What Predicting U.S. Elections and Stock Returns Have in Common.

Matus Mrazik, Investment Manager, systematic equities, Jupiter, and **Jonathan Katz,** Kay Sugahara Professor of Social Sciences and Statistics, Caltech find commonalities between two different fields.

July 2026

Matus Mrazik
Investment Manager, systematic equities, Jupiter

Jonathan Katz
Kay Sugahara Professor of Social Sciences and Statistics, Caltech

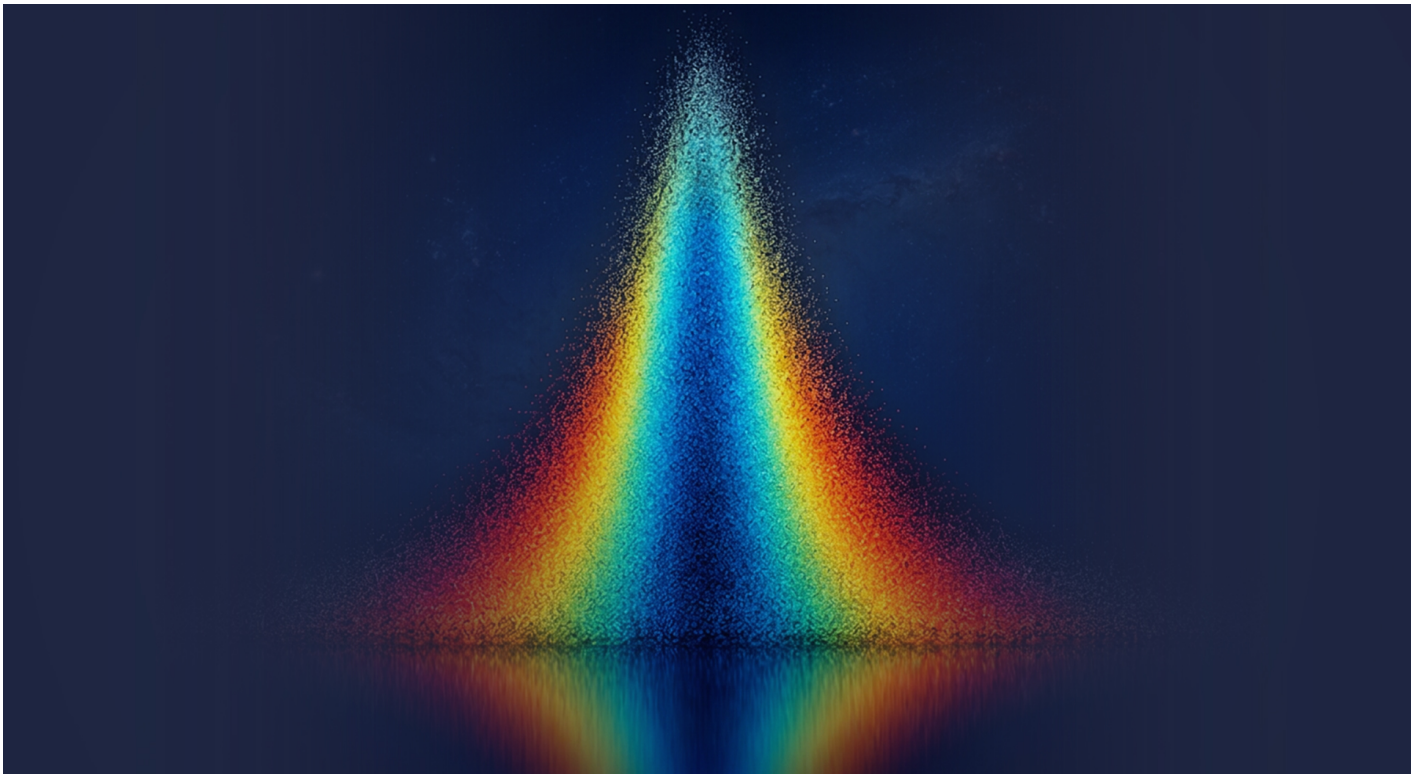

At first glance, election forecasting and quantitative investing seem to have little to do with one another. One is about candidates, voters, districts, parties, and election cycles. The other is about companies, markets, signals, portfolios, and returns. The institutions are different. The participants are different. The language is different.

But underneath the surface, the statistical problem is surprisingly similar.

In both cases, we are trying to extract a small but meaningful signal from a large amount of noisy data. Election models ask why some candidates outperform expectations across districts and election years. Investment models ask why some securities, factors, or strategies outperform others across firms and time. Both settings involve repeated observations of units we have seen before (districts, candidates, firms, markets) and both are shaped by rare but important events that are difficult to predict in advance.

This is where the problem becomes interesting. Modern data sets can be enormous, but the amount of information about the events that really matter can still be quite limited. In systematic investing, we may observe prices and signals every day, and in some cases much more frequently. But the economically important moments, such as factor reversals, regime shifts, liquidity events, credit stress, and momentum crashes, do not happen every day. Measuring more often does not necessarily create more information about these turning points. It can simply give us more observations of roughly the same environment.

The same is true in elections, and for the same reason. What markets and elections really share is a correlated structure: common events move many units at once. In markets, a single macroeconomic announcement, like a U.S. jobs report, an inflation print, or a central-bank decision, can push thousands of stocks up or down on the same day. Elections work the same way. The popularity of the sitting president can move many races together, as can shared economic conditions such as inflation and unemployment, which are the very forces that also move markets. So a model may hold thousands of district-level observations, but when they all respond to the same national environment, they carry far less independent information than their count suggests.

This distinction matters because many classical statistical tools were developed for a world in which sampling error was the main source of uncertainty. In that world, more observations usually meant more precision. That logic is still useful, but it is incomplete. In many modern forecasting problems, the bigger uncertainty is not simply whether we have estimated yesterday's relationship precisely enough. It is whether tomorrow's relationship will look the same.

That is a different kind of uncertainty. It is model uncertainty, regime uncertainty, structural uncertainty. It is the uncertainty that comes from the possibility that the world changes, or that important parts of the world were never fully measured in the first place.

This article is about that shared problem. It explains why standard methods can become overconfident when applied to repeated, noisy, time-varying systems, and why hierarchical modelling offers a useful way to think about the uncertainty that remains. In our research, we are exploring Bayesian hierarchical methods as one way to make these sources of uncertainty more explicit: across firms, across time, and across market conditions. These methods are neither new nor specific to elections. For decades, versions of them have been used across the social and behavioural sciences (public opinion, education, psychology, economics) wherever observations come in groups, repeat over time, or are otherwise correlated rather than independent. They travel so well because they are built to handle exactly that kind of structure, and are flexible enough to fit the problem at hand. Legislative elections are one long-standing application; financial returns, it turns out, are another.

For investors, this is not an abstract statistical concern. Overconfident models do not stay on the page. They become overconfident forecasts, overconfident position sizes, overconfident risk budgets, and overconfident explanations of performance.

From Sampling Uncertainty to Systematic Uncertainty

To see the problem, start with the simplest version of the investment model.

Suppose we want to explain the return of a stock tomorrow using a small number of signals we observe today. The stock might be Nvidia, IBM, or any other company in the investment universe. The signal might be momentum, valuation, profitability, leverage, analyst revisions, or a proprietary measure built by an investment team.

A compact way to write this is:

$$r_{i,t+1} = \alpha_{t+1} + x'_{i,t}\beta_{t+1} + \varepsilon_{i,t+1}$$

Here $r_{i,t+1}$ is the next-period return of stock i. The vector $x_{i,t}$ contains the signals we know at time t. The coefficient β_{t+1} tells us how much those signals are rewarded over the next period. The remaining term, $\varepsilon_{i,t+1}$, is everything the model does not explain.

This is just a regression. In fact, it is close to the kind of regression most people first encounter in an introductory statistics course. The financial version is usually estimated across the cross-section of stocks: at each date, we ask whether stocks with higher exposure to a given signal subsequently performed better or worse than stocks with lower exposure.

For example, imagine a very simple model with only one signal: value. At the end of each day or month, we rank stocks by how cheap or expensive they look. We then ask whether cheaper stocks outperformed more expensive stocks over the next period. The estimated coefficient from that regression is often called the **factor return.** If it is positive, value was rewarded over that period. If it is negative, expensive growth stocks may have outperformed cheap value stocks.

A common practitioner version of Fama–MacBeth[1] does this repeatedly. It estimates a cross-sectional regression at each date, producing a time series of realised factor returns:

$$\hat{\beta}_1, \hat{\beta}_2, \ldots, \hat{\beta}_T$$

Those estimates are then averaged and tested. In plain English, the method asks: "Across many periods, was this signal rewarded on average, and how confident are we about that average?"

This approach has been enormously useful. It is simple, transparent, and closely connected to how many investors already think about factor returns. But it was designed for a world in which the main problem was estimating relationships from limited data.

That is not always the world we face today.

In modern equity data, the cross-section can contain thousands of stocks. As a result, the factor return for any single date can appear to be estimated with great precision. If we run a regression across several thousand firms, the classical standard error can become very small. Mechanically, the model thinks it has learned a lot because it has seen many observations.

But this can be misleading. Observing more firms on the same date does not create more independent information about the state of the world. If there is a market-wide shock, a liquidity event, a central-bank surprise, or a sudden rotation from growth to value, many stocks move together. The model may have thousands of rows of data, but only one market environment.

This is the key distinction. **Sampling uncertainty** is the uncertainty that shrinks when we observe more independent data. It is the uncertainty classical standard errors were designed to measure. If we are estimating the average height of adults in a country, surveying more people genuinely reduces uncertainty.

---

[1] The Fama–MacBeth regression, from Nobel Laureate Eugene F. Fama and James D. MacBeth (1973), is a two-stage procedure. In the first stage, a time-series regression for each asset estimates its exposures (betas) to the factors. In the second stage, a cross-sectional regression in each period relates returns to those estimated betas, producing a series of risk-premium estimates that are averaged over time, with standard errors taken from their period-to-period variation. The version described here is the common practitioner one: it runs that period-by-period cross-sectional regression directly on observed firm signals, rather than on betas from a first-stage time-series regression. See Fama, Eugene F.; MacBeth, James D. (1973). "Risk, Return, and Equilibrium: Empirical Tests". Journal of Political Economy. 81 (3): 607–636. Available at https://www.jstor.org/stable/1831028

A simple example is a factor rotation. Suppose value underperforms growth for several months and then sharply recovers over one quarter. We can observe that quarter using monthly data, daily data, or intraday data. Higher frequency measurement gives us more data points, but it does not necessarily give us more independent regime changes. It may even make the historical relationship look more precise than it really is, because the observations are highly correlated versions of the same underlying event.

This is why the problem is not arithmetic. The regression calculation may be correct on its own terms. The issue is that it is answering too narrow a question. It measures how precisely we estimated a relationship under the assumptions of the model. It does not fully measure how uncertain we should be about whether that relationship will persist, reverse, or behave differently in a new market environment.

The same logic applies to elections. One of us studies U.S. House and Senate elections with two goals that mirror factor investing: not just to forecast outcomes, but to understand the forces that drive them. The data are large (thousands of district results across decades and both chambers) so a pooled regression can again report a reassuringly small sampling error while missing the uncertainty that matters, because the same correlated structure is at work. A good or bad year for a party — a national swing — moves many districts at once, much as a macroeconomic shock moves many stocks. Some incumbents win year after year by larger margins than their districts would lead you to expect: a persistent edge closer to an athlete's skill than to noise. And some of what matters most, such as a scandal, is real but nearly impossible to measure in advance.

Treat all of that as ordinary residual noise and the model misleads in two ways at once. It overstates its own precision, treating correlated observations (such as many districts moving together on a single national swing, many stocks on a single shock) as if each carried independent information. And it understates the risk of a genuine surprise, the kind of scandal or sudden shock the historical record gives little basis to rule out. In elections, the result is a forecast too confident about an unusual outcome; in investing, a model too confident that a factor will behave tomorrow as it did yesterday.

That is the deeper link between the two fields. Both contain many observations but relatively few truly independent shifts in the underlying environment, and each requires a model that distinguishes precision from understanding.

When Residuals Carry Structure

The common problem in both elections and markets is that the residual term, $\varepsilon_{i,t+1}$, in our regression equation above, is often being asked to do too much.

In a simple regression, the residual is meant to capture whatever the model cannot explain. That is fine if the unexplained part is genuinely random: small, independent, short-lived, and centred around zero. But in real forecasting problems, the unexplained part often has structure.

In markets, the residual might contain a firm-specific shock that lasts several months, a liquidity event that affects many stocks at once, a valuation reversal that unfolds over a full cycle, or a broad shock that makes one date very different from another. In elections, it might contain a national swing, a candidate's personal vote, or a scandal that affects one race in a way the measured covariates never anticipated.

None of these are simply "noise". They are not independent white-noise errors that appear for one observation and then vanish. They are persistent and sometimes shared across many units at the same time. If we force all of that structure into a single residual term, the model may still produce estimates, but it will usually sound more confident than it should.

Hierarchical models address this by giving the hidden structure somewhere to live.

A simplified version of the financial model can be written as:

$$r_{i,t+1} = \alpha + x'_{i,t}\beta_{t+1} + a_i + d_{t+1} + \varepsilon_{i,t+1}$$

This looks similar to the earlier regression, but it makes the structure explicit. The date-by-date intercept α_{t+1} of the earlier equation is now written as a common level α (the first term on the right-hand side of the new equation) plus a date effect d_{t+1}, and a persistent firm term a_i is added.

The term a_i (the third term on the right-hand side) captures persistent firm-specific effects. Some firms may repeatedly out- or under-perform what the measured signals imply (carrying a steady premium or discount the signals do not capture) for reasons the model does not directly observe.

The term d_{t+1} captures date-specific effects. Statistically, it allows one date to be different from another, even after accounting for the signals in the model. Some days or months are simply unusual. A liquidity squeeze, a central-bank announcement, a geopolitical shock, an earnings season surprise, or a broad risk-off move may affect many stocks at once. The model does not need to know the exact reason in advance. The date effect says: after accounting for the measured signals, there is still something common to this date that many firms share.

The remaining residual, $\varepsilon_{i,t+1}$, is then no longer forced to carry all of this burden. It is reserved for what is left after the model has accounted for factor payoffs, persistent firm differences, and date-specific common shocks.

This is the key conceptual move. Instead of pretending that every unexplained observation is independent noise, the model separates different types of uncertainty. Some uncertainty belongs to firms. Some belongs to dates. Some belongs to the factor relationship itself. Some remains genuinely idiosyncratic.

The same idea applies naturally to elections. Much of a district's persistent tendency is already measured (through its past election results and its presidential and gubernatorial votes) so it is captured by the variables already in the model, just as a firm's known signals are. What those variables miss is what the extra structure is for: a candidate's personal vote that lasts across several election cycles, or the national swing of a given year that moves many races together. Treating these as ordinary residual noise would miss them. A hierarchical model lets them be estimated separately.

Another way to say this is that the model sits between two extremes.

At one extreme, we could estimate every date completely separately. That is flexible, but it throws away useful information from other periods. At the other extreme, we could force all dates to share one single relationship. That is stable, but too rigid. A hierarchical model lies between these two approaches. It allows relationships to vary across time, firms, or regimes, while still pulling noisy estimates back toward a common centre when the data are weak.

This partial pooling is useful because it reflects how we actually think about markets. We do not believe every month is completely unrelated to every other month. But we also do not believe the payoff to value, momentum, or quality is perfectly constant through time. The model is allowed to learn how much variation there really is.

The quantities that govern this variation become interesting in their own right. For example, the variation in d_{t+1} tells us how much dates differ from one another after accounting for the factors. If this component is large, then calendar time is carrying information that the factor model is missing. In investment terms, that means there are common shocks affecting returns that are not captured by the signals alone.

That is the value of the hierarchical structure. It does not eliminate uncertainty. It organises it. And once uncertainty is organised, we can ask better questions: how much comes from firms, how much comes from dates, how much comes from changing factor relationships, and how much is truly residual.

More Data Is Not Always More Information

A common intuition, especially in a world shaped by high-frequency data and modern computing, is that more frequent measurement should solve the problem. If we can observe prices every day, every hour, or every few minutes, surely uncertainty should shrink. With enough observations, the noise should average out and the underlying structure should become clear.

That intuition is partly right, but only partly.

In classical statistics, more independent observations reduce **sampling uncertainty**. The simplest example is estimating an average.

Suppose there is some true average we care about, call it μ. It might be the average return to a factor, the average effect of a signal, or the average vote share for a party. We do not observe μ directly. Instead, we observe data and calculate an estimate, such as a sample average $\bar{x}$.

The uncertainty around that estimate is measured by its **standard error**. In the simplest case:

$$\mathrm{SE}(\bar{x}) \approx \frac{\sigma}{\sqrt{n}}$$

Here σ is the standard deviation of the observations (how much the individual data points vary) and n is the number of observations. The standard error is different from the standard deviation. The standard deviation describes how spread out the data are. The standard error describes how uncertain we are about our estimate of the true average, μ.

This is why the square root of n matters. If each new observation brings genuinely new information, increasing n makes the estimate more precise. Double the number of observations and the standard error falls a little. Increase the sample size by a factor of one hundred and, all else equal, the standard error falls by a factor of ten.

In the limit, the logic becomes even clearer:

$$\lim_{n\to\infty} \mathrm{SE}(\bar{x}) = \lim_{n\to\infty} \frac{\sigma}{\sqrt{n}} = 0$$

In words: if we can keep adding genuinely independent observations forever, the sampling uncertainty around the average eventually disappears.

This is why the idea is so tempting in modern data-rich settings. With millions, or even billions, of observations, n can feel almost infinite. The formula then seems to tell us that the standard error should be close to zero. The model appears to have learned the parameter with extraordinary precision.

But that conclusion depends on a very strong assumption: that each observation is bringing genuinely new and independent information.

Financial markets generate huge amounts of data, but not all of it is independent information about the structural questions investors care about. Prices can be measured every second, but factor regimes do not change every second. A value rotation, a momentum crash, a credit shock, or a flight to quality may unfold over weeks or months. In a long historical sample, there may be only a handful of truly important episodes of this kind.

More frequent measurement therefore gives us more rows in the dataset, but not necessarily more independent evidence about the events that matter most. If value stocks recover after a long drawdown, observing that recovery daily rather than monthly gives a more detailed picture of the path. It does not turn one rotation into fifty independent rotations.

This matters because many statistical procedures reward the appearance of more data. If a model treats repeated observations from the same environment as though they were independent, uncertainty shrinks mechanically. The estimate begins to look precise, not necessarily because the world has become more predictable, but because the model has counted correlated observations as if they were separate pieces of evidence.

The same issue appears in elections. There may be thousands of district-level observations, but a much smaller number of national election cycles. Observing many districts in one year is valuable, but it does not fully answer the question of what the next unusual election year will look like. A national wave, a political realignment, or an unexpected candidate effect is not replicated thousands of independent times simply because it appears across thousands of districts.

This is what we mean when we say that sampling error can become a second-order problem. The old formulas are not necessarily wrong on their own terms. They answer a specific question: how precisely have we estimated the average behaviour of the system we have already observed?

But in markets and elections, the harder question is often different: what will the system look like next time?

Will value recover or continue to lag? Will momentum survive or reverse? Will a market shock affect many stocks at once? Will a national political wave appear? Will the relationships hold or shift: how much is a factor rewarded in markets, or how strongly does a district's political tendency predict its vote?

These questions cannot be answered simply by increasing the number of observations inside the same environment. They require a model that recognises uncertainty at a higher level: dates, regimes, firms, districts, candidates, and common shocks.

A hierarchical model does not make this uncertainty disappear. Its value is that it makes the uncertainty visible. When dates differ meaningfully from one another, the model allows future dates to differ as well. When firms have persistent behaviour, or an incumbent's personal vote persists across elections, the model can carry that information forward. When relationships vary through time, the model does not pretend they are fixed forever.

This becomes especially important because the final decision is often much more binary than the data. In investing, the question may become: increase exposure or reduce it? Trust momentum or hedge it? Treat the current drawdown as noise or as a possible structural break? In elections, the model may estimate a full distribution of vote shares, but the public question often becomes: who is likely to win?

That gap between rich uncertainty and binary action is where overconfidence becomes dangerous. If the model understates uncertainty, the final decision will look cleaner than it really is. It may say "trade" when the evidence supports only "trade cautiously", or "high conviction" when the right conclusion is "possible, but fragile".

More data is valuable. But more data is not the same thing as more understanding. In systems where the important events are rare, clustered, and slow-moving, the task is not just to collect more observations. It is to model what kind of information those observations actually contain.

What This Means in Practice

The practical question is whether a model is merely producing numbers, or whether it is helping the investor understand uncertainty.

One useful idea here is **coverage.** The concept is simple. If a model says that its 95-percent prediction interval should contain the next realised return, then over time that should happen roughly 95 percent of the time. If it happens much less often, the model is overconfident. It is not just making mistakes; it is understating how uncertain the world really is.

Many standard approaches struggle with this problem. Classical regressions, even when carefully implemented, can produce intervals that are too narrow when the data contain common shocks, persistent firm effects, changing factor payoffs, or heavy tails. There are ways to patch this: clustered standard errors, bootstraps, resampling methods, or distribution-free prediction tools. These can be useful, and in some settings they improve calibration materially.

But there is a difference between widening an interval and understanding why it should be wide.

A black-box correction may tell us that the forecast should be treated cautiously. A structural model can go further. It can help identify whether the uncertainty comes from unstable factor payoffs, persistent differences across firms, date-specific shocks, heavy tails, or market conditions the model has rarely seen before. For an investor, that distinction matters. It is one thing to say, "the interval is wide". It is more useful to say, "the interval is wide because this is the kind of environment in which this signal has historically been fragile".

This is the attraction of hierarchical modelling. It does not simply produce a point forecast and a prediction interval. It organises uncertainty into components that correspond more closely to how investors already think about markets. Some uncertainty comes from the average effect of the signal. Some comes from how that effect varies through time. Some comes from persistent differences across firms. Some comes from date-specific shocks. In richer models, some can also come from explicitly measured market conditions, such as volatility, liquidity, sentiment, or proprietary indicators.

The same logic applies to elections. A forecast that says one candidate has a 60-percent chance of winning is informative. But a forecast is more informative if it also tells us where that uncertainty comes from: a persistent candidate effect, like the personal vote an incumbent carries from one race to the next; the national swing of the year; or an unusual election environment the model has rarely seen. It helps us understand whether the forecast is stable or fragile.

The deeper point is the one we started with. The hard problem in both elections and markets is not simply estimating an average over the past. It is quantifying uncertainty about a future that may not look exactly like the past.

A good model does not pretend this problem away. It does not say that uncertainty has vanished because the dataset is large, or because the equation is elegant, or because the backtest looks clean. It tries to name the uncertainty correctly. It asks which parts of the past should be pooled together, which parts should be kept separate, and where the historical evidence is too thin to justify strong conclusions.

There is a broader lesson here. In life, business, and markets, pretending uncertainty does not exist rarely makes us stronger. "Fake it until you make it" may be useful as a motivational phrase, but it is a dangerous statistical principle. A healthier approach is to face uncertainty directly: to ask what we know, what we do not know, and which unknowns matter most.

Statistics is no exception. At its core, the discipline is not about producing a single confident number: it is about quantifying our uncertainty. Treating difficult structure as harmless noise can make a model look cleaner, but it does not make the world safer or more predictable; it only hides the very thing statistics exists to measure. The more honest approach is to expose the uncertainty, give it structure, and then decide how much confidence it deserves.

That is what hierarchical models are designed to do. They allow the data to decide how much should be shared across firms, dates, regimes, districts, or candidates, and how much should remain separate. They sit between two unhelpful extremes: assuming every situation is unique, and assuming every situation is the same.

For investors, that is not an academic refinement. It is the difference between a model that sounds certain and a model that knows when it is not.

Further reading

Danny Ebanks and Jonathan N. Katz and Gary King, 2025, *If a Statistical Model Predicts That Common Events Should Occur Only Once in 10,000 Elections, Maybe It's the Wrong Model.* Available at https://gking.harvard.edu/10k/

For more papers by Jonathan Katz: https://jkatz.caltech.edu

The value of active minds: independent thinking

A key feature of Jupiter's investment approach is that we eschew the adoption of a house view, instead preferring to allow our specialist fund managers to formulate their own opinions on their asset class. As a result, it should be noted that any views expressed – including on matters relating to environmental, social and governance considerations – are those of the author(s), and may differ from views held by other Jupiter investment professionals.

JAPAN

This document is provided solely in connection with the provision of discretionary investment management services to licensed discretionary investment management companies in Japan. It should not be used for any other purpose or distributed to any other private investors or persons. However anyone in possess of this document should bear in mind that the value of an investment and income can go down as well as up. It may be affected by exchange rate variations and you may not get back the amount invested. Initial charges are likely to have a greater proportionate effect on returns if investments are liquidated in the shorter term. Quoted yields are not guaranteed. Past performance should not be seen as a guide to future performance. Every effort is made to ensure the accuracy of the information but no assurance or warranties are given. Jupiter Asset Management Limited ('JAM') is registered in England and Wales (company number 2036243). The registered office is The Zig Zag Building, 70 Victoria Street, London, SW1E 6SQ. JAM is authorised and regulated by the Financial Conduct Authority in the UK whose address is 25 The North Colonnade, Canary Wharf, London E14 5HS.

US ONSHORE

This document is for informational purposes only and nothing herein should be viewed as an offer, solicitation or recommendation to buy or sell an interest in the private investment fund. Any offer of an interest in a private investment fund will only be made by means of a confidential private offering memorandum as it may be amended, supplemented or restated from time to time (the "Offering Memorandum"), which would contain material information (including certain risks of investing in a private investment fund) not contained in this document and which would supersede and qualify in its entirety the information set forth herein.

Market and exchange rate movements can cause the value of an investment to fall as well as rise, and you may get back less than originally invested. Past performance is no guide to the future. Company examples are for illustrative purposes only and are not a recommendation to buy or sell. Quoted yields are not a guide or guarantee for the expected level of distributions to be received. The yield may fluctuate significantly during times of extreme market and economic volatility.

The views expressed are at a point of time and are subject to change. This is particularly true during periods of rapidly changing market circumstances. Every effort is made to ensure the accuracy of the information provided but no assurance or warranties are given.

Jupiter Asset Management Limited is an SEC registered investment adviser. Registration with the SEC does not imply any level of skill or training. The interests offered in the fund hereby have not and will not be registered under the Securities Act of 1933 act, as amended (the "1933 Act") or any other securities laws, and will be offered and sold for investment only to qualifying United States recipients pursuant to the exemption from the registration requirements of the 1933 Act provided by section 4(a)(2) thereof and in compliance with any other applicable securities laws. There will be no public market for the fund interests, and no such market is expected to develop in the future. Interests in the fund will be subject to restrictions on transferability and resale and may not be transferred or resold except as permitted under the 1933 exchange act and applicable state securities laws pursuant to registration or exemption therefrom.

SOUTH KOREA

Neither Jupiter Asset Management nor any of its associated, affiliated or otherwise connected companies is making any representation with respect to the eligibility of prospective investor to acquire the products referred to in this nor any other document or material in connection therewith under the laws of the Republic of Korea ("Korea") and regulations thereunder. None of such strategies has not been registered under the Financial Investment Services and Capital Markets Act of Korea for public offering, and none of the interests may be offered, sold or delivered, or offered or sold to any person for reoffer or resale, directly or indirectly, in Korea or to any resident of Korea except pursuant to applicable laws and regulations of Korea. This document is not, and under no circumstances is, to be constructed as, a public offering of securities in Korea. Neither the strategies mentioned in this website nor Jupiter Asset Management is making any representation with respect to the eligibility of any recipients of this nor any other document or material in connection therewith to acquire the units therein under the laws of Korea, including but without limitation the Foreign Exchange Transaction law and Regulations thereunder. The units may only be sold in Korea through a local financial investment company registered with the Financial Services Commission on a private placement basis to certain qualified professional investors as defined in the FSCMA.

SINGAPORE

This document (including any data or views given) is for informational purposes only and is not and should not be construed as investment advice. This is not an invitation to subscribe or a recommendation to buy, sell, hold or otherwise transact in any investment including interest in the funds managed by Jupiter. Every effort is made to ensure the accuracy of the information but no assurance or warranties are given. The value of an investment and any income from it can go down as well as up. Any examples of companies, stocks or holdings are for illustrative purposes only. The views expressed are those of the author or presenter at the time of writing or preparation and are not necessarily those of the Jupiter Group as a whole and may change in the future. Initial charges are likely to have a greater proportionate effect on returns if investments are liquidated in the shorter term. Quoted yields are not guaranteed. Past performance is not indicative of future performance. The specific investment objectives, particular needs and financial situation of any specific person has not been considered. If there is any doubt on the suitability of an investment in a particular fund, please contact a financial adviser.

Issued by Jupiter Asset Management (Asia) Private Limited. No part of this document may be reproduced, copied or redistributed in any manner without the prior permission of Jupiter Asset Management (Asia) Private Limited. Jupiter Asset Management (Asia) Private Limited or its affiliates, any directors or employees of Asset Management (Asia) Private Limited or its affiliates shall not be liable for any damages arising from any person's reliance on such information and shall not be liable for any errors or omissions (including but not limited to errors or omissions made by third parties) in such information. The information provided herein is subject to change without further notice.

Jupiter Asset Management (Asia) Private Limited (UEN 200916081Z) is a holder of Capital Markets Services (CMS) Licence issued by the Monetary Authority of Singapore ("MAS") under the Securities and Futures Act 2001of Singapore ("SFA") (CMS License 101788).

This publication has not been reviewed by the Monetary Authority of Singapore.

This document is solely intended for institutional investors as defined in the Securities and Futures Act 2001 of Singapore and is prepared for their information and discussion purposes only. Distribution of the document to any other persons (including any other types of investors) is strictly prohibited.

MALAYSIA

This document does not constitute an offer or an invitation to subscribe for the shares of any securities or investment strategies. It is being made available to you, at your request for information purposes only. No invitation or offer to subscribe for the strategy(ies) is made as the prior recognition by the Malaysian Securities Commission has not been obtained and as this material has not been deposited or registered with the Malaysian Securities Commission under the Capital Markets and Services Act 2007 of Malaysia. Accordingly, neither this website nor any other document or material in connection therewith should be circulated or distributed to any person in Malaysia nor should the strategy(ies) be made available or offered for subscription in Malaysia.

UAE

United Arab Emirates (UAE) (excluding the Dubai International Financial Centre (DIFC) and the Abu Dhabi Global Market (ADGM))

This document, and the information contained herein, does not constitute, and is not intended to constitute, a public offer of securities in the United Arab Emirates ("UAE") and accordingly should not be construed as such. The [Units] are only being offered to a limited number of exempt Professional Investors in the UAE who fall under one of the following categories: federal or local governments, government institutions and agencies, or companies wholly owned by any of them.

The [Units] have not been approved by or licensed or registered with the UAE Central Bank, the Securities and Commodities Authority, the Dubai Financial Services Authority, the Financial Services Regulatory Authority or any other relevant licensing authorities or governmental agencies in the UAE (the "Authorities"). The Authorities assume no liability for any investment that the named addressee makes as a Professional Investor. The document is for the use of the named addressee only and should not be given or shown to any other person (other than employees, agents or consultants in connection with the addressee's consideration thereof)."

HONG KONG

Please ensure you read the Offering Documents for this fund before making an investment decision. These documents contain important information including risk factors and details of charges. This document is intended only for "professional investors" as defined under the Securities and Futures Ordinance (Cap. 571 of the Laws of Hong Kong). This document is issued by Jupiter Asset Management (Hong Kong) Limited and has not been reviewed by the Securities and Futures Commission. Past performance is no guide to the future. The value of investments and the income from them can fall as well as rise and may be affected by exchange rate variations, you may get back less than originally invested. No information in this document should be interpreted as investment advice. You are advised to exercise caution. If you are in any doubt about any of the contents of this document, you should obtain independent professional advice. Jupiter Asset Management Limited acts as investment manager to the Jupiter Global Fund (SICAV) and is registered as distributor to the sub funds of the Jupiter Global Fund (SICAV). This document is for information only and is not an offer to sell or an invitation to buy. In particular, it does not constitute an offer or solicitation in any jurisdiction where it is unlawful or where the person making the offer or solicitation is not qualified to do so or the recipient may not lawfully receive any such offer or solicitation. Any holdings and stock examples are used for illustrative purposes only and should not be viewed as investment advice. The views expressed are those of the presenter at the time of preparation and may change in the future. It is the responsibility of any person in possession of this document to inform themselves, and to observe, all applicable laws and regulations of relevant jurisdictions. The information and any opinions contained herein have been obtained from or are based on sources which are believed to be reliable, but the accuracy cannot be guaranteed. No responsibility can be accepted for any consequential loss from this information.

C2452-EN-07.26

JUPITER
ASSET MANAGEMENT