Morgan Stanley

INVESTMENT MANAGEMENT

Counterpoint Global Insights

BIN There, Done That

How to Reduce the Sources of Forecasting Error

CONSILIENT OBSERVER | March 19, 2020

AUTHORS

Michael J. Mauboussin
michael.mauboussin@morganstanley.com

Dan Callahan, CFA
dan.callahan1@morganstanley.com

Introduction

Anticipating revisions in market expectations is the key to generating long-term excess returns as an investor. This process requires you to understand current expectations and why they are likely to change. A company's fundamental results are the primary catalyst in expectations revisions. These include drivers such as sales growth, operating profit margins, and return on invested capital. Successful long-term investors see where expectations are headed based on their forecasts for fundamental results.

Ample research shows that most experts do not make great forecasts.[1] This might appear to be a problem if you are in the business of making predictions. But it turns out that the ability to explain what happened after the fact, often in a way that flatters your faulty prediction, is an incredibly effective coping mechanism. Barbara Mellers, a professor of psychology at the University of Pennsylvania, says, "We find prediction really hard, but we find explanation fairly easy."[2] We tell stories to ourselves and others to paper over our poor predictions.

The news is not so gloomy. The U.S. intelligence community sponsored a forecasting tournament that allowed scientists to measure the accuracy of predictions pertaining to social, political, and economic events. That analysis revealed that one-in-fifty forecasters, dubbed "superforecasters," consistently made better predictions than did the other participants.[3]

Part of the success of the superforecasters reflects their qualities, including numeracy, intellectual curiosity, probabilistic thinking, and open-mindedness. But part of their success relies on habits that can be "learned and cultivated by any intelligent, thoughtful, determined person."[4] While the superforecasters had a clearer crystal ball than the other forecasters, there are ways the rest of us can sharpen what we see.

Statisticians and psychologists rolled up their sleeves and delved into the data to find out what distinguished the best from the rest.[5] They explain their results using what they call the "BIN Model," where "B" refers to bias, "I" to information, and "N" to noise. Most investors are very familiar with how various forms of bias can degrade the quality of decisions.[6] Investors are also acutely aware that they operate in a world of incomplete information.

Noise, on the other hand, tends to get less attention. Daniel Kahneman, who won the Nobel Prize in Economics based in part on his work on bias, wrote an article with some collaborators where he stated, "Where there is judgment, there is noise—and usually more of it than you think."[7]

This leads to the most striking finding the data revealed: the difference between superforecasters and regular forecasters is "due more to noise than bias or lack of information." The researchers conclude that "reducing noise is roughly twice as effective as reducing bias or increasing information."[8]

We now define the elements of the BIN Model and examine how they apply to investment organizations. We then review the BIN Model in a world of prices. We finish with a discussion of how to reduce noise and bias, and to increase information.

Unpacking the BIN Model

Notwithstanding the model's acronym, we shuffle the order of the components and discuss noise, bias, and then information. This rearrangement reflects the importance of noise and the relative lack of familiarity with the concept and its implications. Bias is better understood in the investment community but must be applied carefully. Information is the easiest to appreciate yet remains a source of potential forecasting edge.

Noise. Noise is "the chance variability of judgments."[9] It is relevant in a number of settings. First is when interchangeable professionals make a judgment based on the same set of facts. For example, *Money* magazine asked 50 accountants to calculate the taxes due for a hypothetical family of 4 with an income of $132,000. Exhibit 1 shows the difference between each of their estimates and the actual taxes due. The accountants all deal with the same tax code, often use the same software, and have the same facts in front of them. However, they commonly rely on their individual judgment to arrive at the tax due. Noise measures the difference between the answers that two accountants provide based on the same case.

For instance, say Accountant A provides a figure of $10,000 and Accountant B says $14,000. Noise is calculated as the difference between the numbers divided by the average. In this case, the noise index would be 33 percent ($4,000/$12,000). Managers find a noise index of around 10 percent acceptable when judgment is necessary.

The accountants in the test by *Money* magazine came up with a range of taxes due from $9,806 to $21,216 and an average noise index of 20 percent, a "depressing" result.[10] Noise audits in the insurance and finance industries reveal even higher average indexes, in the range of 40-60 percent.

Exhibit 1: Estimates by 50 Accountants of Taxes Due

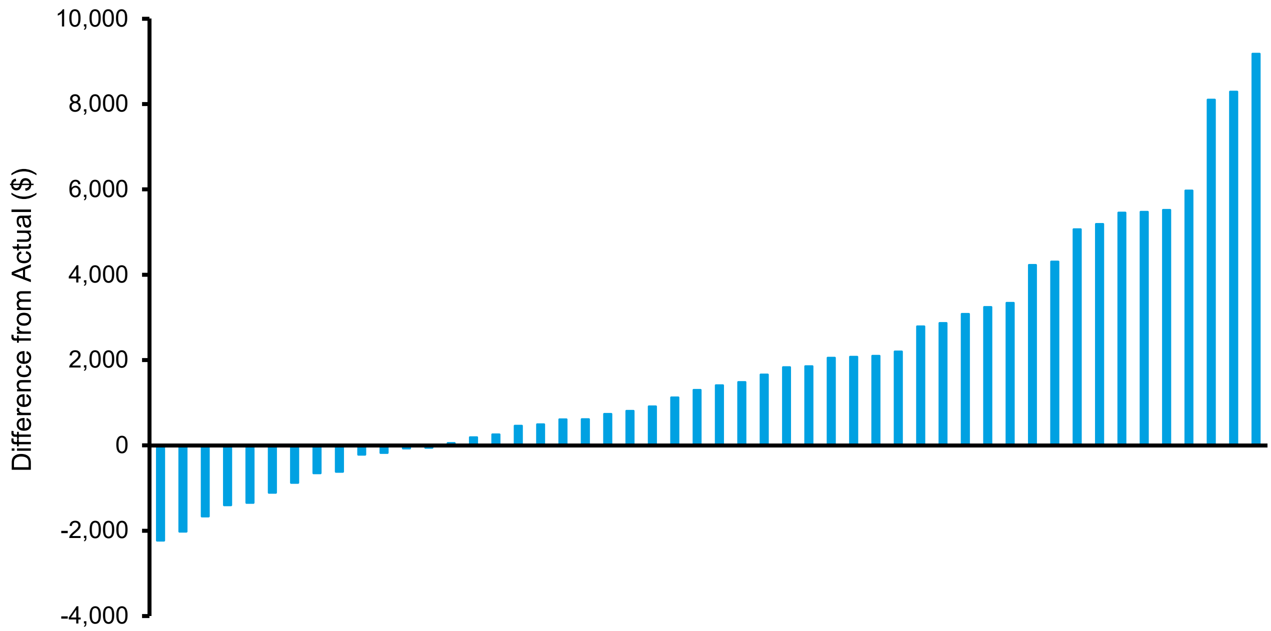

Source: Denise M. Topolnicki, "The Pros Flub Our Third Annual Tax-Return Test," Money*, March 1990, 90-98.*

An organization that has a group of analysts who assess various investment opportunities has to consider the role of noise. What would happen if multiple analysts were assigned to appraise the stock of the same company? Might they come to different or even opposite conclusions? We have done noise audits in the classroom and in the field with ranges similar to other domains.

Another setting where noise is relevant is when the same person evaluates similar decisions over time.[11] You might treat the same decision differently from one moment to the next based on factors such as your mood, hunger, fatigue, or recent experience.

In a classic study, researchers showed pathologists, experts who study body tissue to identify disease, the same biopsy slides at two different times. The average correlation between their overall judgments of disease was 0.63.[12] The correlation would have been 1.0 had the pathologists judged the biopsies the same each time. Wine judges do even worse. Their ratings of the same wine over a short span had an average correlation of just 0.50.[13]

You can imagine analysts presenting ideas to a portfolio manager. That portfolio manager may choose to act or not based on details that have little to do with the merit of the investment idea. For example, research reveals that individuals willingly pass on an opportunity with a positive net present value if they have recently lost money.[14] Kahneman and his colleagues note that "often noise is far above the level that executives would consider tolerable—and they are completely unaware of it."[15]

Noise is also important to consider when a group makes a unique strategic decision that relies on the evaluation of a lot of complex information. Examples include hiring a new employee, proceeding with a costly acquisition, or funding a start-up. In these cases, a lot of information has to be refined into a go or no-go decision.

A key feature of noise is that it is not systematic. You can imagine shots at a target that are scattered randomly around the bullseye (see the right panels of exhibit 2). As a result, it is not possible to anticipate how any particular forecast will deviate from the true signal. Another feature of noise is that you can calculate it without knowing the right answer.

Exhibit 2: Bias and Noise

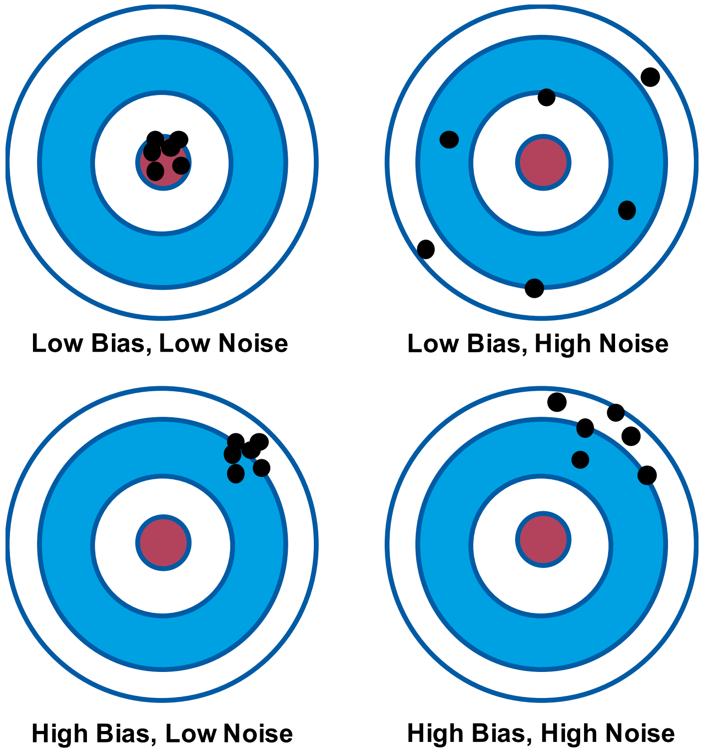

Source: Counterpoint Global.

Bias. Most people use rules of thumb, more formally called heuristics, to cope with the high information demands of decision making. For example, the availability heuristic describes when a decision maker assesses the likelihood of an event by how easy it is to remember similar instances. The availability heuristic explains why people fear flying more than normal after hearing about a plane crash.

Heuristics can be valuable because they save time in making decisions.[16] A bias is the result of an inappropriate application of a heuristic to come to a decision. The refusal to fly after news about a plane crash illustrates the point.

Two biases that are widespread in investing are overconfidence and confirmation. The forms of overconfidence include overestimation, overplacement, and overprecision.[17] Overestimation means you are overconfident in your abilities and overplacement relates to the confidence that you are better than those around you. These forms are relevant when certain conditions are met.

Overprecision is particularly important in investing and business. People tend to be overconfident about how right they are when answering questions that are difficult. This is relevant for forecasting.

We presented 50 true-false questions to more than 10,000 subjects and asked them to answer either true or false and to indicate how confident they are in their answer. Exhibit 3 displays the results. Fifty percent confidence is equivalent to a guess and 100 percent confidence suggests certainty.

We saw that overall, people assign an average confidence of 70 percent but are correct 60 percent of the time. Error also tends to be greatest when the subjects are most confident. This finding is well established in the psychological literature.[18] One manifestation of overconfidence is unprofitable excess trading.[19]

Exhibit 3: Confidence and Calibration for More Than 10,000 Subjects

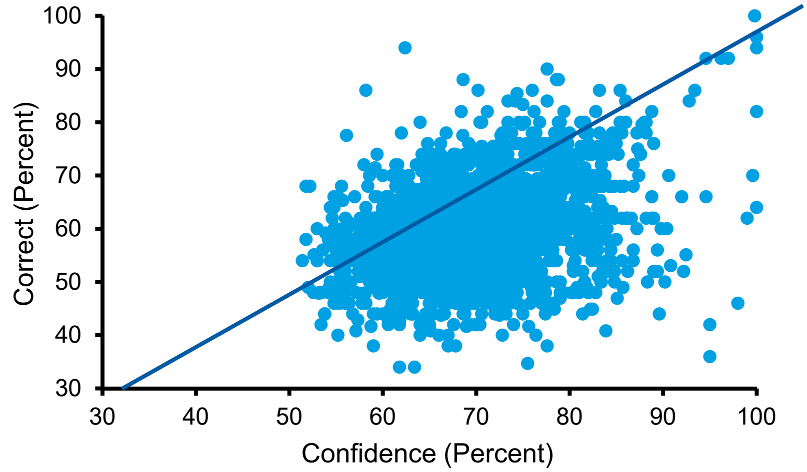

Source: http://confidence.success-equation.com.

Confirmation bias is the tendency to dismiss or discount information that is inconsistent with your prior view. You also interpret ambiguous information in a way that confirms your belief.[20] The ability to update your views effectively is crucial to making sure your beliefs accurately reflect the world. That investors sometimes underreact to information in the short term is consistent with slow belief updating.[21]

We have mentioned two prominent biases investors must overcome, but other biases exist as well.[22]

Heuristics can lead to decisions that have systematic bias. They are predictable. You can now imagine shots at a target that miss the bullseye by similar amounts in the same direction (see the bottom panels of exhibit 2). But to calculate bias you need to know the answer. The nature of the error, nonsystematic for noise and systematic for bias, is a crucial difference and is what allows researchers to separate the impact of each.

Information. Full information allows forecasters to predict with complete accuracy. Information in this context measures the subset of signals that forecasters use relative to full information. There are a few factors that can lead to disparate levels of information among forecasters, including unearthing new information, the speed and accuracy of updating views, placing different weights on the components of information, and being more skillful at dealing with signal complexity.

When the researchers tallied why the superforecasters were so much better than the regular forecasters, they found the difference was roughly 50 percent from reducing noise, 25 percent from reducing bias, and 25 percent from increasing information. It's important to recognize that because noise and bias are independent sources of error, reducing either one of them improves forecasts.

Exhibit 4 summarizes the key characteristics of the BIN Model.

Exhibit 4: The BIN Model

BIN Component	Description	Nature of Error	Contribution to Superforecaster Edge
Bias	Inappropriately applying a rule of thumb	Systematic	~25 percent
Information	Forecasters have incomplete information		~25 percent
Noise	Chance variability of judgments	Nonsystematic	~50 percent

Source: Counterpoint Global.

The BIN Model in a World with Prices

The BIN Model clearly captures concepts that are important in decision making and forecasting. But here's an interesting intellectual exercise to consider: Does noise, bias, or a lack of information matter in an efficient market where stock prices reflect information about future prospects with perfect accuracy?

Noise suggests that two interchangeable analysts might place very different values on the stock of the same company. But what if there is already a price? Noise is irrelevant because the stock price is the correct answer. Bias does not matter for the same reason. And while you may not have full information, the market does and it is in the price.

Markets are not perfectly efficient, of course, but large and liquid markets are very good.[23] The impact of noise and bias is much smaller on stock market investors than it is for accountants or insurance underwriters because there is a baseline of price determined by the collective of investors. Those prices tend to capture a lot of information in a market comprised of investors with diverse models and points of view.[24] That said, we believe the elements of the BIN Model offer investors useful lessons in assessing an investment process.

Financial economists have used the concepts of "noise traders" and "noise allocators" to explain markets for decades.[25] Noise traders "make random trades in financial markets," trade "on noise as if it were information," and "would be better off not trading." Noise allocators "make random allocations to asset managers."[26] Noise traders provide liquidity and earn below-market returns that allow informed investors to generate excess returns to offset the cost of gathering information.

Some models that include noise traders portray the market more accurately than models that rely solely on rational investors.[27] The impact that noise traders have on markets is a function of arbitrage costs, which include identifying and verifying mispricing, implementing and executing trades, and financing and funding securities. When those costs are low, informed investors quickly correct mispricings and markets converge rapidly to fair value. When arbitrage costs are high, gaps between price and value can persist.

The obvious recommendation is to avoid being a noise trader or noise allocator. That means investing based on signals, which requires gathering and properly assessing information. Understanding the investor on the other side of any investment is crucial to having a clear sense of your edge and the motivations of your counterparty.

Another area where noise can creep in is position sizing. Effective investing has two components. The first is finding edge, where a security is mispriced and hence offers the prospect of excess returns. The second is position sizing, or how much money to invest in the idea.

The investment industry spends a lot of time on edge and little time on sizing, and it is likely that sizing decisions are noisy for many portfolio managers. Research shows that investment managers leave returns on the table by failing to follow the position sizes suggested by their own processes.[28]

Individual bias, systematic departures from ideal decisions, are hard to pinpoint in a market that is mostly efficient. But there is some evidence that it exists. For example, recent research suggests that portfolio managers make buy decisions that add value but sell decisions that are poor. A study of more than 780 portfolios and 4.4 million trades from 2000 through 2016 found that portfolio managers would have been better off selling a random position in their portfolio than the security they did sell. The poor selling resulted in raw portfolio returns that were 70 basis points lower per year.[29]

Consultants studied the portfolios of their clients and found that risk-adjusted excess returns of individual stocks follow an inverted-U pattern, the "alpha lifecycle," and that portfolio managers tend to hold stocks too long, on average.[30]

Most major market dislocations are not the result of vagaries in individual forecasting but rather in collective behavior, where investors converge to the same set of beliefs. The dot-com bubble of the late 1990s is a good example. In these cases, the wisdom of crowds flips to the madness of crowds and substantial opportunity presents itself. However, the psychological, professional, and business toll of going against the crowd in these cases can be very high.[31]

How to Avoid the BIN Sin

The goal of training is to improve results. The scientists involved in recruiting competitors for this tournament provided some of the forecasters with training and withheld it from others. This allowed them to measure the impact of training by comparing the group who received it to the control group who did not. Training improved forecasting accuracy by about 10 percent, as captured by a measure of the accuracy of probabilistic predictions called a Brier Score.[32]

Training comes in various categories. Teaching subjects about judgment errors and biases is cheap but useless. Teaching subjects how to check for bias is more effective, as is providing timely feedback. There is also value in learning how to ask good questions and reframe problems.[33] The main point here is that knowing about how you are likely to go wrong isn't as valuable as adopting methods to improve decisions.

When the researchers examined the impact of training on overall accuracy, they found something that surprised them: The training improved accuracy more by reducing noise than by tamping down bias. This finding raises provocative questions about how exactly training affects the various aspects of decision making.

We now turn to methods to improve decisions for each component of the BIN Model.

Noise. The three primary ways to reduce noise are combining judgments, using algorithms, and adopting the "Mediating Assessments Protocol."

Because noise is nonsystematic, the first way to reduce it is to combine forecasts.[34] This is the core idea behind the wisdom of crowds.[35] While each individual may be off the mark in a random way, combining the forecasts reduces the error and increases accuracy. Jack Treynor, a luminary of the investment industry, offered this as a means to understand market efficiency.[36]

Abraham de Moivre produced the math to support this principle in the early 1700s. De Moivre's equation says that the standard error of the mean is equal to the standard deviation of the sample divided by the square root of the sample size. In plain words, the equation says that the error decreases as the sample size increases.[37] A key underlying assumption, which is consistent with the very definition of noise, is that errors are independent and follow a normal, bell-shaped distribution around the correct answer. Forecasting accuracy is compromised when there is no effective way to aggregate the information, or if the errors become correlated and noise transitions to collective bias.

Markets are a mechanism to aggregate views. The involvement of money means there is an incentive to figure out and express a correct view. Negative feedback, where departures from equilibrium are pushed back toward equilibrium, tends to prevail in public markets. Arbitrageurs buy what's cheap and sell what's dear to take advantage of small mispricings. This makes prices more accurate.

But from time to time markets flip from negative to positive feedback, which reinforces a trend. There are instances when positive feedback is useful and necessary. In the context of markets it is often a signal that everyone is thinking the same way, which raises the specter of inefficiency.

Businesses and individuals often don't use aggregation because of the cost. Think of the example of the accountants calculating the taxes due. The average of all the guesses was much closer to the correct answer than the average guess of the individual tax preparers, but retaining 50 accountants to do your taxes is very expensive.

If you are dealing with noisy forecasts and have a cost-effective way to combine them, do so. The forecast is very likely to be more accurate than a single forecast.

A second means to reduce noise is to use an algorithm, which is simply a set of rules or procedures that allow you to achieve a goal. For example, a cake recipe is an algorithm. If you follow the procedures for how to combine the ingredients and bake the batter, you will end up with a tasty cake.

In 1954, Paul Meehl, a clinical psychologist and professor at the University of Minnesota, wrote a book claiming that statistical methods outperformed clinical methods in making patient prognoses.[38] This was not a popular finding, as clinicians placed great value on the judgments based on their training and experience. However, his finding has proven to be robust. One meta-study concluded that the "[s]uperiority for mechanical-prediction techniques was consistent, regardless of the judgment task, type of judges, judges' amounts of experience, or the types of data being combined."[39]

Notwithstanding the superiority of algorithms in many domains, many professionals are still slow to embrace them and are quick to blame them when things go wrong.[40] Even research that shows individuals are willing to defer to algorithms notes that "experienced professionals, who make forecasts on a regular basis, relied less on algorithmic advice than lay people did, which hurt their accuracy."[41]

There are many situations where it is impractical to use an algorithm because the cause and effect is too loosely linked. The advice here is to consider carefully if any aspect of your investment process is best handled systematically. One example is the use of checklists to ensure consistency and thoroughness.[42]

Another example is the size of positions within a portfolio. We saw earlier that many portfolio managers do not stick to their own rules in sizing their positions, to the detriment of their performance.

Portfolio construction is based on inputs, constraints, and objectives.[43] Inputs include expected returns, volatility, and the correlation between assets. Constraints reflect issues such as position, industry, and sector limits, trading volume, and leverage. Objectives include single- or multi-period evaluation horizons. Once a portfolio manager specifies these factors, position sizing can be algorithmic.

Research has found that one way to increase the acceptance of algorithms is to allow humans to adjust the final answer just a bit.[44] Daniel Kahneman calls this "disciplined intuition."[45] The idea is to approach a problem systematically and then give your own judgment or intuition a role in the final decision. The evidence shows that allowing that wiggle room improves the overall quality of the decision.

The final approach to reducing noise is to adopt the Mediating Assessments Protocol (MAP).[46] The idea is to create intermediate ratings based on what's critical in order to come to a final decision that is more informed than relying solely on intuition.

MAP has three parts. First, you "define the assessments in advance." You figure out what's important to arrive at a thoughtful result. For instance, if you are interviewing a candidate to be a financial analyst, you consider the attributes that would lead to success, including skills in valuation, strategy assessment, critical thinking, and communication.

The second step is to use facts to evaluate the candidate based on those criteria. If more than one person is involved with the assessment, each should work independently. The result should be a numerical score for each attribute. To continue with the example of our analyst, each interviewer would evaluate and score the candidate on the relevant criteria. It is also important for each interviewer to ask the same questions in the same order for every candidate.[47]

Finally, there should be a discussion about a final decision only after all of the assessments are complete and scored. The power of the method is that the decision-making approach is standardized.

For investors in public markets, MAP may be more useful in guiding organizational decisions than investment choices. But investors in private markets will find the tool valuable, if only to ensure the investment process is consistent, rigorous, and thorough.

Exhibit 5 summarizes the three primary ways to reduce noise.

Exhibit 5: Techniques for Reducing Noise

Method	Description	Mechanism
Combine judgments	Gather forecasts from individuals who are operating independently	Reduces error by offsetting nonsystematic individual errors
Use algorithms	A set of rules for performing a task	Evidence-based algorithms outperform experts in many domains
Adopt Mediating Assessments Protocol (MAP)	1. Define in advance the attributes you will use to assess something 2. Gather facts to evaluate based on those criteria 3. Make final decision only after all assessments are complete and scored	Standardizes the process

Source: Counterpoint Global.

Bias. We review a few methods to reduce bias. The first is to incorporate base rates into a forecast, which addresses overconfidence. It is common for forecasts to be too optimistic, whether it's the cost and time to remodel a home, how long it will take to complete a task, or the growth rate of a company. Introducing base rates often tempers and grounds an estimate.

Most of us follow a similar pattern when asked to make a forecast: we gather information, combine it with our own experience and views, and project into the future. Psychologists call this approach, which relies on a large dose of intuition, the "inside view."[48]

Integrating base rates, the "outside view," with the inside view generally improves the accuracy of forecasts. The outside view is a "simple statistical analysis of analogous efforts completed earlier."[49] The outside view asks, "what happened when others were in this situation before?" Your specific forecast is recast as an instance of a larger reference class rather than a one-off forecast that relies heavily on the experience and perceptions of the forecaster.

Here's an example to make the concept more concrete. Say a financial analyst forecasts as a base case that a company with annual sales of $10 billion will grow sales 20 percent per year for the next 5 years. In all likelihood, that analyst will have a detailed model substantiating that projection. That is the inside view.

The outside view would examine how many companies of that size have achieved a compound annual growth rate of sales at that level or higher. The answer is 3.5 percent.[50] Would you be inclined to modify a base case that assumes a 50 percent probability of 20 percent growth if the base rate reveals a frequency of 3.5 percent?

There are four steps to using the outside view. The first is to select the reference class. You want the reference class to be broad enough to be robust but narrow enough to be useful in your application. There are good data for many domains, including corporate performance and investing.[51]

The next is to assess the distribution of outcomes for that reference class. Some follow a normal, bell-shaped distribution ("mild randomness"). This is largely true for measures of corporate and investor performance. But others follow a distribution with lots of skew, and in some cases a power law ("wild randomness").[52] Base rates are straightforward to apply for mild distributions and less informative for wild distributions, but we have found that introducing base rates improves the forecaster's thinking. Further, base rates are woefully underutilized even in contexts where their application is clearly beneficial.

The third step is to make a forecast. An understanding of the base rates for the situation should inform both a point forecast and a range of possible outcomes.

The final step is to consider how much to regress your estimate toward the mean. The rule of thumb is that when outcomes are close to random, you should regress your forecast substantially. When outcomes are driven by skill and are therefore persistent, you need not regress much at all.

To continue with our example of a company, sales growth rates are much less persistent than operating profit margins. As a consequence, very high or low expectations for sales growth are better candidates for regression toward the average than are very high or low margins. This recognizes that operating profit margins are to a degree a function of sales growth.[53]

Learning about base rates was the most significant contributor to accuracy for those forecasters who received training. Considering base rates reduces noise and bias.

A second method to remove bias is to create formal mechanisms to open the mind to alternative possibilities. One technique to do this is a premortem.[54] With a premortem, you have a group imagine that a decision was made and that the outcome was a disaster. Each team member then writes down why the result was poor and reviews the reasons. Premortems rely on a mechanism called "prospective hindsight," which consistently generates more scenarios than simply projecting into the future.[55]

Studies show that the premortem method more reliably reduces bias than other techniques do, including getting individuals to consider pros and cons or just cons.[56] Conducting a premortem is not difficult, time consuming, or costly, but there are common mistakes that forecasters make that are avoidable.[57]

Structured approaches to considering other views are also helpful.[58] These include "red teaming," where some individuals are assigned to advocate a view that differs from the consensus.[59] Both premortems and red teams help offset overprecision, one of the manifestations of overconfidence.

A final method to reduce bias is to create a rigorous process to provide accurate and timely feedback. One of the challenges within the investment industry is that the ultimate measure of results is portfolio returns, adjusted for risk. But stock price movements can be very noisy, which means that feedback from the market itself is unreliable in the short-term.

One way to address the problem is to identify signposts and assign numerical probabilities that they occur. An investor trying to generate an excess return buys or sells a stock because he or she has a variant perception, a view that is different than what the market implies.[60] An analyst can distill the thesis that supports the variant perception into specific departures from the consensus. For example, an investor might believe that sales or margins will be higher than what's priced in. As information related to those specific differences is revealed, you can think of yourself as passing a signpost that tells you whether you are on the right path.

A useful signpost has three features. It has an outcome that can be objectively agreed upon, within a specified date, and is important to the overall thesis. For instance, if the consensus is that a company will sell 100 widgets this year and your variant perception is that they will sell more than that, the signpost might be, "There is an 80 percent probability that the company will sell 110 or more widgets in the year."

One crucial point is that forecasts should use numerical probabilities instead of words, which can be interpreted to represent a wide range of probabilities.[61] Numbers allow for an ability to score forecasts precisely and to avoid ex post justifications.

Note that the use of signposts allows for accurate intermediate feedback that leads to the ultimate goal of generating excess returns. By breaking down the process, a forecaster receives more opportunities to learn from mistakes and cultivate the skill of calibration. Research shows that forecasters improve when they receive feedback of this quality.[62]

Debiasing methods introduce base rates to ground forecasts, premortems and red teams to open minds to alternative outcomes, and signposts to create a way to learn from mistakes. The methods produce spillover effects that help manage noise as well.

Information. The first and most obvious way to improve forecasts is to gain access to information before others do. This is especially difficult in markets because of Regulation Fair Disclosure, which was implemented in 2000 and prevents companies from disclosing material privately without simultaneously revealing the same information to the public.

There are some ways, however costly, to access relevant information. These include acquiring and analyzing proprietary data sets, working with top-flight lawyers to interpret laws or regulations, or hiring consultants to accurately assess storm damage.

For example, an alternative data company that tracks the tail numbers of private jets, among other things, observed that a jet owned by Occidental Petroleum landed at the Omaha, Nebraska airport in the spring of 2019. Omaha is the home of Warren Buffett, chief executive officer of the cash-rich conglomerate, Berkshire Hathaway. Occidental announced a large investment from Berkshire Hathaway a couple of days later. The data company's clients included hedge funds thirsty for information.[63] Most investors cannot pursue this source of edge because the cost is high.

A framework including signposts requires a forecaster to explicitly articulate a view that is different than what the market reflects. As new information becomes available, the forecaster must be able to overcome confirmation bias and revise his or her view to reflect that information. Analysis of the forecasters in the tournament revealed that the frequency of belief updating was “the strongest single behavioral predictor of accuracy.”[64]

The ability to weight information effectively is also crucial. Psychologists distinguish between the strength of evidence for a hypothesis and its weight, or “predictive validity.”[65] Use a coin toss as an example. The strength is the ratio of flips that come up heads to those that come up tails. The weight is the sample size, or number of flips.

Forecasters tend to be overconfident when the strength is high and the weight is low. For example, when heads shows up 7 times after 10 flips, the forecaster is likely to be overconfident in an assessment that the coin is biased (this will happen in 1 in 8 times with a fair coin).

Forecasters tend to be underconfident when the strength is low and the weight is high. For example, when heads shows up 5,100 times after 10,000 flips, the forecaster is likely to be underconfident that the coin is biased (this will happen about in 1 in 50 times with a fair coin).

Using information effectively means getting an informational advantage when you can, updating your view accurately and promptly to reflect new information, and placing the proper weight on the information at your disposal.

Conclusion

You can use the BIN Model to assess your decision-making process, with a heightened awareness of the large and potentially overlooked role of noise. The techniques we describe to address noise, including averaging, using algorithms, and adopting the Mediating Assessments Protocol, don’t ensure accuracy but have been shown to help improve predictions. Further, the training techniques designed to reduce bias appear to be effective in dealing with noise as well.

You can also use the model to evaluate the decision-making processes of others, including companies, organizations, and competitors. Because noise tends to be overlooked or ignored, improving consistency can add a great deal of value. For example, when Kahneman and his colleagues asked executives of a large global firm for estimates of the cost of noise in their companies, they suggested it would be “measured in the billions” and that even small reductions in noise would be worth “tens of millions.”[66] Identifying companies that adopt practices to reduce noise may be the source of excess returns.

Please see Important Disclosures on pages 17-18

Endnotes

[1] Philip E. Tetlock, *Expert Political Judgment: How Good Is It? How Can We Know?* (Princeton, NJ: Princeton University Press, 2005).
[2] "Want Better Forecasting Skills? Silence the Noise," *Knowledge@Wharton*, November 26, 2019.
[3] Philip E. Tetlock and Dan Gardner, *Superforecasting: The Art and Science of Prediction* (New York: Crown Publishers, 2015).
[4] Ibid., 18.
[5] Ville A. Satopää, Marat Salikhov, Philip E. Tetlock, and Barbara Mellers, "Bias, Information, Noise: The BIN Model of Forecasting," *Working Paper*, November 2019.
[6] If anything, the complaint now is that there is too much of a focus on bias. See Brandon Kochkodin, "Behavioral Economics' Latest Bias: Seeing Bias Wherever It Looks," *Bloomberg*, January 13, 2020.
[7] Daniel Kahneman, Andrew M. Rosenfield, Linnea Gandhi, and Tom Blaser, "Noise: How to Overcome the High, Hidden Cost of Inconsistent Decision Making," *Harvard Business Review*, Vol. 94, No. 10, October 2016, 38-46 and "Daniel Kahneman's Strategy for How Your Firm Can Think Smarter," *Knowledge@Wharton,* June 8, 2016.
[8] Satopää, Salikhov, Tetlock, and Mellers, "Bias, Information, Noise: The BIN Model of Forecasting."
[9] Kahneman, Rosenfield, Gandhi, and Blaser, "Noise: How to Overcome the High, Hidden Cost of Inconsistent Decision Making."
[10] Denise M. Topolnicki, "The Pros Flub Our Third Annual Tax-Return Test," *Money*, March 1990, 90-98.
[11] Astute readers will see that this is similar to the ergodicity problem. A system is ergodic when the ensemble average and the time average are the same. For instance, the average of 1,000 people flipping a coin (ensemble) will be the same as you flipping it 1,000 times in a row (time). Decision making is non-ergodic. See Ole Peters, "The Ergodicity Problem in Economics," *Nature Physics*, Vol. 15, December 2019, 1216-1221.
[12] Hillel J. Einhorn, "Expert Judgment: Some Necessary Conditions and an Example," *Journal of Applied Psychology*, Vol. 59, No. 5, October 1974, 562-571.
[13] Robert H. Ashton, "Reliability and Consensus of Experienced Wine Judges: Expertise Within and Between?" *Journal of Wine Economics*, Vol. 7, No. 1, May 2012, 70-87.
[14] Baba Shiv, George Loewenstein, Antoine Bechara, Hanna Damasio, and Antonio R. Damasio, "Investment Behavior and the Negative Side of Emotion," *Psychological Science*, Vol. 16, No. 6, June 2005, 435-439.
[15] Kahneman, Rosenfield, Gandhi, and Blaser, "Noise: How to Overcome the High, Hidden Cost of Inconsistent Decision Making."
[16] Max H. Bazerman and Don Moore, *Judgment in Managerial Decision Making, 9th Edition* (Hoboken, NJ: John Wiley & Sons, 2016).
[17] Don A. Moore and Paul J. Healy, "The Trouble with Overconfidence," *Psychological Review*, Vol. 115, No. 2, April 2008, 502-517 and Don A. Moore, *Perfectly Confident: How to Calibrate Your Decisions Wisely* (New York: Harper Business, 2020).
[18] Keith E. Stanovich, Richard F. West, and Maggie E. Toplak, *The Rationality Quotient: Toward a Test of Rational Thinking* (Cambridge, MA: MIT Press, 2016), 141-175.
[19] Kent Daniel and David Hirshleifer, "Overconfident Investors, Predictable Returns, and Excessive Trading," *Journal of Economic Perspectives*, Vol. 29, No. 4, Fall 2015, 61-88 and Brad M. Barber and Terrance Odean, "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment," *Quarterly Journal of Economics*, Vol.116, No. 1, February 2001, 261-292.
[20] Roland G. Fryer, Jr., Philipp Harms, and Matthew O. Jackson, "Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization," *Journal of the European Economic Association*, Vol. 17, No. 5, October 2019, 1470-1501 and Gary Charness and Chetan Dave, "Confirmation Bias with Motivated Beliefs," *Games and Economic Behavior,* Vol. 104, July 2017, 1-23.
[21] George J. Jiang and Kevin X. Zhu, "Information Shocks and Short-Term Market Underreaction," *Journal of Financial Economics*, Vol. 124, No. 1, April 2017, 43-64.
[22] There are purportedly more than 100 biases. See https://en.wikipedia.org/wiki/List_of_cognitive_biases.
[23] Michael J. Mauboussin, "Who Is on the Other Side?" *BlueMountain Investment Research*, February 12, 2019, Eugene Fama and Kenneth French, "Luck versus Skill in the Cross-Section of Mutual Fund Returns," *Journal of*

Finance, Vol. 65, No. 5, October 2010, 1519-1547, and Burton G. Malkiel, *A Random Walk Down Wall Street: The Time-Tested Strategy for Successful Investing, 12th Edition* (New York: W. W. Norton & Company, 2020).
[24] Scott E. Page, *The Model Thinker: What You Need to Know to Make Data Work for You* (New York: Basic Books, 2018), 27-42.
[25] Robert J. Shiller, "Stock Prices and Social Dynamics," *Brookings Papers on Economic Activity*, Vol. 2, 1984, 457-510. For evidence that these investors exist, see Table 1 in Nicolae Gârleanu and Lasse Heje Pedersen, "Efficiently Inefficient Markets for Assets and Asset Management," *Journal of Finance*, Vol. 73, No. 4, August 2018, 1689.
[26] Nicolae Gârleanu and Lasse Heje Pedersen, "Efficiently Inefficient Markets for Assets and Asset Management," *Journal of Finance*, Vol. 73, No. 4, August 2018, 1663-1712 and Fischer Black, "Noise," *Journal of Finance*, Vol. 41, No. 3, July 1986, 529-543.
[27] Andrei Shleifer and Lawrence H. Summers, "The Noise Trader Approach to Finance," *Journal of Economic Perspectives*, Vol. 4, No. 2, Spring 1990, 19-33.
[28] Cameron Hight, "Alpha Theory 2019 Year in Review," *Alpha Theory Blog*, March 6, 2020. In related work, portfolio managers tend to underweight their best positions. See Alexey Panchekha, CFA, "The Active Manager Paradox: High-Conviction Overweight Positions," *CFA Institute; Enterprising Investor*, October 3, 2019.
[29] Klakow Akepanidtaworn, Rick Di Mascio, Alex Imas, and Lawrence Schmidt, "Selling Fast and Buying Slow: Heuristics and Trading Performance of Institutional Investors," *SSRN Working Paper*, September 2019.
[30] Chris Woodcock, Alesi Rowland, and Snežana Pejić, Ph.D., "The Alpha Lifecycle," *Essentia Analytics White Paper,* 2019.
[31] For a model of investor bias see, Aydoğan Alti and Paul C. Tetlock, "Biased Beliefs, Asset Prices, and Investment: A Structural Approach," *Journal of Finance*, Vol. 69, No. 1, February 2014, 325-361. For a discussion of the wisdom and madness of crowds, see Michael J. Mauboussin and Dan Callahan, "Animating Mr. Market: Adopting a Proper Psychological Attitude," *Credit Suisse Global Financial Strategies*, February 10, 2015.
[32] Barbara Mellers, Lyle Ungar, Jonathan Baron, Jaime Ramos, Burcu Gürçay, Katrina Fincher, Sydney E. Scott, Don Moore, Pavel Atanasov, Samuel A. Swift, Terry Murray, Eric Stone, and Philip E. Tetlock, "Psychological Strategies for Winning a Geopolitical Forecasting Tournament," *Psychological Science*, Vol. 25, No. 5, May 2014, 1106-1115.
[33] Welton Chang, Eva Chen, Barbara Mellers, Philip Tetlock, "Developing Expert Political Judgment: The Impact of Training and Practice on Judgmental Accuracy in Geopolitical Forecasting Tournaments," *Judgment and Decision Making*, Vol. 11, No. 5, September 2016, 509-526. For feedback in the investment industry, see Joseph A. Cerniglia and Philip E. Tetlock, "Accelerating Learning in Active Management: The Alpha-Brier Process," *Journal of Portfolio Management*, Vol. 45, No. 5, July 2019, 125-135.
[34] J. Scott Armstrong, ed., *Principles of Forecasting: A Handbook for Researchers and Practitioners* (New York: Springer, 2001), 417-440 and Véronique Genre, Geoff Kenny, Aidan Meyler, Allan Timmermann, "Combining Expert Forecasts: Can Anything Beat the Simple Average?" *International Journal of Forecasting*, Vol. 29, No. 1, January-March 2013, 108-121.
[35] James Surowiecki, *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations* (New York: Doubleday and Company, 2004).
[36] Jack L. Treynor, "Market Efficiency and the Bean Jar Experiment," *Financial Analysts Journal*, Vol. 43, No. 3, May-June, 1987, 50-53.
[37] Howard Wainer, "The Most Dangerous Equation," *American Scientist*, May-June 2007, 249-256.
[38] Paul E. Meehl, *Clinical versus Statistical Prediction: A Theoretical Analysis and a Review of the Evidence* (Minneapolis, MN: University of Minnesota Press, 1954).
[39] William M. Grove, David H. Zald, Boyd S. Lebow, Beth E. Snitz, and Chad Nelson, "Clinical Versus Mechanical Prediction: A Meta-Analysis," *Psychological Assessment*, Vol. 12, No. 1, March 2000, 19-30.
[40] Robyn M. Dawes, "The Robust Beauty of Improper Linear Models in Decision Making," *American Psychologist*, Vol. 34, No. 7, July 1979, 571-582 and Berkeley J. Dietvorst, Joseph P. Simmons, and Cade Massey, "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err," *Journal of Experimental Psychology: General,* No. 144, Vol. 1, February 2015, 114–126.

[41] Jennifer M. Logg, Julia A. Minson, and Don A. Moore, "Algorithm Appreciation: People Prefer Algorithmic to Human Judgment," *Organizational and Human Decision Processes*, Vol. 151, March 2019, 90-103.
[42] Atul Gawande, *The Checklist Manifesto: How to Get Things Right* (New York: Metropolitan Books, 2009).
[43] For example, see Andrew W. Lo, *Hedge Funds: An Analytical Perspective* (Princeton, NJ: Princeton University Press, 2008), 217-236.
[44] Berkeley J. Dietvorst, Joseph P. Simmons, and Cade Massey, "Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them," *Management Science*, Vol. 64, No. 3, March 2018, 1155-1170.
[45] See "Michael Mauboussin and Daniel Kahneman: In Conversation," *Santa Fe Institute Applied Complexity Network Meeting*, October 13, 2015.
[46] Daniel Kahneman, Dan Lovallo, and Olivier Sibony, "A Structured Approach to Strategic Decisions," *MIT Sloan Management Review*, Vol. 60, No. 3, Spring 2019, 67-73.
[47] Iris Bohnet, *What Works: Gender Equality by Design* (Cambridge, MA: Belknap Press, 2016).
[48] Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011), 245-254.
[49] Dan Lovallo and Daniel Kahneman, "Delusions of Success: How Optimism Undermines Executives' Decisions," *Harvard Business Review*, July 2003, 56-63.
[50] Michael J. Mauboussin, Dan Callahan, and Darius Majd, "The Base Rate Book: Integrating the Past to Better Anticipate the Future," *Credit Suisse Global Financial Strategies*, September 26, 2016, 24.
[51] Ibid.
[52] Benoit B. Mandelbrot, *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk* (New York: Springer, 1997), 117-145.
[53] Mauboussin, Callahan, and Majd, "The Base Rate Book: Integrating the Past to Better Anticipate the Future," 40-55.
[54] Gary Klein, *Intuition at Work: Why Developing Your Gut Instincts Will Make You Better at What You Do* (New York: Currency, 2003), 88-91; Gary Klein, "Performing a Project Premortem," *Harvard Business Review*, September 2007, 18-19; and Gary Klein, Tim Koller, and Dan Lovallo, "Bias Busters: Premortems: Being Smart at the Start," *McKinsey Quarterly*, April 2019.
[55] Deborah J. Mitchell, J. Edward Russo, and Nancy Pennington, "Back to the Future: Temporal Perspective in the Explanation of Events," *Behavioral Decision Making*, Vol. 2, No. 1, January/March 1989, 25-38.
[56] Beth Veinott, Gary A. Klein, and Sterling Wiggins, "Evaluating the Effectiveness of the PreMortem Technique on Plan Confidence," *Proceedings of the 7th International ISCRAM Conference – Seattle, USA*, May 2010.
[57] Gary Klein, Paul D. Sonkin, and Paul Johnson, "Rendering a Powerful Tool Flaccid: The Misuse of Premortems on Wall Street," *Working Paper*, February 2019.
[58] For a good summary of techniques, see Bryce G. Hoffman, *Red Teaming: How Your Business Can Conquer the Competition by Challenging Everything* (New York: Crown Business, 2017).
[59] Micah Zenko, *Red Team: How to Succeed By Thinking Like the Enemy* (New York: Basic Books, 2015).
[60] Michael Steinhardt, *No Bull: My Life In and Out of the Markets* (New York: John Wiley & Sons, 2001), 129.
[61] Andrew Mauboussin and Michael J. Mauboussin, "If You Say Something Is 'Likely,' How Likely Do People Think It Is?" *Harvard Business Review*, July 3, 2018 and Jeffrey A. Friedman, *War and Chance: Assessing Uncertainty in International Politics* (New York: Oxford University Press, 2019).
[62] Allan H. Murphy and Harald Daan, "Impacts of Feedback and Experience on the Quality of Subjective Probability Forecasts: Comparison of Results from the First and Second Years of the Zierikzee Experiment," *Monthly Weather Review*, Vol. 112, No. 3, 1984, 413-423 and Danny Hernandez, "How Our Company Learned to Make Better Predictions About Everything," *Harvard Business Review*, May 15, 2017.
[63] Robin Wigglesworth, "Stockpickers Turn to Big Data to Arrest Decline," *Financial Times*, February 11, 2020.
[64] Barbara Mellers, Eric Stone, Pavel Atanasov, Nick Rohrbaugh, S. Emlen Metz, Lyle Ungar, Michael M. Bishop, Michael Horowitz, Ed Merkle, and Philip Tetlock, "The Psychology of Intelligence Analysis: Drivers of Prediction Accuracy in World Politics," *Journal of Experimental Psychology: Applied*, Vol. 21, No. 1, March 2015, 1-14 and Barbara Mellers, Eric Stone, Terry Murray, Angela Minster, Nick Rohrbaugh, Michael Bishop, Eva Chen, Joshua Baker, Yuan Hou, Michael Horowitz, Lyle Ungar, and Philip Tetlock, "Identifying and Cultivating Superforecasters as a Method of Improving Probabilistic Predictions," *Perspectives on Psychological Science,* Vol. 10, No. 3, May 2015, 267-281. There is also evidence that "Robo-Analysts" revise their views more frequently than their human

counterparts. See Braiden Coleman, Kenneth Merkley, and Joseph Pacelli, "Man Versus Machine: A Comparison of Robo-Analyst and Traditional Research Analyst Investment Recommendations," *Working Paper*, January 21, 2020.

[65] Dale Griffin and Amos Tversky, "The Weighing of Evidence and the Determinants of Confidence," *Cognitive Psychology*, Vol. 24, No. 3, July 1992, 411-435.

[66] Kahneman, Rosenfield, Gandhi, and Blaser, "Noise: How to Overcome the High, Hidden Cost of Inconsistent Decision Making."

DISTRIBUTION

This communication is only intended for and will only be distributed to persons resident in jurisdictions where such distribution or availability would not be contrary to local laws or regulations.

EMEA:
This marketing communication has been issued by MSIM Fund Management (Ireland) Limited. MSIM Fund Management (Ireland) Limited is regulated by the Central Bank of Ireland. MSIM Fund Management (Ireland) Limited is incorporated in Ireland as a private company limited by shares with company registration number 616661 and has its registered address at The Observatory, 7-11 Sir John Rogerson's Quay, Dublin 2, D02 VC42, Ireland.

U.S.:
NOT FDIC INSURED | OFFER NO BANK GUARANTEE | MAY LOSE VALUE | NOT INSURED BY ANY FEDERAL GOVERNMENT AGENCY | NOT A BANK DEPOSIT

Hong Kong: This document has been issued by Morgan Stanley Asia Limited for use in Hong Kong and shall only be made available to "professional investors" as defined under the Securities and Futures Ordinance of Hong Kong (Cap 571). The contents of this document have not been reviewed nor approved by any regulatory authority including the Securities and Futures Commission in Hong Kong. Accordingly, save where an exemption is available under the relevant law, this document shall not be issued, circulated, distributed, directed at, or made available to, the public in Hong Kong. **Singapore:** This document may not be circulated or distributed, whether directly or indirectly, to persons in Singapore other than to (i) an accredited investor (ii) an expert investor or (iii) an institutional investor as defined in Section 4A of the Securities and Futures Act, Chapter 289 of Singapore ("SFA"); or (iv) otherwise pursuant to, and in accordance with the conditions of, any other applicable provision of the SFA. This publication has not been reviewed by the Monetary Authority of Singapore. **Australia:** This publication is disseminated in Australia by Morgan Stanley Investment Management (Australia) Pty Limited ACN: 122040037, AFSL No. 314182, which accept responsibility for its contents. This publication, and any access to it, is intended only for "wholesale clients" within the meaning of the Australian Corporations Act.

Japan: This document may not be circulated or distributed, whether directly or indirectly, to persons in Japan other than to (i) a professional investor as defined in Article 2 of the Financial Instruments and Exchange Act ("FIEA") or (ii) otherwise pursuant to, and in accordance with the conditions of, any other allocable provision of the FIEA. This document is disseminated in Japan by Morgan Stanley Investment Management (Japan) Co., Ltd., Registered No. 410 (Director of Kanto Local Finance Bureau (Financial Instruments Firms)), Membership: the Japan Securities Dealers Association, The Investment Trusts Association, Japan, the Japan Investment Advisers Association and the Type II Financial Instruments Firms Association.

IMPORTANT INFORMATION

Investing involves risk including the possible loss of principal. There is no guarantee that any investment strategy will work under all market conditions, and each investor should evaluate their ability to invest for the long-term, especially during periods of downturn in the market.

The views and opinions are those of the author as of the date of preparation of this material and are subject to change at any time due to market or economic conditions and may not necessarily come to pass. Furthermore, the views will not be updated or otherwise revised to reflect information that subsequently becomes available or circumstances existing, or changes occurring, after the date of publication. The views expressed do not reflect the opinions of all investment teams at Morgan Stanley Investment Management (MSIM) or the views of the firm as a whole, and may not be reflected in all the strategies and products that the Firm offers.

Certain information herein is based on data obtained from third party sources believed to be reliable. However, we have not verified this information, and we make no representations whatsoever as to its accuracy or completeness. The views expressed in the books and articles referenced in this whitepaper are not necessarily endorsed by Morgan Stanley Investment Management.

The information herein is a general communications which is not impartial and has been prepared solely for information and educational purposes and does not constitute an offer or a recommendation to buy or sell any particular security or to adopt any specific investment strategy. The material contained herein has not been based on a consideration of any individual client circumstances and is not investment advice, nor should it be construed in any way as tax, accounting, legal or regulatory advice. To that end, investors should seek independent legal and financial advice, including advice as to tax consequences, before making any investment decision.

Charts and graphs provided herein are for illustrative purposes only. Any securities referenced herein are solely for illustrative purposes only and should not be construed as a recommendation for investment.

This communication is not a product of Morgan Stanley's Research Department and should not be regarded as a research recommendation. The information contained herein has not been prepared in accordance with legal requirements designed to promote the independence of investment research and is not subject to any prohibition on dealing ahead of the dissemination of investment research.

MSIM has not authorised financial intermediaries to use and to distribute this document, unless such use and distribution is made in accordance with applicable law and regulation. Additionally, financial intermediaries are required to satisfy themselves that the information in this document is appropriate for any person to whom they provide this document in view of that person's circumstances and purpose. MSIM shall not be liable for, and accepts no liability for, the use or misuse of this document by any such financial intermediary.

Morgan Stanley Investment Management is the asset management division of Morgan Stanley.

This document may be translated into other languages. Where such a translation is made this English version remains definitive. If there are any discrepancies between the English version and any version of this document in another language, the English version shall prevail.