

Non-normality of market returns

A framework for asset allocation decision-making

May 2009



Author

Abdullah Z. Sheikh, FIA, FSA
Director of Research
Strategic Investment Advisory Group
abdullah.z.sheikh@jpmorgan.com



Author

Hongtao Qiao, FRM
Strategic Advisor
Strategic Investment Advisory Group
hongtao.j.qiao@jpmorgan.com

Executive summary

In this paper, the authors investigate non-normality of market returns, as well as its potential impact on portfolio efficiency and the asset allocation process. The main findings are:

- Extreme negative events—due to “non-normality” of asset returns—are observed with much higher frequency than current risk frameworks allow for.
- As a result, traditional asset allocation frameworks—which are based on assumptions of normality in asset returns—can significantly understate portfolio downside risk.
- Using advanced statistical methods, however, risk frameworks can be constructed to better incorporate and account for many types of non-normality.
- In such a “non-normal” risk framework, the authors argue that Conditional Value at Risk (CVaR₉₅) is a more helpful risk quantifier than standard deviation.
- The results of the analysis indicate that asset allocation frameworks incorporating “non-normality” can improve portfolio efficiency in terms of both CVaR₉₅ and Sharpe ratio.

Non-normality of market returns

Market stress and disruptions can have serious consequences for investors, as we have seen during the past year. Admittedly, such events are rare and unpredictable. We believe, however, that risk management frameworks can be modified to better capture the long-term, downside risk associated with these kinds of anomalies.

In broad scope, two specific weaknesses in conventional risk assessment lead to quantifiable underestimation of portfolio risk:

- **Frameworks assuming “normality”¹:** In broad scope, conventional asset allocation frameworks make a range of assumptions about the “normality” of asset returns, the most problematic of which are that returns are independent from period to period² and normally distributed. In reality, we can observe that in many cases returns are not independent, and in all cases they are not normally distributed.
- **Inadequate risk measures:** If one adopts an asset allocation framework that incorporates non-normality, then standard deviation becomes ineffective as the primary quantifier of portfolio risk.

With the latest statistical methods, however, these shortcomings can be addressed. And based on the results presented here, we argue that a modified risk framework may help investors improve portfolio efficiency and resiliency.

Non-normality and its impact on portfolio risk

As sudden market disruptions go, the events of the past year – the sub-prime mortgage crisis and ensuing global financial meltdown—are not without precedent. Just in the last three decades, investors have been faced with a number of financial crises:

- Latin American debt crisis in the early 1980s
- Stock market crash of 1987
- US Savings & Loan crisis in 1989 - 1991
- Western European exchange rate mechanism crisis in 1992
- East Asian financial crisis of 1997
- Russian default crisis and the LTCM Hedge Fund crisis in 1998
- Bursting of US Technology bubble in 2000-2001

¹ The normal distribution – recognisable by its bell-shaped probability density function – is a statistical distribution commonly used to model asset class returns in traditional Mean-Variance Optimisation frameworks.

² Though “independence” is not in strict statistical terms a form of “normality”, we include it here because the assumption of “independence” is one of the central tenets of conventional asset allocation frameworks built around the concept of “normality” of asset returns.

³ In this context, the converse is also true i.e. empirical observations of non-normality are the result of extreme market scenarios.

While these anomalous events are rare, we observe such extreme “non-normality” in real-world markets more frequently than current risk management approaches allow for. Said another way, we believe that conventionally derived portfolios carry a higher level of downside risk than many investors believe, or current portfolio modeling techniques can identify.

The primary reason for this underestimation of risk lies in the conventional approach to applying mean-variance theory, which was pioneered by Harry Markowitz in 1952. Traditional mean-variance frameworks have become the bedrock of top-down asset allocation decision-making. As suggested above, a standard assumption in the mean-variance framework, and indeed many other holistic asset allocation frameworks, is that future asset class returns will be independent from period to period and normally distributed.

Despite being widely recognised as overly simplistic, such broad assumptions of “normality” have appeal due to the ease with which they can be implemented. To implement an asset allocation framework based on normal asset return distributions, practitioners need only make two assumptions for each asset class (mean and standard deviation) and one assumption for each pair (co-variance). The latter applies because one assumes the relationship between each pair of asset classes is linear (another problematic issue discussed later in the paper).

But what if asset returns are not normally distributed?

In fact, in the real world, we can observe empirically that returns are not normally distributed. This leads us to ask: How would incorporating non-normality affect the strategic asset allocation process?

Identifying non-normality

Each of the market events mentioned above was due to an incidence of non-normality of one sort or another³. And the practical implication is that incorporating non-normality into the asset allocation process would force recognition of greater downside risk to the portfolio, precisely from such extreme, unexpected negative events.

Non-normality of market returns

Our focus here is on capturing the impact of non-normality on downside portfolio risk—as well as the asset allocation/optimisation process—rather than identifying the specific sources of non-normality itself (although we do discuss economic or behavioral factors where relevant).

Specifically, we test seven asset classes⁴ and confirm three primary categories of non-normality:

- **Serial correlation:** A critical pillar of many traditional asset allocation frameworks (i.e. frameworks built on a premise of “normality”), is the assumption that asset returns from period to period are independent and identically distributed. However, if one month’s return is ‘influenced’⁵ by the previous month’s return, then there may be a need to account for this effect in future asset projections. Typically, traditional asset allocation frameworks do not allow for serial correlation, but we identify serial correlation in four of the seven asset classes we model.

Serial correlation is often a consequence of the illiquidity and hard-to-price nature of the underlying assets. For example, certain alternative investment strategies—such as Hedge Fund of Funds and Private Equity—show evidence of serial correlation. The difficulty in valuing the underlying assets at regular intervals requires managers or administrators to estimate prices (for example, with reference to the closest marketable security or based on certain economic indicators).

If current asset prices are derived, for example, by updating last month’s asset prices (after allowing for changes in the economic environment since the last valuation), then serial correlation reflects a gradual (rather than instantaneous) recognition of the true underlying value of the asset.

Serial correlation, if not adjusted for in the underlying data, masks true asset class volatility and biases risk estimates downwards, leading to underestimation of overall portfolio risk.

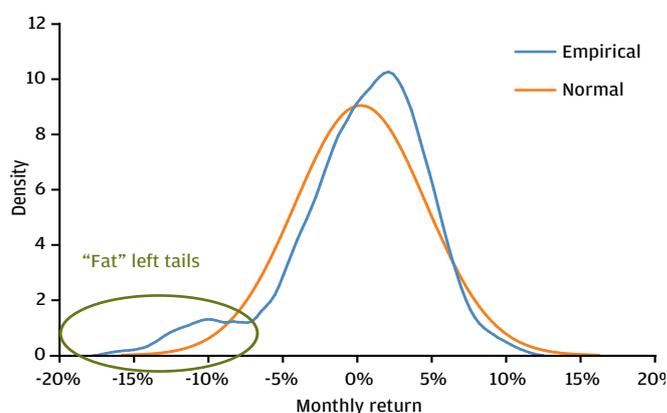
⁴ Asset classes include US Aggregate bonds, US Large Cap Equity, International Equity, Emerging Markets Equity, Real Estate Investment Trusts, Hedge Fund of Funds and Private Equity.

⁵ Influenced in this context refers to a statistically significant coefficient when one month’s return is regressed against the previous month’s return.

- **“Fat” left tails (negative skewness and leptokurtosis):** Our second form of non-normality relates to observing negative returns in greater magnitude and with a higher probability than implied by the normal distribution. This phenomenon is commonly referred to as “fat” left tails.

Exhibit 1 illustrates this phenomenon for monthly dollar-hedged International Equity returns over the ten years to October 2008. The chart plots the empirical (i.e. observed) probability density function of the data relative to a normal distribution.

EXHIBIT 1: INTERNATIONAL EQUITIES—“FAT” LEFT TAILS IN HISTORICAL RETURNS



Source: J.P. Morgan Asset Management. For illustrative purposes only.

One can see that the observed return series (blue line) is more peaked, has a higher density at the extreme left, and leans further to the right than the normal distribution (orange line).

The rightward lean is its “negative skewness.” The consequence of this skewness is that the left slope of the blue line is longer than the left slope of the orange line—i.e. it has a longer tail, which indicates a greater magnitude of extreme negative events.

In addition, the blue line is taller at its apex and shows a higher density at the extreme left end (i.e. leptokurtosis). In particular, the higher density at the left tail indicates a higher probability of extreme negative events.

Non-normality of market returns

An asset allocation framework based on the normal distribution will understate both the frequency and magnitude of extreme negative events, as well as their potential effects on portfolio returns and efficiency.

- **Correlation breakdown:** The simple correlations often used in traditional asset allocation models assume a linear relationship between asset classes—i.e. assume that the relationship between the variables at the extremes is similar to their relationship at less extreme values. Using simple linear correlations is the equivalent of assuming that the ‘joint’ distribution of asset returns is (multivariate) normal. ‘Joint’ distributions capture how asset classes behave together rather than individually.

However, we find that in many cases correlations under extreme conditions are quite different than under normal conditions. In other words, the expected linear correlations breakdown and asset classes exhibit quite different joint behavior. The relationships, in fact, are not linear and the assumption of linearity (by using linear correlation matrices) underestimates the probability of ‘joint’ negative returns under extreme conditions.

Relying on linear correlation matrices tends to overestimate the benefits of portfolio diversification during periods of high market volatility. This leads to a systematic underestimation of downside portfolio risk.

Statistical approaches for incorporating non-normality

Fortunately, we have at our disposal sophisticated statistical tools that allow us to correct for these types of non-normality.

- **Unsmoothing serial correlation:** Serial correlations can be “unsmoothed”. That is, we can correct for the influence of prior-period returns and restore independence to single-period returns. To arrive at our ‘new’ adjusted return series, we apply a variation of Fisher-Geltner-Webb’s well-established ‘unsmoothing’ methodology⁶.

⁶ For a more detailed treatment, please refer to Fisher, J., D. Geltner, and B. Webb. 1994. Value Indices of Commercial Real Estate: A Comparison of Index Construction Methods. Also, Fisher, J. and D. Geltner. 2000. De-Lagging the NCREIF Index: Transaction Prices and Reverse-Engineering.

⁷ Copulas have been applied extensively to the pricing of Collateralised Debt Obligations—in addition to other areas in finance. For a detailed treatment, please refer to “An Introduction to Copulas” by Nelson R. R. 1999.

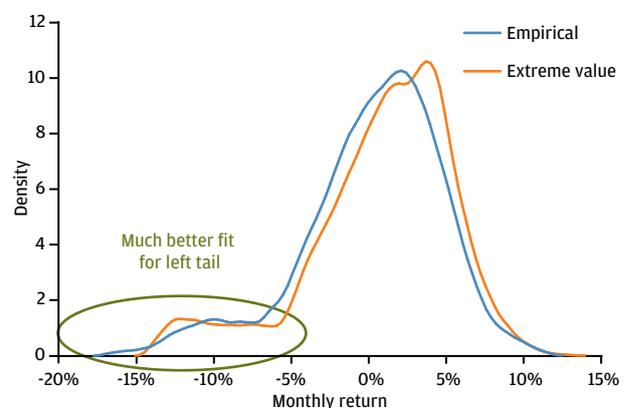
- Our ‘new’ adjusted return series is better reflective of the risk characteristics of the asset class. Notably, the new unsmoothed return series has the same mean as our original return series, but shows a higher volatility, and thus higher downside risk.

- **Modeling “fat” left tails (negatives skewness and leptokurtosis):** “Fat” left tails can be addressed using Extreme Value theory—a body of work specifically designed to look at the probability of high-risk, but low-probability, events such as floods, earthquakes and large insurance losses. In other words, its focus is estimating ‘tail’ risk.

By applying Extreme Value theory we can create asset return distribution models that are a closer “fit” to the return series we observe, much more similar than normal distributions.

Exhibit 2 shows a probability density function for dollar-hedged International Equities, calculated using Extreme Value Theory (orange line).

EXHIBIT 2: INTERNATIONAL EQUITIES—APPLYING EXTREME VALUE THEORY FOR A BETTER FIT



Source: J.P. Morgan Asset Management. For illustrative purposes only.

It is a much closer approximation to empirically observed performance in terms of negative skewness (rightward tilt) and leptokurtosis (“fat” left tails). This process can be applied to all assets in the portfolio, the overall result being an increase in the portfolio’s downside risk profile.

- **Simulating correlation breakdown:** Finally, we can address the issue of correlation breakdown using ‘copula’ theory⁷—a body of work that explicitly looks at the impact of contagion or converging correlations at the total portfolio level.

Non-normality of market returns

Copulas are mathematical functions that allow us to model the joint distribution of asset returns separately from the marginal (i.e. individual asset class) distributions.

By considering ‘joint’ distributions, we turn our focus to how asset classes behave together rather than individually. In particular, copulas allow us to model an increased incidence of joint negative returns (i.e. the “fatter” joint left tails) in our simulation results, just as we observe empirically in real-world market data. And again, using a more accurate proxy for observed data results in recognising higher downside portfolio risk, specifically due to increased dependence of asset returns during periods of market stress.

A better risk quantifier: conditional value at risk

We believe that empirical evidence suggests an imperative to incorporate various types of “non-normality” into the asset allocation and portfolio modeling process, specifically to better understand and model downside portfolio risk. Yet if we take this step, we have to ask whether or not our conventional risk measure (i.e. standard deviation) is up to the new task.

We would argue that, in a framework based on non-normality, standard deviation may not be investors’ most appropriate measure of portfolio risk because it rewards the desirable upside movements as hard as it punishes the undesirable downside movements. This is generally inconsistent with investor risk preferences—primarily as observed in the field of behavioral finance⁸.

Conditional Value at Risk (CVaR₉₅) overcomes many of the drawbacks of standard deviation as a risk measure. Primarily, as it only measures risk on the downside, it captures both the asymmetric risk preferences of investors and the incidence of “fat” left tails induced by skewed and leptokurtic return distributions. Further, given the widespread use by major institutional investors and regulators of its first cousin—Value at Risk—we judge it to be the most appropriate risk measure to incorporate it into our framework.

We define Conditional Value at Risk (CVaR₉₅) as the average real⁹ portfolio loss (or gain) relative to the starting portfolio value in the worst five percent of scenarios, based on our 10,000 Monte Carlo simulations. It is simply the average real loss (or gain) in the worst 500 (5% of 10,000) scenarios i.e. the left tail of the portfolio loss (gain) distribution.

Incorporating non-normality: the impact on portfolio risk

As discussed above, with the right statistical tools, one can develop an asset allocation framework that incorporates non-normality. The next question is: How would this affect the asset allocation and optimisation process?

To illustrate the impact, we apply our non-normal framework to a hypothetical US domiciled investor with a well-diversified portfolio: initial value equals USD 1 billion, with allocations across our seven major asset classes (See **Exhibit 3** below).

EXHIBIT 3: HYPOTHETICAL PORTFOLIO ALLOCATION

Asset class	Current allocation
Total bonds	30%
US Aggregate bonds	30%
Total equity	55%
US Large Cap Equity	40%
International Equity (hedged)	10%
Emerging Markets Equity	5%
Total alternatives	15%
Real Estate Investment Trusts (REITs)	5%
Hedge Fund of Funds	5%
Private Equity	5%
Key statistics	
Expected arithmetic return	9.1%
Expected volatility	10.0%
Expected compound return	8.7%
Sharpe ratio	0.51

Source: J.P. Morgan Asset Management. For illustrative purposes only. Sharpe ratio calculated based on expected risk free return of 4.0% per year as per J.P. Morgan Asset Management. Long Term Capital Market Assumptions (please see Appendix for details).

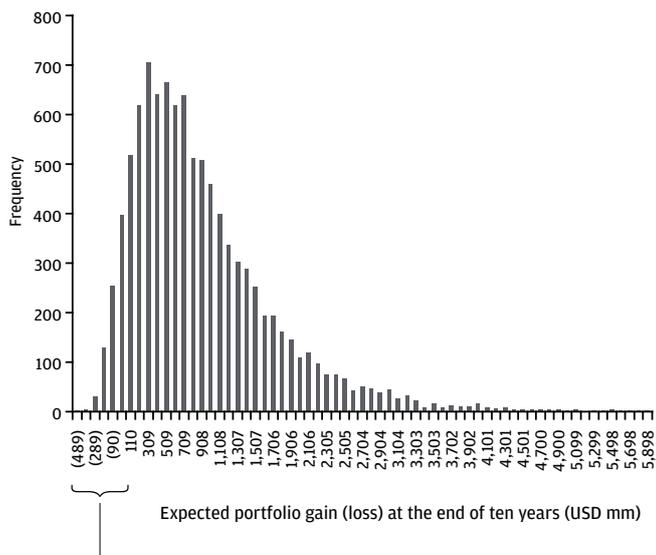
⁸ A key tenet of behavioral finance is the idea of loss aversion i.e. a tendency of investors to prefer avoiding losses than making gains. This translates to risk preferences that are asymmetric in nature.

⁹ Our model calculates real portfolio value by discounting the nominal portfolio value using projected inflation. Inflation itself is projected stochastically in our framework.

Non-normality of market returns

Using our revised framework—which incorporates non-normality—we calculated the $CVaR_{95}$ for the hypothetical portfolio. **Exhibit 4** shows the projected frequency of the portfolio’s real gains and losses at the end of ten years.

EXHIBIT 4: HISTOGRAM OF PROJECTED CUMULATIVE PORTFOLIO GAIN (LOSS) AT THE END OF TEN YEARS ASSUMING NON-NORMALITY



$CVaR_{95}$ is defined as the average (real) cumulative loss in the worst 5% or 500 simulations. This is equal to USD 168 million for the current portfolio.
Source: J.P. Morgan Asset Management. For illustrative purposes only.

The $CVaR_{95}$ of the current allocation, based on our new methodology, is USD 168 million. In other words, the portfolio can expect to lose (on average) USD 168 million in the worst five percent of cases (based on our simulation results). This risk is significant. It indicates a real return (i.e. after allowing for inflation) of -16.8% on the portfolio over an extended time horizon—a result our investor is unlikely to be very happy with.

For the same portfolio, however, risk calculations that assume normality¹⁰ would result in a $CVaR_{95}$ figure of USD 74 million. *Incorporating non-normality more than doubles our prior estimate of $CVaR_{95}$.* In absolute terms, the risk underestimation is USD 94 million or 9.4% of the portfolio’s initial value.

We should note that while the increased risk associated with modeling non-normality is striking, incremental risk in itself

is not necessarily a reason for changing a plan’s asset allocation (unless an absolute threshold value has been breached). More specifically, if one assumes an arbitrarily (but equally) higher downside risk for all asset classes, this in itself, would not impact the efficiency of the portfolio—i.e. its efficiency would not change at all, albeit the $CVaR_{95}$ would now be higher.

The reason non-normality can impact asset allocation is that the downside risk associated with various asset classes is very different. Most obviously, equity and equity type asset classes entail greater degrees of downside risk than, for example, fixed income-type investments. Hence, because the downside risk characteristics of various asset classes are different—and cannot be accounted for using traditional modeling techniques or risk measures such as standard deviation—the efficient allocations in our $CVaR_{95}$ motivated non-normal framework must also be different from a traditional framework.

For this reason—not merely due to higher $CVaR_{95}$ figures—we believe investors need to quantitatively incorporate the impact of non-normality into the asset allocation process.

Incorporating non-normality: the impact on asset allocation

To assess the potential impact that a revised framework would have on asset allocation, we used the revised framework—incorporating non-normality—to create an optimised portfolio. The optimiser minimizes $CVaR_{95}$ for an equivalent target return (9.1%), the same seven asset classes, and no external constraints.

Exhibit 5 compares the initial hypothetical portfolio with the portfolio optimised using the revised framework incorporating non-normality. For illustrative purposes, we also show the optimised allocation derived from a traditional mean-variance framework.

Our results indicate that the ‘optimal’ portfolio using a traditional mean-variance framework actually *increases* (rather than decreases) risk by 22.6% relative to the current allocation – as defined by $CVaR_{95}$. As a traditional framework minimises standard deviation¹¹—which we argue is an inadequate risk measure – it inadvertently exposes our investor to even worse scenarios on the downside than the current allocation.

¹⁰ Our estimates of $CVaR_{95}$ under a traditional framework were derived with an identical asset allocation, except that asset returns were assumed to be individually and jointly normally distributed. Risk and correlations were derived based on the same ten year historical period November 1998 to October 2008.

¹¹ Minimizing standard deviation for a given expected return target is the equivalent of maximizing Sharpe ratio.

Non-normality of market returns

However, the more significant issue by far—and the biggest drawback of the traditional approach – is that it produces a highly concentrated and impractical asset allocation. This is because in the absence of formal constraints, it over allocates to asset classes based on small differences in assumptions. This limits our ability to draw useful insights into the portfolio construction process, using such a framework.

On the other hand, our non-normal CVaR₉₅ based framework produces a diversified solution with allocations across the asset class spectrum. No single asset class significantly dominates the portfolio. Despite the fact that our investor already holds quite a diversified portfolio, our framework suggests that there is still scope for our investor to improve portfolio efficiency further.

EXHIBIT 5: OPTIMISATION RESULTS

	Current allocation	Optimised, Unconstrained Normal allocation	Optimised, Unconstrained Non-normal allocation
Total bonds	30.0%	0.0%	34.5%
US Aggregate Bonds	30.0%	0.0%	34.5%
Total equity	55.0%	18.3%	36.0%
US Large Cap Equity	40.0%	8.9%	21.7%
International Equity	10.0%	7.6%	5.8%
Emerging Markets Equity	5.0%	1.9%	8.5%
Total alternatives	15.0%	81.7%	29.5%
Real Estate Investment Trusts REITs	5.0%	11.9%	11.1%
Hedge Fund of Funds	5.0%	64.1%	7.0%
Private Equity	5.0%	5.6%	11.5%
Total	100.0%	100.0%	100.0%
Key statistics			
Target expected arithmetic return	9.1%	9.1%	9.1%
Expected volatility	10.0%	8.6%	9.5%
Expected compound return	8.7%	8.6%	8.7%
Sharpe ratio	0.51	0.59	0.54
CVaR ₉₅ (\$ million) allowing for non-normality	USD 168 million	USD 206 million	USD 148 million
∆CVaR ₉₅ vs current allocation	-	23% ↑	12% ↓
Return per unit of CVaR ₉₅	0.30	0.25	0.35

Source: J.P. Morgan Asset Management. For illustrative purposes only. Sharpe ratio calculated based on expected risk free return of 4.0%.

The “optimal” portfolio identified by the non-normal framework improves *both the expected Sharpe ratio and reduces the CVaR₉₅* relative to the current allocation. This signifies that our optimal portfolio is more efficient (in Sharpe ratio and CVaR₉₅ space) than the current hypothetical portfolio. Based on our Long Term Capital Market Return assumptions, the allocations to fixed income and alternatives increase, while the allocation to equities decreases.

Incorporating non-normality: a way toward more efficient portfolios

Until recently, investors have been constrained in their ability to incorporate non-normality into the asset allocation process. But now, with the availability of sophisticated statistical tools, we can meet this challenge. Why should we change?

The most straightforward answer is that this is how the world really works—i.e. we observe non-normality with much greater frequency than current frameworks allow for. The more important answer is that ignoring non-normality in equity (and equity-type) return distributions significantly understates downside portfolio risk—in the worst of the worst-cases, potentially posing a solvency risk for the investor.

We also believe that investors need to allow for downside risk in a more robust fashion than standard deviation measures have traditionally assumed. For this reason we recommend CVaR₉₅. We believe this measure is a better fit for investors’ asymmetric risk preferences, as well as the “fat” left tails recognised by non-normal asset allocation frameworks.

Finally, an asset allocation incorporating non-normality has the benefit of reducing the need for external constraints. Investors impose such constraints in an effort to get normal frameworks to provide non-normal solutions—i.e. to better reflect the non-normality we see in the real world. A framework that builds in non-normality up front, however, provides a much more direct, efficient, and elegant way of solving the problem.

Non-normality of market returns

Ultimately, we believe the quantitative results illustrate the point best, and speak for themselves: incorporating non-normality may reduce the portfolio's volatility, improve its efficiency (Sharpe ratio), and reduce its risk relative to unpredictable, extreme negative events.

So, we argue for a revised asset allocation framework because beyond its pure statistical merit, there lies a significant, practical benefit for investors: the potential to improve portfolio efficiency and resilience, in light of a clearer understanding of portfolio risk.

Limitations of reliance

It should be noted that a quantitative framework is only one input into the asset allocation process and cannot replace the professional skill and judgment necessary to arrive at an appropriate strategy. The importance of allowing for subjective—and often qualitative—factors in decision making remains. Further, there is always an explicit need to account for the investor's specific circumstances, including liabilities, when arriving at an appropriate portfolio allocation.

IMPORTANT DISCLAIMER

This document is intended solely to report on various investment views held by J.P. Morgan Asset Management. All charts and graphs are shown for illustrative purposes only. Opinions, estimates, forecasts, and statements of financial market trends that are based on current market conditions constitute our judgment and are subject to change without notice. We believe the information provided here is reliable but should not be assumed to be accurate or complete. The views and strategies described may not be suitable for all investors. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not intended to be, and should not be interpreted as, recommendations. Indices do not include fees or operating expenses and are not available for actual investment. The information contained herein employs proprietary projections of expected returns as well as estimates of their future volatility. The relative relationships and forecasts contained herein are based upon proprietary research and are developed through analysis of historical data and capital markets theory. These estimates have certain inherent limitations, and unlike an actual performance record, they do not reflect actual trading, liquidity constraints, fees or other costs. References to future net returns are not promises or even estimates of actual returns a client portfolio may achieve. The forecasts contained herein are for illustrative purposes only and are not to be relied upon as advice or interpreted as a recommendation. The value of investments and the income from them may fluctuate and your investment is not guaranteed. Both past performance and yield are no guarantee of future results. Please note current performance may be higher or lower than the performance data shown. Please note that investments in foreign markets are subject to special currency, political, and economic risks. Exchange rates may cause the value of underlying overseas investments to go down or up. Investments in emerging markets may be more volatile than other markets and the risk to your capital is therefore greater. Also, the economic and political situations may be more volatile than in established economies and these may adversely influence the value of investments made.

JPMorgan Asset Management (Europe) Société à responsabilité limitée, European Bank & Business Centre, 6 route de Trèves, L-2633 Senningerberg, Grand Duchy of Luxembourg, R.C.S. Luxembourg B27900, corporate capital EUR 10.000.000.

Issued in the UK by JPMorgan Asset Management (UK) Limited which is authorised and regulated in the UK by the Financial Services Authority. Registered in England No. 01161446. Registered address: 125 London Wall, London EC2Y 5AJ.

© 2009 JPMorgan Chase & Co. | JPM2866_Non-Normality_Short