Journal Of Investment Management, Vol. 21, No. 3, (2023), pp. 1–3

JOIM
www.joim.com

PRACTITIONER'S DIGEST

The "Practitioner's Digest" emphasizes the practical significance of manuscripts featured in the "Insights" and "Articles" sections of the journal. Readers who are interested in extracting the practical value of an article, or who are simply looking for a summary, may look to this section.

EFFICIENT GOAL PROBABILITIES: A NEW FRONTIER PAGE 4

Sanjiv Das, Daniel Ostrov, Anand Radhakrishnan and Deep Srivastav

When an investor has multiple goals spread out over time that they hope to fulfill, there are two key questions they want to answer: 1) What is the optimal way to move investments around over time? and 2) What is the optimal way to determine which goals to fully fulfill, partially fulfill, or not fulfill given financial contraints? These questions fundamentally depend on the relative importance of the goals (both full goals and partial goals) to each other. This means attaching a relative weight, that is a "utility weight", to each of these full and partial goals. This is an impossible task for most investors to do directly. After all, who would be able to say "I value giving \$100,000 towards my daughter's college costs sixteen years from now 1.78 times more than I value buying a \$30,000 car seven years from now?"! This paper, however, shows how this weighting can be accomplished instead by using the investor's desired probability for attaining each full or partial goal. This is intuitive for investors and has a simple geometric interpretation through what we call the "Efficient Goal Probability Frontier" (EGPF).

The partial answer to the first key question above (regarding optimally moving investments) is to stay on the Markowitz Efficient Frontier (EF), which is Pareto-optimal. That is, we know that any investment portfolio above the frontier is not possible, any investment portfolio below the frontier is sub-optimal, and so the first key question becomes how to best move the investment portfolio back and forth on the EF over time. Dynamic programming allows us to determine the EGPF, which, like the EF, is a Pareto-optimal frontier. The EGPF is an $(n-1)$-dimensional surface in n-dimensional space, where n is the number of goals. Each point on the frontier contains the complete optimal answer to both key questions above for a given set of utility weights. Just like the EF, anything above the EGPF is not feasible, anything below the EGPF is sub-optimal, and so the question becomes which point on the EGPF should be selected? Once chosen, this generates the needed utility weights.

When an investor states their desired probabilities for attaining each full or partial goal, it corresponds to a point, called the "desired point," in the n-dimensional probability space. If there are insufficient funds to attain the probabilities comprising this point, this paper shows how to generate a "proximity point" on the EGPF that looks to get as close as possible to the desired point, given funding constraints. That is, we are able to generate for the investor the optimal investment strategy and optimal goals-taking strategy that looks to fit their desired probabilities as best as possible given their financial constraints. Alternatively, we also show how to determine the minimal additional initial investment or the minimal additional constant annual infusions over time in either nominal or real dollars that enable the investor to attain all their desired goal probabilities. Further, we explore the sensitivity of these results to increases to the expected returns or to the covariance matrix of the underlying investments. Finally, we show that constraining these investments so they contain no short positions can change the underlying weights considerably, but the resulting change to the EF is small, so the end results are largely unaffected.

THE DETERMINANTS OF INFLATION PAGE 29

William Kinlaw, Mark Kritzman, Michael Metcalfe and David Turkington

In this paper, the authors use a statistical model called a Hidden Markov Model to identify four distinct inflation regimes in the United States from 1960 to present. The model identifies four states: disinflation (where the inflation rate is declining), stable inflation (where the inflation rate is relatively flat), rising steady inflation (where the rate is rising smoothly) and rising volatile inflation (where the rate is accelerating in a very volatile manner). Disinflation periods tend to be associated with economic recessions and rising volatile periods are associated with historical inflation episodes such as the 1970s and the recent inflation surge of 2022.

A simple way to define inflation regimes would be to simply classify historical periods based on whether the level of inflation exceeded a particular threshold. The HMM approach offers two advantages over this approach. First, it accounts for the fact that inflation regimes are persistent. For example, if we are in a rising volatile regime, there is a high likelihood that we will remain in that regime during the next period. Second, it accounts for the fact that regimes differ not only in terms of the mean but also in terms of the standard deviation of inflation. Together, these advantages mean that we capture more persistent and informative inflation regimes.

There are several practical observations from the results. First, we observe that major asset classes have starkly different performance depending on the inflation regime. In nominal terms, stocks perform best during "steady" and "rising stable" regimes whereas bonds perform best during "disinflation" regimes when real interest rates and bond yields are typically declining. Cash performs best during rising volatile regimes when short-term interest rates are typically rising. In real terms, the "rising volatile" regime is by far the worst for both stocks and bonds.

Second, the analysis looks across a range of economic variables to identify the most important determinants of inflation. As of February 2022, the most important determinant of the recent spike in inflation was spending by the federal government. This analysis offers insights to both investors and policy makers regarding both the state of inflation and its key drivers.

ASSET ALLOCATION WITH NON-PECUNIARY ESG PREFERENCES: EFFICIENTLY BLENDING VALUE WITH VALUES PAGE 42

Douglas M. Grim, Giulio Renzi-Ricci and Anna Madamba

While the average investor builds a multi-asset portfolio based only on financial expectations and risk tolerance, it is well-documented that some investors derive non-financial, positive social and/or private benefits from the ESG-related features of assets. Incorporating this nonfinancial preference can rationally influence their portfolio construction process. Unfortunately, the most common method used to determine the portfolio allocation for this type of investor only assesses one dimension of risk and arbitrarily applies weight constraints for assets with a certain ESG profile. Doing so subordinates financial goals and can be difficult for an advisor to justify.

In this paper, we propose a practical and comprehensive framework that efficiently integrates an investor-specific, accessible, and quantifiable measure of ESG-driven non-pecuniary utility into a portfolio optimization process that simultaneously considers different dimensions of investor risk preferences. Using sensitivity analysis, we demonstrate that the optimal portfolio choice for an investor is a function of risk preferences, the financial expectations and ESG attributes of funds considered, and the degree of non-pecuniary desire that an investor has for such attributes. This demonstrates the practical importance of being able to estimate a variety of financial and non-financial inputs and provides for more transparent trade-off discussions between the investor and their advisor.

The framework can accommodate any type of ESG strategy with concessionary or non-concessionary asset return expectations. Overall, it is unique in its practical ability to enable explicit, customized, and intuitive investor choices for ESG taste and systematic, factor, and active risk aversion. This ensures that an appropriate portfolio can be found for investors with different pecuniary and non-pecuniary preferences and expectations.

INVESTMENT MANAGEMENT LESSONS LEARNED FROM THE MANAGEMENT AND MISMANAGEMENT OF IMPENDING BANK RUNS PAGE 61

Seoyoung Kim

Long-horizon total returns are reasonably predictable using simple models of expected return. Under certain assumptions, the classic Gordon model produces equity expected returns equal to the earnings yield plus expected inflation. The random-walk model for bond yields implies that a bond's expected return is equal to its yield. Although these basic formulations of expected return by no means perfectly predict future returns, rigorous statistical analyses indicate that these models contain real information and therefore should prove useful to asset allocators as a starting point when constructing optimal portfolios.