the CGF.

By matching the powers of v_t on the LHS and the RHS, we obtain two ODEs:

$$\begin{aligned} a_t &= -\lambda\bar{v}b, \\ b_t &= \tfrac{1}{2}\beta - \tfrac{1}{2}\beta^2 + (\lambda - \rho\sigma\beta)b - \tfrac{1}{2}\sigma^2 b^2. \end{aligned} \tag{26}$$

The second ODE is a general Ricccati equation, which can be solved in a standard way using the boundary condition for the particular solution. By substituting the solution in the first ODE, one then obtains $a(\beta, t, T)$. The final solutions are:

$$\begin{aligned} a(\beta, t, T) &= -\frac{\lambda\bar{v}\phi}{\sigma^2}\left(\phi - (\lambda - \rho\sigma\beta)\right)(T-t) + 2\log\frac{\phi + (\lambda - \rho\sigma\beta) + (\phi - (\lambda - \rho\sigma\beta))\mathrm{e}^{-\phi(T-t)}}{2\phi} \\ b(\beta, t, T) &= (\beta^2 - \beta)\frac{1 - \mathrm{e}^{-\phi(T-t)}}{\phi + (\lambda - \rho\sigma\beta) + (\phi - (\lambda - \rho\sigma\beta))\mathrm{e}^{-\phi(T-t)}}, \end{aligned} \tag{27}$$

where

$$\phi = \sqrt{(\lambda - \rho\sigma\beta)^2 + \sigma^2(\beta - \beta^2)}. \tag{28}$$

They satisfy the boundary condition at $t = T$. Then, the CGF of the log-return (hence skipping x_t as a parameter and subtracting $\log\mathrm{e}^{x_t\beta}$ from the MGF of the log-factor price) is:

$$\begin{aligned} c(\beta, v_t, t, T) &= \log G(\beta, x_t, v_t, t, T) - \log\mathrm{e}^{x_t\beta} \\ &= \mu\beta(T-t) + a(\beta, t, T) + b(\beta, t, T)v_t. \end{aligned} \tag{29}$$

Now let us we evaluate these expressions for $\beta = 1$ since we need the CGF at $\beta = 1$.

Since $b(1, t, T) = 0, \phi(\beta = 1) = \lambda - \rho\sigma$, and $a(1, t, T) = -\frac{\lambda\bar{v}\phi}{\sigma^2}0 + 2\log\frac{2\phi}{2\phi} = 0$, $a(1) = b(1) = 0$. Take $t = 0$. Then:

$$c_T(1) - c_T^*(1) = (\mu - r_f)T \tag{30}$$

This result suggests that the FRP in the Heston model does not depend on variance. In other words, even if volatility is stochastic, the FRP captures just the difference between the physical drift (μ) and the risk-neutral one r_f. The difference in risk-neutral and physical parameters of the Heston model is irrelevant for the FRP since for $\beta = 1$, $b(1, t, T) = b^*(1, t, T) = 0$ and the multiplier of the stochastic volatility v_t in $c(1, v_t, t, T)$ is zero. However, the FRP for a general leveraged asset is different from zero since these assets load on the variance risk premium through their leveraged exposure.

A.5. CRP with power utility

The compensation for higher-order cumulant risk measured by the CRP can be computed in standard economic models, for example in a setting with power utility (over the log-factor return instead of consumption growth). With a risk-aversion parameter γ, the risk-neutral cumulant generating function can be expressed as a function of the physical one: $c^*(\beta) = c(\beta - \gamma) - c(-\gamma)$ (see, e.g., Backus et al. (2011) for the derivation). Then, we can write the CRP as:

$$CRP_T = c_T(1) + c_T(-\gamma) - c_T(1-\gamma) - \left(\mathrm{E}[\log R_T] - r_{f,T}\right). \tag{31}$$

In the particular case when $\gamma = 1$ (log-utility), CRP is:

$$CRP_T = c_T(1) + c_T(-1) - \left(\mathrm{E}[\log R_T] - r_{f,T}\right) = 2 \sum_{n=2,\text{even}}^{\infty} \frac{\kappa_{n,T}}{n!} - \left(\mathrm{E}[\log R_T] - r_{f,T}\right), \tag{32}$$

which is closely related to the returns on the short-both strategy for $\beta = 1$ that give $2\sum_{n=2,\text{even}}^{\infty} \frac{\kappa^*_{n,T} - \kappa_{n,T}}{n!}$.

A.6. Measuring the CRPO

We can also construct a bet on implied vs. realized odd-order cumulants by buying an ETF and selling its opposite ETF. This strategy extracts the β^n-weighted CRPO plus the log risk premium LRP ($\mathrm{E}[\log R_T] - \mathrm{E}^*[\log R_T] = \kappa_{1,T} - \kappa^*_{1,T}$), since the exposure to even-order cumulants cancels out. The returns on this "short-one" strategy are:

$$r_{\mathrm{SO},T} = 2 \sum_{\mathrm{n} \geq 1,\ \mathrm{odd}}^{\infty} \frac{\beta^n (\kappa_{n,T} - \kappa^*_{n,T})}{n!}. \tag{33}$$

Let us denote $CRPO_T(\beta) = \frac{1}{2} r_{\mathrm{SO},T}$. Table B.5 shows the summary stats for the returns on the short-one strategy and the estimate of the $CRPO_T(\beta)$. The results show that $CRPO_T(\beta)$ is negative for natural gas, oil, currencies and high yield bonds, just like the $CRPE_T(\beta)$. However, in contrast to the $CRPE_T(\beta)$, $CRPO_T(\beta)$ is positive for the S&P 500, most equity sectors, VIX, and emerging market equities. These facts show that the β^n-weighted mixture of cumulants can be positive or negative for assets with different loadings on the same factor: higher-order cumulants matter not only through the CRP of the factor but also through the sign of β.

8. Appendix B – Additional tables and figures

Table B.1
$CRP(\beta)$ as a share of the factor risk premium (FRP). $CRP(\beta)$ is estimated as α from the regression $r_{\text{ETF},t}(\beta) = \alpha + \beta r_{bmk,t} + (1-\beta) r_{f,t} + \epsilon_t$ for several markets and leverages, where $r_{\text{ETF},t}(\beta)$ is the return on an ETF with leverage β, $r_{bmk,t}$ is the return on the ETF benchmark, and β is the ETF leverage. We estimate $r_t = \log \text{E}[R_t]$ by first calculating $\text{E}[R_t]$ as the average daily return, and then running monthly regressions of $\log \text{E}[R_{\text{ETF},t}(\beta)]$ on $\log \text{E}[R_{bmk,t}]$. The numbers in the table are the ratios of the $CRP(\beta)$ to the FRP.

			$\frac{CRP(\beta)}{FRP}$		
Leverage= β	-3	-2	-1	2	3
S&P 500	-1.28	-0.95	-0.55	0.19	0.20
Nasdaq	-0.54				-0.00
Russell 2000	-0.98				-0.05
Financials	-1.33				-0.21
Consumer services		-0.73		-0.00	
Basic materials		-1.01		-0.10	
Technology	-0.14				-0.36
Utilities		-3.49		0.68	
Industrials		-1.07		0.04	
Real Estate	0.11	-1.32		0.28	-1.47
Emerging Markets	-0.83	-1.19	1-.01	0.47	-0.48
VIX			1.05	0.89	
Treasuries 7-10 yr	-0.75	-0.51	-0.35	-0.50	-1.17
Treasuries more 20 yr	-0.42	-0.38	-0.29	-0.32	-0.73
High Yield			-0.48	-0.25	
Gold	-0.83	-0.79		-0.20	-0.09
Silver	2.00	2.41		0.88	0.87
Nat gas	0.82	0.45		0.24	0.64
Oil	3.49	0.75		-2.80	-6.29
Euro/US Dollar		1.36		1.11	
Yen/US Dollar		-1.33		-1.90	

Table B.2
Starting dates of the ETFs used in the short-both strategy. The table shows the first date when a long and inverse ETF with leverage of β and $-\beta$, respectively become available in a given asset.

Asset	β	Starting date
S&P 500	1	22/06/2006
S&P 500	3	06/11/2008
Nasdaq	3	12/02/2010
Russell 2000	3	06/11/2008
Financials	3	07/11/2008
Consumer services	2	02/02/2007
Basic materials	2	02/02/2007
Technology	3	18/12/2008
Utilities	2	02/02/2007
Industrials	2	02/02/2007
Real estate	3	17/07/2009
Emerging markets	1	02/11/2007
Emerging markets	3	18/12/2008
VIX	1	05/01/2011
Treasuries 7-10 yr	1	05/04/2011
Treasuries 7-10 yr	3	17/04/2009
Treasuries 20 yr+	1	21/08/2009
Treasuries 20 yr+	3	17/04/2009
High yield	1	23/03/2011
Gold	2	04/12/2008
Silver	3	01/01/2009
Nat gas	3	08/02/2012
Oil	3	06/01/2017
Euro/US Dollar	2	26/11/2008
Yen/US Dollar	2	26/11/2008

Table B.3
Short-both strategy with fees. The table shows the short-both strategy using before-fees returns. Columns 2–4 use log-returns, 5–7 simple returns. Columns 2 and 5 are in basis points, whereas columns 3, 4, and 6 in %.

	β (1)	Mean log daily (2)	Mean log, annual (3)	Mean CRPE, annual (4)	Mean simple, daily (5)	Mean simple, annual (6)	Sharpe ratio (7)
S&P 500	1	2.09	5.23	-2.61	0.49	1.23	0.06
S&P 500	3	14.03	35.08	-17.54	1.35	3.43	0.73
Nasdaq	3	13.92	34.80	-17.40	0.66	1.66	0.41
Russell 2000	3	23.17	57.93	-28.96	2.35	6.05	1.12
Financials	3	30.46	76.15	-38.08	3.88	10.18	1.27
Consumer services	2	7.33	18.33	-9.16	1.15	2.92	0.09
Basic materials	2	14.5	36.25	-18.13	2.28	5.86	0.36
Technology	3	17.05	42.63	-21.31	1.77	4.52	0.77
Utilities	2	9.14	22.85	-11.43	2.99	7.76	0.38
Industrials	2	10.03	25.08	-12.54	1.75	4.47	0.18
Real Estate	3	19.07	47.68	-23.84	2.63	6.80	0.55
Emerging Markets	1	4.4	11.00	-5.50	0.74	1.87	0.24
Emerging Markets	3	21.9	54.75	-27.38	1.85	4.73	0.71
VIX	1	32.08	80.20	-40.10	7.2	19.71	0.49
Treasuries 7-10 yr	1	0.59	1.48	-0.74	0.44	1.11	0.16
Treasuries 7-10 yr	3	3.16	7.90	-3.95	1.58	4.03	0.27
Treasuries more 20 yr	1	1.71	4.28	-2.14	0.91	2.30	0.66
Treasuries more 20 yr	3	8.8	22.00	-11.00	1.2	3.05	0.71
High Yield	1	2.59	6.48	-3.24	2.29	5.89	1.87
Gold	2	4.76	11.90	-5.95	0.55	1.38	0.31
Silver	3	31.13	77.83	-38.91	4.53	11.99	0.25
Nat gas	3	60.85	152.13	-76.06	6.42	17.40	1.07
Oil	3	45.3	113.25	-56.63	6.16	16.64	0.26
Euro/US Dollar	2	1.01	2.53	-1.26	-0.35	-0.87	-0.27
Yen	2	1.14	2.85	-1.43	-0.3	-0.75	-0.18

Table B.4

$CRP(\beta)$ with fees. The table shows the annualized $CRP(\beta)$ in %, estimated as α from the regression $r_{\text{ETF},t}(\beta) = \alpha + \beta r_{bmk,t} + (1-\beta) r_{f,t} + \epsilon_t$ for several markets and leverages, where $r_{\text{ETF},t}(\beta)$ is the return on an ETF with leverage β, $r_{bmk,t}$ is the return on the ETF benchmark, and β is the ETF leverage. We estimate $r_t = \log \mathrm{E}[R_t]$ by first calculating $\mathrm{E}[R_t]$ as the average daily return (before fees), and then running monthly regressions of $\log \mathrm{E}[R_{\text{ETF},t}(\beta)]$ on $\log \mathrm{E}[R_{bmk,t}]$. All estimates are significantly different from zero at the 5% level except those in *italics*. Here and in the subsequent analysis $r_{f,t}$ is the 1-month Treasury rate and standard errors are computed using the Newey and West (1987) estimator with lag selection based on the Bartlett kernel (e.g., Andrews (1991)). Daily frequency, from the first leveraged ETF inception date in a given market to April 2021 (February 2018 for VIX, June 2020 for gold and gas since some long and inverse ETFs were delisted).

	Annualized $CRP_T(\beta)$ in %				
Leverage= β	-3	-2	-1	2	3
S&P 500	-8.05	-5.88	-3.08	2.24	2.38
Nasdaq	-4.59				0.88
Russell 2000	-9.32				*0.53*
Financials	-11.82				*-1.11*
Consumer services		-5.77		*0.96*	
Basic materials		-7.95		*0.04*	
Technology	*-1.28*				-4.87
Utilities		-12.47		3.50	
Industrials		-7.53		*1.25*	
Real Estate	*2.08*	-12.18	-4.20	3.67	-13.67
Emerging Markets	*-4.80*	-7.51	-6.26	4.24	*-2.20*
VIX			-35.84	-29.67	
Treasuries 7-10 yr	-1.78	-1.02	*-0.41*	-1.75	-3.33
Treasuries more 20 yr	-1.42	-1.93	-0.72	*-0.91*	-3.11
High Yield			-1.63	*-0.41*	
Gold	-2.69	-2.53		*-0.05*	*0.85*
Silver	-7.32	-9.38		*-3.00*	*-2.28*
Nat gas	-11.21	-5.64		*-2.51*	-8.38
Oil	31.81	7.34		-23.18	-52.70
Euro/US Dollar		*-0.18*		*0.05*	
Yen/US Dollar		*0.27*		*-0.08*	

Table B.5

Returns of the short-one strategy with log-returns. The table shows summary statistics for the short-one strategy $r_{\mathrm{SO},T}$. The last column is the average $r_{\mathrm{SO},T}/2 = CRPO_T(\beta)$ minus column 7 LRP_T ($= \kappa_{1,T} - \kappa^*_{1,T}$ calculated as $\mathrm{E}[\log R_T] - \mathrm{E}\left[\log R_{f,T}\right]$). The numbers in the table are in basis points. Daily frequency, from the first leveraged ETF inception date in a given market to April 2021 (February 2018 for VIX, June 2020 for gold and gas since some long and inverse ETFs were delisted).

Asset	β (1)	Mean (2)	S.d. (3)	Median (4)	Min (5)	Max (6)	LRP_T (7)	Mean $CRPO_T(\beta) - LRP_T$ (8)
S&P 500	1	8.52	251.13	9.08	-2248.35	2500.37	2.56	1.70
S&P 500	3	33.95	712.51	32.22	-7049.16	5613.44	4.19	12.79
Nasdaq	3	47.11	727.52	54.87	-7058.27	5723.65	6.82	16.74
Russell 2000	3	36.89	913.52	41.10	-8109.58	5337.28	4.23	14.22
Financials	3	37.54	1033.85	27.55	-9078.44	9461.06	3.55	15.22
Consumer services	2	20.86	472.89	18.57	-3915.47	3973.73	3.64	6.79
Basic materials	2	16.4	693.48	22.31	-5101.35	5223.2	1.54	6.66
Technology	3	48.27	781.52	49.79	-7308.02	6901.42	7.26	16.88
Utilities	2	14.90	482.89	18.78	-3972.98	5569.37	1.18	6.27
Industrials	2	19.44	559.89	20.47	-4653.39	5461.23	2.68	7.04
Real estate	3	34.28	809.29	40.53	-9316.34	5945.92	1.10	16.04
Emerging markets	1	5.08	381.36	6.26	-3625.65	4638.08	-0.17	2.71
Emerging markets	3	23.62	896.3	27.74	-7249.14	5545.77	3.38	8.43
VIX	1	-24.08	992.36	-85.77	-2645.25	25674.21	-28.36	16.32
Treasuries 7-10 yr	1	2.53	73.45	1.39	-441.35	424.8	1.47	-0.21
Treasuries 7-10 yr	3	7.68	237.88	2.46	-1573.4	1554.27	1.38	2.46
Treasuries more 20 yr	1	4.53	177.89	4.77	-1300.79	1510.68	2.46	-0.20
Treasuries more 20 yr	3	13.27	550.82	21.77	-3812.67	4162.12	2.20	4.44
High yield	1	2.76	108.9	2.33	-1191.9	1382.29	1.95	-0.57
Gold	2	10.65	409.87	4.60	-3536.27	2404.64	2.04	3.29
Silver	3	-0.41	1010.5	0.00	-7891.06	4843.32	-2.11	1.91
Nat gas	3	-38.83	1475.96	-13.93	-13273.22	12788.54	-10.09	-9.33
Oil	3	-40.14	1206.98	0.00	-16344.26	8073.39	-8.17	-11.90
Euro/US Dollar	2	-1.64	231.11	0.00	-1283.22	1532.62	-0.40	-0.42
Yen/US Dollar	2	-2.34	235.53	0.00	-1362.25	1589.58	-0.47	-0.70

Fig. B.1. Trading against assets with opposite β-s: extracting the $CRPE(\beta)$. The figure shows the profit dynamics of liquidity provision to assets with opposite β-s using a binomial tree example. The figure illustrates the dynamics of the factor and the corresponding profits for a market-maker who sells short a pair of assets with opposite β-s ($\beta = 2$ and $\beta = -2$). For each period, the parameters of the tree are $u = 1.05$ (gross return in the up-state) and $ud = 1$, where d is gross return in the down-state. Red areas indicate nodes where the market-maker loses money, and green ones show where she makes profit. More color-intense nodes indicate larger losses or profits.

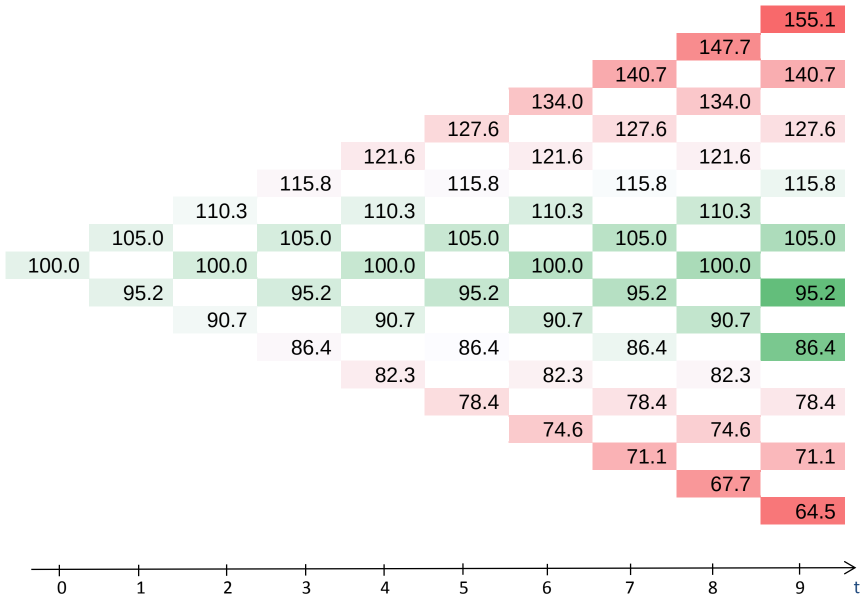

Fig. B.2. Under-performance of the ETF with simple returns for S&P 500 ($\beta = 3$), Nasdaq ($\beta = 3$), Small cap stocks ($\beta = 3$), Emerging market stocks ($\beta = 3$), Treasuries 20yr+ ($\beta = 3$), Japanese Yen/US Dollar ($\beta = 2$), Natural gas ($\beta = 3$), Oil ($\beta = 2$) and Gold ($\beta = 3$). The graphs show 30-day cumulative differences between the return on the ETF and the sum of the return on the leveraged benchmark and the risk-free rate: $r_{ETF} - (\beta r + (1-\beta) r_f)$. Blue lines are long ETFs ($\beta > 0$), red lines are inverse ETFs ($\beta < 0$).

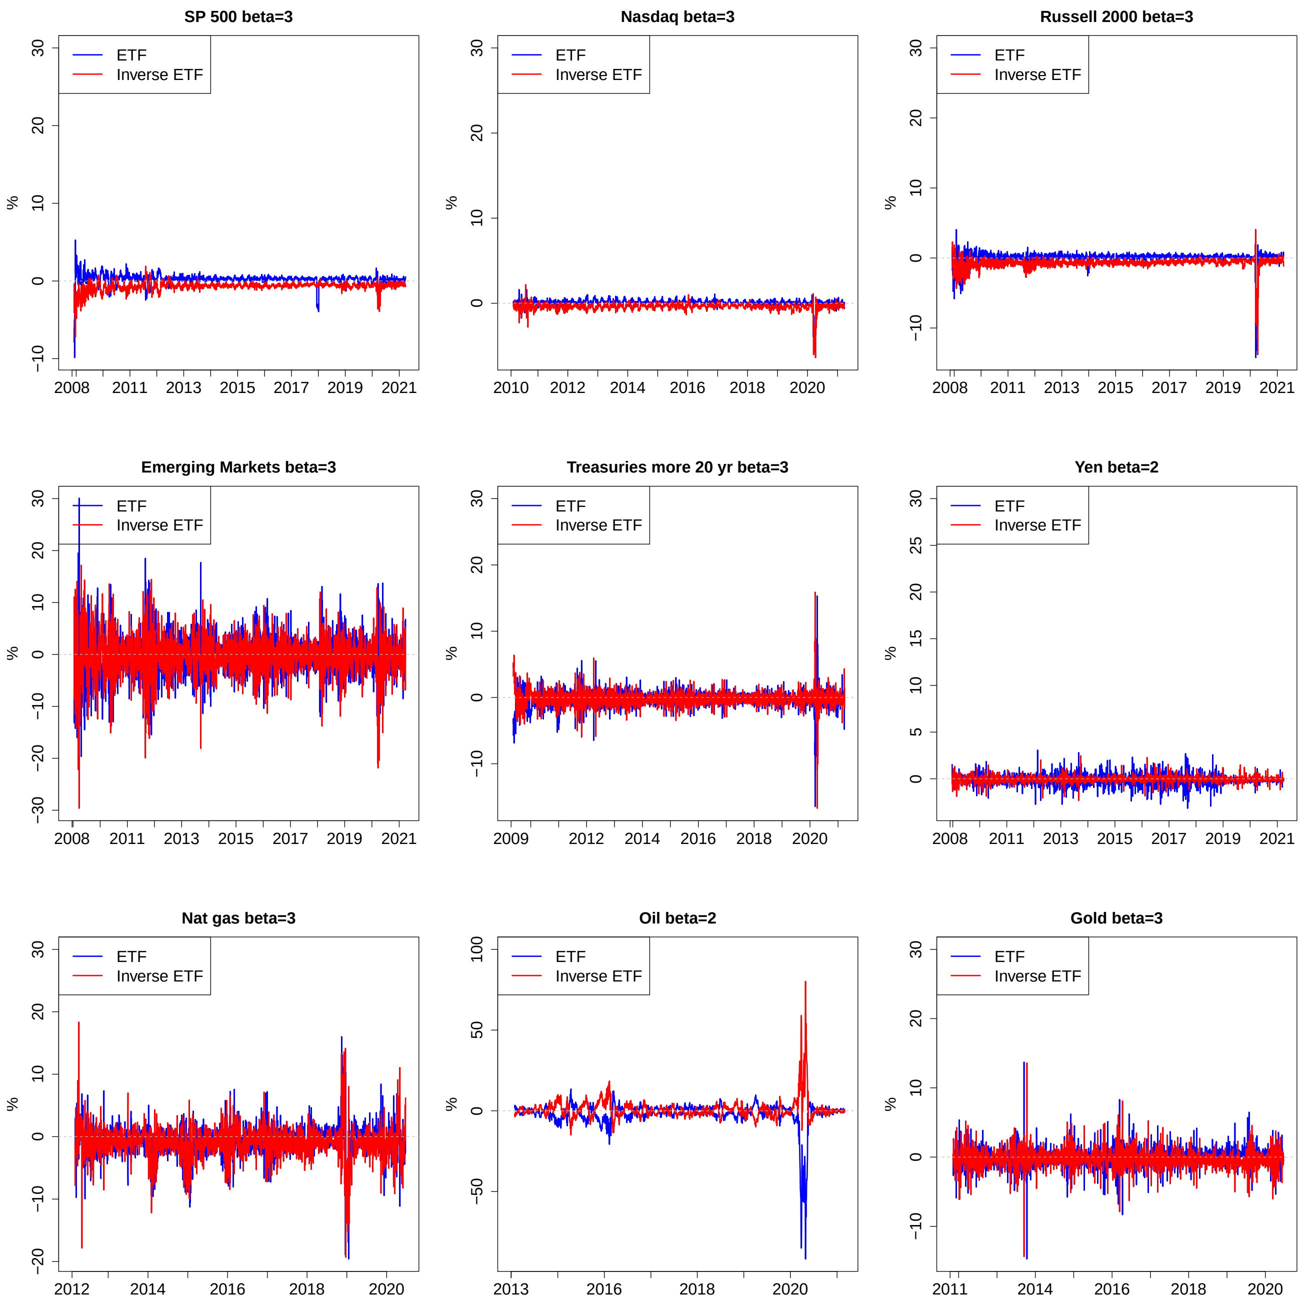

Fig. B.3. The returns on the short-both strategy with daily simple returns for S&P 500 ($\beta = 3$), Nasdaq ($\beta = 3$), Small cap stocks ($\beta = 3$), Emerging market stocks ($\beta = 3$), Treasuries 20yr+ ($\beta = 3$), Japanese Yen/US Dollar ($\beta = 2$), Natural gas ($\beta = 3$), Oil ($\beta = 2$) and Gold ($\beta = 3$). Plots are 30-day cumulative returns (not annualized).