



ABSTRACT

Economic Regime Management® (ERM) is a framework to help investors understand how economic conditions affect their overall portfolios and individual investments. While intuitively investors know that economic factors are intimately related to future returns, no clear method exists to incorporate that information into actionable knowledge. This remains the case despite the fact that the relationship between many of the factors we will discuss and asset returns have been widely studied through time. While insightful and innovative research on this topic is plentiful, the ERM methodology focuses on being truly implementable in real-world scenarios, in real-time. This focus on implementation and a few key principles will allow ERM to avoid some of the common pitfalls in the use of economic factors. While we do not endeavor to cover all of the process details completely in this paper, we will give an intuitive explanation for the usefulness of the ERM framework.

This paper will show how they can isolate individual sources of risk at a macro-level. Utilizing this information, investors can more easily find solutions that provide the best risk and return profiles while diversifying the sources of risk and return effectively. Finally, using insights about economic regimes, an investor can conduct scenario analyses that illustrate how different economic scenarios may affect investor's portfolios.

MAKING THE CONNECTION

In this first part of a three-part series, we will focus on understanding why connecting economic fundamentals with asset returns proves to be a difficult task. We will then propose a simple method of measurement that can be used for general insight for long term investors. This method of measurement will highlight the insights that drive the ERM framework and lay a foundation that we will build on later in the series.

The economic environment to which investors are subject may have a dominating influence on investment outcomes. Investors intuitively understand this fact, but turning this deceptively simple concept into actionable knowledge has proven to be a difficult challenge.

To understand why this is the case, and before a framework can be built, investors need to understand the problems that exist and how that leads to common pitfalls in implementation.


MAKING THE CONNECTION PART I:

DATA PROBLEMS, A STORY OF LAGS AND REVISIONS

Perhaps the biggest problem when dealing with economic data is the fact that the real-time data is often not the same as the historical data. There are two main reasons for this disparity. The first is that most economic data is released with some sort of lag. For instance, Real GDP growth is released each quarter in the U.S., but the final estimate of that number is not ready until one quarter after the end of the quarter the data is covering. So even when trying to forecast real growth over the next quarter, analysts are not yet sure about the *current* level of real growth.

The second part of this problem is that economic data is often revised several times, and is released in different 'vintages' through time. This means that if an analyst examines Real GDP (gross domestic product) during the year and then again in the next year, it is likely that some changes due to revisions will be reflected in the data. For example, in July of 2014, Research and Development spending was accounted for using a different methodology, and with a simple change of an accounting rule,¹ the U.S. economy became approximately 3% bigger.

The net effect of these two issues is the data that investors must deal with in real-time is not of the same quality as the data available for historical analysis. This means that any analysis using the historical relationship between most economic factors and asset prices is very likely not truly implementable. Keeping with the Real GDP example where these issues are particularly acute, we can see in the chart below that, with the benefit of the revised historical data, the positive relationship between Real GDP growth and U.S. Equity Returns² is noticeable. Contrasting that with the data that would be available in real-time, the relationship disappears, as can be seen in the flat regression line.



This type of problem is prevalent in forecasting and brings to mind a famous quote attributed to Nils Bohr, a Nobel laureate in Physics.

"Prediction is very difficult, especially if it's about the future"
- Nils Bohr

¹ http://www.bea.gov/faq/index.cfm?faq_id=1028

² U.S. Equity returns are represented by the MSCI USA Total Return Index.


MAKING THE CONNECTION PART II:

ALL ABOUT EXPECTATIONS

SECTION A: EXPECTATIONS AND SURPRISES

The difficulties in understanding the relationship between economic factors and asset prices only grow in complexity from here. The market is constantly using its current view of economic projections to inform its assessment of the 'fair' market price for an asset. For example, holding all else equal, if the market was projecting 3% growth in Real GDP, and then 3% growth was realized, an investor should not expect asset prices to change in response to the growth in GDP. Instead, prices would only change if, for example, the market was expecting 3% Real GDP growth, and instead 5% growth was realized. Therefore, changes, or perhaps more accurately *surprises*, in economic factors are only useful relative to those expectations.


This dynamic is the most clear when using the effects of inflation as an example. Below is a chart of annualized rolling 3-year U.S. inflation as measured by the Consumer Price Index (CPI) relative to its historical average. The high inflation period of the 1970s is easily identifiable, and it is clear that inflation has been falling persistently over time. A first thought then may be that an inflation-linked asset was more likely to return higher amounts during the 1970s than it would have recently.



However, when we overlay inflation with the corresponding rolling 3-year return from U.S. TIPS³ (Treasury Inflation Protected Securities) it can be easily seen that there are many instances when the returns from TIPS do not correspond with changes in inflation in such a straightforward manner. This is in large part due to the fact that changes in real rates will have just as large an impact as changes in inflation on the performance of TIPS.

³ In retrospect we know what the changes in inflation and rates were over this period. This allows us to build a highly accurate simulated TIPS return stream that allows us to extend our analysis back in time. See Appendix for details.


Annualized Rolling 3-Year Inflation vs. U.S. TIPS Return



This simple illustration has a powerful implication because it shows that even if a factor like Inflation, Growth, etc., is high (or low) it does not necessarily mean that investment returns in assets linked to those factors will be correspondingly high (or low).

It is worth a quick aside at this point to discuss why a seemingly natural division of the inflation data into “low” and “high” periods can lead to erroneous conclusions. The primary shortfalls of the high/low division approach result from differing perspectives. For example, what an investor might consider to be high inflation today would have ranked as low inflation in the 1970’s (see the following chart).


U.S. CPI in Subperiods




An important takeaway from the chart above is that if we were to use any kind of sub-period, an overwhelming majority of the “high” observations would come from the 1970s, and we believe that this can lead to bias due to the fact that many characteristics of that time period may not repeat themselves in a future scenario.

Another important example of perspective is in making cross-country comparisons. The following chart illustrates the year-over-year Real GDP growth in both the U.S. and China.


Again, if some arbitrary definition of “high” versus “low” was made, these definitions would not apply well to both the cases of China and the U.S.



The ERM framework uses two insights revolving around market expectations that will help explain why returns fluctuate as we have seen. As mentioned earlier, more important than the absolute level of any economic metric like Real GDP growth and/or Inflation is the unexpected changes or surprises in that data relative to the market’s expectations. This brings up an immediate problem, as now market expectations must be defined somehow. Luckily, some of the most widely used economic data includes market forecasts. For instance, the Philadelphia Federal Reserve Bank published a Survey of Professional Forecasters dataset that covers many important economic factors. The chart below shows the performance of these forecasts relative to realized inflation over time.



Although this is a great start to understanding how realized inflation compared to market expectations, the reader may notice that this series only begins in 1981. To extend this analysis, we created our own representative agent forecaster, as seen in the chart below.



The idea of a representative agent forecaster allows for the extension of the analysis back in time before there was a survey of forecasters available. This history is important especially for the CPI inflation series, because some of the most interesting times were during the high inflation 1970s, which pre-dates the survey.

A representative agent is a real time forecaster that uses the ‘vintage’ of economic data that was available at the time to estimate consensus forecasts. This is a realistic way to analyze economic data because of the constant revisions, and it is working with the same data as the Survey of Forecasters.

This representative agent provides time-period flexibility, and when comparing its performance with the Survey of Professional Forecasters, it is remarkably similar. The errors or surprises in forecasts are ~89% correlated, which provides confidence that it is truly a good representation of market expectations.⁴


SECTION B: UNCERTAINTY ABOUT THE FUTURE

Another insight that the ERM framework exploits is that the level of the market’s expectation for the future is not the only important piece of information to consider when evaluating the market’s current economic view. The market could have a very positive future economic expectation but, if its confidence about that expectation is very low, then the positive nature of that outlook may not count for much. Thus, ERM focuses as well on the *uncertainty* of future expectations, which can be an overriding component of market expectations.


Sticking with the example of inflation, we can again look at the professional forecast data from the Philadelphia Federal Reserve. In the chart below, we can see that over time this group of forecasting professionals has had varying amounts of agreement within their forecasts. The series below has clear periods of relatively high and low certainty about Inflation forecasts. During periods where forecast uncertainty was relatively low, it would

⁴ Please see the appendix for more detail.

make sense to have higher confidence in the inflation forecasts versus the periods when forecast uncertainty was relatively high.



Uncertainty is an important consideration when evaluating the effects of economic factors. As we have stated, the market has 'priced-in' some expectations of the future economic environment. When the range of forecasts, otherwise known as the *dispersion*, is high, there is greater uncertainty. Unsurprisingly, as we see below in the 1st decile (*lowest*) of dispersion/uncertainty, we see a closely bunched group of small forecast surprises, but as we go toward the bottom of the chart to the 10th decile (*highest*), these surprises flatten out across a larger range in both positive and negative directions. This type of effect strongly impacts on the risk of assets exposed to the underlying economic fundamental.



MEASURING THE IMPACT OF ECONOMIC FUNDAMENTALS ON FINANCIAL ASSETS

So far in this paper we have provided examples in only two economic factors, Inflation and Growth. Going forward, we will focus on four factors that describe the largest amount of price variation across assets, though the ERM framework is also capable of incorporating other factors.


The first additional factor many readers may have already guessed will be interest rates. This important economic factor has a direct effect on many assets. Normally, interest rate sensitivity is measured using duration. However, duration, which is the normal measure for rate-sensitive assets, it is not commonly used for most asset classes. Therefore, to have one measure that can be used seamlessly across all asset classes, we focus on interest rates in the same way we considered inflation and growth. Specifically, we compare changes in the 10-Year Treasury yield to the market forecast of that rate.

The macro-economic factors of growth, inflation, and interest rates are all macro-factors that have been well studied. The final factor we will introduce is focused on helping us understand that since we are studying assets in a financial marketplace, it is not enough to understand the macro-economic environment, but we also need to have a good grasp of the macro-financial environment.

The metric we use to capture the macro-financial conditions is called Systemic Risk. As the name implies, Systemic Risk is a measure of the risk 'system-wide'. In other words, it is the amount of risk that is not specific to equities, bonds, or any other individual asset or asset class, but rather shared across all markets.

Armed now with these four macro-factors and the Philadelphia Federal Reserve Survey of Professional Forecasters, we can propose a simple measurement system. In this initial model, we use the one year forward forecast for each of the macro-economic factors, then measure the realized change over that same year. Any difference between the realized change in the factor and the forecasted change will be considered to be a surprise in the factor. In the case of the macro-financial factor of Systemic Risk, it is by definition a random variable (if it could be known it would not be risky); therefore, any change in that metric is measured as a surprise.

In the chart below, we illustrate the asset sensitivity to changes in each macro-factor. At this point we encourage the reader not to get too bogged down in interpreting the factor sensitivity, as we will address this issue later. That said, factor sensitivity measures the average return expectation relative to a change in the factor. For example, if the sensitivity to a growth surprise was 1, and a 2% surprise occurred, we would expect an increase in our return over that period to be 2% ($2\% * 1 = 2\%$).



In the chart above, there are several interesting initial conclusions. First, when looking at Interest Rate Surprise sensitivity, we see that for fixed income assets the relative sensitivities are extremely similar to what we would suspect with the traditional Duration measure but again, interest rate surprise sensitivity can be used across all assets.

Using Surprises instead of actual changes can be a bit difficult at first, as we don't yet have a good sense of what kind of magnitude in Surprise we can expect (we will address this shortly), but the benefit is that we can apply this same metric across all asset classes, which becomes a powerful point of comparison.

The second big conclusion highlights that this method of measurement isolates each effect in the context of all four factors. For example, in focusing on TIPS, we first notice that although TIPS hedge out Inflation Surprise, TIPS are not the most inflation sensitive asset class. This is because TIPS as an asset are not only linked to Inflation Surprise, but also to Interest Rate Surprises, and the magnitude of that link is much stronger for interest rates. This can be seen by comparing longer dated TIPS, which exhibit a very small positive Inflation Surprise sensitivity but very large negative Interest Rate Surprise sensitivity, to shorter dated TIPS, which exhibit the reverse.

It is likely that Inflation Surprise and Interest Rate Surprise will move in related ways over some periods. Some methods of measuring inflation's effect on asset classes over the long term will combine these two factors because it can be justified that if inflation is rising, interest rates should rise as well, on average. Our view is that although that relationship may have broadly been true, on average, in the past, there are many periods when it breaks down. For instance, coming out of the 2008 Financial Crisis the Federal Reserve continued to bring rates down but, by implementing the Quantitative Easing (QE) program, inflation expectations began to rise. Without separating interest rates and inflation, it would have been difficult to understand those dynamics and the resulting asset class performance.

Finally, when comparing Growth Surprise Sensitivity and Systemic Risk Sensitivity, we see that they also tend to be mirror images. This relationship makes sense because almost all asset classes are positively exposed to Growth Surprises, when the economy is growing faster than expected, so are the values of assets. The only exceptions are assets like high grade bonds with short durations and cash that act as the "flight to safety" asset when Systemic Risk is on the rise. These two factors can be intimately related.

THE DISTRIBUTION AND SIZE OF SURPRISES AND ECONOMIC REGIMES

THE HISTORICAL POINT OF VIEW

To get a sense of the size of surprises, we again begin with the Philadelphia Federal Reserve Survey of Professional Forecasters. Keeping with our simple measurement methodology, we consider any realized difference in a factor relative to the forecast of that factor one year prior to be a surprise. The only exception, of course, is Systemic Risk, where any change is considered to be a surprise.

Using these data and definitions, we can see in the table below the historical size of surprises from each macro factor over one year.

	Inflation Surprises	Growth Surprises	Interest Rate Surprises	Changes in Systemic Risk
-2 Standard Deviations	-3.8%	-4.3%	-0.7%	-16.1%
-1 Standard Deviation	-1.7%	-2.2%	-0.2%	-8.2%
Average	0.3%	-0.2%	0.4%	-0.3%
1 Standard Deviation	2.4%	1.9%	1.0%	7.5%
2 Standard Deviations	4.4%	3.9%	1.6%	15.4%

There are a few things to notice from this table. The first is that the size and distribution of Inflation and Growth Surprises are extremely similar, and this relationship is convenient because when comparing sensitivity between these two factors, the raw factor sensitivities are equivalent to one another. In other words, a factor sensitivity of 1 equates to basically the same amount in both Growth and Inflation Surprises. Perhaps more importantly, the distribution of Growth and Inflation Surprises shows that the sizes of surprises are likely to

happen at roughly the same probability. For example, for both Growth and Inflation Surprises, close to 68% of the surprises are likely to fall in between the range of -2% to 2%.

Jumping a column to Systemic Risk, we see that the scale of this risk is much larger. This means that the factor sensitivity to this financial macro factor carries with it a much larger significance than the other factors. In practice this means a change in the level of Systemic Risk has much more impact on portfolio return than surprises in other factors.

Moving to Interest Rate Surprises, we see that the size and probability of rates surprising to the positive side are much larger than those to the negative side in the historical period. This is largely due to the fact that the Philadelphia Forecast data does not begin until 1981 for the 10-year interest rate, and for this entire period rates have been generally falling. As a result, any reversal in this trend has tended to catch markets off-guard.


Lastly, notice that all of the factors have average surprises that are non-zero. As was highlighted in the case of Interest Rate Surprise, we believe this is a consequence of the time period under study. It would be hard to justify, and would not make much rational sense, for the market to persistently over- or underestimate any of these factors. As a result when we apply these insights into a forward looking analysis we do not assume that these non-zero averages (sometimes called drifts) will exactly repeat themselves in the future.

In closing this section, historical information is often the best way to build a foundation of knowledge. However, ERM is a forward-looking framework and, therefore, it needs to build on this historical foundation. To be more explicit, if we were to stop at this point, we would be assuming that the historical distribution of surprises was the 'true' distribution. Some approaches have argued this view, however the ERM perspective is that to take that approach is analogous to trying to win the last war instead of the next one. In the next section we describe how ERM and Economic Regimes can make sense out of the historical distribution of surprises and provide the insight necessary for forward-looking investors to make informed decisions and understand their exposures to these factors.

ECONOMIC REGIMES AND FACTOR SURPRISES


Before jumping into economic regimes and how they affect factor surprises, we first examine the distribution of monthly changes in U.S. Inflation. In the chart below, we can see the historical monthly changes of U.S. Inflation as a histogram overlaid with the normal distribution that we would expect, given the summary statistics. We can easily see that the normal distribution is not a good fit for these observations, as the peak goes well beyond what we would expect, there are fat tails, and there seems to be a positive skew.

Histogram of Monthly Inflation Changes vs. Normal Distribution




The implication of this poor fit to the normal distribution is that by focusing on simple monthly changes in inflation, we have a poor and lumpy metric to compare to asset returns. As we discuss below, when we switch to surprises in inflation as opposed to simple changes in inflation, we have a much better and smoother metric for comparison.

The chart below shows that by focusing on surprises and controlling for the remaining ERM factors of real growth surprises, interest rate surprises, and systemic risk, we substantially improve the clarity of the relationship between changes in the TIPS yield and inflation.



The ERM methodology suggests we can make still further improvements by considering different regimes of surprises. Trying to fit all of these surprises into a single distribution that is constant through time is not consistent with economic rationale. Economies and markets are not static objects. Instead, as we might expect, economies and markets go through periods with very different characteristics. As we saw with the Survey of Professional Forecasts data, we are likely to see persistent periods, or “Economic Regimes”, when the factor surprises are likely to be positive, negative, or widely dispersed (i.e., uncertain market expectations). Perhaps then, instead of looking at a single distribution, we should combine distributions from different regimes with varying frequency. We can use the chart below to help get a better sense of the implications.

Monthly Inflation Surprises by Regime



The chart above again shows a histogram overlaid by three normal distributions. However, in this chart the histogram is of Inflation Surprise as opposed to simple changes. Notice that the histogram of Surprises still has many of the same problems as the histogram of changes, but that the scale of those issues like fat tails, etc., is decreased. In other words, the distribution is already a bit smoother.


Now overlaying the histogram are distributions that represents the three types of regimes that we discussed previously. Those are 1) a positive surprise regime where market expectations are likely to be too low and the realized change in inflation will be higher, 2) a negative surprise regime where the market’s expectations are likely to be too high, and 3) an uncertain regime where both very large positive and negative surprises are likely to be seen.

Although we will not spend much time on how these regimes are determined in this first part of our three part series, the idea of three different distributions is at the core of the ERM methodology. Instead of trying to fit a single distribution to the data, ERM is asking: If these are the three distinct regimes (distributions) from which surprises can come, with which regime (distribution) is any given observation most likely to be associated?

In practice, this means that the ERM methodology has fundamentally reversed the question. Usually, risk analysis assumes that a single normal distribution fits the data and then predictions would be attempted based on that assumption. In contrast, the ERM methodology approaches this problem by tying what we have learned in the earlier sections on expectation, surprises, and uncertainty to a rationale of why the distribution we observe has the distinct non-normal characteristics that we have seen.

By making the link between surprises, the different regimes (distributions) that may have produced them and asset prices, we can make a strong tie from economic theory to an intuitive rationale. The rationale is that asset prices adjust to surprises in economic factors because any expected change is already priced in. Surprises can come from one of three regimes (distributions) that represent forecasters being overly optimistic (leading to negative surprises), overly pessimistic (vice-versa), and having a large amount of uncertainty (leading to a wider dispersion of surprises).

The ERM methodology states that because the observed Surprises come from one of these three distributions in varying frequency, we can explain why the shape of the distribution of surprises is non-normal. Briefly, notice that the three distributions overlap, and that where there is overlap, the histogram is larger than you otherwise might expect. To get a better sense of how this methodology works, we examine the following chart.



The chart above shows the same information as the previous histogram. We can think of the chart above as looking down on the previous chart. The chart directly above contains some important additional information. First, the numbers on the chart represent the historical

frequency of the observations. Specifically, the Negative Surprise Regime occurred most often during the available sample, at 41% of the time. We extended the timespan here by using the Representative Agent Forecaster mentioned above, and this allows us to extend the analysis back to 1968. We can see that the Positive Surprise Regime and Uncertain Regime have occurred 37% and 23% of the time, respectively. The second piece of additional information we can see are the actual observations represented as dots from within each Regime and then added together to form the final 'All' distribution, which was represented as a histogram of Surprises previously. By combining these different Regimes, we can make sense of how the distribution of Surprises we see today came to be, and rationalize what we see with the changing economic environment.

Understanding the dynamics of this Regime addition is an important piece of the ERM methodology. We have built a good understanding of history, but looking forward it is not likely that inflation will continue to be falling, and hence we are not likely to be in a Negative Surprise Regime at the same frequency. If we were to use only the historical data and not adjust for this fact, we would be over-emphasizing Negative Inflation Surprises and falling inflationary periods relative to Positive Inflation Surprises.


PUTTING ECONOMIC REGIME MANAGEMENT TO WORK

In the last few sections we have measured the sensitivity of financial assets to macro-factor Surprises and shown that these factor Surprises come from a combination of Regime distributions that add up to what we have observed historically. In this first part of the series, we will focus on how long-term investors can make use of this information by examining their portfolios' exposures to the macro-financial factors. This can lead to important insights such as where there may be macro-level risk concentrations in the portfolio. Because the level of macro-financial sensitivity can be measured for each asset, if an investor finds that her portfolio is overly exposed to one macro-factor, the most effective change to the portfolio can be easily found. For example, considering the graph on page 9, we can see that if investors are concerned about their exposure to growth, reducing Public REITs would be the most effective course. The process can become considerably more complicated when trying to get the most effective use from the four macro-factors simultaneously combined with other investment objectives. However, the ERM methodology puts a clear structure around this complicated process and can help guide investors towards a more effective use of these macro-level exposures.


The most obvious insight is that, similar to looking at factor sensitivities for individual asset classes, we can measure the long-term sensitivity of a total portfolio to different macro-factors. In the chart below, we see the same simple measurement system used earlier to measure asset class sensitivity at the total portfolio level. Focusing specifically on the portfolio exposure to surprises in growth, the chart below shows ~4.5X sensitivity to growth. This implies that if a growth surprise of 2% was realized, then the change in the portfolio total return would be 9% ($4.5 * 2\% = 9\%$) higher than expectations. For example, the portfolio represented below is a diversified portfolio with an 8% return target.⁵ Holding

⁵ Please see appendix.

everything else constant in the previous example, the annual return over that one-year period would be estimated at 17%.




This kind of analysis can be a powerful tool to see the source of portfolio risk in terms of macro-factors. Although this view can make it easy to think of examples of events where growth surprises by 1%, interest rates by 2%, and so forth as we saw in the historical and regime specific distributions, the range of these shocks in Surprise can differ substantially between factors and Regimes. One adjustment we can make to get a better sense based on the likelihood of the scale of these Surprises is to show the portfolio effect at a one standard deviation event. The advantage of this adjustment is two-fold. First, instead of using factors sensitivity, we can see the effect on the Y-axis in terms of percentage change in portfolio return. The second advantage is that at one standard deviation, each Surprise would take place at the same level of probability.




After making this adjustment there are a few key conclusions. Most obviously, Growth Risk and Systemic Risk make up the majority of the macro-factor risk within this portfolio.

When comparing a traditional 60% equity and 40% bond benchmark to our sample “Diversified Portfolio,” a distinct concentration of risk can be seen. In the chart below, we can see that the Diversified Portfolio has reduced exposure to the largest sources of risk and is much more evenly diversified among sources of risk.




We have highlighted some simple metrics that can help investors understand the sensitivity of their portfolios to Positive or Negative Surprises. These metrics can help clarify the sources of macro-factor risk, decompose that risk, and show the composition of total risk from that perspective. One insight from ERM discussed, but not yet shown, is that uncertainty can have a large effect on Surprises and the performance of assets and portfolios, as well as variance. To get a sense of how to use this information, let us revisit an earlier chart, but swap Growth for Inflation to avoid repetition.



We can see the historical frequency of the different Regimes of Growth Surprises and how the uneven addition of those regimes led to the non-normal distribution of Surprises that we can observe today.

As previously stated, it is important to learn from this historical foundation, but it is unlikely that the future will look exactly like the past. Therefore, when deciding to be forward looking, we can adjust these frequencies to describe not only a Neutral future scenario but also long-term scenarios where the Surprise Regimes tend to be more Negative, Positive, or Uncertain. We can see this illustrated in the chart below, where the probabilities show the frequency of each Regime in the four future scenarios and the resulting differences in the average and the 95th percentile of Surprises. In contrast to the earlier regime addition charts, we have annualized these numbers so that they represent expectations over a single year.

Growth Scenario Regime Addition



As illustrated in the chart above, we can describe different future scenarios by varying the amount of time Growth Surprises are in the Negative, Positive, or Uncertain Regime. This changes the final distribution of Surprises that a long-term investor will face. Note that the average expected Surprise is different in each scenario. We show the Surprise at the 95th percentile, and the chart shows that if the Uncertain Regime dominates, the size of this Surprise can increase substantially.


Changes in the distributions of Surprises are directly tied to future return expectations. In the chart on the following page, our future return expectations vary in each scenario compared to the base case expectations in response to these varying regime dynamics.

Focusing first on the Neutral scenario, notice that although the average return expectation is the same as the base case, several characteristics are changed. Accounting for the exposure to Growth Surprises introduces a positive skewness to the distribution. This causes the tail risk,⁶ or expected loss from an 'extreme' scenario, to be substantially worse than the base expectation for that same event. This 'fat' left tail is more closely aligned with investors' past experiences, and we believe that this type of behavior (especially for heavily Growth-exposed portfolios) will persist in the future. It is worth noting that in each of the scenarios that we constructed, this 'fat' left tail effect is present and is worse than the base expectation would lead investors to believe. The Uncertain scenario is on average worse

⁶ We chose to use (95th Percentile) VaR as the primary metric for our analysis as it represents a single number with an intuitive, standard meaning that can be easily compared across different examples and other types of analysis.

than the Negative scenario, but this averaging hides the fact that some of the largest positive surprises are during Uncertain Regimes. The takeaway is focused more on risk and dispersion than any particular average. In short, we can see that even in cases where bad news is more likely, this tends to be better than times that are dominated by uncertainty; however, if investors sought to avoid uncertainty altogether, they would miss some of the largest and most positive outcomes for their portfolio.

Return Expectations by Growth Scenario



In closing, we have seen that combining a straightforward system of sensitivity measurement and general scenario analysis through varying regime frequency can provide valuable descriptive insight. These insights can help investors understand their portfolio sensitivities to different economic factors and to shed some light on how those portfolios might perform in some future economic scenarios. In the chart above we are only examining different scenarios in Growth. We could examine any of the macro-financial factors and even combinations of two or all four macro-financial factors.

In this first part of our three part series we suggest that investors begin by looking at portfolio exposures in isolation. This focus on individual macro-financial factors - one at a time - can add significant clarity to a complex problem and aid investors in diversifying sources of macro-factor risk. Building long-term scenarios to obtain a better understanding of their risk exposures in different economic scenarios can add to this clarity and put these insights into action for the portfolio. Through that process, investors should have a good sense of what sources of macro-risk are driving their portfolios. If investors can also have a good understanding of what they risk losing in extreme scenarios, this can be an important insight to help build better portfolios.

In the next two parts of this series, we will introduce more specific regime definitions, scenario construction, optimal factor diversification, and exposure trade-off analysis. At their core, each of these new ideas is an extension of an idea that we have put forth in this first paper. We suggest that investors begin with an analysis of the type we have introduced above.

In summary, understanding the sensitivity of a portfolio to macro-factors is extremely valuable to long-term investors. If these tools can help investors distribute their sensitivities to macro-factor risk and set more informed thresholds about the type of risk they are willing to take on to meet their long-term return objectives, then the insights that ERM can offer will lead to improved investment outcomes.

APPENDIX

EXTENDING TIPS BACK IN TIME

As stated in retrospect since we know what the changes in inflation and rates were over this period. This allows us to build a highly accurate simulated TIPS return stream that allows us to extend our analysis back in time. We use a combination of the fundamental formula for the TIPS yield below:

$$r_{TIPS} = \left(\rho + \gamma_{\rho} + \tau \left(\frac{\pi}{1+\pi} + \gamma_{\pi} \right) \right) / (1 - \tau)$$

Where:

- r_{TIPS} is the real yield on TIPS
- ρ is the real rate of interest
- γ_{ρ} is the real rate of interest risk premium
- τ is the tax rate
- π is the inflation rate
- γ_{π} is the inflation rate risk premium

Additionally, we use the nominal bond yield, the trade weighted dollar, the real yield on cash, the Bloomberg Commodities Index, and realized volatility of the Consumer Price Index as predictors of the model above to build a maximum likelihood estimation of the real TIPS yield.

REPRESENTATIVE AGENT MODEL

In this construction of the representative agent we have used a simple extension of what is called an Auto-Regressive process of order one. These class of models attempt to forecast the next observation by looking at the previous observation and they tend to perform reasonable well in tests. Our extension is to use a rolling average input as well. The intuition for this is that in the market we have two types of forecasters 1) those looking at the long term trend and 2) those with a short term view. Our model captures exactly that and weights between the two.

$$X_i = \varphi_1 * X_{i-1} + \varphi_2 * \frac{1}{12} * \sum_{n=1}^{12} X_{i-n}$$

DIVERSIFIED PORTFOLIO WEIGHTS

Asset Name	Weight (%)
Bank Loans	2.0
Commodities	4.0
Core Bonds	9.0
Core Infrastructure	1.0
Core Real Estate	4.0
EAFE Equity	10.0
EM Debt	6.0
EM Equity	11.0
Frontier Markets	1.0
High Yield	4.0
Natural Resources	3.0
Non-Core Infrastructure	1.0
Opportunistic RE	1.0
Private Equity	7.0
Public NR	2.0
Public REITs	1.0
RE Debt (HY)	1.0
TIPS	9.0
U.S. Equity	21.0
Value-Added RE	2.0