GMO White Paper

Table of Contents

The Idolatry of Interest Rates
Part II: Financial Heresy

James Montier

In many ways this is perhaps my most personal essay, not because I'm about to share some deep and dark personal revelation (I can almost hear the collective sigh of relief), but rather because this essay reflects my views and mine alone. Others at GMO should not be tarred with the brush of my beliefs, and I have no doubt that many will disavow any association with the views I express here.[1] In fact, my colleague Ben Inker's "rebuttal" follows this piece.

In this essay I want to build on some of the ideas that were developed in my last essay[2] specifically as they pertain to thinking about asset markets. The most obvious place to start is with the idea that the natural rate of interest is a myth. Accepting this idea has many ramifications for the way in which one conducts asset pricing.

An unanchored system: the world without a natural rate of interest

Perhaps the clearest implication from my previous essay with respect to investing is that the cash rate is potentially unanchored. That is to say, without a natural rate it isn't obvious what cash rate one should expect to see in the long term. Or, in the parlance of GMO, what is the long-term level of cash to which we should mean revert, or, perhaps, should cash even be considered to mean revert?

In a previous piece[3] I questioned the logic of mean reversion when it comes to cash rates (and hence bonds). If there isn't a natural rate, then we are left trying to guess what a group of people (the Federal Reserve, for example) will think is appropriate seven years[4] from now. Because we can't even know for sure who is going to be on the FOMC in seven years' time, this is an exceedingly difficult task. Pythia herself may have struggled with such a challenge.

As Exhibit 1 shows, the real rate is probably best thought of as a policy variable, set by the central bank. The breaks in the real rate series (statistically estimated) track incredibly well with the tenure of various Federal Reserve Chairmen.
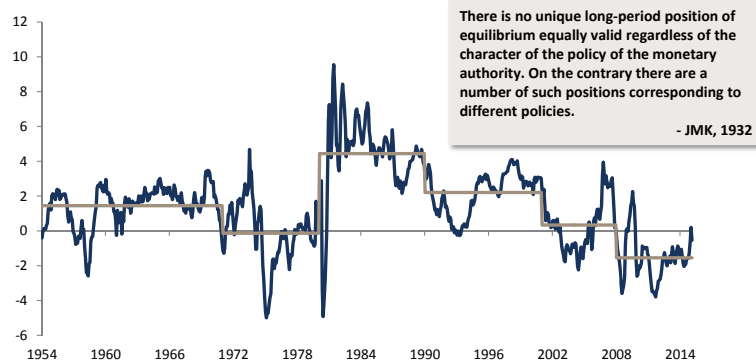
---

[1] It is this open disagreement that makes GMO a wonderfully stimulating place to work. As the late, great Sir Terry Pratchett put it, "The presence of those seeking the truth is infinitely to be preferred to the presence of those who think they've found it."

[2] See "The Idolatry of Interest Rates, Part I: Chasing Will-'o-the-Wisp," May, 2015, available at www.gmo.com.

[3] See "The Purgatory of Low Returns," GMO 2Q 2013 Quarterly Letter, available at www.gmo.com.

[4] GMO predicts asset class returns by assuming that profit margins and price/earnings ratios will move back to long-term average in seven years from whatever level they are today.

Exhibit 1:

Real Fed Funds Over Time (%)

> There is no unique long-period position of equilibrium equally valid regardless of the character of the policy of the monetary authority. On the contrary there are a number of such positions corresponding to different policies.
>
> - JMK, 1932

Source: U.S. Federal Reserve, GMO

So what is one do to? One could turn to the tender mercies of history to see if any meaningful answers can be gleaned from a careful reading of the past. However, and perhaps unsurprisingly, history provides us with little comfort (consistent with my view of real interest rates as a policy variable). Exhibit 2 shows the average, median, and standard deviation across 21 countries[5] over different time periods.

Exhibit 2:

Panel Data on Real Interest Rates Across 21 Countries
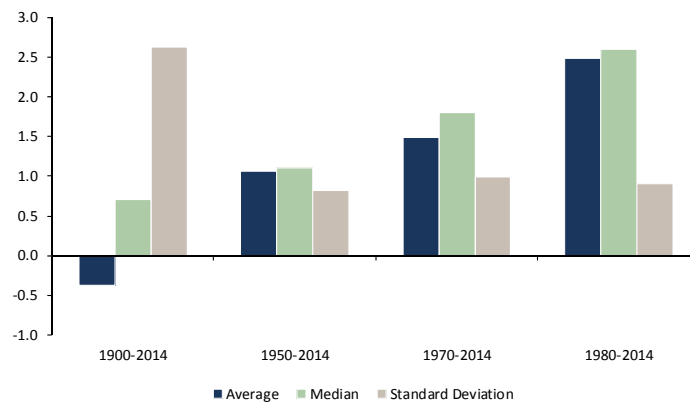
Source: Dimson, Marsh and Staunton, GMO

The answers obtained from history depend massively on the sample period chosen. Using the longest sample period reveals an average real interest rate of -0.4% p.a., but this is coupled with a median real rate of 0.7% and a very wide distribution. I would also caution that mixing differing monetary regimes may further muddy the already opaque waters. For instance, the prevailing real rate under a system such as the gold standard is likely to be very different from that seen under a fiat currency system.[6]

---

[5] Countries include Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, South Africa, Spain, Sweden, Switzerland, United Kingdom, and United States.

[6] The failure to distinguish between different monetary regimes is a common failure amongst economists. For instance, Wray and Nersisyan (2010) point out that all of the defaults that Rogoff and Reinhart document occurred under systems of fixed exchange rates. Yet their results are often generalized (very inappropriately) to all monetary regimes.

GMO
North America | Europe | Asia-Pacific

To help illustrate this point, Exhibit 2 breaks out three sub samples: 1950-2014 represents the post war period; 1970-2014 captures the period of fiat currencies; and 1980-2014 attempts to determine whether the high inflation experience of the 1970s significantly impacted the previous selection. The results reduce the standard deviations across all of the samples, but we can still see that the average ranges from 1% to 2.5% and each is very sample-specific. History really isn't much of a guide as to a reasonable level of real rates to assume for the long term.

In a minority of cases one could turn to the predictions of the central bank itself. For instance, the Fed produces its now infamous "dots" diagram, which shows the interest rate forecasts (albeit in nominal terms) from voting members of the FOMC. This can be combined with information from, say, the inflation swap market (or the inflation forecasts from the Fed) to derive the "Fed's view" of long-term "neutral" real interest rates. Doing such an exercise today shows the median Fed expectation for the longer term to be 3.75% nominal; using an inflation "forecast" of 2% would give a long-term real rate expectation of 1.75%.

There are three problems I can think of with doing this. The first is that the Fed is no doubt completely aware of people like me doing things like this. Thus, we can't disentangle what the Fed wants us to believe from what the Fed really believes. That is to say, the dots may not represent the "true" belief of the Fed, but rather the belief that the Fed would like the market to think it is their true belief. Second, this estimate will not be independent of the model used by the various FOMC members, almost all of whom are seemingly wedded to so-called Dynamic Stochastic General Equilibrium models. Some of the problems inherent in such models were discussed in Part I of this series. Suffice to say, they assume what they seek to prove: a natural rate of interest based on time preference. Third, this kind of process assumes that the Fed knows more than anyone else does – a highly debatable point. Remember, this is the institution that, despite its army of PhD economists, failed to warn of either of the two most recent major economic disasters – the TMT bubble and the housing bubble.

Alternatively, one could turn to the market and ask what is implied for future cash rates from market pricing. With the help of data and models from the New York Fed (the ACM model) and the Cleveland Fed (expected inflation), one can back out the market's implied real rate expectations. Exhibit 3 shows the average real rate expected over the short term (0 to 5 years) and the longer term (5 to 10 years). A clear difference of opinion between the market and the Fed can be observed. The market is implying a long-term real rate of 1% versus the Fed's view of 1.75%.

Exhibit 3:

**Market Implied Real Rate Expectations (U.S.)**

Source: FRB New York, FRB Cleveland, GMO

GMO
North America | Europe | Asia-Pacific
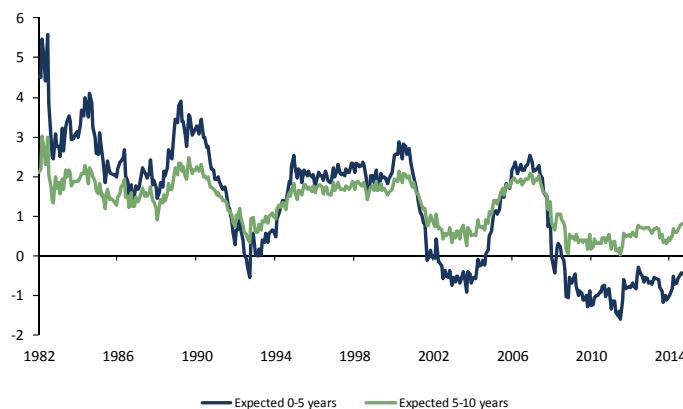
Of course, this raises the question as to whether the market implied view has been of any use in "forecasting" the future real rates. The evidence doesn't provide much by way of succor. Exhibit 4 shows the average expected real rate over the coming 10 years, and compares it with the actual delivered real rate over the same horizon. Thus, the reading for 2004 represents the market implied view of the real interest rate over the next 10 years (around 1% real) compared with the actual outturn over those 10 years (not far off -1% real).

During the early part of our sample, the market implied view and the actual outturn coincided fairly well. However, since the Greenspan years (1987 onwards), the implied view and the outturn have borne little relationship to one another.
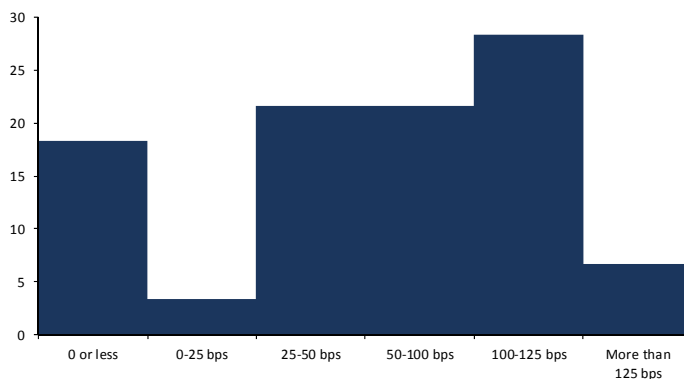
Exhibit 4:

Implied Real Rate vs. Actual Rate over 10 Years
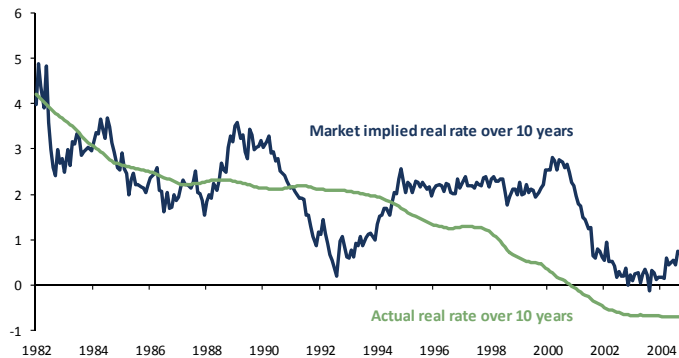
Source: FRB New York, FRB Cleveland, GMO

Out of curiosity, last year my colleague Sam Wilderman and I asked our in-house research conference audience (around 70 investment professionals) what they thought the terminal cash rate in seven years' time would be. The results are displayed in Exhibit 5. The average turned out to be 80 bps, but what was more interesting to me was the spread of opinions. The standard deviation of the responses was 0.5%, hence we can say that GMO is 95% confident that the real rate of interest in seven years' time will be somewhere between -20 bps and +180 bps. As Ygritte from Game of Thrones would say, "You know nothing, Jon Snow."

Exhibit 5:

GMO Investment Professionals' View of the Ending Cash Rates in 7 Years' Time

Source: GMO

So we have a range of potential views on what might serve as an anchor for cash in the absence of any natural rate of interest, as set out in Exhibit 6. The average of these various estimates is just over 1.25% (oddly enough this is the number we use to construct our baseline forecasts for cash and bonds in the 7-Year Asset Class Forecast framework). Whilst the range in Exhibit 6 may not look high, when translated into a confidence interval, this data is equivalent to saying that we can be 95% sure that the long-term rate lies in the range of 0.4% to 2.1%!
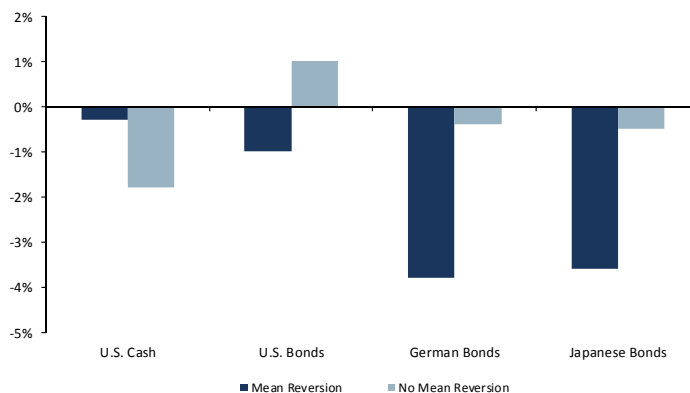
Exhibit 6:

Summary of Views on the Likely "Long-term" Real Rate

Source	Estimate
History	1.50%
Federal Reserve	1.75%
Market	1.00%
Investment Professionals at GMO	0.80%

Source: GMO

What do we actually do? As noted above, when we construct our baseline forecasts we use a terminal cash rate of 1.25% real. Based on the evidence presented above, this doesn't seem wildly out of line with a range of views. However, we also run versions of the cash and bond forecasts wherein we don't assume any mean reversion at all: Cash rates stay where they are, inflation stays where it is, and the term premia (for bonds) stay at current levels. This allows us to gauge the impact of our baseline assumption upon our perception of value. A comparison of the forecasts with the assumption of mean reversion and without mean reversion is shown in Exhibit 7.

Exhibit 7:

Select Fixed Income Forecasts With and Without Mean Reversion
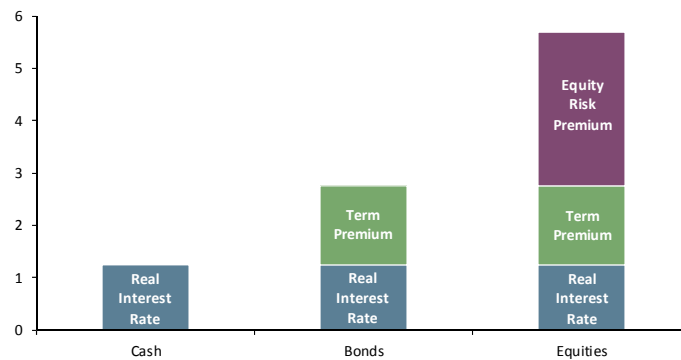

Source: GMO

Perhaps the most striking aspect of Exhibit 7 is just how unattractive certain bond markets are, even under the assumption of no mean reversion. This echoes the point made by Ben Inker in his recent quarterly missive[7] – you have to believe some really extreme things to think that several major bond markets look anything except appalling from an investment point of view.

---

[7] See Ben Inker, "Breaking Out of Bondage," GMO 1Q 2015 Quarterly Letter, available at www.gmo.com.

GMO
North America | Europe | Asia-Pacific

However, I think there is a deeper lesson we need to take from the above work as to the appropriate level of the real cash rate. We need to be honest with ourselves and admit that no one truly knows what the future real rate will be. An obvious corollary to this statement is that it would be dangerous to build an approach to asset pricing based on something very uncertain, or as Keynes put it, "It would be foolish, in forming our expectations, to attach great weight to matters which are very uncertain."

Yet an approach that puts the real cash rate at its very heart is surprisingly common. It is the risk premia building block approach to thinking about asset valuation. Exhibit 8 is a stylized, simple example of this kind of approach. In essence, the approach starts with the real interest rate and then builds on various risk premia (i.e., term premium for government bonds, plus credit risk premium for corporate bonds, and plus equity risk premium to get to equities). As one user of this approach wrote recently, "The expected returns of both mainstream equities and bonds are heavily dependent on the equilibrium real rate of interest, which serves as a key building block for all long-term investment returns in an economy."[8]

Exhibit 8:

Building on Quicksand – the Risk Premia Approach to Asset Pricing

Source: GMO

To me this is the intellectual equivalent of building on quicksand. The more blocks that get added, the more precarious the edifice becomes. Bonds and cash are clearly related, and I obviously have no quibble there, but when it comes to equities, I get very nervous about the approach.

The equity risk premium – a block too far

This is where this essay gets really personal. I recently presented some of the ideas that follow at another of our internal conferences. I started by asking the audience (around 80 investment professionals) to raise their hands if they thought that interest rates mattered for equity valuation. It transpired that over 99% of those present thought they did, leaving me in a minority of one. (I beg you to recall Orwell's words that perhaps the lunatic was simply a minority of one, although I suspect any number of colleagues and readers would suggest the reverse causality.) Louis Nizer once opined, "I know of no higher fortitude than stubbornness in the face of overwhelming odds." So, undaunted,

---

[8] Shane Shepherd, "Slow Growth: A Tale of Two Theories," Research Affiliates, May, 2015. I quote this not in any way to ridicule Research Affiliates, whom I hold in great esteem, but rather to highlight the prevalence of this kind of thinking, even amongst the very smart. Indeed, a number of my colleagues (also very smart people) subscribe to this viewpoint.

I want to make the case that given the uncertainty over the real cash rate in the long term, we are better off ignoring it and valuing equities independently.[9]

The idea that interest rates should matter for equity valuations seems blindingly obvious to anyone familiar with the general dividend discount approach. After all, the present value of an asset should be the discounted sum of its future cash flows.

$$P_0 = \frac{D_1}{(1+k)} + \frac{D_2}{(1+k)^2} + \frac{D_3}{(1+k)^3} + \cdots \frac{D_T}{(1+k)^T}$$

However, this general approach says nothing about the appropriate discount rate to use, other than that it should be the required return on equity. In effect, the decision to make K (the discount rate) a function of the interest rate is an implicit assumption. The required return to equity investment could be conceived to be independent of the interest rate if one chose.

A brief history of the equity risk premium

As far as I can establish, the equity risk premium (ERP) was born of ex post studies of the returns to investing in equities and bonds (e.g., Edgar Lawrence Smith's seminal 1924 book, Common Stocks as Long Term Investments). That is to say, although not known by the name ERP until considerably later, the ERP was simply a term for the relative performance of equities vis-à-vis bonds in a given sample.

The first instance of the ERP being used in a forward-looking (ex ante) valuation sense that I've come across is in John Burr Williams' The Theory of Investment Value. This is the first book to outline the dividend discount model (or, as he termed it, "evaluation by the rule of present worth"). At the start of the book, Williams talks about the discount rate as the required rate of return to equity investors.

However, when talking later in the book about constructing a forward-looking discount rate for use in empirical work, he notes, "The customary way to find the value of a risky security has always been to add a 'premium for risk' to the pure interest rate, and then use the sum as the interest rate for discounting future receipts." Indeed, Williams provides a table of "Interest Rates, Past, Present and Future," which takes the riskless rate as the long-term government bond rate of 4% and the expected return to "good stocks" as 5.5%. This is, of course, exactly the same approach as the risk premia building blocks approach I mentioned above.

Importantly, it should be noted that Williams pays attention to the danger in using cyclically-low bond yields as the basis for the discount rate. He notes (referring to yields in 1938), "Present long-term yields…are too small…and should be larger to correspond with what seems to be a reasonable expectation for the real rate of interest in the future." Herein lies the major problem, as per the above discussion: How on Earth do you know what the appropriate real rate of interest actually is?

When you look at the ex post data (see Exhibit 9), the basic lack of a strong anchor for cash rates again becomes clear. However, the real equity return has been mean-reverting. Of course, it has displayed significant volatility over time, but historically has reverted to around 6% p.a.

---

[9] An oft-heard version of this argument is that the lack of decent alternatives suggests that we should own equities because they are "relatively attractive," even if not absolutely so.

This accords nicely with my mental model of how the world works. As already mentioned, the cash rate is best thought of as a policy variable that shifts around over time and doesn't display any real tendency to mean revert. However, as I will explore a little later, I think I can provide a framework for thinking about the return on capital – and the determinants of its mean reversion.

Exhibit 9:

Ex Post U.S. 10-year Equity and Cash Real Returns (% p.a.)
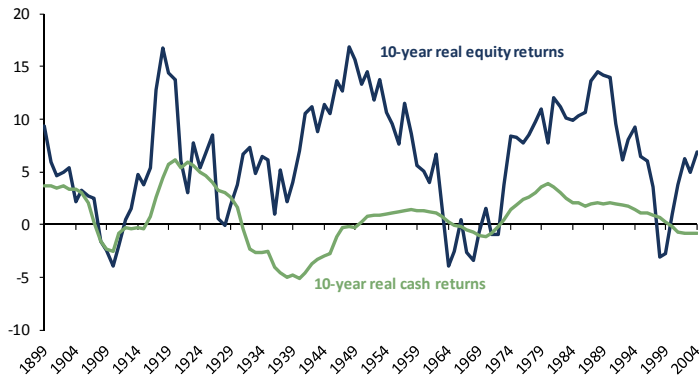

Source: Dimson, Marsh and Staunton, GMO

Given an equity return that has tended to mean revert, and a cash return that hasn't particularly, the concept of the equity risk premium is not likely to be of very much use. Any properties that is has to mean revert are likely to be driven by the equity side of the equation. Thus, using the ERP in valuation is, to me at least, akin to taking a glass of water (the likely return on equities) and adding some mud (the interest rate), and then finding one's self with a glass of muddy water.

Exhibit 10:

Using the ERP Is Equivalent to this...
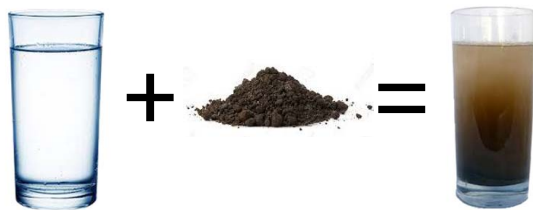
For the sake of clarity, I want to explore a couple of toy examples that may help to make my position a little more plain. One of my colleagues is fond of the following way to pose a question: What if an angel appeared and told you that cash would return 5% real forevermore?

First, it would lead to a pretty radical reassessment of some of my religious beliefs or, more likely, a trip to an asylum because listening to angels might have me questioning my own sanity. But let's assume that I haven't lost my marbles, and that somehow I now know for sure that real cash rates will be 5% forevermore. How would this affect my view of equities? Obviously, and I think in keeping with every sane person, I'd require more than 6% real from equities.

GMO
North America | Europe | Asia-Pacific

Conversely, if the angel revealed that future cash rates would be -10% forevermore, I'd happily admit that I would require less than 6% real for equities.

Haven't I just admitted that real cash rates matter for equity valuation? Yes, in extremis, but in neither of the two cases above do I really believe that the real cash rate would be sustainable at such levels in perpetuity; the real cash rate is probably bounded and can't just wander off all over the place, barring events like hyperinflations. In effect, I'm arguing that given the range of probable outcomes, the epistemological limits are binding. That is to say that we can't really tell what the real rate will be in a range that matters for our investments.

Nor is this ignorance limited to the real interest rate. Let's now turn to the state of our knowledge of another of the building blocks in the risk premium approach: the ERP.

What do we know about the ERP?

If we don't know much about the real interest rate, what do we know about the ERP? Once again we run into trouble with regards to our stock of knowledge. Let's start with the academic work. As the classic paper by Mehra and Prescott[10] showed, there was no way that standard economic models could account for the ERP. Indeed, such a model (calibrated to fit history in regard to consumption growth) suggested a very high risk-free rate (around 13%) and a very low equity risk premium (around 1.4%), almost the exact opposite of the world we observe, while in the period of their original sample, the risk-free rate was 0.8% and the ERP was 6.9%. In a published retrospective on his own paper, Mehra notes, "The 1.4% value represents the maximum equity risk premium that can be obtained in this class of models" given "sensible" constraints on the parameters.

Mehra continues, "The puzzle cannot be dismissed lightly because much of our economic intuition is based on the very class of models that fall short so dramatically when confronted with financial data. It underscores the failure of paradigms central to financial and economic modeling…Hence, the viability of using this class of models…is thrown open to question."[11]

Of course, since that seminal work of 30 years ago, a micro industry has sprung up, seeking to explore the equity risk premium puzzle via "relaxing the assumptions" made by Mehra and Prescott. These "fixes" range from changing the measure of the risk-free rate to using alternative preference structures (i.e., habit formation), from modelling rare disasters to understanding borrowing constraints, and just about everything in between.

However, the problem with the plethora of "fixes" that have been offered up is not only that, as Mehra states, "None has fully resolved the anomalies,"[12] but collectively they provide an embarrassment of riches for those of us in the debunking game. By selecting your favoured explanations, you can come up with just about any risk premium for equities that you would like. So, unfortunately, theory has very little to tell us about the sources of the ERP, and hence very little to say about what we might expect it to be in the future and how it might change over time.

---

[10] Rajnish Mehra and Edward C. Prescott, "The Equity Premium: A Puzzle," Journal of Monetary Economics 15 (1985), 145-161.

[11] Mehra, "The Equity Premium: Why Is it a Puzzle?" *Financial Analysts Journal*, 2003.

[12] For instance, the rare disaster literature has much in common with the way a number of my colleagues express their thoughts (i.e., equities are paid because of a chance of the 1930s happening). However, as Nakamura et al. (2013) show, these models require a high elasticity of intertemporal substitution (that is, consumption to be sensitive to changes in the real interest rate). The empirical literature, as per the Havranek (2015) meta survey, puts the estimate "deep below 1" with a point estimate of around 0.3, nowhere near large enough for the rare disaster literature to be correct.

There are other sources of views on the ERP that we might consult. However, without any form of theoretical underpinning, we must remember that all of these are little more than guesses and numbers. Ivo Welch of Brown University has occasionally surveyed finance professors for their views on the long-term ex ante ERP. In his most recent survey (2008), the average long-term ERP across some 400 finance professors was 5% with a standard deviation of 1.7%. Thus, finance professors in aggregate are 95% certain that the ERP lies somewhere between 1.6% and 8.4%!

John Graham of Duke University regularly surveys chief financial officers, and one of the questions he asks concerns their expectation for the long-term ERP. The estimate of the average expected ERP over the next 10 years is around 3.73%, accompanied by a large standard deviation (2.63%).

The last group one can turn to is investors themselves (probably the most useful group because they collectively actually help determine the ERP). In a previous existence, I did exactly that and asked investors what they expected the ERP to be over the long term. The average response was a number surprisingly close to the CFO group's response at 3.5%, although the standard deviation of estimates was markedly narrower than that of the CFO group. All three sets of average responses and 95% confidence intervals are shown in Exhibit 11.

Various Estimates of the Ex Ante ERP and Their 95% Confidence Intervals
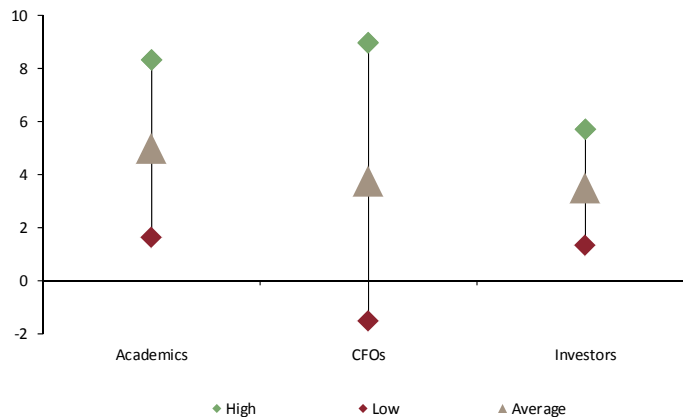
Source: GMO

As with the real rate cash estimates shown earlier, whilst the average values may be reasonably close, the range of outcomes is huge. Hence, when they are joined together à la the building blocks approach, we are effectively compounding ignorance. Remember, the real cash estimates ranged from 0.4% to 2.1%. Couple these estimates with ERP estimates that range from around 1.5% to 8% (I'm ignoring the negative tail on the CFO distribution), and perhaps you can begin to appreciate why I don't have very much faith in this kind of approach (and that is without even beginning to consider the role of the term premium).

Examples of ERP models

Let's look at the performance of a simple ERP model versus an alternative model based on simple Shiller P/E. I've constructed the ERP model in the following way. Let's assume that we have perfect foresight on real interest rates. That is to say that at any given point in time we can look 10 years

forward and know with certainty exactly what the real interest rate will have averaged over those 10 years. To this I've added a constant ERP figure of 4.1% (the average of the averages across the three groups shown in Exhibit 11). The sum of these two figures gives us our expected equity return, which we then invert to give a justified P/E. We can then see if the predictions derived from this approach are more accurate than those from one that simply assumes a constant "fair value" P/E.

As soon as one begins to operationalize the ERP approach, a road block is hit. At some points in time, the discount rate calculated becomes negative because the perfect foresight 10-year cash rate is below -4.1%. A negative discount rate blows up the model completely. This obviously renders the building blocks approach extremely problematic.

To be as kind as I could be to this approach, I've looked only at the post war period, when this problem doesn't arise. Exhibit 12 shows a comparison between the perfect foresight ERP approach (PF ERP), a standard constant Shiller P/E-based model, and the actual outturns.

Exhibit 12:

Comparison of Forecasting Approaches
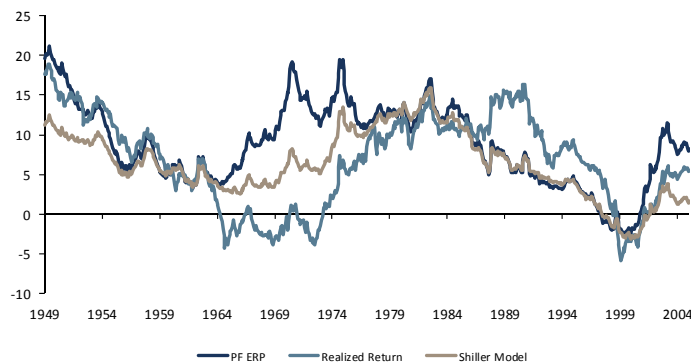
Source: GMO

On average, the standard Shiller P/E model has had a smaller forecast error than the PF ERP model. That is to say that you were better off ignoring the impact of real rates, even if you had perfect knowledge of what they would be! When combined with the standard Shiller model's more robust nature (i.e., not having the problem of negative discount rates), to me at least, this adds up to some hefty weight in favour of the standard Shiller-style model.

Another way of looking at the power of an ERP-based approach is to take the forecast errors from the standard Shiller approach and see if the real interest rate can explain them. In effect, is the standard Shiller "forecast" error positive when rates are low and negative when rates are high? (Or, does the real interest rate help explain the forecast error?)

Let's start by looking at the forecast error itself (and its sources). As Exhibit 13 shows, the forecast error over time is substantial, and that shouldn't be a surprise to anyone who has ever used a value approach, which tells you nothing about timing. Cheap assets can always get cheaper, and expensive assets can always get more expensive.

I've decomposed the forecast error into the contribution from three sources: the fundamentals, the P/E, and the "margin." You can see that there is no particular pattern to the sources of the forecast error. The fact that the "margin" and the P/E contributions often offset each other is why one should use a cyclically-adjusted P/E. You can see the impact of the poor fundamentals associated with, say, the Great Depression by the significant contribution they made to the overall forecast error in that period.

When we look at the TMT period, we can see the model underpredicting reality and, more importantly, see that it was really all about the P/E, in keeping with the mania-like nature of the episode. More recently, we see a similar (although far less pronounced) experience: The model has been too conservative when compared to the outcome, driven once again largely by the P/E.

Exhibit 13:

Shiller Forecast Error and Its Decomposition
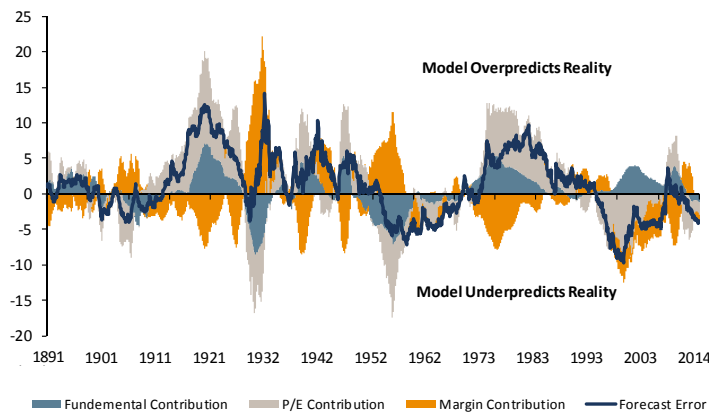

Source: Shiller, GMO

This last fact may seem like prima facie evidence of the idea that interest rates matter for valuation. However, when I compare the forecast error (or the contribution from the P/E – it makes no difference to the conclusion) to the real interest rate I simply don't see any evidence of a relationship at all, as seen in Exhibit 14.[13]

There have been times when the real interest rate was negative and the model overpredicted reality – the complete opposite of the ERP approach claim. Similarly, there have been periods of relatively high real interest rates when the model underpredicted reality. There simply doesn't appear to be any consistent pattern.[14]

---

[13] These charts start in the 1930s because the earliest interest rate (at least at the short end) data begins in the 1920s, so a 10-year perfect foresight means that we need to start in the 1930s.

[14] I have also looked at only the valuation-driven component of the total error, and exactly the same pattern holds. This isn't just a case of the total forecast error obscuring the devil hidden in the details.

Forecast Error and the Perfect Foresight Real Interest Rate
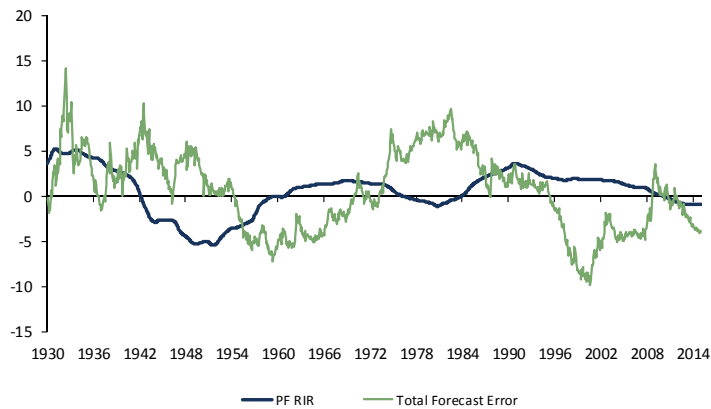
Source: GMO

To me this provides further evidence of the idolatry of interest rates: Not only is their impact in real economic terms pretty damn limited as per my previous essay, it also appears that their impact is at best unpredictable (and at worst unfathomable) when it comes to equity valuations.

Alternative Approach

As noted above, there is an alternative to assuming that the cost of equity is tightly tied to the interest rate as per the ERP approach. It is, of course, to assume that the cost of equity is a set number (that is to say that your expectation of equilibrium return doesn't vary with the real rate of interest). This naturally raises the $64 million dollar question as to what that number should be.

We know from history that the "fundamentals" (yield plus growth) have given around 6% p.a. (at least in the U.S.), and because in equilibrium investors shouldn't expect valuation changes, it isn't obviously wrong to think that 6% is a reasonable rate of return over the long term.

Indeed, Graham and Dodd[15] reached much the same conclusion in 1934 when they wrote, "It is difficult to see how average earnings less than 6% upon the market price could ever be considered as vindicating that price…16 times average earnings is taken as the upper limit of prices for an investment purchase." It is probably worth noting that in later editions (including the 1942 edition) the authors suggested that 5%, and hence 20x, was the maximum valuation for an investment purchase.

However, it would be nice to have some theoretical understanding of what determines the equilibrium return on capital (ROC). In order to provide this, we need to return to some papers written in the mid 1950s and early 1960s and largely ignored by the mainstream in finance and economics (strangely enough). This understanding goes by the name of the Kaldor-Pasinetti Theorem, which, whilst it may sound like it belongs in an episode of the Big Bang Theory, may actually be a useful framework for understanding the long-run ROC. For the wonks, nerds, and geeks out there I've outlined a simple version of the theorem in the appendix.

---

[15] Benjamin Graham and David Dodd, Security Analysis, McGraw-Hill, 1934.

The theorem (in the form we use in our empirical work) states that the equilibrium ROC (P/K) is

$$\frac{P}{K} = \frac{g_{n\,(1-x)}}{sc}$$

Where gn is the rate of potential growth defined as the sum of labour force growth and labour productivity, x is the amount of investment financed by new equity issuance, and sc is the savings by firms.

We can also estimate the long-run ROC using this approach and compare it with the actual observed ROC. This is shown in Exhibit 15. Remember that the long-run ROC isn't calculated using the actual ROC, yet it does a pretty good job of showing the "gravitational attractor" that the realized ROC tends toward. It is also pretty smooth in the general scheme of things, thus operating with a constant ROC assumption seems to have worked pretty well.

Of course, there is nothing to say that going forward the driving forces between the long-run equilibrium ROC won't change, but at least this framework helps us to think about how the rate of ROC may evolve over time, and what factors are important in driving those changes.[16]

Exhibit 15:

Non-financial NIPA ROC
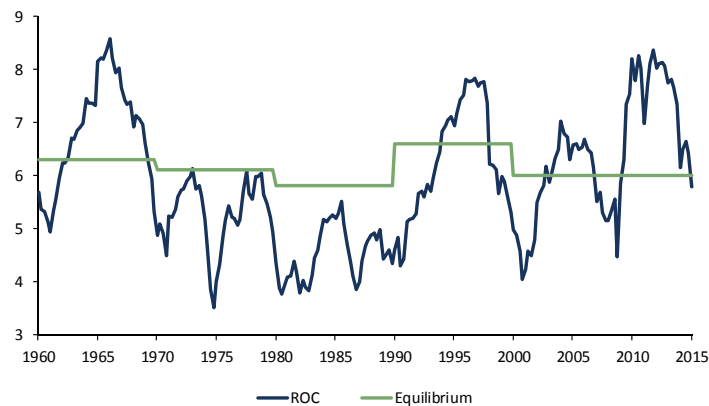
Source: GMO

Even if you don't believe a word I've written

If, like any number of my colleagues, you don't believe a word I've written (and quite possibly think that I either have escaped from or belong in an asylum), then let me leave you with one last chart. Even if you believe that real interest rates do matter for equity market valuations, and, hence, that lower real rates justify higher sustainable levels of P/E, you still need to ask yourself the following question: Would I need to believe that today's U.S. equity market valuations represent fair value?

Exhibit 16 tries to answer this question. It assumes that equity market valuations are a function of the real rates of interest, and asks a very simple question: How long would you need to see real interest rates stay at -2% p.a. in order to be able to call today's cyclically-adjusted P/E fair value?

---

[16] For instance, in the data on Japan, omitted from this note because it was already very long, the same Kaldor-Pasinetti framework shows a steadily declining long-run rate of ROC. The reasons for this are beyond the scope of this paper, but I may well return to this topic at some point in the future.

The answer is a staggering 60 years! According to data from Reinhart and Sbrancia (2011), the average length of a period of financial repression is 22 years plus or minus 12. That puts 60 years of negative rates at a 3-standard deviation event, and twice as long as any period of financial repression ever actually observed! Seems like a pretty extreme belief to have to hold to me.
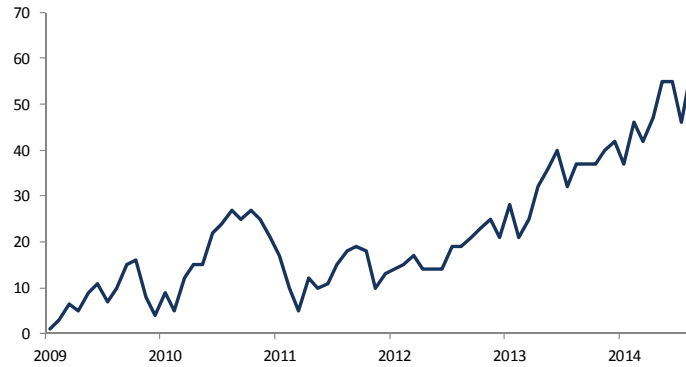
Source: GMO

James Montier is a member of GMO's Asset Allocation team. Prior to joining GMO in 2009, he was co-head of Global Strategy at Société Générale. Mr. Montier is the author of several books including "Behavioural Investing: A Practitioner's Guide to Applying Behavioural Finance; Value Investing: Tools and Techniques for Intelligent Investment"; and "The Little Book of Behavioural Investing." Mr. Montier is a visiting fellow at the University of Durham and a fellow of the Royal Society of Arts. He holds a B.A. in Economics from Portsmouth University and an M.Sc. in Economics from Warwick University.

Appendix: The Kaldor-Pasinetti Theorem

Consider a simple closed private economy, total net income (Y) can be split into two broad categories, Wages (W) and Profits (P)

$$Y \equiv W + P$$

And Savings can also be divided into two categories. (Here Kaldor and Pasinetti use slightly different terminology. Kaldor refers to savings out of profits and savings out of wages, whilst Pasinetti refers to capitalists' savings [Sc] and workers' savings [Sw]). We will follow Pasinetti for convenience here, but one should remember both possible interpretations as they will be relevant in the empirical section below.

$$S \equiv Sc + Sw$$

Let's assume simple proportional savings functions Sw = swW and Sc = scP (where sw and sc lie between 0 and 1 and represent the propensity to save of the workers and capitalists, respectively).

Remember that I = S

So we can write:

$$I = swW + scP = swY + (sc - sw)P$$

We can re-arrange this to give an expression for the profit share (P/Y)

$$\frac{P}{Y} = \frac{1}{sc - sw}\frac{I}{Y} - \frac{sw}{sc - sw}$$

We can derive an expression for the profit share (by dividing by capital, K)

$$\frac{P}{K} = \frac{1}{sc - sw}\frac{I}{K} - \frac{sw}{sc - sw}\frac{Y}{K}$$

If we assume that workers don't save, sw=0[17] and recalling from the Harrod growth model that in equilibrium I/K = gn (which is labour force growth + labour productivity, which Harrod labelled as the natural rate of growth), then the above reduces to:

$$\frac{P}{Y} = \frac{g_n}{sc}k$$

Where k is capital to output ratio

And

$$\frac{P}{K} = \frac{g_n}{sc}$$

When we use this framework, we actually use a version closer to Kaldor (1966), which accounts for net issuance by companies – a materially important element given the role of buybacks over the last few decades. But in essence the result is very similar.

---

[17] I'm sympathetic to this assumption as it captures the essence of Kalecki's viewpoint that "workers spend what they earn, and capitalists earn what they spend." However, you don't have to make this assumption. Pasinetti (1962) has shown the results hold as long as you assume that the interest rate and the profit rate are the same thing. I personally dislike this assumption. Moore (1974) has shown that the Kaldor-Pasinetti Theorem still holds even if you don't like either of these assumptions as long as you define equilibrium as having a constant distribution of wealth, and a constant return on capital, which seems like a pretty good definition of equilibrium. Furthermore, Fazi and Salvadori (1981) show that as long as the rate of interest is below the rate of profit, then the Kaldor conclusions are correct, even if workers save some of their income.

The Idolatry Companion: Potential Utility in an Equity Risk Premium Framework

Ben Inker

When I first approached James to join GMO about six years ago, I told him that we had been big fans of his work for many years. I went on to say that four out of five times we agreed with what he was saying, and the fifth time always led to an interesting conversation where we learned something useful. I'd say that parts of James' paper "The Idolatry of Interest Rates, Part II: Financial Heresy" is in the fifth out of five camp. In it, James is making a number of good points and taking them to a logical conclusion. It's not the same logical conclusion that many of the rest of us on the Asset Allocation team arrive at, but it is an important enough challenge to how both we at GMO and investors generally tend to think about the role of interest rates in equity valuations that it seemed a very worthwhile white paper to have him write. What follows is not so much a rebuttal as an explanation for why I find the construct of an equity risk premium (ERP) to be a useful one, despite some of the problems that James so ably points out.

Perhaps the central reason why a number of us consider that the ERP is a useful construct is that it allows us to explore a pretty central question: Under what circumstances should we expect the long-run pricing of equities to shift? James notes that the return to equities (in the U.S. at least) has averaged around 6% real and that the return on equity capital has likewise been fairly stable at 6%. But there is another way to read the data that suggests things have not been quite so stable.

The blue line in Exhibit 1 is the return consistent with the CAPE version of the earnings yield for the S&P 500, the red line is the 20-year moving average of that series, and the green line is a regression fit for the series. While there is no particular reason the world is best described with a linear regression, a falling required return to equities looks like a better fit to the historical data than a constant 6% real.

GMO
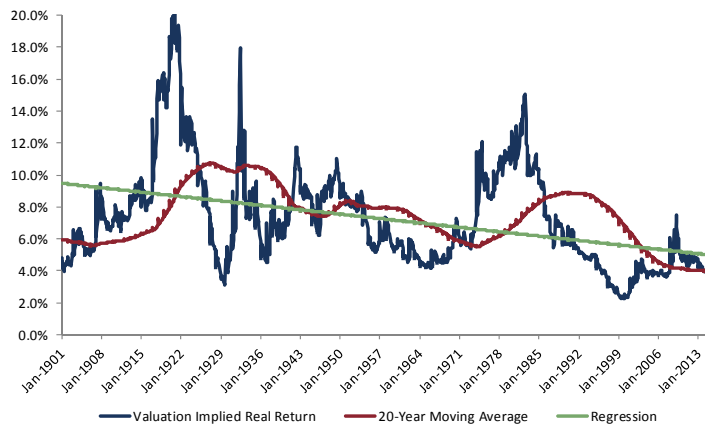
Data from 1901-2015
Source: Robert Shiller, GMO

We could continue the green line into the future and assume that the required return to equities will fall forever, but it seems far more helpful to try to understand why the cost of equity capital has been falling and whether it is plausible to imagine that it will continue or not. To that end, the ERP seems to me a helpful framework. If we separate the required return into a return to a low-risk asset – either cash or high quality bonds – and an ERP, we can then ask what plausible levels are for the ERP over a low-risk asset, and what plausible levels are for the expected return on the low-risk asset. Even if we wind up coming to the conclusion that we cannot forecast with any certainty the expected return to the low-risk asset in the long run, that at least injects some useful uncertainty about our required return to equities. To my mind, therefore, it is a useful framework even if it does not decrease our uncertainty about the future required return to equities.

And with regard to the inherent difficulty in predicting the future required return to low-risk assets, James may possibly be guilty of overstating the hopelessness of the task. Exhibit 2 shows the ex ante real yield on U.S. bonds going back to 1900.[1] I'm using bond yields instead of cash yields here because the cyclical nature of cash yields both creates a "phantom" mean reversion as cash yields rise and fall across the economic cycle, and because bonds can be thought of as related to the market's expectation for cash rates over the medium term, making them a natural discount rate to think of for equities.

---

[1] There is something of an art to calculating an ex ante real bond yield, as it involves estimating what investors believed inflation was likely to be at the time. Since the 1970s we have survey-based data on long-term inflation expectations, but before then we had to estimate an expected inflation based on previous inflation and priors that investors may have had. I can't swear that I have done it all perfectly, but this is unlikely to be profoundly wrong.

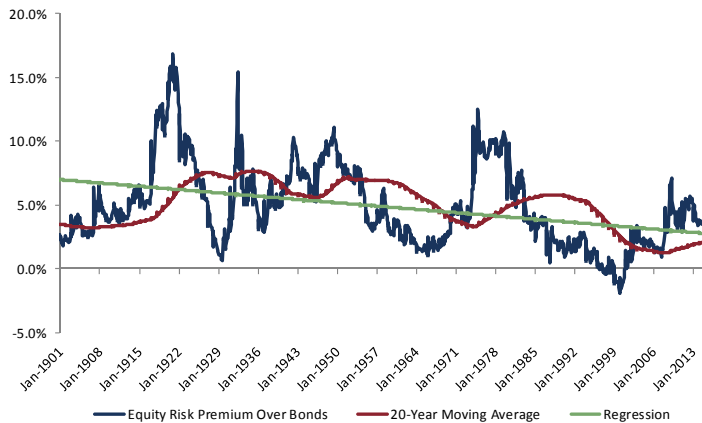## Exhibit 2:
Ex Ante Real Bond Yield

Data from 1901-2015
Source: Robert Shiller, Federal Reserve, Survey of Professional Forecasters, GMO

While the periods in which bond yields have been continually higher or lower than trend are somewhat longer than the comparable equity chart, there is less evidence of a significant trend here than there was for equities. The yearly "force" of mean reversion looks weaker than equities, but there is not obviously a lot of drift to the series. It is worth pointing out, however, that the recent period has been about as extreme an event as anything we have seen since 1900.

As a slight caricature, you can look at the two exhibits above and say that the falling trend to required equity returns (as embodied in Shiller P/E) has for most of the last 115 years been best thought of as a case of a declining ERP. That process must have some natural limitation, assuming the ERP must stay meaningfully positive. The most plausible case today for a continued fall in the required return to equities from here would therefore be driven by a fall in the required return to the low-risk asset. Such a permanent fall would be consistent with the current level of bond yields, but just because the bond market is predicting something is far from a guarantee that it will occur.

Such a fall in bond yields is the essence of the "Hell" scenario I have written about. And you can look at the difference between the embedded return in stock valuations versus bond valuations by subtracting the one from the other in Exhibit 3.

Exhibit 3:
"Equity Risk Premium" Over Real Bond Yields

Data from 1901-2015
Source: Robert Shiller, Federal Reserve, Survey of Professional Forecasters, GMO

This exhibit is a bit tricky to interpret. While it does look to be the case that the current "equity risk premium" is a bit higher than both the trend and the 20-year moving average, that is driven by the fact that current real bond yields are very far below trend, and well below even our estimated required yield in the "Hell" scenario. So while you could say that equities look mildly cheap versus bonds today, this only leads you to believe that equities are a buy if you are convinced that today's low bond yields are "correct." As James put it in his paper, in order for equities to be fairly valued today, you need to believe that real cash yields will be well below zero for the next 60 years. Exhibit 3 is effectively assuming that real cash yields will be well below zero forever. But the point of thinking in terms of an ERP is not to collapse everything we think about stocks and bonds into a single chart. Stocks and bonds are capable of being simultaneously overvalued or undervalued. When both are expensive, the very long duration of stocks makes them particularly dangerous, and when both are cheap, equities are particularly compelling. But just because stocks and bonds are not completely comparable as assets does not mean that comparing them cannot be enlightening.

And this brings up the second reason why I like thinking in terms of an ERP over a low-risk asset. It reminds us, as asset allocators, that all portfolios at all times are fully invested in something. It is no good to complain that equities viewed in isolation are overvalued. The question is how attractive equities are, or any other asset, relative to the other assets in your opportunity set. Today we are staking out a position that both stocks and bonds are overvalued, which pushes us to own a significant slug of cash. But we cannot ignore the very low yield available on cash today when determining that allocation. I worry that James' methodology can encourage us to ignore the low yield on cash, which I believe is dangerous.

But ignoring James' important challenge to the ERP framework would be more dangerous. The historical evidence for a stable ERP is no stronger, and probably less strong, than the evidence for a stable total required return to equities. Modifying your equity forecasts on the basis of cash yields or bond yields has given a worse fit than ignoring them. Thoughtful investors would be well-served by recognizing these facts in building their models and portfolios, even if they believe that the future may turn out to be different from the past on this front, as I tend to.

Ben Inker is co-head of GMO's Asset Allocation team and member of the GMO Board of Directors. In addition, he oversees the Developed Fixed Income team. He joined GMO in 1992 following the completion of his B.A. in Economics from Yale University. In his years at GMO, Mr. Inker has served as an analyst for the Quantitative Equity and Asset Allocation teams, as a portfolio manager of several equity and asset allocation portfolios, as co-head of International Quantitative Equities, and as CIO of Quantitative Developed Equities. He is a CFA charterholder.