



Quantifying the non-Gaussian gain

David Allen^{1,2} · Stephen Satchell³ · Colin Lizieri⁴

Revised: 6 October 2023 / Accepted: 9 October 2023 / Published online: 28 November 2023
© The Author(s) 2023

Abstract

In this paper, we quantify the economic gain from accounting for departures from normality for the mean-variance (MV) investor. We provide two models that account for the key empirical regularities of financial returns: stochastic volatility, asymmetric returns, heavy tails and tail dependence. We show that accounting for departures from normality leads to significant gains in expected utility commensurate with or exceeding typical active management fees. The majority of the uplift in expected utility derives from accounting for stochastic volatility.

Keywords Asset allocation · Non-normality · Asymmetry · Tail dependence · Expected utility

Introduction

The 2008–2009 and more recent financial crises have spurred renewed interest in risk management and approaches to capital allocation. The Markowitz (1952) approach, for long the mainstay of MBA textbooks, has been criticised for the maintained hypotheses that returns are well-described by a Gaussian distribution or that investors have mean-variance (MV) utility.^{1,2} With this in mind, we provide closed-form solutions for the gain in expected utility from accounting for non-normality in a mean-variance framework. Our analytical approach circumvents the pitfalls that can thwart empirical studies. We make assumptions about return and volatility

generating models that change from very strong assumptions to versions closer to empirical data. Readers should be aware that we do not claim that each model is an accurate representation of the returns faced by asset managers; rather, such an approach allows us, at least partially to decompose the changes in expected utility that arise. We quantify the gain in expected utility for both international and domestic investors. First, we show that there are economically significant gains in mean-variance utility from accounting for non-normality commensurate with typical mutual fund fees. Second, we find that most of the gains in expected utility derive from accounting for stochastic volatility. Fleming et al. (2001, 2003, 2012), Gomes (2007), Han (2006), Busse

✉ Stephen Satchell
ses999gb@yahoo.co.uk

¹ Plato Investment Management, 35/60 Margaret Street, Sydney, NSW 2000, Australia

² Honorary Associate of the University of Technology, Sydney, Australia

³ Trinity College, University of Cambridge, Cambridge CB2 1TQ, UK

⁴ Department of Land Economy, University of Cambridge, 19 Silver Street, Cambridge CB3 1EP, UK

¹ Ingersoll (1987) also shows that if returns are multivariate normal, expected utility is again a function of the mean and the variance under mild restrictions of the utility function (4.46–4.49)

² While the Markowitz approach and mean variance are often used interchangeably, a careful reading of Markowitz (1952) reveals no mention of the Gaussian distribution, nor does it claim that investors have mean-variance utility. See also Kritzman and Markowitz (2017).



(1999), and Kim and In (2012) all provide empirical evidence that accounting for stochastic volatility adds economic value. Our analytical model identifies the potential origin of this uplift in performance. Our result provides justification for an increased use of conditional volatility models, showing real economic gains for investors.

The remainder of this paper is organised as follows. In the section “**Non-Normality and Investment**”, we survey literature relating to investor preferences and the non-Gaussian characteristics of financial data. In the section “**Analytical Framework**”, we set out our stochastic return models and analytical approach. In the section “**Derivation of Propositions**”, we propose two distributions for our state variables that control asymmetry and stochastic volatility. Next, we derive closed-form solutions for the gain in expected utility from accounting for non-normality and then provide two empirical applications for a domestic and an international investor, before concluding.

Non-normality and investment

It has long been established that the returns of financial assets depart from normality. Mandelbrot (1963) documented extreme departures from normality for commodity prices, “warranting a radically new approach to the problem of price variation”. Analysing large cap equities, Fama (1963) reported that 5-standard deviation events occur 2000 times more often than the Gaussian distribution predicts.³ Further, financial data tend to contain more extreme negative returns than extreme positive returns: the return distribution is asymmetric or skewed. In common parlance, “busts” are more common than “booms” (Beedles 1979; Alles and Kling 1994). Mandelbrot (1963) also identified volatility clustering for commodities where “large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes”. More recent research reveals asymmetric dependence structures between assets. In the international context, Karolyi and Stulz (1996) find that the dependence between US and Japanese stocks increases during large shocks to the respective markets. Within the US market, Ang and Chen (2002) show that the dependence between individual stocks and the aggregate market index is significantly higher for extreme downside moves than for extreme upside moves of the index.⁴

³ We estimate that in the 2010-2015 period, 17 three standard deviation returns have occurred during the 1,000 most recent daily returns for the S&P 500 index against an expected 2.7 for a normal distribution.

⁴ For a more extensive review of recent literature, see Alcock & Satchell (2018).

Throughout our analysis, we consider a single investor with mean-variance utility. While the mean-variance assumption is largely normative in practice, our use of the mean-variance utility assumption is motivated in three ways. First, despite it being well established that returns are non-Gaussian and that investors are concerned with higher moments, the mean-variance approach remains ubiquitous. In a survey of investment managers, Amenc et al. (2011) find that most investors do not employ extreme risk measures, instead relying on the Gaussian distribution or not quantifying tail events whatsoever (see also Fabozzi et al. 2007). Second, the use of the mean-variance approximation is also justified through a second-order Taylor series expansion of expected utility. Samuleson’s (1970) *Fundamental Approximation Theorem* shows that for “compact” probabilities involving less and less risk, the mean-variance approach becomes increasingly accurate. Third, the mean-variance assumption is consistent with exponential utility when returns are normally distributed. Finally, in addressing the claims of DeMiguel et al (2009) that 1/N portfolios outperform optimised ones, Allen et al. (2019) show that M-V approaches can outperform more passive 1/N approaches with minimal forecasting ability. Although we formally compare the expected utility of “distinct” informed and uninformed investors, this is consistent with treating them as the same entity with and without information. This is important as comparing utility across different individuals leads to a number of additional requirements, which we do not wish to assume.

Analytical framework

Stochastic return models

The relationship between conditional returns and conditional variance intuitively plays an important role in determining the effect of accounting for stochastic volatility. If, for example, conditional returns and conditional variance are unrelated, it should be possible to increase utility by increasing exposure when risk is low, thereby capturing higher returns, and decreasing exposure when risk is high. Under the efficient market hypothesis, agents must be compensated for systematic risk and hence it follows that the conditional variance of market returns should be positively related to conditional excess returns. Empirical evidence, however, is inconclusive. Pindyck (1984), French et al. (1987), and Ghysels et al. (2005) find a positive relationship, while Campbell (1987), Glosten et al. (1993), Whitelaw (1994), (2003) and Brandt and Kang (2004) document negative relationships. Given this lack of a clear relationship between risk and return, we use two stochastic return models with different assumptions on the relationship between conditional



Table 1 Unconditional and conditional moments under Model 1 with skew and stochastic volatility

	Return vector	Covariance matrix
Unconditional	μ	$\Omega = \mu\mu' \text{Var}(s_1) + \Sigma$
Investor type		
Uninformed unconditional	$E_U[r] = \mu$	$E_U[\Omega] = \Sigma$
Informed conditional	$E_I[r] = \mu$	$E_I[\Omega s_2] = \mu\mu' \text{Var}(s_1) + s_2\Sigma$

Unconditional and conditional moments of the uninformed, informed and omniscient investors in the presence of skew and stochastic volatility under the first stochastic representation

variance and conditional excess returns. Both stochastic representations accommodate skew, heavy tails, tail dependence and stochastic volatility. Although the term unconditional expected mean variance may sound as if it has a redundancy, we use it to remind the reader that we have accounted for all variation in state variables. We also assume, in effect, three information sets for each of our different models. Knowing the full set and taking expectations is denoted by a subscript O (“Omniscient”), repeating the exercise with “possible” information is denoted by the subscript I (Informed), and for knowing less than the informed set, we use the subscript U (Uninformed). The precise details of each information set depend upon the model structure that we assume. This will change throughout the paper. Readers may wish to peruse Table 1 which sets out a sense of how this works.

In the first model, we assume that conditional mean returns and conditional variance are independent. The first stochastic representations is given by

$$\text{Model 1 : } r_t = s_{1,t}\mu + s_{2,t}^{\frac{1}{2}}Hz_t \quad (1)$$

where r_t is a vector of returns, μ , is the vector of unconditional mean returns, H is the Cholesky decomposition of the covariance matrix, Σ , s_1 and s_2 are state variables, and Ω is the unconditional covariance matrix.

The vector z_t consists of iid normal variables. Hereinafter, we drop the time subscript for the state variables. For clarity, we initially drop information subscripts for expectations, showing them in later models. In both models, the expected values of the two state variables are: $E[s_1] = E[s_2] = 1$ and the two state variables are independent of each other, $E[s_1s_2] = 1$, and independent of z_t . In Model 1, when $s_1 > 1$, the conditional expected return is greater than μ . Similarly, when $s_2 > 1$, the conditional covariance matrix is increased in the sense that the difference between the new and old matrices is positive semi-definite. The first model subsumes many notable cases. If $s_1 = 1$ and $s_2 = 1$, we recover the Gaussian distribution; if $s_1 = s_2 \sim N^-(\lambda, \chi, \psi)$ where N^- denotes the generalised inverse Gaussian distribution,⁵ we have the central generalised hyperbolic distribution;⁶ if

$s_1 = s_2 \sim IG(v/2, v/s)$, where IG denotes the inverse gamma distribution, and v is the degrees of freedom parameter, we have the central multivariate skewed-t distribution of DeMarta and McNeil (2005).

The first state variable, s_1 , introduces skew, heavy tails and tail dependence to the unconditional distribution with jumps occurring simultaneously across all assets. The second state variable, s_2 , introduces time-varying volatility and heavy tails to the unconditional distribution. In this way, we capture the four stylised facts of financial assets: heavy tails, negative skew, volatility-clustering and tail dependence.

In the second model, we assume that conditional returns and the conditional variance of the market are positively dependent, in line with standard equilibrium theory. The second stochastic model is given by

$$\text{Model 2 : } r_t = s_2^{\frac{1}{2}}(s_1\mu + Hz_t) \quad (2)$$

Under the second model, assets react differently to information depending on the state of the market. When volatility is high, a given negative shock, $s_1 < 1$, will result in a larger price decline than when volatility is low.

Informed and uninformed investors

To quantify the benefits from accounting for non-normality, we assume the existence of informed and uninformed investors (or, equivalently, consider an investor with and without information, given our prior caveat). The informed investor is aware that the return distribution is skewed and that volatility is stochastic, whereas the uninformed investor assumes that returns are i.i.d. multivariate Gaussian. We therefore assume that informed investors are aware of the state variable, s_1 , but are unaware of the value of s_1 at any point in time. This is consistent with the rational expectations equilibrium model of Veronesi (1999). Intuitively, this means that the informed investor accounts for the probability of a crash but cannot predict when that crash will occur. Additionally,

⁵ McNeil, Frey and Embrechts (2005), Appendix 2.5

⁶ McNeil, Frey and Embrechts (2005), 3.2.3



we assume that the informed investor can observe the conditional level of variance, s_2 .⁷

Derivation of propositions

We compare the objective expected utility of the mean-variance investor who ignores non-normality, with the objective expected utility of the informed mean-variance investor who takes non-normality into account. Our focus is on the objective conditional expected utility as known by the “omniscient being” that observes the true underlying probability distribution. We denote this with the subscript O. We also refer to the subjective expected utility that the investor expects to receive given her conditional expectations of return and risk which may or may not be equal to the objective conditional utility. We assume that we have no parameter uncertainty, no forecasting ability⁸ and no budget constraint. We assume:

1. Returns are skewed, and volatility is stochastic
2. The investor has mean-variance utility
3. The uninformed investor assumes the mean is non-stochastic: $\text{var}(s_1) = 0$
4. The uninformed investor assumes volatility is non-stochastic: $\text{var}(s_2) = 0$ and sets $s_2 = 1$
5. The informed investor is aware that the overall mean is stochastic: $\text{var}(s_1) > 0$
6. The informed investor is aware that volatility is stochastic: $\text{var}(s_2) > 0$
7. The informed investor is unaware of the level of the mean in each period.
8. The informed investor conditions on the level of volatility s_2

The unconditional and subjective investor moments are shown in Table 1.

For the informed investor, returns are skewed and volatility is stochastic under Model 1. The utility function is:

$$\omega' \mu - \lambda \omega' \Omega \omega / 2 \quad (3)$$

The optimal weights are given by

$$\hat{\omega} = \frac{\Omega^{-1} \mu}{\lambda} = \frac{(\mu \mu' \text{Var}(s_1) + s_2 \Sigma)^{-1} \mu}{\lambda} \quad (4)$$

⁷ An alternative and equally valid interpretation is that the state variable, s_2 , represents a form of estimation error in the covariance matrix, see, for example, Coles, Loewenstein and Suay (1995). Whereas Kan and Zhou (2007) and Allen, Lizieri and Satchell (2019) evaluate the impact of estimation error in the individual terms of the covariance matrix, the current work can be viewed as evaluating errors in the overall level of estimated market volatility.

⁸ See Allen et al. (2019) for an analysis of the impact of forecasting ability in MV allocations.

Substituting the optimal weight of the informed investor into the expected utility function gives the following:

Proposition 1 *The unconditional expected mean-variance utility of the informed mean-variance investor under the assumptions of stochastic volatility and skew under Model 1 and where $\alpha = \mu' \Omega^{-1} \mu$ is given by*

$$E_O[U_I] = E_I[U_I] = \frac{\alpha}{2\lambda} \left(E \left[\frac{1}{s_2 + \text{Var}(s_1)\alpha} \right] \right) \quad (5)$$

Proof: See “Appendix A”

For the uninformed investor, returns are skewed and volatility is stochastic under Model 1

The optimal weights of the uninformed investor are again given by

$$w = \frac{\Sigma^{-1} \mu}{\lambda}$$

The conditional expected return of the uninformed investor is given by

$$E_O[r_p | s_1] = \frac{s_1 \mu' \Sigma^{-1} \mu}{\lambda}$$

The conditional expected risk is then:

$$\begin{aligned} E_O[\sigma_p^2 | s_1, s_2] &= \omega' \Omega \omega \\ &= \frac{1}{\lambda^2} (\Sigma^{-1} \mu)' (\text{Var}(s_1) \mu \mu' + s_2 \Sigma) (\Sigma^{-1} \mu) \\ &= \frac{1}{\lambda^2} (s_2 \alpha + \text{Var}(s_1) \alpha^2) \end{aligned}$$

The expected utility of the uninformed investor is hence given by

$$E_O[U_U | s_1, s_2] = \left(\frac{s_1 \alpha}{\lambda} - \frac{1}{2\lambda} (s_2 \alpha + \text{Var}(s_1) \alpha^2) \right)$$

Now, since $E_O[s_1] = 1$ and $E_O[s_2] = 1$, we have:

Proposition 2 *The unconditional expected mean-variance utility of the uninformed mean-variance investor under the assumptions of stochastic volatility and skew under Model 1*

$$E_O[U_U] = \frac{\alpha}{2\lambda} (1 - \text{Var}(s_1) \alpha) \quad (6)$$

It is again clear that this is less than the subjective expected utility. The gain in utility is given by:

$$\begin{aligned} E_O[U_I] - E_O[U_U] &= \frac{\alpha}{2\lambda} \left(\left(E_O \left[\frac{1}{s_2 + \text{Var}(s_1)\alpha} \right] \right) - (1 - \text{Var}(s_1)\alpha) \right) \end{aligned}$$

Now from Jensen’s inequality, we know



Table 2 Unconditional and conditional moments under Model 2 with skew and stochastic volatility

	Return vector	Covariance MATRIX
Unconditional	μ	$\Omega = (\mu\mu'Var(s_1) + \Sigma)$
Investor type		
Uninformed unconditional	$E_U[r] = \mu$	$E_U[\Omega] = \Sigma$
Informed conditional	$E_I[r s_2] = s_2^{1/2}\mu$	$E_I[\Omega s_2] = s_2(\mu\mu'Var(s_1) + \Sigma)$

Unconditional and conditional moments of the uninformed and informed investors in the presence of skew and stochastic volatility under the second stochastic representation

$$\frac{\alpha}{2\lambda} \left(\left(E_O \left[\frac{1}{s_2 + Var(s_1)\alpha} \right] \right) \right) > \frac{\alpha}{2\lambda} \left(\frac{1}{1 + Var(s_1)\alpha} \right)$$

Hence,

$$E_O[U_I] - E_O[U_U] > \frac{\alpha}{2\lambda} \left(\left(\frac{1}{1 + Var(s_1)\alpha} \right) - (1 - Var(s_1)\alpha) \right) > 0$$

The effect of skew and stochastic volatility under Model 2

We now turn to our second model where conditional returns are positively related to conditional variance, and the size of the crash depends on the prevailing level of volatility. Rewriting, for convenience:

$$Model\ 2 : \quad r_t = s_2^{1/2}(s_1\mu + Hz_t)$$

The objective conditional mean-variance utility functions for our second model are given by:

$$E_O[U_I|s_1, s_2] = \omega' \mu s_1 s_2^{1/2} - \frac{\lambda s_2}{2} \omega' (\mu\mu'Var(s_1) + \Sigma) \omega \quad (7)$$

where λ is the risk aversion parameter, and, as before, the subscript ‘‘O’’ refers to the expectation of the ‘‘omniscient being’’, with full information.

As in Model 1, the expected values of the two state variables are $E_O[s_1] = E_O[s_2] = 1$ and the two state variables are independent, implying that $E_O[s_1 s_2] = 1$. As well as capturing the characteristics of market behaviour, this approach yields convenient mathematical properties. The objective mean return is determined by the interaction of the two state variables as follows:

$$E_O[r|s_1, s_2] = s_1 s_2^{1/2} \mu$$

The objective conditional covariance matrix is defined as follows

$$\Omega_s = Cov[r_t|s_1, s_2] = s_2(\mu\mu'Var(s_1) + \Sigma)$$

Assuming that the elements of μ are positive, the first state variable increases each element of the covariance matrix due to the effect of stochastic mean returns, while the second state variable scales the resulting covariance matrix. Under this model, the informed investor is aware that mean

returns are stochastic and that volatility is also stochastic. Again, the informed investor is unaware of the value of the state variable s_1 . The informed investor is unaware that the true mean return is a function of the second state variable s_2 . Again, the informed investor can condition volatility on s_2 and thus can forecast the future level of volatility without error. The uninformed investor uses the Gaussian i.i.d. assumption.

The unconditional and conditional moments of the investors are shown in Table 2:

For informed investors, returns are skewed and volatility is stochastic under Model 2.

Model 2

The conditional expected return of the informed investor is given by:

$$E_I[r|s_2] = s_2^{1/2} \mu$$

The informed covariance matrix is defined as follows:

$$E_I[\Omega|s_2] = s_2(\mu\mu'Var(s_1) + \Sigma)$$

The optimal weights are then given by:

$$\hat{\omega} = \frac{(\mu\mu'Var(s_1) + \Sigma)^{-1} \mu}{s_2^{1/2} \lambda}$$

Substituting the optimal weights of the informed investor into the objective utility function gives:

Proposition 3 *The unconditional expected mean-variance utility of the informed mean-variance investor under the assumptions of stochastic volatility and skew under Model 2*

$$E_O[U_I] = E_I[U_I] = \frac{\alpha}{2\lambda} \left(\frac{1}{1 + Var(s_1)\alpha} \right) \quad (8)$$

For the uninformed investor, returns are skewed and volatility is stochastic in Model 2.

The full results are derived in Appendix B. Substituting the weights of the uninformed investor as before into the objective utility function gives:



Proposition 4 *The unconditional expected mean-variance utility of the uninformed mean-variance investor under the assumptions of non-stochastic volatility and skew under Model 2*

$$E_o[U_U] = \frac{\alpha}{2\lambda} \left(2E_o[s_2^{1/2}] - (1 + \text{Var}(s_1)\alpha) \right) \quad (9)$$

It is easily seen that this is less than the subjective expected utility as $s_2^{1/2}$ is a concave function so

$$E_o[s_2^{1/2}] < 1 \text{ as } E_o[s_2] = 1.$$

Using Propositions 3 and 4, the gain in utility is given by:

$$\begin{aligned} E_o[U_I] - E_o[U_U] &= \frac{\alpha}{2\lambda} \left(\frac{1}{1 + \text{Var}(s_1)\alpha} \right) - \frac{\alpha}{2\lambda} \left(2E_o[s_2^{1/2}] - (1 + \text{Var}(s_1)\alpha) \right) \\ &= \frac{\alpha}{2\lambda} \left(\frac{1}{1 + \text{Var}(s_1)\alpha} + (1 + \text{Var}(s_1)\alpha) - 2E_o[s_2^{1/2}] \right) \\ &= \frac{\alpha}{2\lambda} \left(\frac{1 + (1 + \text{Var}(s_1)\alpha)^2}{1 + \text{Var}(s_1)\alpha} - 2E_o[s_2^{1/2}] \right) \end{aligned} \quad (10)$$

Taking the first term with the brackets of Eq. (10), we have

$$\frac{1 + (1 + \text{Var}(s_1)\alpha)^2}{1 + \text{Var}(s_1)\alpha} = \frac{1}{1 + \text{Var}(s_1)\alpha} + (1 + \text{Var}(s_1)\alpha)$$

Now,

$$\frac{1}{1 + \text{Var}(s_1)\alpha} + (1 + \text{Var}(s_1)\alpha) > 2$$

and

$$\frac{1}{1 + \text{Var}(s_1)\alpha} > 1 - \text{Var}(s_1)\alpha \quad \text{as } 1 > 1 - \text{Var}(s_1)^2 \alpha^2$$

Thus, the first term within the brackets of equation (10) is greater than two, and the second term is strictly less than two by Jensen's inequality and the gain from accounting for non-normality is unambiguously positive under both of our models.

Closed-form expressions for expected utility under non-normality

Modelling the state variables

In order to quantify the gains from accounting for non-normality, we need to model the distributions of the two state variables. For the first state variable, we require a

distribution that can accommodate skew and has a mean of one. We use a modified Bernoulli distribution defined as follows

$$s_1 = a + (b - a)\text{Bernoulli} \quad (11)$$

where the steady-state scaling factor, $a > 1$, with a probability, p , and a crash scaling factor, $b < 1$ with probability, $1 - p$. Multiplying the mean return vector, μ , by our modified Bernoulli variable creates negative skew and a heavy left tail while leaving the unconditional return unchanged on average. This is a further example of a mean-preserving spread (Rothschild and Stiglitz 1970). The variance of the Bernoulli distribution is given by

$$V(s_1) = (b - a)^2(1 - p)p$$

Since the expected value of the second state variable is 1,

$$E(s_1) = a + (b - a)p = 1$$

and we can express the variance solely in terms of the crash probability, p , and the crash scaling factor, b .

$$V(s_1) = \frac{1}{p}(1 - b)^2(1 - p) \quad (12)$$

For a fixed probability, $1 - p$, we can see that the variance of our modified Bernoulli distribution increases as the absolute magnitude of the crash, b , increases.

The second state variable that we use to capture stochastic volatility must be strictly positive to prevent the variance from becoming negative and also have a mean of one to ensure the unconditional variance is unchanged on average. We use the scaled chi-square distribution to capture time-varying volatility. This distribution has the attractive feature that the inverse moment can be calculated which allows for explicit expressions for expected utility. Like the standard chi-square, the scaled chi-square is bounded by zero from below and is generated by summing squared random normal variables, consistent with the calculation of variance. Unlike the standard chi-square, the mean is one for all values of k , corresponding to the steady-state market volatility level. The distribution of s_2 is given by

$$s_2 = \frac{\chi^2(k)}{k} \quad (13)$$

where χ^2 defines a chi-square variable with $k \geq 3$ degrees of freedom. The m -th moment of the distribution is given by:

$$E[X^m] = \frac{2^m \Gamma\left(m + \frac{k}{2}\right)}{k^m \Gamma\left(\frac{k}{2}\right)} \quad (14)$$



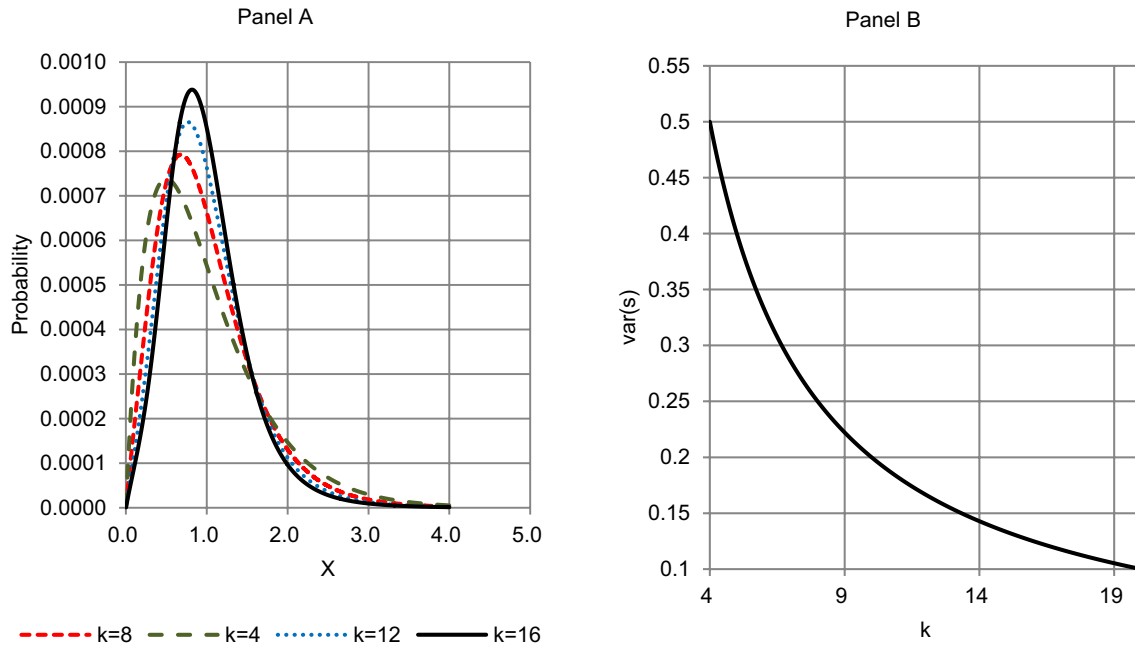


Fig. 1 Probability density of the scaled chi-square distribution (panel A). Variance of scaled chi-square distribution vs. degrees of freedom, k (panel B). The left-hand panel shows the probability density func-

tion of the scaled chi-square distribution for 4, 8, 12 and 16 degrees of freedom. The right-hand panel shows the variance of the scaled chi-square distribution vs. k as given in equation 18

The parameter k controls the level of variability in the volatility level, s_2 . As k increases, the variability in the level of volatility approaches zero. The second moment of the scaled chi-square distribution is given by

$$E[X^2] = \frac{4\Gamma(2+\frac{k}{2})}{k^2\Gamma(\frac{k}{2})} = 1 + \frac{2}{k} \quad (15)$$

$$\text{var}(X) = \frac{2}{k} \quad (16)$$

Panel A of Fig. 1 shows the density function of the scaled chi-square distribution for four values of k . Panel B shows the variance of the distribution conditional on k .

We also draw upon the inverse moment of our modified chi-square distribution.

$$E[X^{-1}] = \frac{k\Gamma(\frac{k}{2}-1)}{2\Gamma(\frac{k}{2})} = \frac{k}{k-2} \quad (17)$$

We now use our modified Bernoulli and scaled chi-square distributions to give closed-form solutions for the expected gain in utility from accounting for non-normality.

Incorporating stochastic volatility and skew under Model 1

We now look to derive closed-form solutions under Model 1 incorporating skew and stochastic volatility. For the informed investor, it is possible to derive a closed-form solution for expected mean-variance utility for particular distributions of the state variable, s_2 using the following relationship:

$$\begin{aligned} E_O[U_I] &= \frac{\alpha}{2\lambda} \left(E \left[\frac{1}{s_2 + \text{Var}(s_1)\alpha} \right] \right) \\ &= \frac{1}{2\lambda \text{Var}(s_1)} \left(E \left[\sum_{j=0}^{\infty} \left(\frac{-s_2}{\text{Var}(s_1)\alpha} \right)^j \right] \right) \end{aligned}$$

As far as we are aware, however, it is not possible to derive a closed-form solution for expected utility of the informed investor using the scaled chi-square distribution. Accordingly, in the section ‘‘Empirical Application’’ we use numerical integration to give the expected utility of the informed investor. For the uninformed investor, we substitute the expected variance of the modified Bernoulli distribution into Proposition 2 to derive a closed-form expression for expected utility as follows:



Corollary 1 *The unconditional expected mean-variance utility of the uninformed mean-variance investor under the assumptions of stochastic volatility and skew under Model 1 using the expected first moment of the modified Bernoulli distribution is*

$$E_O[U_U] = \frac{\alpha}{2\lambda} \left(\frac{p+(1-b)^2(1-p)\alpha}{p} \right) \quad (18)$$

For our second stochastic representation, where risk and return is known, we can provide closed-form solutions for the gain in expected utility for both investors as we now show by substituting the variance of the modified Bernoulli distribution (14) into Proposition 3:

Corollary 2 *The unconditional expected mean-variance utility of the informed mean-variance investor under the assumptions of non-stochastic volatility and skew under Model 2 using the expected moments of the scaled chi-square and the modified Bernoulli distribution is:*

$$E_O[U_I] = \frac{\alpha}{2\lambda} \left(\frac{p}{p+(1-b)^2(1-p)\alpha} \right) \quad (19)$$

Substituting the definition of the k th moment of the scaled chi-square distribution (16) with $m=0.5$ and the variance of the modified Bernoulli distribution (14) into Proposition 4 gives:

Corollary 3 *The unconditional expected mean-variance utility of the uninformed mean-variance investor under the assumptions of non-stochastic volatility and skew under Model 2 using the expected moments of the scaled chi-square and the modified Bernoulli distribution is:*

$$E_O[U_U] = \frac{\alpha}{2\lambda} \left(\frac{2\sqrt{2}\Gamma\left(\frac{k+1}{2}\right)}{\sqrt{k}\Gamma\left(\frac{k}{2}\right)} - \left(\frac{p+(1-b)^2(1-p)\alpha}{p} \right) \right) \quad (20)$$

The gain in expected utility is then given by the difference between Corollaries 2 and 3.

$$E_O[U_I] - E_O[U_U] = \frac{\alpha}{2\lambda} \left(\frac{p^2+(p+(1-b)^2(1-p)\alpha)^2}{p(p+(1-b)^2(1-p)\alpha)} - \frac{2\sqrt{2}\Gamma\left(\frac{k+1}{2}\right)}{\sqrt{k}\Gamma\left(\frac{k}{2}\right)} \right) \quad (21)$$

Again, we see that this is unambiguously nonnegative as the first term within the brackets is strictly greater than or equal to 2 and the second term is strictly less than or equal to 2.

Corollary 4 *The unconditional expected mean-variance utility of the informed mean-variance investor under the assumptions of non-stochastic volatility, skew and no estimation error under Model 2 using the expected moments of the modified Bernoulli and scaled chi-square distributions is:*

$$E_O[U_I] = \frac{\alpha}{2\lambda} \left(\frac{p}{p+(1-b)^2(1-p)\alpha} \right) \quad (22)$$

Corollary 5 *The unconditional expected mean-variance utility of the uninformed mean-variance investor under the assumptions of non-stochastic volatility and skew under Model 2 using the expected moments of the modified Bernoulli and scaled chi-square distributions is:*

$$E_O[U_U] = \frac{\alpha}{2\lambda} \left(\frac{2\sqrt{2}\Gamma\left(\frac{k+1}{2}\right)}{\sqrt{k}\Gamma\left(\frac{k}{2}\right)} - \left(\frac{p+(1-b)^2(1-p)\alpha}{p} \right) \right) \quad (23)$$

Empirical application

Data

In this section, we quantify the gain in expected utility from accounting for non-normality by calibrating our models to empirical data. Now, we have stepped away from the general case, and it is a valid criticism that our results are now subject to the vagaries of the data sets we have employed. To mitigate this risk and the scope for data-mining, we use asset classes that span significant periods of time that are commonly used in the literature. We consider the cases of both an international and a domestic investor. The opportunity set of our international investor is comprised of the MSCI equity indices of the G7 countries, Canada, France, Germany, Italy, Japan, the UK and the USA. To calibrate the two models, we use the MSCI total return indices for the period February 1999 to March 2013. Table 3 provides the sample summary statistics. All of the asset returns are negatively skewed and display excess kurtosis. Using the Jarque–Bera test, we can reject normality for each of the assets at the 1% level. We also reject normality in an average of 28% of overlapping 1,000 day sub-periods. The autocorrelation in the absolute value of returns is large and positive, indicative of heteroskedasticity.

The investment universe of our domestic investor contains the five value-weighted US sectors using the K.R. French data for the period 1/1983-12/2012. The universe contains all NYSE, AMEX and NASDAQ stocks. Sector definitions are based on the four-digit SIC codes. The 30-year interval spans multiple market regimes including the October 1987 crash, the Russian debt default in and collapse of LTCM in 1998, the Tech bubble and ensuing correction beginning in March, 2001, the Financial Crisis of 2007-8 and the ensuing recovery. Table 4 provides the summary statistics. Again, we can reject the hypothesis of normality at the 1% level for each of the indices.



Table 3 G7 MSCI indices summary statistics

	MSCI Canada	MSCI France	MSCI Germany	MSCI Italy	MSCI Japan	MSCI UK	MSCI USA	Mean
Mean return	0.24	0.11	0.13	0.39	0.43	0.11	0.09	0.21
Standard deviation	3.43	3.46	3.73	3.69	2.83	2.98	2.66	3.26
Skewness	-0.68	-0.57	-0.46	-0.67	-0.10	-0.71	-0.48	-0.52
Kurtosis	6.2	4.4	3.6	5.3	1.6	9.3	5.0	5.1
Maximum	19.4	14.8	16.4	19.9	12.0	17.7	12.3	16.1
Minimum	-22.9	-23.4	-22.9	-23.6	-14.8	-24.1	-18.2	-21.4
Jarque-Bera	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
% of Jarque-Bera failures	21	21	21	36	29	36	29	28
% of statistically significant skewness	14	14	7	36	14	7	14	15
% of statistically significant kurtosis	7	21	21	29	7	21	29	19
$\rho(r_t, r_{t-1})$	-0.10	-0.05	-0.01	-0.02	-0.06	-0.10	-0.08	-0.06
$\rho(r_t , r_{t-1})$	0.44	0.18	0.26	0.28	0.13	0.24	0.30	0.26

The summary statistics for the G7 countries for the period 2/1999-3/2013 using weekly data for the MSCI total return indices. The mean return refers to the weekly arithmetic return. The standard deviation is also weekly. Kurtosis is defined as excess kurtosis such that the normal distribution has a kurtosis of zero. The % of Jarque-Bera failures refers to the proportion of 1000 day overlapping sub-periods where normality is rejected at the 5% level. The % of statistically significant skewness and % of statistically significant kurtosis refers to the proportion of sub-periods where the skewness and kurtosis are statistically significant at the 5% level. $\text{corr}(r_t, r_{t-1})$ refers to the autocorrelation in returns at one lag. $\text{corr}(|r_t|, |r_{t-1}|)$ refers to autocorrelation of the absolute value of returns at one lag.

Modelling the state variables with the method of simulated moments

To quantify the gain in utility from accounting for non-normality, we first need to estimate the parameters of the two return models. We employ the Bernoulli and scaled chi-square distributions for the two state variables; other choices are also valid, as discussed above. Estimating the models allows us to decompose the unconditional distribution into a Gaussian i.i.d. component employed by the uninformed investor and a non-Gaussian component that is employed by the informed investor. The stochastic representations of our two models are given by

$$\begin{aligned} \text{Model 1} : r_t &= s_1 \mu + s_2^{1/2} H z_t \\ \text{Model 2} : r_t &= s_2^{1/2} (s_1 \mu + H z_t) \end{aligned} \quad (24)$$

where s_1 and s_2 are described in equations (11) and (13).

We require estimates of the parameters, \hat{b} , and \hat{p} , for the first state variable, the parameter, k for the second state variable, the vector $\hat{\mu}$ and the upper triangular matrix, \hat{H} where \hat{H} is the Cholesky decomposition of the covariance matrix of the uninformed investor, $\hat{\Sigma}$. By simultaneously estimating the parameters, we can decompose the unconditional distribution into a skew-related component and a standard covariance-related component. It is then strictly true that the unconditional covariance matrix, Ω , is greater than Σ over all elements. In total, we have $3 + 2n + n(n-1)/2$ terms to estimate. This is a formidable challenge made more difficult by

the lack of a closed-form multivariate density function, ruling out standard maximum likelihood estimation. McFadden (1998) introduced the method of simulated moments (MSM) to estimate the parameters of multivariate functions for problems where the density function may be difficult to estimate, and it is this approach that we use here. Prior examples of employing the MSM for fitting return generating distributions in an asset allocation context include Brandt (1999), Ait-Sahalia and Brandt (2001) and Das and Uppal (2004).

We are interested in replicating the non-Gaussian characteristics of our data, and we select the parameters, θ , to minimise weighted squared deviation between the second, third and fourth simulated and empirical moments as described below in (17). To ensure the covariance structure is also captured, we also employ the cross-moments of order two. The i -th empirical moment is given by the time-series average as follows:

$$\hat{m}_{simkT,k} = \frac{1}{T} \sum_{t=1}^T m_i(z_t^{emp}), \quad i = \dots, n_m \quad (25)$$

where T is the number of data points, $m_i(\cdot)$ is a function that gives the moment of order i , and z_t^{emp} is the observed return in period t .

The simulated counterpart of (15) is given by

$$\hat{m}_{simkT,i} = \frac{1}{T} \sum_{t=1}^{kT} m_i(z_t^{sim}(\theta)), \quad i = \dots, n_m \quad (26)$$



Table 4 Value-weighted US sectors: summary statistics

	Cons.	Manuf.	Tech.	Health.	Other.	Mean
Mean return	0.26	0.26	0.23	0.26	0.22	0.25
Standard deviation	2.29	2.28	3.12	2.50	2.80	2.60
Skewness	-0.52	-0.82	-0.35	-0.37	-0.55	-0.52
Kurtosis	4.2	5.8	3.3	3.5	4.7	4.32
Maximum	11.0	10.3	16.8	15.8	11.8	13.17
Minimum	-17.1	-16.1	-18.6	-17.0	-18.4	-17.43
Jarque–Bera	0.00	0.00	0.00	0.00	0.00	0.00
% of Jarque–Bera failures	28	41	17	38	31	31
% of statistically significant skewness	24	28	17	31	14	23
% of statistically significant kurtosis	28	31	10	24	24	23.45
$\rho(r_t, r_{t-1})$	-0.05	-0.08	-0.08	-0.08	-0.06	-0.07
$\rho(r_t , r_{t-1})$	0.18	0.22	0.15	0.14	0.22	0.18

The summary statistics for the 5 US value-weighted sectors for the period 1/1983-12/2012 using weekly data. The mean return refers to the weekly arithmetic return. The standard deviation is also weekly. Kurtosis is defined as excess kurtosis such that the normal distribution has a kurtosis of zero. The % of Jarque–Bera failures refers to the proportion of 1000 day overlapping sub-periods where normality is rejected at the 5% level. The % of statistically significant skewness and % of statistically significant kurtosis refers to the proportion of sub-periods where the skewness and kurtosis are statistically significant at the 5% level. $\text{corr}(r_t, r_{t-1})$ refers to the autocorrelation in returns at one lag. $\text{corr}(|r_t|, |r_{t-1}|)$ refers to autocorrelation of the absolute value of returns at one lag.

where $k \geq 1$ is a scaling factor that determines the length of the simulation. We select k to give 100,000 simulated periods per asset.

The vector of differences between the empirical and simulated moments is then given by

$$g_{s,T}(\theta, z_t^{\text{emp}}) = \hat{m}_{\text{simkT}} - \hat{m}_{\text{simpT}}$$

To obtain the optimal values of the parameter vector, θ , we minimise the error term, $Q_{k,T}(\theta)$, as follows.

$$\hat{\theta} = \arg \min Q_T(\theta) = g_{s,T}(\theta)' W g_{s,T}(\theta) \quad (27)$$

where W is a covariance matrix that we discuss below.

We perform k -step MSM (Hall 2005) as follows. To obtain an initial estimate of $\hat{\theta}$, we minimise the squared differences between the simulated and observed moments with respect to the identity matrix. This of course gives equal weight to each of the moment and cross-moment conditions. We can improve on this result by repeating the minimisation with respect to an optimal weighting matrix \hat{W} that provides the smallest asymptotic covariance of the estimator. We can estimate the optimal weighting matrix, \hat{W} , as the inverse of the estimated asymptotic covariance matrix of the moment conditions using the simulated data series. The most elementary estimator we could construct to capture all of the autocovariance is given by

$$\hat{\Omega}(\hat{\theta}) = \hat{\Gamma}_0 + \sum_{i=1}^{T-1} (\hat{\Gamma}_i + \hat{\Gamma}_i') \quad (28)$$

where

$$\hat{\Gamma}_i = T^{-1} \sum_{T=i+1}^{T-1} g(\hat{\theta}, z_t^{\text{emp}}) g(\hat{\theta}, z_{t-i}^{\text{emp}})' \quad (29)$$

We follow common practice and use a Newey–West (1987) covariance estimator to weight the autocovariances. The result is a heteroskedasticity and autocovariance consistent (HAC) estimator that is guaranteed to be positive definite and has a single solution. The Newey–West estimator computes the asymptotic covariance matrix as if the moment process is a vector moving average and uses the sample autocovariances up until a given lag, l , as follows.

$$\hat{\Omega} = \hat{\Gamma}_0 + \sum_{l=1}^l \frac{l+1-i}{l+1} (\hat{\Gamma}_i + \hat{\Gamma}_i') \quad (30)$$

While there are a number of complicated formulae that can be used to calculate the optimal number of lags, l , we satisfy ourselves with the smallest integer that satisfies $l = T^{1/4}$.⁹ The optimal weighting matrix for the r th run (iteration) is then given by

$$\hat{W}_r = \hat{\Omega}_r^{-1}(\hat{\theta}_{r-1}) \quad (31)$$

We then solve equation (21) iteratively, updating the optimal weighting covariance, \hat{W}_r , using the most recent parameter estimates, $\hat{\theta}_{r-1}$, until we achieve convergence.

⁹ See Franke (2008)



Table 5 Parameter estimates for G7 MSCI indices and value-weighted sectors

	Model 1		Model 2	
	$r_t = s_1\mu + s_2^{1/2}Hz_t$	$r_t = s_2^{1/2}(s_1\mu + Hz_t)$	Countries	Sectors
\hat{k}	6	4	6	4
\hat{p}	0.996152 (0.0045)	0.991531 (0.0351)	0.998369 (0.0040)	0.991531 (0.0311)
\hat{b}	-84.76 (51.49)	-84.67 (115.72)	-84.75 (76.12)	-84.67 (126.99)
$corr(var_{emp}, var_{sim})$	0.9965	0.9989	0.9994	0.9997
$corr(covar_{emp}, covar_{sim})$	0.9989	0.9982	0.9998	0.9993
$corr(skew_{emp}, skew_{sim})$	0.8732	0.9788	0.8569	0.9777
$corr(kurt_{emp}, kurt_{sim})$	0.9941	0.9938	0.9917	0.9663
degrees of freedom	4	2	4	2
J-critical	9.49	5.99	9.49	5.99
J-statistic	7.58	2.37	8.56	2.03
p-value	0.11	0.31	0.07	0.36

The method of simulated moments parameter estimates for the G7 MSCI Indices for the period 2/1999 to 3/2013 and the 5 value-weighted US sectors for the period 1/1983 to 12/2012. \hat{k} refers to the estimated degrees of freedom of the scaled chi-square distribution; \hat{p} and \hat{b} refer the probability and scale parameters of the modified Bernoulli distribution. Standard errors are shown in brackets. The J-statistic tests whether the model is consistent with the empirical moments and converges asymptotically to a chi-square distribution

Parameter estimates

The optimal parameters for the two models for the G7 countries and the five value-weighted sector indices are shown in Table 5.

Both models describe the data well. Using the J test, we cannot reject the hypothesis that the empirical returns are generated by either of the stochastic return models. The correlation coefficients between the observed historical moments and the estimated moments of the simulated models are very high for each model and for each data set. In particular, the variance, covariance and kurtosis terms map on extremely well with the correlation coefficients approaching unity. For the sake of brevity, we do not report the t-statistics comparing the empirical and simulated moments.

Consistent with the high correlation coefficients, t-statistics for the moment conditions exceed the 5% critical value for both of the models. It appears that both models are satisfactory in reproducing the key moments of the two data sets. Further, the parameter estimates appear plausible. For example, in the case of the first model for the G7 MSCI index data, the crash probability is $1 - p = 0.38\%$, and the crash state value is -84.76. If we take the average weekly asset return across the G7 countries of 0.21% and substitute into the stochastic return equation of Model 1, we have a steady-state return of 0.26% per week and a crash return of -17.8% per week occurring once every 263 weeks or 5 years. The standard errors of the estimated crash probabilities, \hat{p} ,

are small, particularly for the G7 data set. We recognise, however, that the period used to estimate the models for the G7 MSCI data set included the dot-com crash and the global financial crisis and thus *may* overstate the long-term crash probability and potentially overstate the benefits of accounting for departures from normality, although the possibility of sharp downward corrections in response to major shocks remains a valid concern.

Expected utility and the effect of non-normality under Model 1

We now use Proposition 1 and Corollary 3 to quantify the gain in expected utility from accounting for non-normality using the first stochastic representation. In Fig. 2, we show the expected utility of the informed and uninformed mean-variance investors conditional on the variability of volatility for three different levels of skew. The horizontal axis shows the variability of the second state variable given by the variance of the scaled chi-square distribution in equation (19). We increase the magnitude of negative skew by increasing the scaling factor, b while keeping the crash probability constant. The left-hand panel pertains to the G7 MSCI indices, and the right-hand panel pertains to the value-weighted US sectors. In the left panel, we show weekly crash sizes of 0, -8.9% and -17.8%, all at the 0.4% probability level consistent with the best fit parameters in the first column of Table 5. In the right panel, we show weekly crash sizes of 0, -10.6%



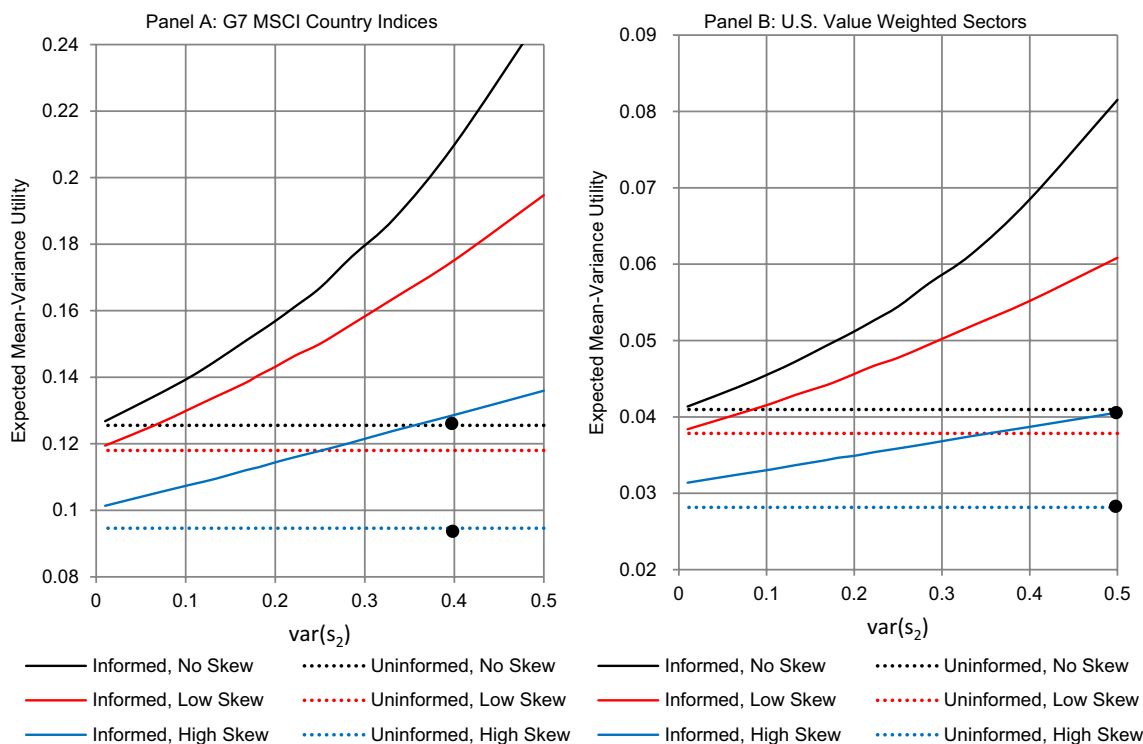


Fig. 2 Model 1: Expected utility for the informed and uninformed investors for the G7 and US sectors allocation problems. The expected mean-variance utility of the informed and uninformed investors using Proposition 1 and Corollary 3. The left-hand panel relates to the G7 MSCI country indices for the period 2/1999 to 3/2013. The right-hand panel relates to the US value-weighted sectors for

the period 1/1983 to 12/2013. The models are calibrated using the method of simulated moments (MSM). We use Monte Carlo simulation to derive the expected values for Proposition 1 and direct substitution for Corollary 3. The parameters used in Proposition 1 and Corollary 3 are given in Table 5. We use a risk aversion $\lambda = 0.05$.

and -21.2%, at the 0.85% probability level. For the best-fit parameters, we have marked the expected utility of the uninformed and informed investors with a black dot. In each case, we show a zero skew, low skew and high skew case. The high negative skew case in each panel has been selected to correspond with the best fit parameters in Table 5.

The effect of increasing the magnitude of negative skew is straightforward, decreasing the expected utility of the informed and uninformed investors. The loss in expected utility is not very large, consistent with Das and Uppal (2004). When volatility is non-stochastic, the expected utility of the informed investor shifts downwards by more than the expected utility of the uninformed investor, because the informed investor is accounting for the additional component of risk due to extreme events, whereas the uninformed investor is not. The introduction of extreme joint returns increases both the correlation and the variance of the unconditional distribution relative to the naïve expectations of the uninformed investor. A parallel can be drawn with the normal mean-variance mixing family where the variance is comprised of a standard variance component and a component driven by the variance in the mean.

The effect of stochastic volatility requires more interpretation. For the informed investor, increasing the level of variability in volatility actually increases expected utility. This is perhaps counter-intuitive and requires more explanation. The reason lies in the convexity of expected utility with respect to volatility. To make this concrete, consider the case of a mean-variance investor with a risk aversion of one in a two-state world investing in a single asset. In the two equally likely states, the variance is equal to 0.5 units in the first state and 1.5 units in the second, and the mean return is 1 unit in both states. The unconditional variance is therefore equal to 1 unit. Assuming that the informed investor can forecast the level of volatility without error, the expected utility is 1 in the first state and 0.33 in the second giving an unconditional expected utility of 0.66. If volatility is non-stochastic, it is easily seen through the standard relation

$$E[U] = \frac{\mu' \Sigma^{-1} \mu}{2\lambda} \quad (32)$$

that the expected utility equals 0.5. Thus, the informed investor gains more when volatility is low than she loses when volatility is high and, paradoxically, a departure from Gaussian i.i.d. returns leads to an increase in expected utility



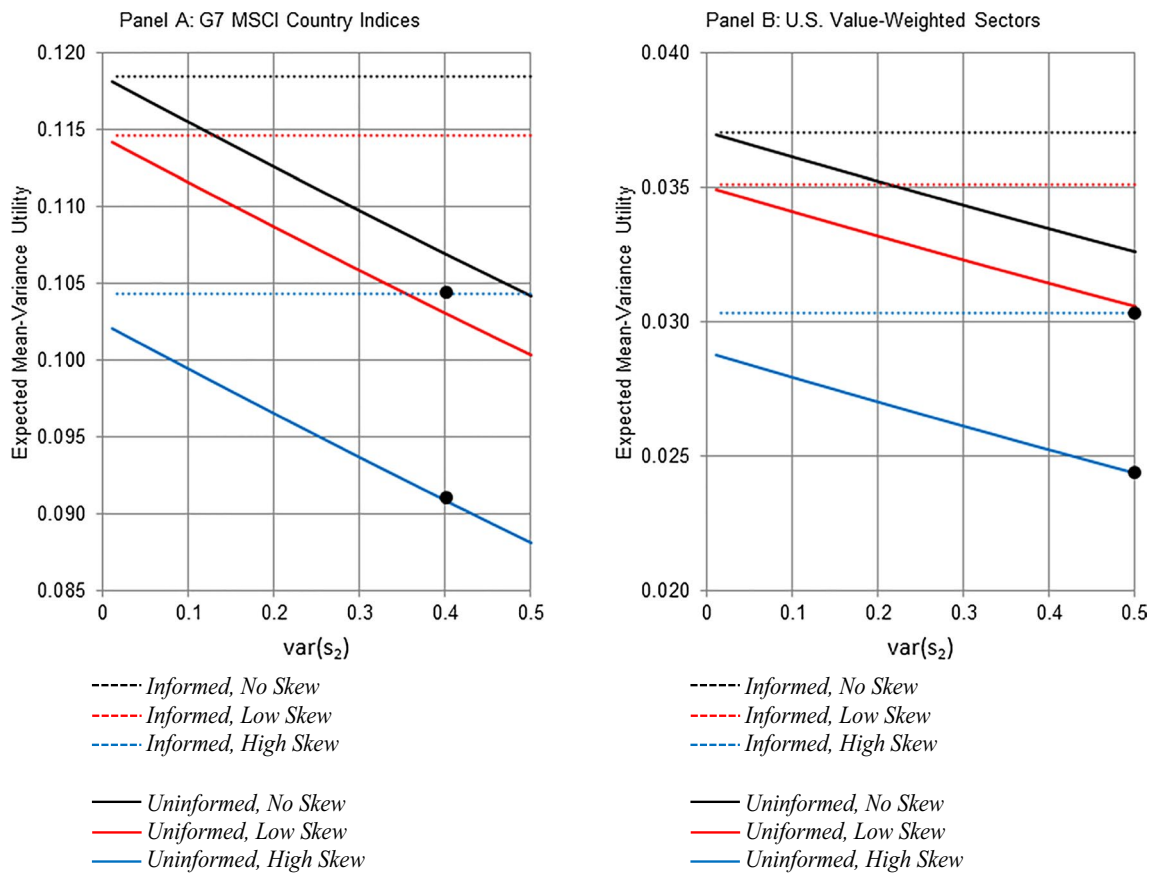


Fig. 3 Model 2: Expected utility for the informed and uninformed investors for the G7 asset allocation problem. The expected mean-variance utility of the informed and uninformed investors using corollaries 6 and 7. The left-hand panel relates to the G7 MSCI country indices for the period 2/1999 to 3/2013. The right-hand panel relates

to the US value-weighted sectors for the period 1/1983 to 12/2012. The models are calibrated using the method of simulated moments (MSM). The parameters used in corollaries 6 and 7 are given in Table 5.

over the non-stochastic case. This is perhaps an unexpected result and an under-appreciated source of utility. For the uninformed investor, the objective expected utility in state one is 0.75 and 0.25 in state two, giving an unconditional expected utility of 0.5 units. The reason the uninformed investor underperforms the informed investor in the presence of non-stochastic volatility is more intuitive. The investor takes too little risk when risk is low and too much risk when risk is high. Put another way, in a world where you can model and account for changes in risk, stochastic volatility represents an opportunity.

The gain in expected utility from accounting for non-normality is given by the vertical distance between the black dots in each panel. For the international investor, the gain in expected mean-variance utility is approximately 3 bps per week, or crudely 1.5% per year. For the domestic investor, the gain is approximately 1.25 bps per week, or 0.65% per year. It is common in the practitioner literature to equate the gain in certainty equivalence to the incremental fee the asset

manager could charge by applying a given technique. Malkiel (2013) finds that the average mutual fund fees are approximately 0.9%.¹⁰ In this context, a gain in certainty equivalence of 0.65% to 1.5% per year is significant. A potential criticism of quantifying the gains in this way is that the gain in expected utility is a function of risk aversion which we have had to estimate. To break the link between the level of risk aversion and the benefits of accounting for non-normality, we consider the percentage gain in expected utility. Dividing Corollary 4 by Proposition 1, we can see that the percentage gain in expected utility is independent of the level of risk aversion. The percentage gain in utility is approximately 38% for the international investor, and 44% for the domestic investor. The majority of the uplift in each case comes from accounting for stochastic volatility. This is a highly significant uplift suggesting that investors *should* account for non-normality.

¹⁰ Excluding Exchange Traded Funds (ETFs) and index funds



Expected utility and the effect of non-normality under Model 2

In Fig. 3, we show the expected utility of the informed and uninformed investors using corollaries 6 and 7, for the second stochastic representation. While in the first model, mean returns and covariance are independent, in the second model, the return vector and the covariance matrix move in lock-step.

As in Model 1, the introduction of skew leads to a decrease in the expected utility for both the informed and the uninformed investors; the uninformed investor experiences a larger decrease in expected utility than the informed investor. In contrast to Fig. 2, as the variability in volatility increases moving left to right, the expected utility of the informed investor remains constant. This is because when volatility is high (low), exposure is low (high) coinciding with when expected returns are high (low). Thus, the investor is not able to exploit low volatility periods and benefit from the variability in volatility in the same way as under our first model. What about the uninformed investor? Increasing the variability in volatility in this case leads to a decrease in expected utility. Whether volatility is above or below the steady-state level, the utility of the uninformed investor is less than the utility of the informed investor. Again, we have highlighted the expected utility for the informed and uninformed investors for the optimal MSM estimates with black dots. For the international investor, the gain in expected utility is approximately 1.3 bps per week, or 0.7% per year. In percentage terms, this translates to a gain of 15%. For the domestic investor, the gain is approximately 0.6 bps per week, or 0.30% per year. The percentage gain for the domestic investor, however, is higher at 25%. Again, the majority of the utility gain comes from accounting for stochastic volatility. Under both our models, whether mean returns are independent of volatility or not, there are economically significant gains from accounting for both stochastic volatility and skew. The gains, however, are more pronounced in the first model where conditional mean returns and conditional variance are independent. This makes sense in that the informed investor can benefit during low volatility periods by scaling up portfolio weights and capturing higher returns. Given the lack of clear empirical evidence showing that conditional mean returns are positively related to conditional variance, it could be argued that the conclusions of our first model should carry more weight.

Conclusion

In this paper, we believe we have been the first to analytically quantify the economic gains that can be captured by the mean-variance investor from accounting both for the skew and

stochastic volatility known to be present in asset returns. These gains are economically significant, commensurate with typical mutual fund management fees. The percentage uplift in certainty equivalent is also highly significant at approximately 40%. While accounting for skew is important, the majority of the gains are due to accounting for non-stochastic volatility. This finding aligns with the empirical work of Fleming et al. (2001, 2003), Kirby, and Ostdiek (2012), Gomes (2007) and Han (2006) on volatility timing and the mutual fund work of Busse (1999). The utility gains from incorporating the effect of stochastic volatility into the asset allocation decision are perhaps under-appreciated. In particular, if expected mean returns and volatility are independent, we show that the expected utility of the informed investor actually exceeds the expected utility of the mean-variance investor when returns are non-stochastic. Thus, paradoxically, violations of the i.i.d. Gaussian assumption can increase expected utility relative to the non-stochastic case. This finding also provides theoretical support for the use of conditional volatility models including exponentially weighted moving averages, option implied volatility and GARCH models for portfolio construction. The mean-variance approximation is ubiquitous in the finance industry today and often goes unquestioned by practitioners.

Appendix A: Derivation of propositions: the effect of stochastic volatility and skew under Model I

Proof of Proposition 1

The unconditional expected mean-variance utility of the informed mean-variance investor under the assumptions of stochastic volatility and skew under Model 1.

The conditional moments of the informed investor are

$$E_I[r] = \mu$$

$$Cov_I[r|s_2] = \mu\mu'Var(s_1) + s_2\Sigma$$

The optimal weights are given by

$$\hat{\omega} = \frac{\Omega^{-1}\mu}{\lambda} = \frac{(\mu\mu'Var(s_1)+s_2\Sigma)^{-1}\mu}{\lambda} \quad (33)$$

Now,

$$(Var(s_1)\mu\mu' + s_2\Sigma)^{-1} = \frac{\Sigma^{-1}}{s_2} - \frac{Var(s_1)}{s_2^2} \frac{\Sigma^{-1}\mu\mu'\Sigma^{-1}}{1 + \frac{Var(s_1)}{s_2}\mu'\Sigma^{-1}\mu}$$

Making use of $E[s_1] = 1$, we have



$$E_O[U_I|s_2] = \frac{1}{2\lambda} \mu' \left(\frac{\Sigma^{-1}}{s_2} - \frac{\text{Var}(s_1)}{s_2^2} \frac{\Sigma^{-1} \mu \mu' \Sigma^{-1}}{1 + \frac{\text{Var}(s_1)}{s_2} \mu' \Sigma^{-1} \mu} \right) \mu \quad (34)$$

$$E_O[U_I|s_2] = \frac{1}{2\lambda} \left(\frac{\alpha}{s_2} - \frac{\text{Var}(s_1)}{s_2^2} \frac{\alpha^2}{1 + \frac{\text{Var}(s_1)}{s_2} \alpha} \right)$$

$$E_O[U_I|s_2] = \frac{1}{2\lambda} \left(\frac{\alpha}{s_2} - \frac{\text{Var}(s_1) \alpha^2}{s_2^2 + s_2 \text{Var}(s_1) \alpha} \right)$$

$$E_O[U_I] = \frac{1}{2\lambda} \left(E \left[\frac{\alpha}{s_2} \right] - E \left[\frac{\text{Var}(s_1) \alpha^2}{s_2^2 + s_2 \text{Var}(s_1) \alpha} \right] \right)$$

In turn, we have

$$E_O[U_I] = \frac{1}{2\lambda} \left(E \left[\frac{1}{s_2} \right] \alpha - E \left[\frac{1}{s_2^2 + s_2 \text{Var}(s_1) \alpha} \right] \text{Var}(s_1) \alpha^2 \right)$$

$$E_O[U_I] = \frac{\alpha}{2\lambda} \left(E \left[\frac{1}{s_2} \right] - \left(E \left[\frac{1}{s_2} \right] - E \left[\frac{1}{s_2 + \text{Var}(s_1) \alpha} \right] \right) \right)$$

Proposition 1

$$E_O[U_I] = \frac{\alpha}{2\lambda} \left(E \left[\frac{1}{s_2 + \text{Var}(s_1) \alpha} \right] \right) \quad (35)$$

Proof of Proposition 2

The unconditional expected mean-variance utility of the uninformed mean-variance investor under the assumptions of stochastic volatility and skew under Model 1

The conditional expected return of the uninformed investor is given by:

$$E_O[r_p|s_1] = \frac{s_1 \mu' \Sigma^{-1} \mu}{\lambda}$$

The expected risk is then

$$E_O[\sigma_p^2|s_2] = \omega' \Omega \omega$$

$$= \frac{1}{\lambda^2} (s_2 \alpha + \text{Var}(s_1) \alpha^2)$$

The expected utility of the uninformed investor is hence given by

$$E_O[U_U|s_1, s_2] = \left(\frac{s_1 \alpha}{\lambda} - \frac{1}{2\lambda} (s_2 \alpha + \text{Var}(s_1) \alpha^2) \right)$$

Since $E[s_1] = 1$, we thus prove the result.

$$E_O[U_U] = \frac{\alpha}{2\lambda} (1 - \text{Var}(s_1) \alpha)$$

Appendix B: Derivation of propositions: the effect of stochastic volatility and skew under Model 2

Our second stochastic representation is given by

$$r_t = s_2^{\frac{1}{2}} (\mu s_1 + H z_t)$$

where μ is the unconditional mean, Σ is the unconditional covariance matrix, s_1 and s_2 are stochastic state variables, and H is as previously defined.

The objective conditional mean-variance utility function for our second model is given by

$$E_O[U_I|s_1, s_2] = \omega' \mu s_1 s_2^{1/2} - \frac{\lambda s_2}{2} \omega' (\mu \mu' \text{Var}(s_1) + \Sigma) \omega \quad (36)$$

where λ is the risk aversion parameter, and the subscript "O" refers to the expectation of the omniscient being.

The objective mean return is determined by the interaction of the two state variables as follows

$$E_O[r|s_1, s_2] = s_1 s_2^{\frac{1}{2}} \mu$$

The objective covariance matrix is defined as follows

$$\Omega_s = \text{Cov}[r|s_1, s_2] = s_2 \Sigma$$

Proof of Proposition 3

The unconditional expected mean-variance utility of the informed mean-variance investor under the assumptions of stochastic volatility and skew under Model 2

The conditional expected return of the informed investor is given by

$$E_I[r|s_2] = s_2^{1/2} \mu$$

The conditional covariance matrix is defined as follows

$$\text{Cov}_I[r|s_2] = s_2 (\mu \mu' \text{Var}(s_1) + \Sigma)$$

The optimal weights are given by

$$\hat{\omega} = \frac{(\mu \mu' \text{Var}(s_1) + \Sigma)^{-1} \mu}{s_2^{1/2} \lambda} \quad (37)$$

$$E_O[r_p|s_2] = \frac{\mu' (\mu \mu' \text{Var}(s_1) + \Sigma)^{-1} \mu}{\lambda}$$

Thus, the expected return is independent of the second state variable, s_2 .

Now,



$$\begin{aligned}\Omega^{-1} &= \frac{1}{s_2} (\text{Var}(s_1) \mu \mu' + \Sigma)^{-1} \\ &= \frac{1}{s_2} \left(\Sigma^{-1} - \text{Var}(s_1) \frac{\Sigma^{-1} \mu \mu' \Sigma^{-1}}{1 + \text{Var}(s_1) \mu' \Sigma^{-1} \mu} \right)\end{aligned}$$

Thus, we have

$$\begin{aligned}E_O[U_I] &= \frac{1}{2\lambda} \mu' \left(\Sigma^{-1} - \text{Var}(s_1) \frac{\Sigma^{-1} \mu \mu' \Sigma^{-1}}{1 + \text{Var}(s_1) \mu' \Sigma^{-1} \mu} \right) \mu \\ E_O[U_I] &= \frac{1}{2\lambda} \left(\alpha - \text{Var}(s_1) \frac{\alpha^2}{1 + \text{Var}(s_1) \alpha} \right)\end{aligned}$$

Hence Proposition 3:

$$E_O[U_I] = \frac{\alpha}{2\lambda} \left(\frac{1}{1 + \text{Var}(s_1) \alpha} \right)$$

Proof of Proposition 4

The unconditional expected mean-variance utility of the uninformed mean-variance investor under the assumptions of stochastic volatility and skew under Model 2.

The optimal weights of the uninformed investor are again given by

$$w = \frac{\Sigma^{-1} \mu}{\lambda} \quad (38)$$

The conditional expected return of the uninformed investor is given by

$$E_O[r_p | s_1, s_2] = \frac{s_1 s_2^{1/2} \mu' \Sigma^{-1} \mu}{\lambda}$$

The expected risk is

$$\begin{aligned}E_O[\sigma_p^2 | s_2] &= \omega' \Omega \omega \\ &= \frac{s_2}{\lambda^2} (\Sigma^{-1} \mu)' (\text{Var}(s_1) \mu \mu' + \Sigma) (\Sigma^{-1} \mu) \\ &= \frac{s_2}{\lambda^2} (\alpha + \text{Var}(s_1) \alpha^2)\end{aligned}$$

The expected utility of the uninformed investor is hence given by

$$E_O[U_U | s_1, s_2] = \left(\frac{s_1 s_2^{1/2} \alpha}{\lambda} - \frac{s_2}{2\lambda} (\alpha + \text{Var}(s_1) \alpha^2) \right)$$

Since $E[s_1] = 1$, we have Proposition 4:

$$E_O[U_U] = \frac{\alpha}{2\lambda} \left(2E[s_2^{1/2}] - (1 + \text{Var}(s_1) \alpha) \right)$$

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agarwal, V., and N.Y. Naik. 2004. Risks and portfolio decisions involving hedge funds. *Review of Financial Studies* 17: 63–98.
- Ait-Sahalia, Y., and M.W. Brandt. 2001. Variable selection for portfolio choice. *Journal of Finance* 56: 1297–1351.
- Alcock, J., and S. Satchell, (eds.). 2018. *Asymmetric dependence in finance: diversification, correlation and portfolio management in market downturns*. Hoboken: Wiley.
- Alfarano, S., T. Lux, and F. Wagner. 2005. Estimation of agent-based models: The case of an asymmetric herding model. *Computational Economics* 26: 19–49.
- Allen, D., C. Lizieri, and S. Satchell. 2019. In defense of portfolio optimization: What if we can forecast?. *Financial Analysts Journal* 75 (3): 20–38.
- Alles, L.A., and J.L. Kling. 1994. Regularities in the variation of skewness in asset returns. *Journal of Financial Research* 17: 427–438.
- Amenc, N., F. Goltz, and A. Lioui. 2011. Practitioner portfolio construction and performance measurement: Evidence from Europe. *Financial Analysts Journal* 67: 39–50.
- Andersen, T.G., and T. Bollerslev. 1998. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39: 885–905.
- Andersen, T.G., T. Bollerslev, F.X. Diebold, and H. Ebens. 2001. The distribution of realized stock return volatility. *Journal of Financial Economics* 61: 43–76.
- Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys. 2003. Modeling and forecasting realized volatility. *Econometrica* 71: 579–625.
- Ang, A., and G. Bekaert. 2002. International asset allocation with regime shifts. *Review of Financial Studies* 15: 1137–1187.
- Ang, A., and J. Chen. 2002. Asymmetric correlations of equity portfolios. *Journal of Financial Economics* 63: 443–494.
- Baillie, R.T., and T. Bollerslev. 1989. The message in daily exchange rates: A conditional-variance tale. *Journal of Business & Economic Statistics* 7: 297–305.
- Basak, Suleyman, and Georgy Chabakauri. 2010. Dynamic mean-variance asset allocation. *Review of Financial Studies* 23: 2970–3016.
- Beedles, W.L. 1979. Asymmetry of market returns. *Journal of Financial and Quantitative Analysis* 14: 653–660.
- Bloomfield, T., R. Leftwich, and J.B. Long. 1977. Portfolio strategies and performance. *Journal of Financial Economics* 5: 201–218.
- Brandt, M.W. 1999. Estimating portfolio and consumption choice: A conditional Euler equations approach. *Journal of Finance* 54: 1609–1645.
- Brandt, M.W., A. Goyal, P. Santa-Clara, and J.R. Stroud. 2005. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Review of Financial Studies* 18: 831–873.



- Brandt, M.W., and Q. Kang. 2004. On the relationship between the conditional mean and volatility of stock returns: A latent VaR approach. *Journal of Financial Economics* 72: 217–257.
- Brennan, M.J., E.S. Schwartz, and R. Lagnado. 1997. Strategic asset allocation. *Journal of Economic Dynamics & Control* 21: 1377–1403.
- Busse, J.A. 1999. Volatility timing in mutual funds: Evidence from daily returns. *Review of Financial Studies* 12: 1009–1041.
- Campbell, J.Y. 1987. Stock returns and the term structure. *Journal of Financial Economics* 18: 373–399.
- Chacko, G., and L.M. Viceira. 2005. Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets. *Review of Financial Studies* 18: 1369–1402.
- Coles, J.L., U. Loewenstein, and J. Suay. 1995. On equilibrium pricing under parameter uncertainty. *Journal of Financial and Quantitative Analysis* 30: 347–364.
- Cremers, J.H., M. Kritzman, and S. Page. 2005. Optimal hedge fund allocations - do higher moments matter?. *Journal of Portfolio Management* 31: 70.
- Das, S.R., and R. Uppal. 2004. Systemic risk and international portfolio choice. *Journal of Finance* 59: 2809–2834.
- DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal. 2009. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?. *Review of Financial Studies* 22: 1915–1953.
- Dittmar, R.F. 2002. Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *Journal of Finance* 57: 369–403.
- Duchin, R., and H. Levy. 2009. Markowitz versus the Talmudic portfolio diversification strategies. *Journal of Portfolio Management* 35: 71.
- Engle, R.F. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of united-kingdom inflation. *Econometrica* 50: 987–1007.
- Fabozzi, F.J., S. Focardi, and C. Jonas. 2007. Trends in quantitative equity management: Survey results. *Quantitative Finance* 7: 115–122.
- Fama, E.F. 1963. Mandelbrot and the stable paretian hypothesis. *Journal of Business* 36: 420–429.
- Fama, E.F. 1965. The behavior of stock-market prices. *Journal of Business* 38: 34–105.
- Fleming, J., C. Kirby, and B. Ostdiek. 2003. The economic value of volatility timing using “Realized” Volatility. *Journal of Financial Economics* 67: 473–509.
- Franke, Reiner. 2009. Applying the method of simulated moments to estimate a small agent-based asset pricing model. *Journal of Empirical Finance* 16: 804–815.
- French, K.R., G.W. Schwert, and R.F. Stambaugh. 1987. Expected stock returns and volatility. *Journal of Financial Economics* 19: 3–29.
- Ghysels, E., P. Santa-Clara, and R. Valkanov. 2005. There is a risk-return trade-off after all. *Journal of Financial Economics* 76: 509–548.
- Gilovich, T., R. Vallone, and A. Tversky. 1985. The hot hand in basketball - on the misperception of random sequences. *Cognitive Psychology* 17: 295–314.
- Glosten, L.R., R. Jagannathan, and D.E. Runkle. 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48: 1779–1801.
- Gomes, Francisco J. 2007. Exploiting short-run predictability. *Journal of Banking & Finance* 31: 1427–1440.
- Hall, A.R. 2005. *Generalized method of moments*. Oxford: Oxford University Press.
- Han, Y.F. 2006. Asset allocation with a high dimensional latent factor stochastic volatility model. *Review of Financial Studies* 19: 237–271.
- Harvey, C.R., and A. Siddique. 2000. Conditional skewness in asset pricing tests. *Journal of Finance* 55: 1263–1295.
- Hwang, S.S., and S.E. Satchell. 1999. Modelling emerging market risk premia using higher moments. *International Journal of Finance & Economics* 4: 271–296.
- Ingersoll, J.E., 1987. *Theory of financial decision making* (Rowman & Littlefield Publishers).
- Jagannathan, R., and T.S. Ma. 2003. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *Journal of Finance* 58: 1651–1683.
- Jorion, P. 1991. Bayesian and capm estimators of the means - implications for portfolio selection. *Journal of Banking & Finance* 15: 717–727.
- Kan, Raymond, and Guofu Zhou. 2007. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis* 42: 621–656.
- Kim, Sangbae, and Francis In. 2012. False discoveries in volatility timing of mutual funds. *Journal of Banking & Finance* 36: 2083–2094.
- Kirby, C., and B. Ostdiek. 2012. It’s all in the timing: Simple active portfolio strategies that outperform naive diversification. *Journal of Financial and Quantitative Analysis* 47: 437–467.
- Kritzman, M., and H. Markowitz. 2017. An Interview with Nobel Laureate Harry M. Markowitz. *Financial Analysts Journal* 73 (4): 16–21.
- Kroll, Y., H. Levy, and H.M. Markowitz. 1984. Mean-variance versus direct utility maximization. *Journal of Finance* 39: 47–61.
- Levy, H., and H.M. Markowitz. 1979. Approximating expected utility by a function of mean and variance. *American Economic Review* 69: 308–317.
- Liu, J., F.A. Longstaff, and J. Pan. 2003. Dynamic asset allocation with event risk. *Journal of Finance* 58: 231–259.
- Lizieri, Colin, Stephen Satchell, and Qi. Zhang. 2007. The underlying return-generating factors for reit returns: An application of independent component analysis. *Real Estate Economics* 35: 569–598.
- Longstaff, F.A. 2001. Optimal portfolio choice and the valuation of illiquid securities. *Review of Financial Studies* 14: 407–431.
- Mandelbrot, B. 1963. The variation of certain speculative prices. *Journal of Business* 36: 394–419.
- Markowitz, Harry. 1952. Portfolio selection. *Journal of Finance* 7: 77–91.
- McNeil, A, R Frey, and P Embrechts, 2005. *Quantitative risk management: Concepts, techniques and tools* (Princeton University Press).
- Merton, R.C. 1971. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 3: 373–413.
- Merton, R.C. 1980. On estimating the expected return on the market - an exploratory investigation. *Journal of Financial Economics* 8: 323–361.
- Mitchell, M.L., and J.H. Mulherin. 1994. The impact of public information on the stock-market. *Journal of Finance* 49: 923–950.
- Mitton, Todd, and Keith Vorkink. 2007. Equilibrium underdiversification and the preference for skewness. *Review of Financial Studies* 20: 1255–1288.
- Nelson, D.B. 1990. Arch models as diffusion approximations. *Journal of Econometrics* 45: 7–38.
- Nelson, D.B. 1992. Filtering and forecasting with misspecified arch models.1. Getting the right variance with the wrong model. *Journal of Econometrics* 52: 61–90.
- Newey, W.K., and K.D. West. 1987. A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance-matrix. *Econometrica* 55: 703–708.
- Pindyck, R.S. 1984. Risk, inflation, and the stock-market. *American Economic Review* 74: 335–351.



- Pulley, L.B. 1981. A general mean-variance approximation to expected utility for short holding periods. *Journal of Financial and Quantitative Analysis* 16: 361–373.
- Rogers, L., S. Satchell, and Y. Yoon. 1994. Estimating the volatility of stock prices: A comparison of methods that use high and low prices. *Applied Financial Economics* 4: 241–247.
- Rosenberg, B., and J.A. Ohlson. 1976. Stationary distribution of returns and portfolio separation in capital-markets - fundamental contradiction. *Journal of Financial and Quantitative Analysis* 11: 393–402.
- Rothschild, M., and J.E. Stiglitz. 1970. Increasing risk.1. *Definition*, *Journal of Economic Theory* 2: 225–243.
- Samuelson, P.A. 1970. Fundamental approximation theorem of portfolio analysis in terms of means, variances and higher moments. *Review of Economic Studies* 37: 537–542.
- Schwert, G.W. 1989. Why does stock-market volatility change over time. *Journal of Finance* 44: 1115–1153.
- Scott, R.C., and P.A. Horvath. 1980. On the direction of preference for moments of higher-order than the variance. *Journal of Finance* 35: 915–919.
- Shiller, R.J. 1981. Do stock-prices move too much to be justified by subsequent changes in dividends. *American Economic Review* 71: 421–436.
- Simaan, Y. 1993. What is the opportunity cost of mean-variance investment strategies. *Management Science* 39: 578–587.
- Svenson, O. 1981. Are we all less risky and more skillful than our fellow drivers. *Acta Psychologica* 47: 143–148.
- Veronesi, P. 1999. Stock market overreaction to bad news in good times: A rational expectations equilibrium model. *Review of Financial Studies* 12: 975–1007.
- Weiss, A. 1984. Arma models with arch errors. *Journal of Time Series Analysis* 5: 129–143.
- Westerfield, J.M. 1977. Examination of foreign-exchange risk under fixed and floating rate regimes. *Journal of International Economics* 7: 181–200.
- Whitelaw, R.F. 1994. Time variations and covariations in the expectation and volatility of stock-market returns. *Journal of Finance* 49: 515–541.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

David Allen has a PhD from the University of Cambridge. He is head of long/short strategies at Plato Investment Management, Sydney, Australia, previously worked in quantitative finance for JP Morgan and is an Honorary Associate of the University of Technology Sydney.

Stephen Satchell has PhDs from London School of Economics and Cambridge University. He has supervised over 50 PhD students and published many books and articles. His research activities are focused on finance and econometrics. He spends much of his time participating in practitioner research projects.

Colin Lizieri has a PhD from the London School of Economics and has taught at City University, London, the University of Reading and, since 2009, the University of Cambridge, where he was Grosvenor Professor of Real Estate Finance. He is the director of the Cambridge Real Estate Research Centre. His research focuses on commercial real estate finance and investment, global capital flows and urban development, publishing extensively in those areas. He has served on the research committees of many professional and industry bodies and chaired a World Economic Forum industry panel on the Future of Real Estate and Urbanization.

