CAMBRIDGE ASSOCIATES LLC

MEAN-VARIANCE ASSUMPTIONS: AN INTRODUCTION

Tom Danielson
Beth Newhouse

Mean-Variance Assumptions: An Introduction

The global capital markets are by nature unpredictable and volatile and making assumptions about the future of the capital markets is a challenging and dynamic endeavor. Capital market assumptions can be developed using many different methodologies drawn from countless schools of thought, each with a different goal or application in mind.

In our case, we develop risk, return, and correlation estimates across asset classes for use in long-term strategic asset allocation planning. The first section of this paper defines our objectives; the second outlines our philosophy, including the key principles that shape our framework; and the third discusses some of the key conceptual dilemmas we face in constructing these assumptions. Appendix A contains a description of various approaches to the development of mean-variance assumptions, parts of which are incorporated into our process. Appendix B discusses taxable considerations for U.S. investors.

Objectives

A given set of assumptions will reflect the specific purpose for which it has been developed. When comparing alternate assumption sets, one should therefore recognize and understand how differences among them reflect their differing objectives.

The assumptions we develop are designed as the inputs to mean-variance analysis used as a tool in long-term strategic investment planning, leading to decisions on such key issues as long-term asset allocation and spending policy. These long-term decisions require analysis that is no less relevant in two years than it is today. To accommodate these requirements, we develop equilibrium assumptions which are independent of beginning and end points.[1]

Philosophy and Principles

In order to be consistent and coherent in our methodology, we start with core principles that serve as the foundation of the assumption framework, providing points of reference and guidance as we assess different options throughout the process.

- **Return and Volatility Are Simplistic.** The mean-variance framework is a crude and limited landscape in which to evaluate investment decisions. Investment decisions should not be made solely on the basis of the average expected returns over time or the likely variability of those returns—many other variables and considerations should be taken into account. In recognition of these inherent flaws, we believe it important to resist the temptation to require an illusory perfection in our input assumptions.

- **No Free Lunch.** In the mean-variance assumption framework, risk is represented by variability of returns. We assume that the greater the risk of an asset class, the higher the returns must be to compensate investors for incurring more risk. While not insisting on a perfect correspondence between relative risk and return across asset classes, we do regard this relationship as a cornerstone of our construct and are skeptical of equilibrium risk/return assumptions that deviate from it to a significant degree.

[1] Note that equilibrium assumptions are long term in nature and as such are not appropriate guides for shorter-term decisions (e.g., tactical asset allocation).

- **Complex and Imperfect.** There is no such thing as a "correct" assumption—all assumptions and all assumption-making methodologies are subject to valid criticisms. Because the process and the results are necessarily flawed and inherently imprecise, any assumption framework reflects many compromises among competing priorities. Consequently, we represent our mean-variance assumption framework as one of many useful but inherently flawed lenses through which investors can view and assess some of the trade-offs they face.

- **Adapting to the Imprecision.** Since the assumption-making process is inherently imprecise, we incorporate this reality into our framework in two principal ways. First, in recognition of the fact that actual returns can diverge significantly from mid-point estimates, even for periods as long as 25 years, we complement our estimates of the long-term mean return of an asset class with an indication of the range within which we estimate returns will fall in half of all 25-year periods. Second, we round all standard deviations to the nearest tenth of a percentage point and returns to the nearest half percentage point. Rounding in this way allows for meaningful distinctions among asset classes without implying too much precision.

- **Internal Consistency.** We place a high priority on internal consistency, meaning that individual assumptions that make up the framework must make sense relative to each other. Methodologies and rules should be reasonably consistent with those applied elsewhere in the framework. The desired result is a consistent, reasonable, and comprehensible assumption set.

- **Simplicity.** A process that has countless dimensions but limited precision presents a challenge in the form of a delicate balance between complexity and simplicity. In our framework we try to avoid complexity that offers marginal improvement, erring in favor of a comprehensible framework to one with confusing layers of analysis that may provide only the illusion of analytical rigor.

- **Role of Data.** The judicious use of data also presents a significant challenge. Every data series has its own strengths and weaknesses. In some cases there is not enough data. In others, the volume of data masks an underlying deficiency in what it represents.[2] The *quality* of a data set also merits attention.[3] The relevance of historical data can also vary, since asset classes constantly evolve, sometimes dramatically.[4] Finally, in the case of illiquid investments, the expression of performance with time-weighted returns is awkward at best, used here only to compare them with more liquid asset classes.

Process and Framework

The mean-variance assumption framework is composed of an array of rules and definitions. This section covers some of the key judgments involved, many of which are predicated on the founding principles from the previous section.

- **Definition of Asset Class.** The definition of "asset class" can be complex, subject to as much scrutiny and dispute as the estimates themselves. Consistent with our preference for simplicity, especially where greater complexity does little to enhance the results, we include only what we would characterize

[2] For example, 30 annual returns are often more valuable than 120 monthly returns, since ten years is typically not long enough to reflect multiple economic environments.
[3] For example, much of the capital market data during World War II is of suspect quality, not suitable for use in deriving equilibrium assumptions.
[4] For example, as the inflation-linked bond market broadens (in number of issues and number of issuer countries), becoming more attractive to investors, the characteristics of the asset class evolve.

as "primary" asset classes, which we would define as asset classes with fundamentally different and distinctive sources of return. Thus, for example, equities and bonds are clearly different asset classes, since they easily meet this definition; however, large- and small-cap equities are not, since they share the same fundamental economic basis of return.

- **Active versus Passive.** Our assumption framework represents passive investment in an asset class. Where no passive alternative is available, the assumptions are a broad representation of the asset class. The framework does not attempt to incorporate the unique style, goals, or capabilities of individual managers or institutions.

- **Risk/Return Line Definition.** The risk/return line is the foundation on which risk/return trade-offs are developed. The line is defined on the risky end by equities and on the risk-free end by cash, with bonds falling in between.

- **Global Perspective.** Our perspective is that of a global investor, subject to global inflation and invested globally in multiple asset classes and currencies. While global assumptions in local currencies do not reflect performance actually experienced by any investor, they do serve as the universal framework for risk/return relationships. The assumptions are then translated into the investor's base currency.

  Historically, marketable and non-marketable alternative assets (i.e., hedge funds and private equity) have been predominantly U.S. oriented. However, the opportunity set in these asset classes is increasingly global and there are many reasons to believe this will continue. Consequently, we treat these as global asset classes, just as we do traditional equities and bonds.

  Therefore, all of our asset classes are global with the exception of some regional and locally defined asset classes which depend specifically on the region or country. For example, local cash, fixed income, and equities are defined for an investor's home market in that market's base currency.

- **Currency Translation.** Our assumptions incorporate the impact of currency risk that is incurred when investing in foreign assets. Given the potential for enormous complexity in modeling currency, the goal of a methodology that is simple and defensible is particularly salient. We assess the volatility of the investor's currency relative to foreign currencies. The currency volatility is incorporated with the volatility of each asset class to reflect the level of volatility specific to the investor's base currency. No adjustment is made to the return because the long-term expected return for currency is zero. Hedged assumptions are derived similarly to unhedged, in global risk/return space, without the addition of currency.

- **Review Frequency.** Our equilibrium assumptions are subject to intensive review every three to five years. While our goal of creating an assumption set that can contribute to long-term strategic decision making does not change, many of the components are reviewed.[5] The strengths and weaknesses of the current approach are considered, including any aspects of the current framework that have attracted significant internal and external discussion. We also assess the merits of ideas found in the recent publications of others researching the area, including newly published data. ■

[5] For example, a new asset class may be added as it becomes increasingly attractive to institutional investors.

Appendix A: Developing Mean-Variance Assumptions: Various Approaches

The main text describes how we approach the challenge of developing input assumptions for use in mean-variance analysis. This appendix discusses various approaches, some of which are incorporated into our process to varying degrees.

Historical Performance

The simplest of all approaches is to use historical data, as is, on the naïve assumption that since we cannot predict the future, we might as well assume it will match the past. Although this has the merit of simplicity, it ignores secular changes in the structure and composition of various asset classes over time, which reflect the evolution of the underlying economic drivers of return. We would argue that raw historical data should never be used as is, without first being subject to an informed evaluation of its continued relevance.

Ground-Up Assumptions

A more complex approach involves building total return expectations from the ground up. For example, the estimation of equity returns would include estimates for several components of total return: real earnings growth, dividend payout/ share buybacks, and multiple expansion/ contraction. This has the considerable virtue of requiring the investor to account for the specific sources of expected return, but it adds additional layers of assumption estimation, which, however rigorously pursued, may not in fact enhance the final results. It is also problematic when applied to some of the more esoteric asset classes, where the composition of total return is more ambiguous. In short, we find this approach appropriate for estimating the likely range of short- to intermediate-term (e.g., five to ten years) returns for asset classes like equities, for which the fundamental bases of return are reasonably well defined, but not appropriate for the development of a framework of equilibrium assumptions across all asset classes.

Risk Premium

A third approach is to base all return assumptions on estimates of the risk premium of each asset class relative to a base asset class. Typically, this involves using Treasury bonds or bills as the base, since the nominal return for a given period is embedded in the current yield-to-maturity of bonds or bills maturing at the end of that period, and then estimating the risk premium of each asset class relative to this base. The core principle here is very similar to the one included in our own approach, which is the assumption that investors must be paid a return premium to induce them to incur greater risk than can be earned from a "riskless" investment.[1] However, we find three major defects in this approach. First, ex-post risk premia have proved unreliable as a means of estimating ex-ante risk premia. Second, no other method of estimating ex-ante risk premia has proved robust over time. Finally, the whole construct is therefore built on unreliable estimates

[1] Of course, neither Treasury bills nor Treasury bonds are without risk. The risk of the former is re-investment risk and of the latter, inflation risk. Since the prospective real rate of return of Treasury bonds cannot be derived from the current yield, the risk-premium approach can only be used to estimate prospective *nominal* returns—which are far less relevant than prospective real returns. We have as yet insufficient history to determine whether the real return of 25- and 30-year Treasury Inflation-Protected Securities (TIPS) will correspond with the yield prevailing at the time of purchase; if it does, then TIPS yields would become the logical base on which to construct a risk-premium approach.

of the risk premia relative to a base return, which itself varies considerably over even quite short periods of time.

Nevertheless, because we agree with the underlying principle of the risk-premium approach, one of the reality tests we apply to our own methodology is to see whether it results in assumptions that incorporate risk premia that seem unrealistic.

Incorporating Valuations/Forecasts

Our approach to the development of capital market assumptions is vulnerable to the criticism that it ignores current valuations. This means, in effect, that our assumption for equity market returns, for example, will be the same whether the current market multiple is 30 or 10, which is tantamount to assuming that a traveler's distance to, say, London will be the same regardless of whether the starting point is New York or Tokyo, which is obviously nonsense. Our answer to this is that our median assumption for an asset class return is only one point in an estimated range, and that, of course, our expectation of where returns will actually fall within that range is highly dependent on both beginning- and end-point valuations for that asset class for any given period. Therefore, we regard using current valuations as more relevant to tactical than to strategic asset allocation, although they are also reflected in strategic decisions to the extent that investors incorporate (as they should) some estimate of the likely distribution of returns in their assumptions.

Specific Implementation

The level of detail and scope of an assumption set can vary widely. One common difference is the degree to which specific investment styles and small distinctions in implementation are incorporated. For example, some investors may want to reflect their strategy weights within an asset class or incorporate alpha for their managers. Although it can be constructive to drill into the assumptions with the goal of making them more relevant, it is difficult to determine just how many of these efforts actually improve rather than degrade the resulting output. The challenge involves balancing slight improvements to an inherently imperfect framework against the greater complexity one has introduced without any great assurance that the results will prove any more robust. ■

Appendix B: Taxable Considerations for U.S. Investors

In developing assumptions for use in our models, we first assume a tax-exempt investor and then we adapt those assumptions to make them more relevant to U.S. taxable investors. To ensure consistency, both the standard and taxable assumption sets are designed to produce what we have characterized as equilibrium estimates.

Every investor's tax situation is unique. Tax laws are anything but static, and investors implement their allocations in a variety of ways. Accordingly, we do not make our taxable assumptions with the goal of precisely representing an investor's current and prospective tax concerns. Rather, we intend to capture the relative tax efficiency of asset classes. Consistent with this approach, our assumptions about the tax characteristics of each asset class are relatively generic.

Tax efficiency is an unstable characteristic for most asset classes, such that different beginning and end points can have significant effects on tax efficiency in both absolute and relative terms. We aim to represent the normative case or the expected tax efficiency of an asset class, noting that tax efficiency, like performance, is quite volatile.

The taxable framework consists of several types of estimates. The first is tax rates, where the highest marginal U.S. federal tax rates are considered. The second involves breaking total return into equilibrium income and capital gains estimates. The third layer is the taxable turnover of an asset class, including how much of the turnover is short term and how much is long term. In estimating turnover costs and the percentage of returns attributable to short-term capital gains, we assume some tax-aware management of traditional equity assets.

While pre-tax returns are rounded to the nearest 1%, after-tax returns are rounded to the nearest 0.5% to better capture the effect of taxes. This does not imply an additional level of precision but is instead needed to portray the cost of taxes that is otherwise blurred with more liberal rounding. ■