*Chapter 2*

# *Evaluating the Historical Record*

Primitive peoples, with no knowledge of modern science, express confidence in the proposition that the sun will rise tomorrow. The reason is that the historical record is unambiguous on this point. Ask whether it will rain tomorrow, though, and doubt arises. Because of random variation in weather, the historical record is a good deal more ambiguous. Rain today does not necessarily mean rain tomorrow.

With respect to the equity premium, the confidence that can be placed in the assumption that the future will be like the past depends on two related characteristics of the historical data: how accurately the historical premium can be measured and the extent to which the measured premium depends on the choice of the sample period. Before those questions can be addressed, however, there is the issue of how the average returns that go into the premium should be computed in the first place.

## Computing the Average Premium: Arithmetic versus Geometric

The historical equity risk premium equals the difference between the average return on equities and the average return on treasury



securities calculated over a specified time period. It can be seen in Table 1.2, for instance, that over the full sample period between 1926 and 1997, the average return on stocks was 13.0% and the average return on treasury bills was 3.8%, so the equity risk premium over bills was 9.2%. Those are arithmetic averages. They are computed in the standard way: Add up all the annual returns and divide by the numbers of years (in this case, 72).

Although it is familiar, the arithmetic average has a peculiar property. As an illustration, suppose that an investor earns returns of 10%, 20%, –25%, and 15% in 4 consecutive years. The arithmetic average of the four returns is 5%. Now consider an investor who starts with \$100. If he or she earns 10%, 20%, –25%, and 15% in each of 4 years, his or her ending wealth will be \$113.85. However, if that investor earns 5% per year for 4 years, he or she will end up with \$121.55. This is a general problem. Investors who earn the arithmetic average of a series of returns wind up with more money than investors who earn the series of returns that are being averaged.

The geometric average solves this problem. By definition, the geometric average is the constant return an investor must earn every year to arrive at the same final value that would be produced by a series of variable returns. The geometric average is calculated using the formula

$$\text{Geometric Average} = (\text{Final Value}/\text{Initial Value})^{1/n} - 1$$

where $n$ is the number of periods in the average. When the formula is applied to the preceding example, the results are as follows:

$$\text{Geometric Average} = (113.85/100)^{1/4} - 1 = 3.29\%$$

An investor who earns 3.29% for 4 years will end up with \$113.85.

There are four properties of arithmetic and geometric averages that are worth noting:

- The geometric average is always less than or equal to the arithmetic average. For instance, in Table 1.2 the arithmetic average stock return is 13.0%, but the geometric average is only 11.0%. (The geometric averages are reported at the bottom of the path of wealth columns in Table 1.2.)
- The more variable the series of returns, the greater the difference between the arithmetic and geometric average. For example, the returns for common stock are highly variable. As a result, the arithmetic average exceeds the geometric average by 200 basis points. For treasury bonds, whose returns are less variable, the difference between the two averages is only 40 basis points.
- For a given sample period, the geometric average is independent of the length of the observation interval.[1] The arithmetic average, however, tends to rise as the observation interval is shortened. For instance, the arithmetic average of monthly returns for the S&P 500 (calculated on an annualized basis by compounding the monthly arithmetic average) over the period between 1926 and 1997 is 13.1%, compared with the 13.0% average of annual returns.
- The difference between the geometric averages for two series does not equal the geometric average of the difference. Consider, for instance, stock returns and inflation. Table 1.2 reveals that the geometric average stock return is 11.0% and the average inflation rate is 3.1%, for a difference of 7.9%. However, Table 1.3 shows that the geometric average real return on common stock was 7.7%. This discrepancy does not arise for arithmetic averages, where the mean difference always equals the difference of the means.

With respect to the equity risk premium, the manner in which the average is calculated makes a significant difference. When compared with treasury bills over the full 1926-to-1997 period, the arithmetic average risk premium is 9.2%, whereas the geometric average premium is only 7.2%. Which average is the more appropriate choice? That depends on the question being asked. Assuming that the returns being averaged are largely independent and that the future is like the past, the best estimate of expected returns over a given future holding period is the arithmetic average of past returns over the same holding period. For instance, if the goal is to estimate future stock-market returns on a year-by-year basis, the appropriate average is the annual arithmetic risk premium. On the other hand, if the goal is to estimate what the average equity risk premium will be over the next 50 years, the geometric average is a better choice. Because the ultimate goal in this book is to arrive at reasonable forward-looking estimates of the equity risk premium, both arithmetic and geometric averages are employed where they are useful.

It is worth reiterating that projection of any past average is based on the implicit assumption that the future will be like the past. If the assumption is not reasonable, both the arithmetic and geometric averages will tend to be misleading.

## How Accurately Can the Historical Risk Premium Be Measured?

The accuracy with which the historical risk premium can be measured depends on the variability of the observations from which the average is calculated. In an assessment of the impact of that variability, the best place to start is with an expanded version of Table 1.2 that includes monthly returns for the four asset classes over the period between 1926 and 1997. Given this expanded data set, one way to assess the variability of the ex-post risk premium, defined as the difference between the observed returns for stocks and the related treasury securities, is to plot one histogram for stocks versus bonds and another for stocks versus bills. Each bar on the histogram represents the fraction of the 864 monthly

[1] This follows immediately from the fact that the geometric average depends only on the initial and final values of the investment.