Exhibit 4.4: Long-Term Corporate Bonds, Real and Nominal Return Terminal Index Value 1926–2020 (Year-end 1925 = $1.00)

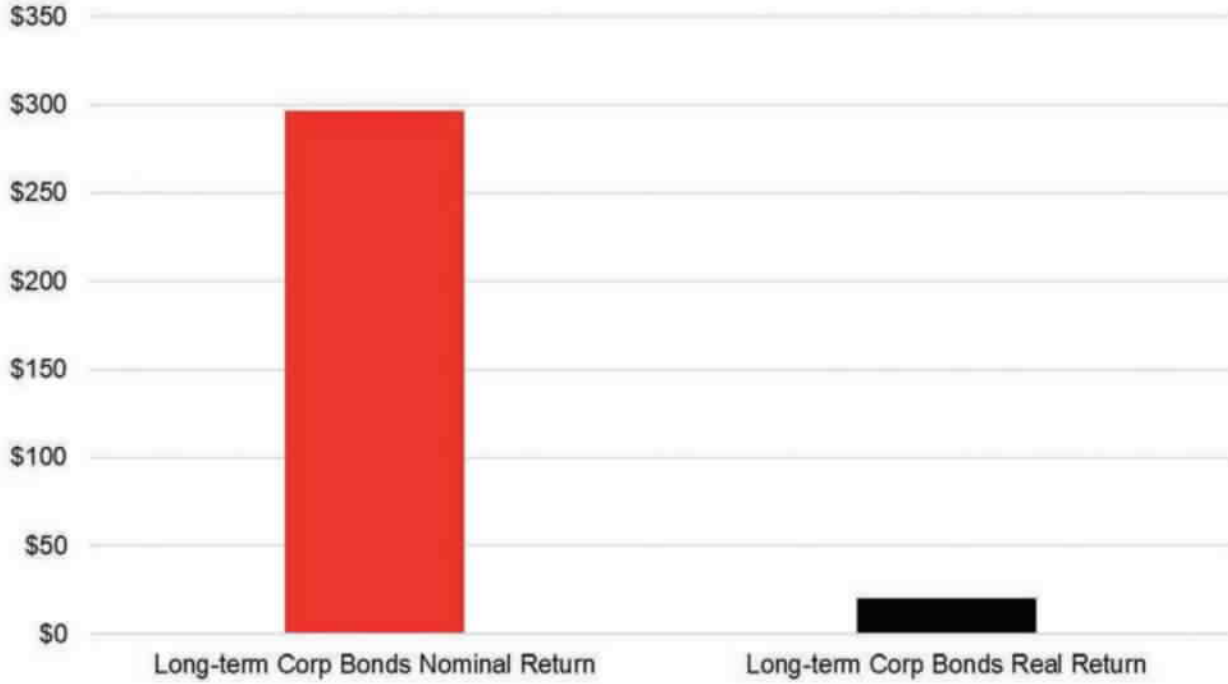

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset class represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Long-term (i.e., 20-year) Corporate Bonds: IA SBBI® US LT Corp TR USD, and (ii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission

Long-term Government Bond Real Returns

Construction

The inflation-adjusted return is a geometric difference and is approximately equal to the arithmetic difference between the long-term government bond total return and the inflation rate. The monthly inflation-adjusted long-term government bond total return is given by:

$$\frac{(1+LT\,Govt\,Bond\,TR)}{(1+Inflation)}-1$$

Because government bond returns are composed of inflation, the real riskless rate, and the horizon premium, the inflation-adjusted government bond returns may also be expressed as:

$$[(1+Real\,Riskless\,Rate)\times(1+Horizon\,Premium)]-1$$

Exhibit 4.5 depicts (i) what $1.00 invested at the end of December 1925 in long-term government bonds would have grown to by the end of 2020 and (ii) what $1.00 invested at the end of December 1925 in long-term government bonds would have grown to by the end of 2020 if long-term government bond returns were adjusted for inflation.

Exhibit 4.5: Long-term Government Bonds, Real and Nominal Return Terminal Index Value 1926–2020 (Year-end 1925 = $1.00)

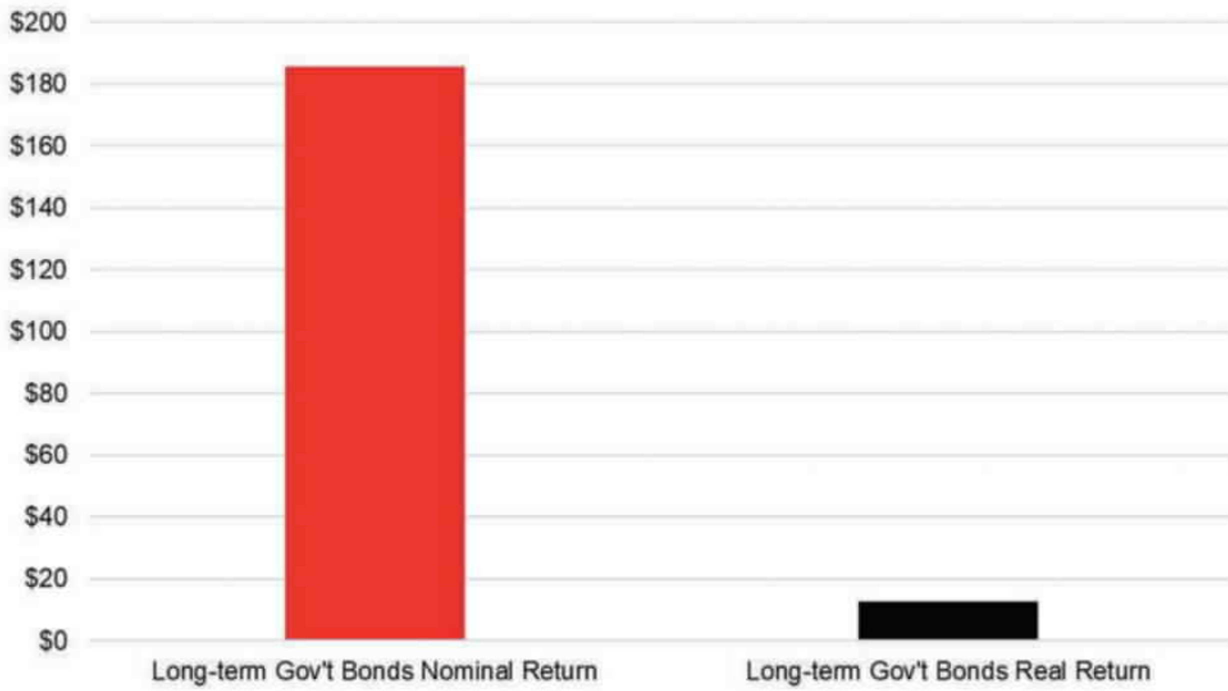

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset class represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, and (ii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission

Intermediate-term Government Bond Real Returns

Construction

The inflation adjusted return is a geometric difference and is approximately equal to the arithmetic difference between the intermediate-term government bond total return and the inflation rate. The monthly inflation-adjusted intermediate-term government bond return is given by:

$$\frac{(1+IT\ Govt\ Bond\ TR)}{(1+Inflation)}-1$$

Exhibit 4.6 depicts (i) what \$1.00 invested at the end of December 1925 in intermediate-term government bonds would have grown to by the end of 2020, and (ii) what \$1.00 invested at the end of December 1925 in intermediate-term government bonds would have grown to by the end of 2020 if intermediate-term government bond returns were adjusted for inflation.

Exhibit 4.6: Intermediate-term Government Bonds, Real and Nominal Return Terminal Index Value 1926–2020 (Year-end 1925 = \$1.00)

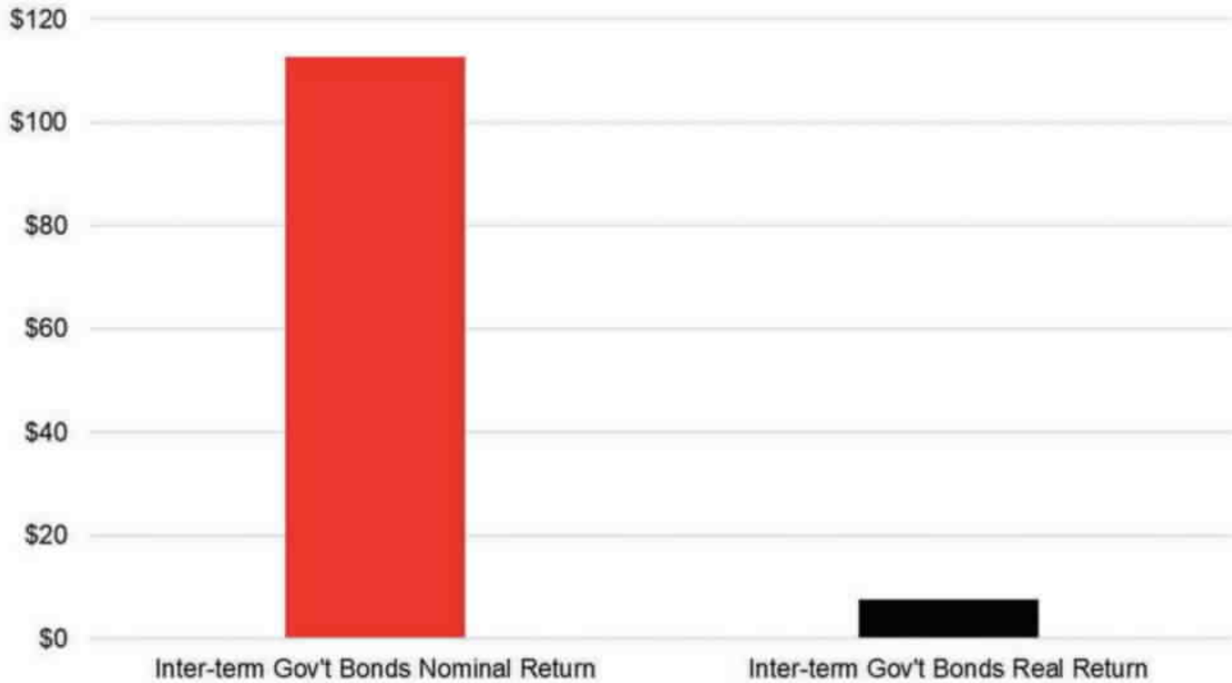

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset class represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Intermediate-term (i.e., 5-year) Government Bonds: IA SBBI® US IT Govt TR USD, and (ii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission

Real Riskless Rates of Return (U.S. T-Bill Real Returns)

Treasury bills returned 3.3% compounded annually over the 1926–2020 period, in nominal terms, but only a 0.4% compound annual return in real (inflation-adjusted) terms. Thus, an investor in Treasury bills would have barely beaten inflation (or retained purchasing power) over the 95-year period.

Construction

The real riskless rate of return is the difference in returns between U.S. Treasury bills and inflation. This is given by:

$$\frac{(1+\textit{Treasury Bill TR})}{(1+\textit{Inflation})}-1$$

Exhibit 4.7 depicts (i) what $1.00 invested at the end of December 1925 in U.S. Treasury bills would have grown to by the end of 2020 and (ii) what $1.00 invested at the end of December 1925 in U.S. Treasury bills would have grown to by the end of 2020 if U.S. Treasury bill returns were adjusted for inflation.

Exhibit 4.7: U.S. Treasury Bills, Real and Nominal Return Terminal Index Value 1926–2020 (Year-end 1925 = $1.00)

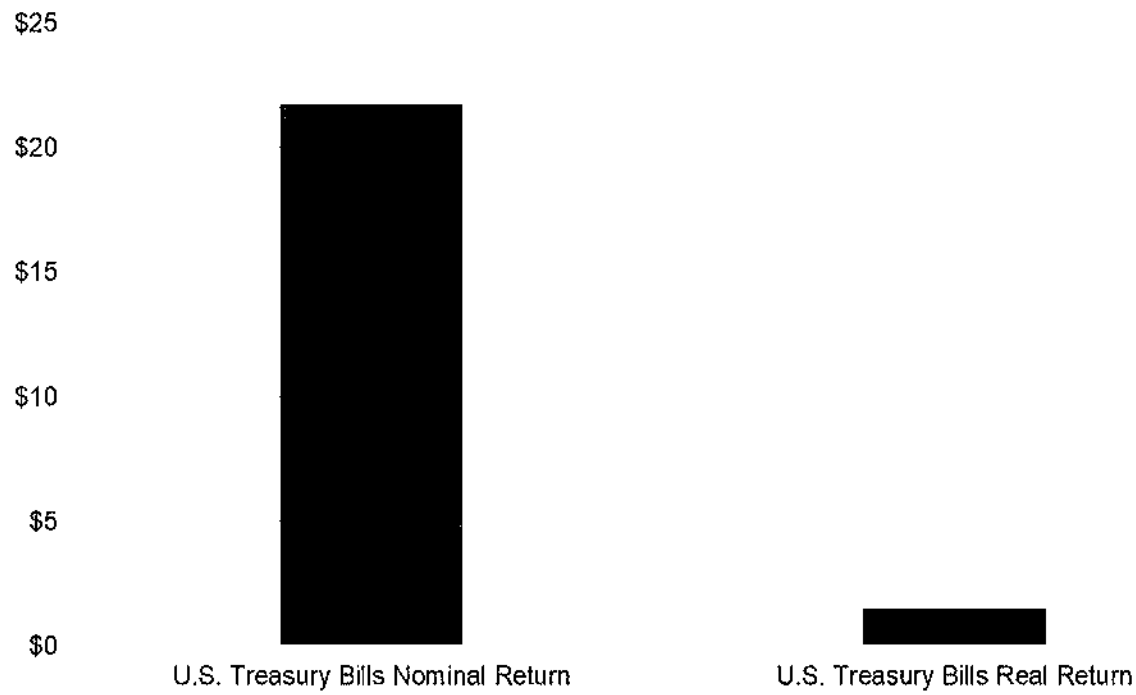

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset class represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD, and (ii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission

Exhibit 4.8 shows the levels, volatility, and patterns of real interest rates (inflation-adjusted U.S. Treasury Bill return) over the last 95 years.

Exhibit 4.8: Annual Real Interest Rate Returns (%) 1926–2020

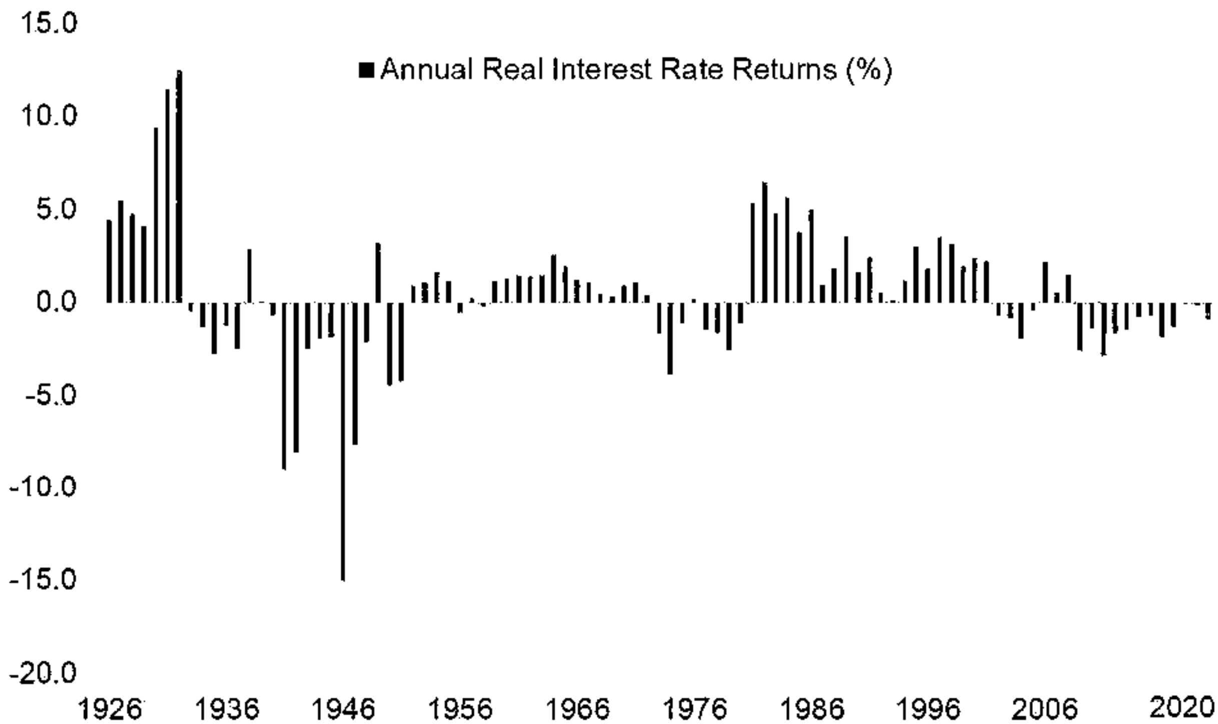

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset class represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD, and (vii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission

Chapter 5
Annual Returns and Indexes

Returns and benchmark indexes are used to measure the rewards investors earn for holding an asset class. Indexes represent levels of wealth or prices while returns represent changes in levels of wealth. Total returns for specific asset classes can be divided into two primary components: income and capital appreciation. The income return measures the cash stream earned by holding the security, such as coupon interest for bonds or dividend payments for stocks. In contrast, the capital appreciation return results from a change in the price of the security. The method for computing a return varies with the nature of the payment (income or capital appreciation) and the period of measure (monthly or annual). Indexes are computed by establishing a base period and base value and increasing that value by the returns in successive periods. Indexes are used to illustrate the cumulative growth of wealth from holding an asset class. This chapter describes the computation of the annual returns and indexes.

The first generation of stock indexes was created to assess the market's general direction. One of the oldest and most recognizable market indexes is the Dow Jones Industrial Average, or DJIA, first published on May 26, 1896. When Charles Dow initially calculated the DJIA, which originally consisted of only 12 stocks, the process was simple: Add up the share prices of the stocks in the index and then divide by the number of stocks in the index.[137,138] In this type of index, known as a price weighted index, higher-priced stocks have a greater influence than lower-priced stocks.

However, most modern indexes are market-capitalization weighted, meaning companies with greater overall market capitalization (share price times number of shares outstanding) have a larger influence than companies with lesser market capitalization. Market-cap weighting has a strong theoretical motivation because the capital asset pricing model, or CAPM, implies in its simplest form that every investor should hold every security in proportion to its market capitalization. In contrast, price weighting lacks any theoretical motivation, so it is rarely used outside of the Dow Jones Averages (S&P Dow Jones Indices uses market-cap weighting for most of its other indexes).

Market-cap weighting is widely considered to be the central organizing principle of good index construction. Its practical advantage is that the weights adjust automatically as share prices fluctuate, eliminating the need for the frequent and expensive rebalancing that can occur with other weighting schemes. So-called strategic beta or smart beta, or other alternative weighting

137 None of the original 12 companies listed in the DJIA remain as a component of the average. The total number of companies listed in the DJIA has not changed since 1928, when the number of companies in the index was increased to 30. For more information on the historical makeup of the DJIA, please visit the S&P Dow Jones Indices website at https://us.spindices.com/indices/equity/dow-jones-industrial-average.

138 In June 2018, General Electric (GE) was the last company (out of the original 12) to be removed from the DIJA. See: https://www.businessinsider.com/dow-original-companies-2016-12.

schemes, seek to outperform market-cap weighted indexes but deviate from CAPM and are not macro consistent.

Market-cap weighting is usually implemented with a “float” adjustment that subtracts the number of closely-held and illiquid shares from the number of shares outstanding. A float-adjusted market-cap weighted portfolio is macro-consistent, meaning that if all investors held such a portfolio, all available shares of its constituent stocks would be held, with none left over. Accepting this is an extension of Roll’s critique, which states in part that no portfolio can capture shares of all assets (e.g., jewelry, fine wine, automobile collections, etc.) and instruments with market value, so indexes can only approximate a true diversified CAPM portfolio.[139]

With all other weighting schemes, it is mathematically impossible for all investors to hold the index portfolio. While there is wide agreement on the general principles of equity index construction, index providers differ in their methodologies that determine which stocks are selected for inclusion, the number of stocks to include, and other details. Exhibit 5.2 (at the end of the chapter) summarizes the construction methodologies of the major broad indexes of the U.S. equity market.

Annual and Monthly Returns

Returns on the Basic Asset Classes

Summary statistics of annual total returns of the seven basic SBBI® asset classes are discussed in Chapter 2.

Calculating Annual Returns

Annual returns are formed by compounding the 12 monthly returns. Compounding, or linking, monthly returns is multiplying together the return relatives, or one plus the return, then subtracting one from the result. The equation is denoted as the geometric sum as follows:

$$r_{year} = \left[(1+r_{Jan})\times(1+r_{Feb})\ldots\times(1+r_{Dec})\right]-1$$

Where:

r_{year}	=	The compound total return for the year
$r_{Jan}, r_{Feb}, \ldots, r_{Dec}$	=	The returns for the 12 months of the year

139 See Roll, R. 1977. “A critique of the asset pricing theory’s tests Part I: On past and potential testability of the theory.” *Journal of Financial Economics*, Vol. 4, No. 2, P. 129.

The compound return reflects the growth of funds invested in an asset. The following example illustrates the compounding method for a hypothetical year:

Month	Return (%)	Return (Decimal)	Return (Relative)
January	1	0.01	1.01
February	6	0.06	1.06
March	2	0.02	1.02
April	1	0.01	1.01
May	-3	-0.03	0.97
June	2	0.02	1.02
July	-4	-0.04	0.96
August	-2	-0.02	0.98
September	3	0.03	1.03
October	-3	-0.03	0.97
November	2	0.02	1.02
December	1	0.01	1.01

The return for this hypothetical year is the geometric sum:

(1.01 x 1.06 x 1.02 x 1.01 x 0.97 x 1.02 x 0.96 x 0.98 x 1.03 x 0.97 x 1.02 x 1.01) – 1 = 1.0567 – 1 = 0.0567

or a gain of 5.67%. One dollar invested in this hypothetical asset at the beginning of the year would have grown to slightly less than $1.06. Note that this is different from the simple addition result, (1 + 6 + 2 + 1 – 3 + 2 – 4 – 2 + 3 – 3 + 2 + 1) = 6%

Calculation of Returns from Index Values

Equivalently, annual returns, r_t, can be formed by dividing index values according to:

$$r_t = \left[\frac{V_t}{V_{t-1}}\right] - 1$$

Where:

r_t = The annual return in period t

V_t = The index value as of year-end t

V_{t-1} = The index value as of the pervious year-end, t-1

The construction of index values is discussed later in this chapter in the section entitled "Calculation of Index Values."

Calculation of Annual Income Returns

The conversion of monthly income returns to annual income returns is calculated by adding all the cash flows (income payments) for the period, then dividing the sum by the beginning period price:

$$r_I = \frac{(I_{Jan} + I_{Feb} \ldots + I_{Dec})}{P_0}$$

Where:

r_I	=	The income return for the year
$(I_{Jan}, I_{Feb} \ldots I_{Dec})$	=	The income payments for the 12 months of the year
P_0	=	The price of the security at the beginning of the year

The following example illustrates the method for a hypothetical year:

Month	Beginnin of Month Price ($)	Income Return (Decimal)	Income Payment ($)
January	100	0.006	0.60
February	102	0.004	0.41
March	105	0.002	0.21
April	101	0.001	0.10
May	99	0.005	0.50
June	103	0.004	0.41
July	105	0.003	0.32
August	103	0.002	0.21
September	105	0.003	0.32
October	103	0.004	0.41
November	106	0.001	0.11
December	105	0.002	0.21

Sum the income payments (not the returns), and divide by the price at the beginning of the year:(0.60 + 0.41 + 0.21 + 0.10 + 0.50 + 0.41 + 0.32 + 0.21 + 0.32 + 0.41 + 0.11 + 0.21) / 100 = 0.0381 or an annual income return of 3.81%.

Annual income and capital appreciation returns do not sum to the annual total return. The difference may be viewed as a reinvestment return, which is the return from investing income from a given month into the same asset class in subsequent months within the year.

Index Values

Index values represent the *cumulative* (i.e., compound) effect of investment returns. For example, in 1926 the total return (i.e., with dividends reinvested) of large-cap stocks was approximately 11.6%. A hypothetical investor investing $1.00 as of December 31, 1925, in large-cap stocks would have seen their investment grow to approximately $1.12 ($1.00 x (1+ 0.116)) by the end of 1926. During the following year (1927), the total return of large-cap stocks was approximately 37.5%. The $1.12 our hypothetical investor began 1927 with would have grown to approximately $1.53 ($1.12 x (1 + 0.375)) by the end of 1927. This can also be calculated as:

$\$1.53 = \$1 \times (1+ 0.116) \times (1 + 0.375)$

Of course, the cumulative effect over the entire 1926–2020 period can also be calculated as follows:

$\$1.00 \times (1+ r_{1926}) \times (1 + r_{1927}) \ldots \times (1 + r_{2020})$

Where "r_{year}" is the total return in a given year.

Following this methodology, the $1.00 invested in large-cap stocks at year-end 1925 by our hypothetical investor would have grown to nearly $11,000 by the end of 2020. Such growth reveals the power of compounding (reinvesting) one's investment returns.

Year-end index levels (based upon $1.00 invested at the end of 1925) for all six SBBI® asset classes plus inflation are displayed in Exhibit 5.1 as of year-end (i.e., December 31) 2010, 2011, and 2012.[140] Exhibit 5.1 also includes year-end index levels for the capital appreciation ("Capital App") component of total return for large-cap stocks, long-term government bonds, and intermediate term government bonds.

Note that the capital appreciation component of total return is generally a small contributor to total return over the longer term. For example, $1.00 invested at the end of 1925 in large-cap stocks would have grown to $3,532.55 by the end of 2012, but capital appreciation contributed only $111.77 to total return, or about 3.2% ($111.77 ÷ $3,532.55). This implies that the other two

[140] 2010–2012 was selected for example purposes only. *Precalculated* monthly and annual index values from January 1926 through December 2019 are presented in table format in the full-version *2020 SBBI® Yearbook* for the seven SBBI® total return series, and for all accompanying SBBI® capital appreciation series, as follows: Large-Capitalization Stocks: Capital Appreciation Index, Small-Capitalization Stocks: Total Return Index, Long-term Corporate Bonds: Total Return Index, Long-term Government Bonds: Total Return Index, Long-term Government Bonds: Capital Appreciation Index, Intermediate-term Government Bonds: Total Return Index, Intermediate-term Government Bonds: Capital Appreciation Index, U.S. Treasury Bills: Total Return Index, Inflation Index. For more information, visit dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook.

components of total return (dividend return and reinvestment return) contributed approximately 96.8% to the total return of large-cap stocks over the 1926–2012 period.

Exhibit 5.1: SBBI® Series Terminal Index Values as of Year- 2010, 2011, and 2012 (Year-end 1925 = $1.00)

	Large-Cap Stocks		**Small-Cap Stocks**	**Long-term Corp Bonds**	**Long-term Gov't Bonds**	
Year	Total Returns	Capital App	Total Returns	Total Returns	Total Returns	Capital App
2010	2,982.23	98.56	16,054.70	133.38	92.94	1.12
2011	3,045.21	98.56	15,532.07	157.32	118.13	1.38
2012	3,532.55	111.77	18,364.60	174.12	122.18	1.39
etc…						

	Inter-term Gov't Bonds		**U.S. Treasury Bills**	**Inflation***
Year	Total Returns	Capital App	Total Returns	Inflation
2010	84.12	1.60	20.55	12.24
2011	91.53	1.71	20.56	12.61
2012	93.05	1.73	20.57	12.83
etc…				

* NOTE: On February 26, 2021 Morningstar revised the "IA SBBI US Inflation" series. The revisions were applied to various dates from February 1926 through December 2020. The revisions were small and did not materially affect long-term averages.

Calculation of Index Values

It is possible to mathematically describe the nature of the index values in Exhibit 5.1 precisely. These indexes are initialized as of December 31, 1925, at $1.00 (represented by V_0 in the equation below). At the end of each month, a cumulative wealth index (V_n) for each of the monthly return series is formed. This index is formed for month n by taking the product of one plus the returns each period, in the following manner:

$$V_n = V_0\left[\prod_{t=1}^{n}(1+r_t)\right]$$

Where:

V_n	= The index value at end of period *n*
V_0	= The initial index value at time 0
r_t	= The return in period *t*

Using Index Values for Performance Measurement

Index values can be used to determine whether an investment portfolio accumulated more wealth over a period than another portfolio would have done, or whether the investment performed as well as an industry benchmark. In the following example, which produced more wealth: the "investor portfolio" or a hypothetical S&P 500 index fund returning exactly the S&P total return?[141]

	Investor Portfolio (%)	**S&P 500 (%)**
January 1990	-5.35	-6.71
February 1990	0.65	1.29
March 1990	0.23	2.65
Accumulated Wealth	\$0.955	\$0.970

Taking December 1989 as the base period (i.e., \$1.00 invested at the end of 1989) and using the computation method described above, the S&P 500 *outperformed* the investor portfolio.

Computing Returns for Non-Calendar Periods

Index values are also useful for computing returns for non-calendar-year periods. For example, using the index values in Exhibit 5.1 for the "Investor Portfolio," the return over the January 1990 to March 1990 period can be calculated by dividing the index value at the end of March 1990 (\$0.955) by the starting index value as of December 1998 (\$1.00), and subtracting 1. This yields:

(\$0.955 ÷ \$1.00) – 1 = -4.5%.

The same calculation can be performed for the S&P 500 Index over the January 1990 to March 1990 period:

(\$0.970 ÷ \$1.00) – 1 = -3.0%.

[141] In this example, each index measures total return and assumes monthly reinvestment of dividends.

Inflation-Adjusted Returns and Indexes

Inflation-adjusted returns and indexes can be used to measure investors' real returns and real changes in wealth over time.[142,143]

For example, a hypothetical investment of $1.00 in SBBI® large-cap stocks over the most recent 10-year period (2011–2020) would have grown to $3.67 by December 31, 2020 in nominal terms, and $3.09 in real (i.e., inflation-adjusted) terms. This demonstrates that investors in large-cap stocks multiplied their nominal wealth over the 2011–2020 period by a factor of 3.67 ($3.67 ÷ $1.00) and their real wealth (i.e., purchasing power) over the 2011–2020 period by a factor of 3.09 ($3.09 ÷ $1.00).

Overview of Major Broad Market U.S. Equity Indexes

The "market" is typically represented by a broad-based equity index that includes a wide range of industries and arguably behaves like the market as a whole. The SBBI large-cap stock series (essentially the S&P 500 index) is one example of this. Exhibit 5.2 provides an overview of some additional major broad market U.S. equity indexes.

142 The calculation of inflation-adjusted returns is discussed in Chapter 4, "Description of the Derived Series".

143 *Precalculated* annual inflation-adjusted index values (Year-end 1925 = $1.00) for all years 1926–2020 are presented in table format in the full-version *2021 SBBI® Yearbook* for the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD, (iii) Long-term (i.e., 20-year) Corporate Bonds: IA SBBI® US LT Corp TR USD, (iv) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, (v) Intermediate-term (i.e., 5-year) Government Bonds: IA SBBI® US IT Govt TR USD, and (vi) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD. For more information, visit: dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook.

Exhibit 5.2: Major Broad Market U.S. Equity Indexes*

	Index Family			
	Morningstar	**MSCI**	**Russell**	**S&P Dow Jones**
Broad Market Index	Morningstar U.S. Market Index	MSCI Investable Market	Russell 3000	S&P Composite 1500**
Percent U.S. Market Cap Coverage	97%	>99%	98%	90%
Total Number of Stocks	1441	2375	3000	1500
Transparent, Rules-Based Methodology	Yes	Yes	Yes	Yes
Eligibility	Stocks of companies domiciled in the U.S. listed on the NYSE, NYSE MKT, or NASDAQ	Stocks of companies domiciled in the U.S. listed on the NYSE, NYSE MKT, NYSE Arca, or NASDAQ	Stocks of the largest 3000 companies domiciled in the U.S. listed on a U.S. exchange	Stocks of companies domiciled in the U.S. listed on the NYSE, NYSE MKT, NYSE Arca, or NASDAQ chosen for market size, liquidity, and industry group representation by the S&P Index Committee
Exclusion Criteria	ADRs Limited Partnerships, Investment Trusts (except REITs), and Holding Companies	ADRs, Fixed Dividend Shares, Convertible Notes, Warrants, and Rights, Limited Partnerships, Master Limited Partnerships, LLCs, Business Development Companies, Pooled Investment Vehicles, Royalty Trusts and Investment Trusts (except REITs)	ADRs Limited Partnerships, Closed-end Mutual Funds, Price < $1, and Royalty Trusts and LLCs	ADRs Limited Partnerships, Investment Trusts (except REITs), Tracking Stocks and Holding Companies, and Royalty Trusts and LLCs
Market Cap Cutoff Method	Market Cap Percent	Fixed Number of Stocks	Fixed Number of Stocks	Fixed Number of Stocks
Unique Cap Cutoff Method	Yes	Yes	Yes	Yes
Unique Style Classification	Yes	No, stocks may be included in more than one style index	No, stocks may be included in more than one style index	No, stocks may be included in more than one style index
Core Style Index	Yes	No	No	No
Reconstitution Frequency	Semiannual	Semiannual	Annual	Ad hoc

*The broad market indices shown in Exhibit 5.3 can be disaggregated into capitalization and style indices. For example, the S&P Composite 1500 can be disaggregated into the S&P 500 (large-cap stocks), S&P 400 (mid-cap stocks), and the S&P 600 (small-cap stocks).

**The market for U.S. large-cap stocks is represented by the S&P 500 throughout the Ibbotson® SBBI® Yearbook series.

Chapter 6
Statistical Analysis of Returns

Statistical analysis of historical asset returns can reveal the growth rate of wealth invested in an asset or portfolio, the riskiness or volatility of asset classes, the co-movement of assets, and the random or cyclical behavior of asset returns. This chapter focuses on arithmetic and geometric mean returns, standard deviations, and serial and cross-correlation coefficients, and discusses the use of each statistic to characterize the various asset classes by growth rate, variability, and safety.

Calculating Arithmetic Mean Return

The arithmetic mean of a series is the simple average of the elements in the series. The arithmetic mean return equation is:

$$r_A = \frac{1}{n}\sum_{t=1}^{n} r_t$$

Where:

r_A = The arithmetic mean return

r_t = The series return in period t, that is, from time t - 1 to time t

n = The inclusive number of periods

Calculating Geometric Mean Return

The geometric mean of a return series over a period is the compound rate of return over the period. The geometric mean return equation is:

$$r_G = \left[\prod_{t=1}^{n}(1+r_t)\right]^{\frac{1}{n}} - 1$$

Where:

r_G = The geometric mean return

r_t = The series return in period t

= The inclusive number of periods

The geometric mean return can be restated using beginning and ending period index values. The equation is:

$$r_G = \left[\frac{V_n}{V_0}\right]^{\frac{1}{n}} - 1$$

Where:

r_G = The geometric mean return

V_n = The ending period index value at time n

V_0 = The initial index value at time 0

n = The inclusive number of periods

The annualized geometric mean return over any period of months can also be computed by expressing n as a fraction. For example, the beginning of 2020 to the end of May 2020 is equivalent to five twelfths of a year, or 0.4167. V_n would be the index value at the end of May 2020, V_0 would be the index value at the beginning of 2020, and n would be 0.4167.

Geometric Mean Versus Arithmetic Mean

A simple example illustrates the difference between geometric and arithmetic means. Suppose \$1.00 was invested in a large-cap stock portfolio that experiences successive annual returns of 50% and negative 50%. At the end of the first year, the portfolio is worth \$1.50 and at the end of the second year, it is worth \$0.75. The annual arithmetic mean is 0.0%, whereas the annual geometric mean is -13.4%. Both are calculated as follows:

$$r_A = \frac{1}{2}(0.50 - 0.50) = 0.00$$

$$r_G = \left[\frac{0.75}{1.00}\right]^{\frac{1}{2}} - 1 = \text{-}0.134$$

The geometric mean is backward-looking, measuring the change in wealth over more than one period. On the other hand, the arithmetic mean better represents a typical performance over a single period.

In general, the geometric mean for any period is less than or equal to the arithmetic mean. The two means are equal only for a return series that is constant (i.e., the same return in every period). For a non-constant series, the difference between the two is positively related to the variability or standard deviation of the returns. For example, in Exhibit 6.1 the difference between the arithmetic and geometric mean is much larger for risky large-cap stocks than it is for nearly riskless Treasury bills. This is because the "variability" (as measured by annual standard deviation) of large-cap stock returns is nearly 20%, while the standard deviation of Treasury bill returns is much lower, at just over 3%.

Exhibit 6.1: Comparison of the Annual Geometric (i.e., Compound) Mean Return and Annual Average Return of Large-Cap Stocks and U.S. Treasury Bills (%)
1926–2020

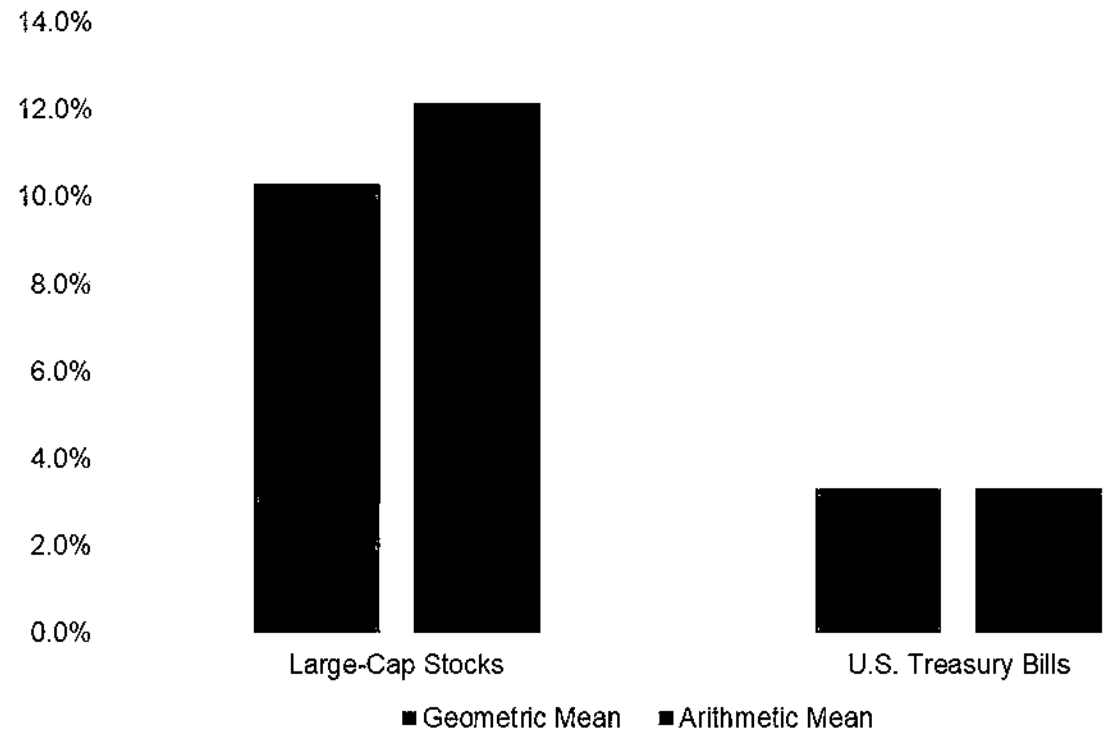

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes are represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, and (ii) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Calculating Standard Deviation

The standard deviation of a series is a measure of the extent to which observations in the series differ from the arithmetic mean of the series. For a series of asset returns, the standard deviation is a measure of the volatility, or risk, of the asset.

In a normally distributed series, about two-thirds of the observations lie within one standard deviation of the arithmetic mean; about 95% of the observations lie within two standard deviations; and more than 99% lie within three standard deviations.

For example, the standard deviation for large-cap stock returns from 1926 to 2020 was 19.7% with an annual arithmetic mean of 12.2%. Therefore, roughly two-thirds of the observations have annual returns between -7.5% and 31.9% (12.2% plus or minus 19.7%); approximately 95% of the observations are between -27.2% and 51.6% (12.2% plus or minus 39.4%).

The equation for the standard deviation of a series of returns (σ_r) is:

$$\sigma_r = \sqrt{\frac{1}{n-1}\sum_{t=1}^{n}(r_t - r_A)^2}$$

Where:

r_t = The return in period t

r_A = The arithmetic mean of the return series r

n = the number of periods

The scaling of the standard deviation depends on the frequency of the data; therefore, a series of monthly returns produces a monthly standard deviation. For example, using the monthly returns for the hypothetical year on page 5-3, a monthly standard deviation of 2.94% calculated as follows:

$$0.0294 = \frac{1}{12-1} \times \begin{bmatrix} (0.01-0.005)^2 + (0.06-0.005)^2 + (0.02-0.005)^2 + \\ (0.01-0.005)^2 + (-0.03-0.005)^2 + (0.02-0.005)^2 + \\ (-0.04-0.005)^2 + (-0.02-0.005)^2 + (0.03-0.005)^2 + \\ (-0.03-0.005)^2 + (0.02-0.005)^2 + (0.01-0.005)^2 \end{bmatrix}^{\frac{1}{2}}$$

It is sometimes useful to express the standard deviation of the series in another time scale. To calculate *annualized* monthly standard deviations (σ_n) one uses the following equation.[144]

$$\sigma_n = \sqrt{[\sigma_1^2 + (1+\mu_1)^2]^n - (1+\mu_1)^{2n}}$$

[144] The equation appears in Levy, H. & Gunthorpe, D. 1993. "Optimal Investment Proportions in Senior Securities and Equities Under Alternative Holding Periods." *Journal of Portfolio Management*, Vol. 19, No. 4, P. 33.

Where:

n = The number of periods per year, e.g., 12 for monthly, 4 for quarterly, etc.

σ_1 = The monthly standard deviation

μ_1 = The months arithmetic mean

Applying this formula to the prior monthly standard deviation of 2.94% results in an annualized monthly standard deviation of 10.78%. The *annualized* monthly standard deviation is calculated as follows:

$$\sqrt{\left[0.0294^2+(1+0.005)^2\right]^{12}-(1+0.005)^{2(12)}}$$

This equation is the exact form of the common *approximation*:

$$\sigma_n \approx \sqrt{n}\sigma_1$$

The "approximation" treats an annual return as if it were the sum of 12 *independent* monthly returns, whereas the "exact form" treats an annual return as the *compound* return of 12 independent monthly returns. While the approximation can be used for "back of the envelope" calculations, the exact formula should be used in applications of quantitative analysis. Forming inputs for mean-variance optimization is one such example. Note that both the exact formula and the approximation assume that there is no monthly autocorrelation.

Limitations of Standard Deviation[145]

Using the statistical measure of standard deviation of returns is clearly the easiest way to mathematically express the concept of risk. However, practitioners and academics alike have noted that standard deviation misses important and essential qualities of risk from the standpoint of an investor of capital.

One limitation of standard deviation as a measure of risk is the tacit assumption that returns can be described by a measure that assumes a normal distribution of returns, while it is empirically acknowledged that many financial market returns exhibit excess kurtosis relative to the normal (Gaussian) distribution. This characteristic is referred to as a leptokurtic, or "fat-tailed," return distribution. Fat-tailed outcomes reflect market movements far larger than one would reasonably expect from a normal distribution of returns. One of the most extreme examples of a fat-tailed return profile occurred on Oct. 19, 1987, when the Dow Jones Industrial Average declined by 22.68%, or more than 20 standard deviations. The magnitude of the deviation from normal returns

[145] The Limitations of Standard Deviation, the Semi-Variance and Semi-Standard Deviation, and the Issues Regarding Semi-Variance sections were written by Erik Kobayashi-Solomon and Philip Guziec.

can be understood when considering that a normal distribution would predict such a move once in more than 4.5 billion years. More recently, 2008 had 11 days with declines greater than 4 standard deviations, and on May 6, 2010, the Dow Jones Industrial Average declined by 9% in a matter of minutes on an intraday basis, a move that on a daily basis would have been among the top 10 declines in recorded history. Clearly, an awareness of the nature of statistical descriptions of market moves beyond standard deviation is helpful in developing a representative profile of market risk. As Laurence B. Siegel wrote in July 2018:[146]

"Although most of classical finance focuses only on risk and expected return, investors differ in their tastes and preferences and assets differ in their characteristics other than risk and return."

An excellent recent book discusses how investor preferences for various assets and premia, broadly characterized as the "popularity" of these characteristics, influence returns. In 2019, Roger G. Ibbotson and colleagues Thomas M. Idzorek, CFA, Paul D. Kaplan, CFA, and James X. Xiong, CFA, published a new CFA Institute Research Foundation monograph entitled *Popularity: A Bridge Between Classical and Behavioral Finance* (available for download at: https://www.cfainstitute.org/en/research/foundation/2018/popularity-bridge-between-classicaland- behavioral-finance).[147]

Semivariance and Semistandard Deviation

Given academic and practitioner concerns about variance, various approaches have been suggested to more appropriately measure risk. We take a moment here to briefly discuss investor perception of risk and to review another measure, semi variance.

One criticism of variance and standard deviation is that an investor is less worried about bidirectional variation in value (the essence of the standard deviation measure) than about an ultimately unrecoverable shortfall in investment capital. In considering risk from this point of view, two cases stand out as the most salient: (i) suffering a realized or mark-to-market loss of capital that prevents fulfillment of a goal or mandate over an investment time frame and (ii) allocating capital in investments that appreciate too little to fulfill a goal or mandate over the investment time frame. The former case involves an excess of variation in an unacceptable direction; the latter case involves a paucity of variation to an acceptable magnitude.

Of these two cases, most academic work has focused on developing a framework to accurately measure and analyze directionally-specific variance. Foremost in this attempt has been the concept of semivariance.

Semivariance characterizes the downside risk of a distribution and focuses on the portion of risk that is below (to the left of) the mean or a specific target. For example, for a 4% target return, the

---

146 Laurence B. Siegel is the Gary P Brinson Director of Research at the CFA Institute Research Foundation. The quote is from the Forward of *Popularity – A Bridge between Classical and Behavioral Finance*, by Roger G Ibbotson, Thomas M. Idzorek, CFA, Paul D. Kaplan, CFA, and James X. Xiong, CFA

147 Or, go to the CFA Institute website at cfainstitute.org and search for "popularity".

semivariance describes the variance of the data points below (to the left of) the return of 4%. The semivariance below the mean uses the mean return as the target return. The semistandard deviation is simply the square-root of the semivariance. The semivariance (semistandard deviation) is always lower than the total variance (standard deviation) of the distribution.

$$SV_t = \frac{1}{n}\sum_{r_t<r_T}^{n}(r_T - r_t)^2$$

$$SV_t = \frac{1}{n}\sum_{r_t<r_T}^{n}(r_T - r_t)^2$$

$$SSTD_m = \sqrt{SV_m}$$

$$SSTD_t = \sqrt{SV_t}$$

Where:

SV_m	=	The semi-variance below mean
SV_t	=	The semi-variance below target
r_A	=	The arithmetic mean return
r_t	=	The series return in period *t*
r_T	=	The target selection return
n	=	The inclusive number of periods
$SSTD_m$	=	The semi-standard deviation below mean
$SSTD_t$	=	The semi-standard deviation below target

Issues Regarding Semivariance

While semivariance seems to intuitively address issues regarding directionality, it does have empirical, theoretical, and practical shortcomings. Empirically, when returns are measured over relatively short time frames, distributions tend to be symmetric. As such, using semivariance for short time frames effectively gives no extra explanatory power (because semivariance simply equates to one half of the variance) and, in fact, limits the data available for analysis (because the calculation of semivariance discards any positive return observations). When returns are measured over relatively longer time frames (on the order of a year or more), asset returns tend to follow a distribution that is positively skewed. As such, for investors with longer time horizons, semivariance has less explanatory power because the data set is limited to the less germane case, while the richer part of the data set is discarded.

From a theoretical standpoint, the assumption implicit in the calculation of semivariance – investors do not care about positive variance – has repercussions regarding investor utility functions. Namely, ignoring positive variation implies that an investor is indifferent when presented with the choice between making an uncertain but positive return bet and making a bet that is certain to generate the expected payoff. For example, investors would, under the assumptions of semivariance, be agnostic between a 50-50 bet of generating either 5% or 10% and a sure bet paying 7.5%.

Practically speaking, ignoring upside variation means that we ignore the second aspect of risk mentioned above – a paucity of magnitude. In other words, if one attempts to minimize semivariance, without regard to the degree to which an asset or allocation has upside potential, one runs the risk of generating returns which, while low in downside variance, are also low in upside variance. In this case, one has protected oneself from one class of risk by taking on yet another. Given these issues, semivariance has met with limited acceptance among academics and practitioners alike.

Volatility of the Markets

The volatility of stocks and long-term government bonds is shown by the bar graphs of monthly returns in Exhibit 6.2. The stock market was tremendously volatile in the first few years studied; this period was marked by the 1920s boom, the crash of 1929–1932, and the Great Depression years. The market seemingly settled after World War II and provided more stable returns in the postwar period. In the 1970s and 1980s, stock market volatility increased, but not to the extreme levels of the 1920s and 1930s. In the 1990s, 2000s, and 2010s, volatility was relatively moderate. Bonds present a mirror image. Long-term government bonds were extremely stable in the 1920s and remained so through the crisis years of the 1930s, providing shelter from the storms of the stock markets. Starting in the late 1960s and early 1970s, however, bond volatility soared; in the 1973–1974 stock market decline, bonds did not provide the shelter they once did. Bond pessimism (i.e., high yields) peaked in 1981 and subsequent returns were sharply positive. While the astronomical interest rates of the 1979–1981 period have passed, the volatility of the bond market remains higher.

Exhibit 6.2: Month-by-Month Returns on Stocks and Bonds (%) 1926–2020

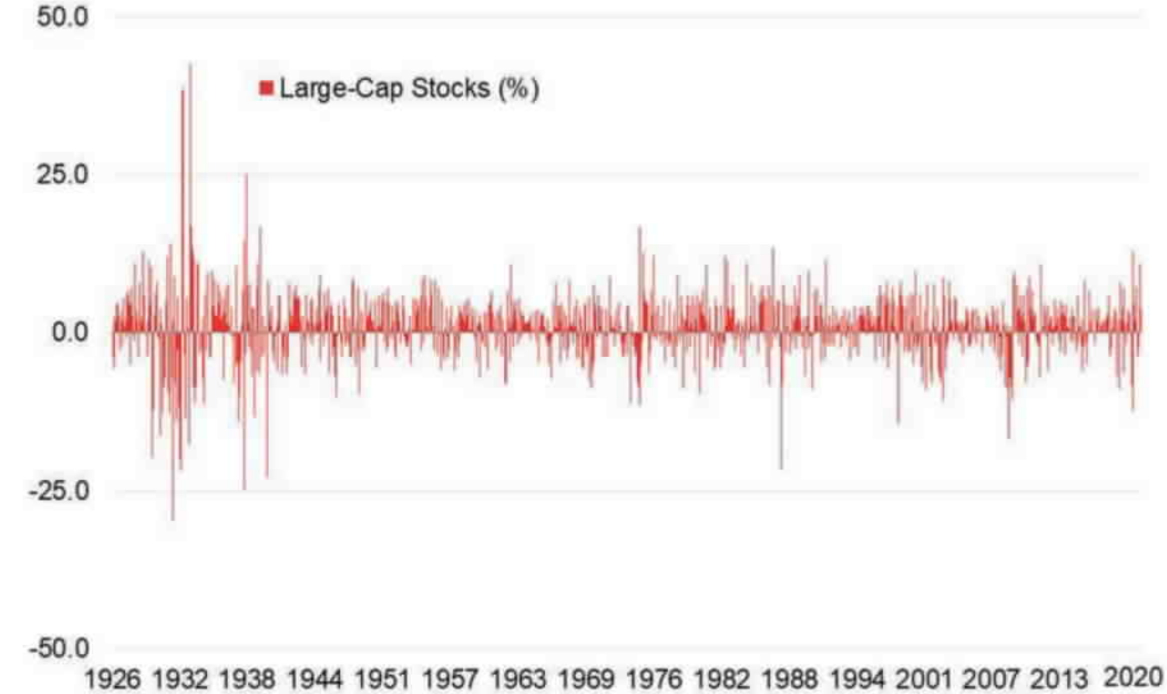

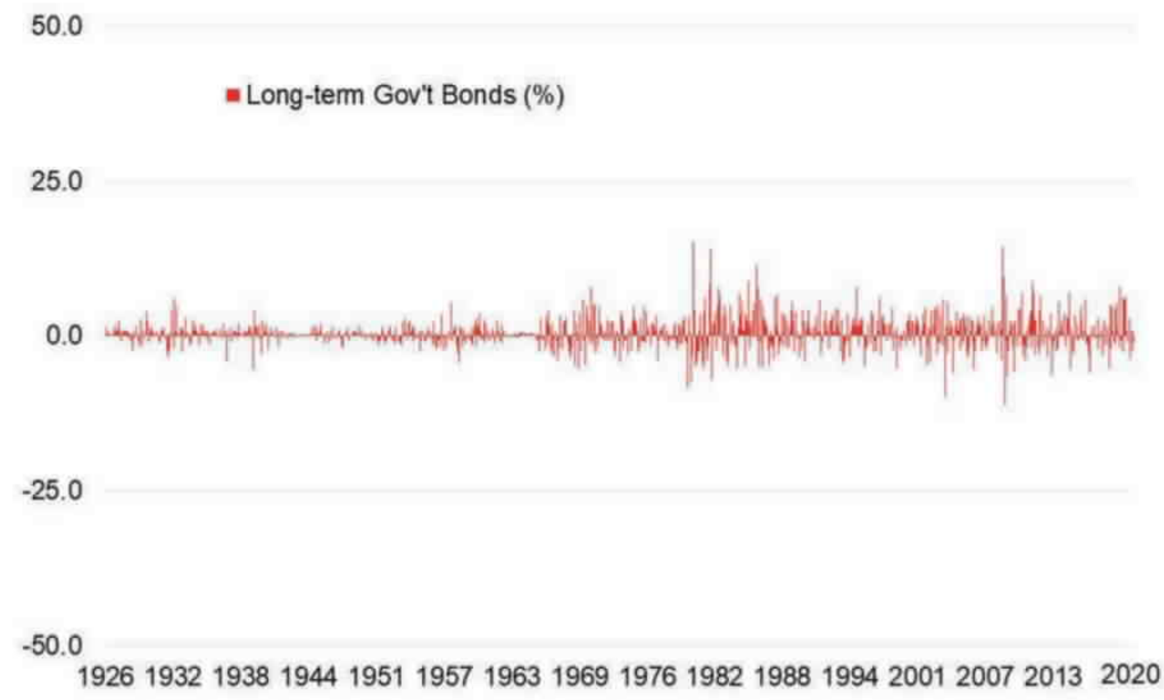

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Changes in the Risk of Assets Over Time

Another time series property of great interest is change in volatility or riskiness over time. Such change is indicated by the standard deviation of the series over different sub-periods. Exhibit 6.3 shows the annualized monthly standard deviations of the basic data series by decade beginning in 1926 and illustrates differences and changes in return volatility. In this exhibit, the 1920s cover the period 1926–1929. Equity returns have been the most volatile of the basic series, with volatility peaking in the 1930s due to the instability of the market following the 1929 market crash. The significant bond yield fluctuations of the 1980s caused the fixed-income series' volatility to soar compared to prior decades. Small cap stocks were the *most* volatile SBBI® asset class in all time periods shown in Exhibit 6.3. Treasury bills were the *least* volatile SBBI® asset class in all time periods shown in Exhibit 6.3.

Exhibit 6.3 displays the *annualized* standard deviation of the *monthly* returns on each of the basic and derived series from January 1926 to December 2020. The estimates in Exhibit 6.3 are not strictly comparable to Exhibits 2.4 or 2.8 where the 95-year period standard deviation of *annual* returns around the 95-year annual arithmetic mean is reported. The arithmetic mean drifts for a series that does not follow a random pattern. A series with a drifting mean will have much higher deviations around its long-term mean than it has around the mean during a particular calendar year.

Exhibit 6.3: Annualized Monthly Standard Deviations by Decade (%)

	1920s*	1930s	1940s	1950s	1960s	1970s	1980s	1990s	2000s	2010s
Large-Cap Stocks	23.9	41.6	17.5	14.1	13.1	17.1	19.4	15.9	16.3	14.1
Small-Cap Stocks	24.7	78.6	34.5	14.4	21.5	30.8	22.5	20.2	26.1	19.6
Long-term Corp Bonds	1.8	5.3	1.8	4.4	4.9	8.7	14.1	6.9	11.7	9.0
Long-term Gov't Bonds	4.1	5.3	2.8	4.6	6.0	8.7	16.0	8.9	12.4	10.9
Inter-term Gov't Bonds	1.7	3.3	1.2	2.9	3.3	5.2	8.8	4.6	5.2	3.2
U.S. Treasury Bills	0.3	0.2	0.1	0.2	0.4	0.6	0.9	0.4	0.6	0.2
Inflation	2.2	2.6	3.1	1.3	0.8	1.3	1.3	0.7	1.6	1.0

*Based on the period 1926–1929

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes and inflation represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD, (iii) Long-term (i.e., 20-year) Corporate Bonds: IA SBBI® US LT Corp TR USD, (iv) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, (v) Intermediate-term (i.e., 5-year) Government Bonds: IA SBBI® US IT Govt TR USD, (vi) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD, and (vii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Large-cap stocks and equity risk premiums have virtually the same annualized monthly standard deviations because there is very little deviation in the U.S. Treasury bill series (the monthly equity risk premium is calculated as [(1+ Large Stock total return) ÷ (1+Treasury Bill total return) – 1], as described in Exhibit 4.1). The series with drifting means (U.S. Treasury bills, inflation rates, and inflation-adjusted U.S. Treasury bills) all tend to have very low annualized monthly standard deviations since these series are quite predictable from month to month. There is much less predictability for these series over the long term. Because it is difficult to forecast the direction and magnitude of the drift in the long-term mean, these series have higher standard deviations over the long term in comparison to their annualized monthly standard deviations.[148]

Equity investors may have the impression that periods of economic turmoil mainly affect equity markets, but this is not true. For example, during the 2008–2009 global financial crisis the price fluctuation of bonds in general and long-term corporate bonds in particular was dramatic. The annualized standard deviation of long-term corporate bonds recorded an all-time high value of 25.5% in 2008, a year that saw the collapse of storied investment banks Lehman Brothers and Bear Sterns. The previous record annualized standard deviation for long-term corporate bonds of 20.2% (measured during 1981) is more than one-fifth lower than the 2008 value, and more than twice the Depression era record of 11.7%, recorded in 1933. Another mark of the severity of bond price fluctuations can be seen by noting that the annualized standard deviations for long-term corporate bonds in 2008 were only three percentage points shy of the annualized standard deviations for the S&P 500 “Large-Cap Stocks” in 2009.

The annualized monthly standard deviation of all of the SBBI® asset classes increased in 2020 compared to 2019, but long-term corporate and U.S. government bond volatility was still significantly lower than the levels reached in 2008. Equity volatility, however, reached levels not seen for decades. The annualized monthly standard deviation of large-cap stocks in 2020 reached 31.6%, a level not seen since 1987. The price volatility of large-cap stocks in 2020 was the fifth highest on record with 1933’s 99.8% being the highest. The annualized monthly standard deviation of small-cap stocks in 2020 reached 43.2%. The price volatility of small-cap stocks in 2020 was the twelfth highest on record with 1933’s 286.6% being the highest.

Correlation Coefficients: Serial and Cross-Correlations

The behavior of an asset return series over time reveals its predictability. For example, a series may be random or unpredictable, or it may be subject to trends, cycles, or other patterns, making the series predictable to some degree. The serial correlation coefficient of a series determines its predictability given knowledge of the last observation. The cross-correlation coefficient (often shortened to “correlation”) between two series determines the predictability of one series, conditional on knowledge of the other.

[148] *Precalculated* annualized monthly standard deviations (1926–2020) are presented in table format in the full-version *2021 SBBI® Yearbook* for the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series and derived series, as follows: Large-Cap Stocks, Small-Cap Stocks, Long-term (i.e., 20-year) Corporate Bonds, Long-term (i.e., 20-year) Government Bonds, Intermediate-term (i.e. 5-year) Gov't Bonds, U.S. (30-day) Treasury Bills, Inflation, Equity Risk Premium, Small Stock Premium, Bond Default Premium, Bond Horizon Premium, Inflation-adjusted T-Bills (i.e., real interest rates), For more information, visit dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook.

Serial Correlations

The serial correlation of a return series, also known as the first order autocorrelation, describes the extent to which the return in one period is related to the return in the next period. A return series with a high (near 1) serial correlation is very predictable from one period to the next, while one with a low (near 0) serial correlation is random and unpredictable.

The serial correlation of a series is closely approximated by the equation for the cross correlation between two series. The data, however, are the series and its “lagged” self. For example, the lagged series is the series of one-period-old returns:

Period	Return Series	Lagged Return Series
1	0.10	undefined
2	-0.10	0.10
3	0.15	-0.10
4	0.00	0.15

Cross Correlations

The cross-correlation between two series measures the extent to which they are linearly related.[149] The correlation coefficient measures the sensitivity of return on one asset class or portfolio to the return of another. The correlation equation between return series X and Y is:

$$p_{X,Y} = \left[\frac{Cov(X,Y)}{\sigma_X \sigma_y} \right]$$

Where:

$Cov(X,Y)$ = The covariance of X and Y, defined below

σ_X = The standard deviation of X

σ_y = The standard deviation of Y

The covariance equation is:

[149] Two series can be related in a nonlinear way and have a correlation coefficient of zero. An example is the function $y = x^2$, for which $p_{x,y} = 0$.

$$Cov(X,Y) = \frac{1}{n-1}\sum_{t=1}^{n}(r_{X,t} - r_{X,A})(r_{Y,t} - r_{Y,A})$$

Where:

$r_{X,t}$ = The return for series X in period t

$r_{Y,t}$ = The return for series Y in period t

$r_{X,A}$ = The arithmetic mean of series X

$r_{Y,A}$ = The arithmetic mean of series Y

n = The number of periods

Correlations of the Basic Series

Exhibit 6.4 presents the annual cross correlations and serial correlations for the seven basic series. Long-term government bond returns and long-term corporate bond returns are highly correlated with each other but negatively correlated with inflation. To the degree that inflation is unanticipated, it has a negative effect on fixed-income securities. In addition, U.S. Treasury bills and inflation are reasonably highly correlated, a result of the post-1951 "tracking" described in Chapter 2. Lastly, both the U.S. Treasury bills and inflation series display high serial correlations.

Exhibit 6.4: Basic Series: Serial and Cross-Correlations of Historical Annual Returns 1926–2020

	Large-Cap Stocks	**Small-Cap Stocks**	**Long-term Corp Bonds**	**Long-term Gov't Bonds**	**Inter-term Gov't Bonds**	**U.S. Treasury Bills**	**Inflation**
Large-Cap Stocks	1.00						
Small-Cap Stocks	0.79	1.00					
Long-term Corp Bonds	0.17	0.05	1.00				
Long-term Gov't Bonds	0.01	-0.10	0.89	1.00			
Inter-term Gov't Bonds	-0.02	-0.11	0.84	0.86	1.00		
U.S. Treasury Bills	-0.02	-0.08	0.14	0.17	0.47	1.00	
Inflation	-0.01	0.05	-0.15	-0.14	0.01	0.42	1.00
Serial Correlation	0.01	0.06	0.02	-0.15	0.15	0.91	0.63

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes and inflation represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD, (iii) Long-term (i.e., 20-year) Corporate Bonds: IA SBBI® US LT Corp TR USD, (iv) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, (v) Intermediate-term (i.e., 5-year) Government Bonds: IA SBBI® US IT Govt TR USD, (vi) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD, and (vii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Correlations of the Derived Series

The annual cross-correlations and serial correlations for the four risk premium series and inflation are presented in Exhibit 6.5. Notice that inflation is negatively correlated with the horizon premium. Increasing inflation causes long-term bond yields to rise and prices to fall; therefore, a negative horizon premium is observed in times of rising inflation.

Exhibit 6.6 presents annual cross-correlations and serial correlations for the inflation-adjusted asset return series. It is interesting to observe how the relationship between the asset returns are substantially different when these returns are expressed in inflation adjusted terms (as compared with nominal terms). In general, the cross-correlations between asset classes are higher when one accounts for inflation (i.e., subtracts inflation from the nominal return).

Exhibit 6.5: Risk Premia and Inflation: Serial and Cross-Correlations of Historical Annual Returns 1926–2020

	Equity Risk Premium	**Small Stock Premium**	**Bond Default Premium**	**Bond Horizon Premium**	**Inflation**
Equity Risk Premium	1.00				
Small Stock Premium	0.26	1.00			
Bond Default Premium	0.31	0.17	1.00		
Bond Horizon Premium	0.03	-0.11	-0.48	1.00	
Inflation	-0.07	0.12	0.00	-0.27	1.00
Serial Correlation	0.02	0.36	-0.31	-0.16	0.63

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes and inflation represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD, (iii) Long-term (i.e., 20-year) Corporate Bonds: IA SBBI® US LT Corp TR USD, (iv) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD, and (vii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Exhibit 6.6: Inflation-Adjusted Series: Serial and Cross-Correlations of Historical Annual Returns 1926–2020

	Inflation-Adjusted						
Inflation-Adjusted Series	Large-Cap Stocks	Small-Cap Stocks	Long-term Corp Bonds	Long-term Gov't Bonds	Inter-term Gov't Bonds	U.S. Treasury Bills*	Inflation
Large-Cap Stocks	1.00						
Small-Cap Stocks	0.79	1.00					
Long-term Corp Bonds	0.23	0.08	1.00				
Long-term Gov't Bonds	0.09	-0.06	0.91	1.00			
Inter-term Gov't Bonds	0.07	-0.06	0.89	0.90	1.00		
U.S. Treasury Bills	0.09	-0.06	0.51	0.49	0.70	1.00	
Inflation	-0.19	-0.07	-0.54	-0.48	-0.58	-0.70	1.00
Serial Correlation	0.00	0.03	0.13	-0.06	0.20	0.66	0.63

*Real Interest Rates

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes and inflation represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD, (iii) Long-term (i.e., 20-year) Corporate Bonds: IA SBBI® US LT Corp TR USD, (iv) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, (v) Intermediate-term (i.e., 5-year) Government Bonds: IA SBBI® US IT Govt TR USD, (vi) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD, and (vii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Is Serial Correlation in the Derived Series Random?

The risk/return relationships in the historical data are represented in the equity risk premium, the small stock premium, the bond horizon premium, and the bond default premium. The real/nominal historical relationships are represented in the inflation rates and the real interest rates. The objective is to uncover whether each series is random or is subject to any trends, cycles, or other patterns.

The one-year serial correlation coefficients measure the degree of correlation between returns from each year and the previous year for the same series, as seen in Exhibit 6.6. Highly positive (near 1) serial correlations indicate trends, while highly negative (near -1) serial correlations indicate cycles. Looking to exhibit 6.7, the analysis suggests that both inflation rates and real riskless rates follow trends. Serial correlations near zero suggest no patterns (i.e., random behavior), so the analysis suggests that the equity risk premium and the bond horizon premium are random variables (although this is less-strongly indicated in the case of the bond horizon premium).

The small stock premium and the bond default premium, however, fall into a middle range that makes it more difficult to determine whether the small stock premium is a trend (or is random) and whether the bond default premium is a cycle (or is random). In these two cases, one could argue that the small stock premium is a *possible* trend, and the bond default premium is a *possible* cycle.

Exhibit 6.7: Interpretation of the Annual Serial Correlations

Series	Serial Correlation	Interpretation
Equity Risk Premium	0.02	Random
Small Stock Premium	0.36	Possible Trend
Bond Default Premium	-0.31	Possible Cycle
Bond Horizon Premium	-0.16	Random/Possible Cycle
Real Interest Rates	0.66	Trend
Inflation	0.63	Trend

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes and inflation represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD, (iii) Long-term (i.e., 20-year) Corporate Bonds: IA SBBI® US LT Corp TR USD, (iv) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD, and (vii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Rolling-Period Standard Deviations

Rolling-period standard deviations are obtained by rolling a window of fixed length along each time series and computing the standard deviation for the asset class for each window of time. They are useful for examining the volatility or riskiness of returns for holding periods similar to those actually experienced by investors. Exhibits 6.8 and 6.9 graphically depict this volatility. Monthly data are used to maximize the number of data points included in the standard deviation computation. The first 60-month (i.e., 5-year) rolling period covered is January 1926–December 1930, so each of the graphs start at 1930.

Exhibit 6.8 examines the 60-month rolling standard deviation for large-cap stocks, small-cap stocks, and long-term government bonds. It is interesting to see the relatively high standard deviation for small- and large-cap stocks in the 1930s with an apparent lessening of volatility for 60-month holding periods during the 1980s. Note also how the standard deviation for long-term government bonds reaches the level of both stock asset classes during part of the 1980s. Exhibit 6.9 examines the 60-month rolling standard deviation for long-term and intermediate-term government bonds and U.S. Treasury bills. Note that the vertical scale (from 0.0% to 25.0%) of Exhibit 6.8 is different than the vertical scale (from 0.0% to 5.0%) of Exhibit 6.9.

Exhibit 6.8: Rolling 60-month Standard Deviations: Large-cap Stocks, Small-cap Stocks and Long-term Government Bonds (%) 1926–2020

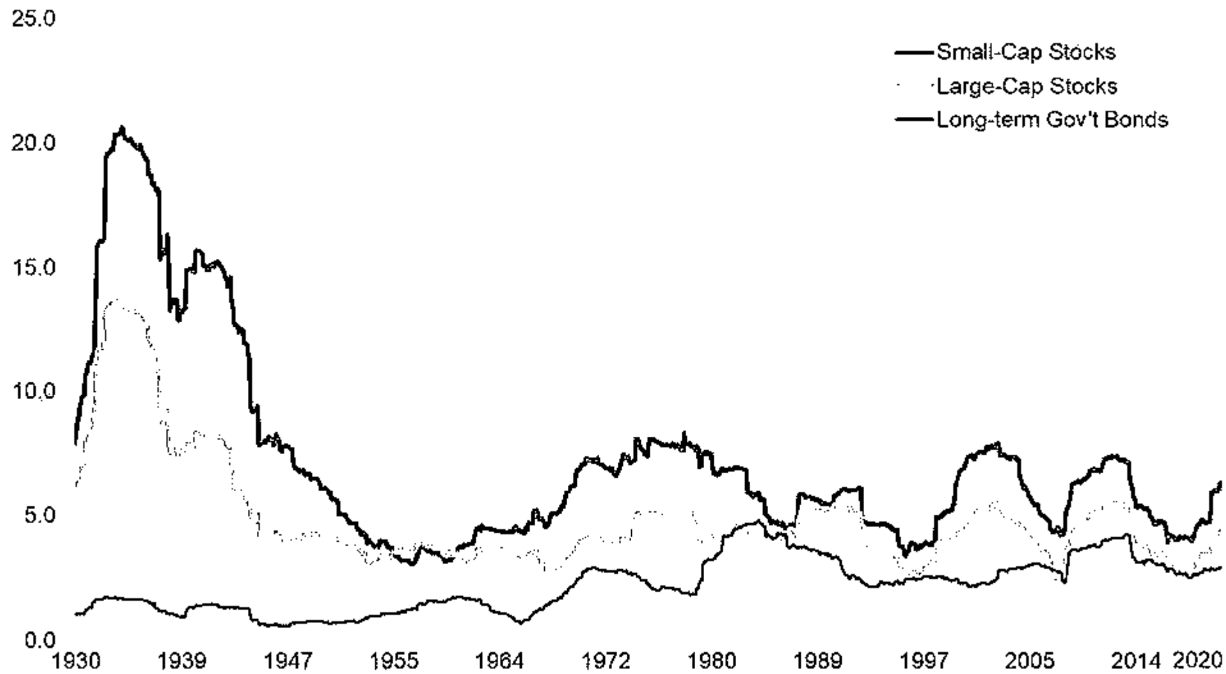

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD, (iii) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Exhibit 6.9: Rolling 60-month Standard Deviations: Long-term Government Bonds, Intermediate-term Government Bonds, and Treasury Bills (%) 1926–2020

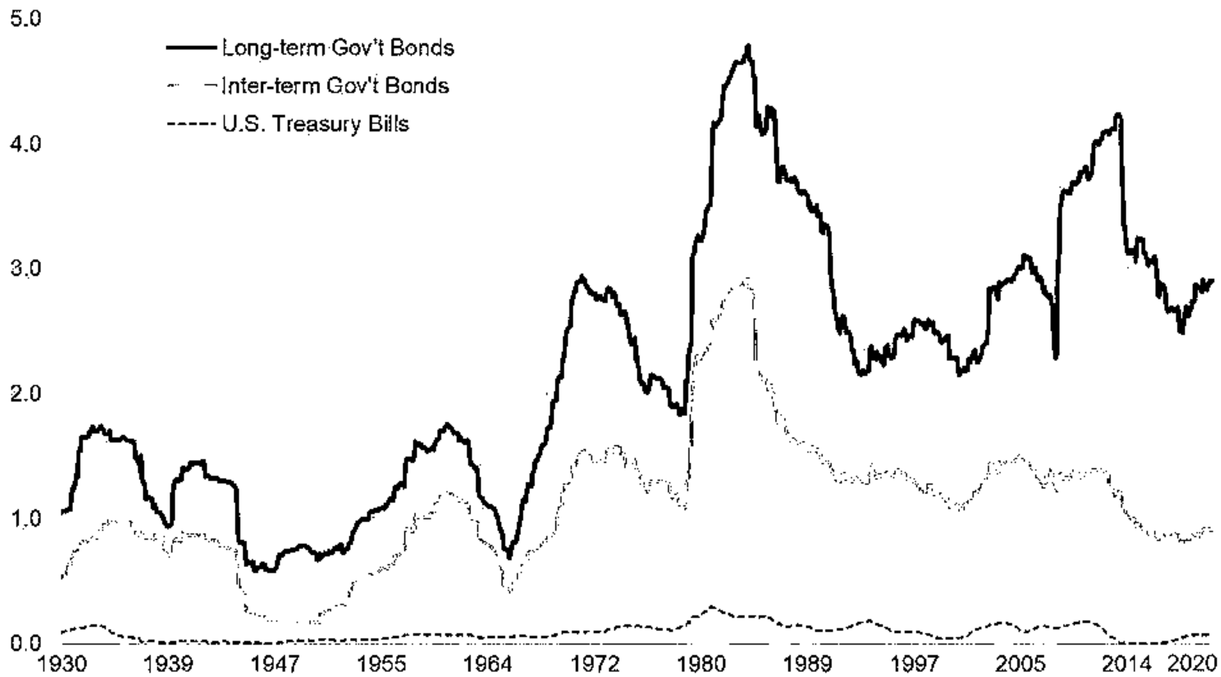

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, (v) Intermediate-term (i.e., 5-year) Government Bonds: IA SBBI® US IT Govt TR USD, (vi) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Rolling-Period Correlations

Rolling-period correlations are obtained by moving a window of fixed length along time series for two asset classes and computing the cross-correlation between the two asset classes for each window of time. They are useful for examining how asset class returns vary together for holding periods similar to those actually experienced by investors. Exhibits 6.12 and 6.13 graphically depict cross-correlation. Monthly data are used to maximize the number of data points included in the correlation computation. The first 60 month (i.e., 5-year) rolling period covered is January 1926– December 1930, so each of the graphs start at 1930.

Exhibits 6.10 and 6.11 show cross-correlations between two asset classes for 60-month holding periods. Exhibit 6.10 shows the volatility of the correlations between large-cap stocks and long-term government bonds. There are wide fluctuations between strong positive and strong negative correlations over the past 95 years. Exhibit 6.11 shows the correlation between Treasury bills and inflation. These asset classes also show wide fluctuations in correlation over the past 95 years.

Exhibit 6.10: Rolling 60-month Correlation: Large-cap Stocks and Long-term Government Bonds 1926–2020

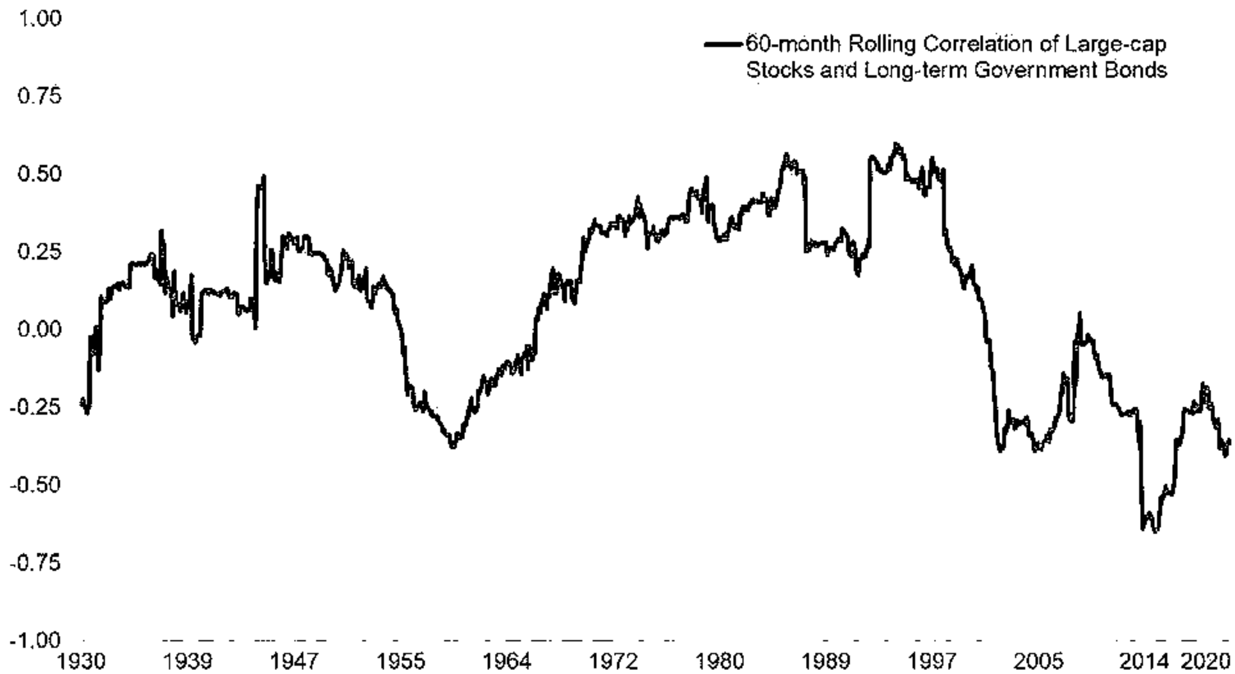

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Exhibit 6.11: Rolling 60-month Correlation: U.S. Treasury Bills and Inflation 1926–2020

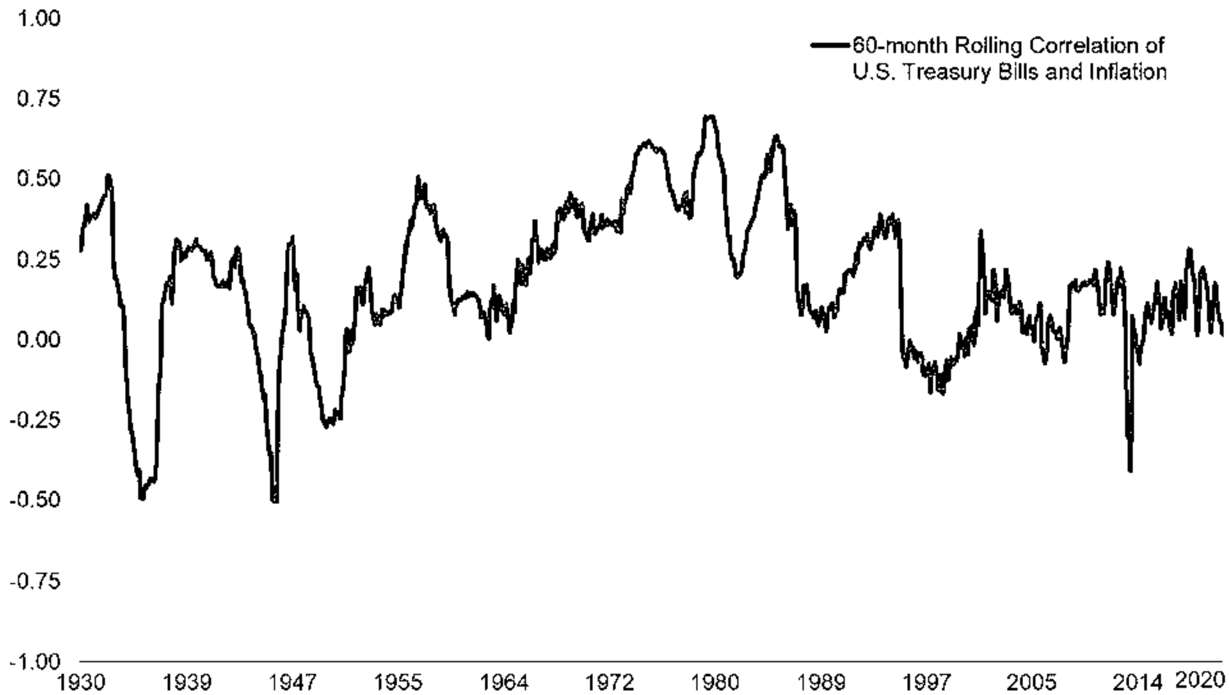

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD, and (vii) Inflation: IA SBBI® US Inflation. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

The True Impact of Asset Allocation on Return

The importance of asset allocation has been the subject of considerable debate and misunderstanding for decades. A May 2010 article pinpoints one of the primary sources of confusion surrounding the importance of asset allocation: How much of a portfolio's level of return comes from a fund's asset-allocation policy?[150]

Note: The following section was written by Thomas Idzorek, president of the Morningstar Investment Management division.

The Debate

The seminal work on the importance of asset allocation, the catalyst of a 25-year debate, and unfortunately the source of what is arguably the most prolific misunderstanding among investment professionals, is the 1986 article, "Determinants of Portfolio Performance," by Gary Brinson,

[150] Idzorek, T. M. "Asset Allocation is King." *Morningstar Advisor*, April/May, 2010, P. 28.

Randolph Hood, and Gilbert Beebower (BHB).[151] BHB regressed the time series returns of each fund on a weighted combination of indexes reflecting each fund's asset-allocation policy. In one of the many analyses that BHB carried out (and probably one of the least important ones), they found that the policy mix explained 93.6% of the average fund's return variation over time (as measured by the R squared of the regression) – the key word being "variation."

Unfortunately, this 93.6% has been widely misinterpreted. Many practitioners incorrectly believe the number means that 93.6% of a portfolio's return level (for example, a fund's 10-year annualized return) comes from a fund's asset-allocation policy. This is not true. The truth is that, in aggregate, 100% of portfolio return levels comes from asset-allocation policy.

Return 'Levels' Versus Return 'Variations'

It is imperative to distinguish between return levels and return variations. In the big picture, investors care far more about return levels than they do return variation. The often-cited 93.6% says nothing about return levels, even though that is what so many practitioners mistakenly believe. It is possible to have a high R-squared, indicating that the return variations in the asset class factors did a good job of explaining the return variations of the fund in question, yet see the weighted-average composite asset-allocation policy benchmark produce a significantly different return level from the fund in question. This is the case in BHB's study. Despite the high average 93.6% R-squared of their 91 separate time-series regressions, the average geometric annualized return of the 91 funds in their sample was 9.01% versus 10.11% for the corresponding policy portfolios.

So even though 93.6% is the number that seems to be stuck in everyone's mind, 112% (10.11% divided by 9.01%) of return levels in the study's sample came from asset-allocation policy. To put it bluntly, when it comes to return levels, asset allocation is king. In aggregate, 100% of return levels come from asset allocation before fees and somewhat more after fees. This is a mathematical truth that stems from the concept of an all-inclusive market portfolio and the fact that active management is a zero-sum game. This fundamental truth is somewhat boring; therefore, it is often lost.

[151] Brinson, G.P., Hood, L.R., & Beebower, G.L. 1986. " Determinants of Portfolio Performance." *Financial Analysts Journal,* Vol. 42, No. 4, P. 39.

Chapter 7
Company Size and Return[152]

In previous versions of the Stocks, Bonds, Bills, and Inflation® (SBBI®) Yearbook the discussions in this chapter used data from the Center for Research in Security Prices (CRSP) to demonstrate various concepts about company size and return (i.e., the "size effect"). Starting with the 2020 SBBI® Yearbook®, these concepts are demonstrated using the Morningstar/Ibbotson Associates SBBI Large Stock series and SBBI Small Stock series.[153,154,155]

One of the most remarkable discoveries of modern finance is the finding of a relationship between company size and return, generally referred to as the "size effect." The size effect is based on the empirical observation that companies of smaller size tend to have higher returns than do larger companies.

A 1981 study by Rolf Banz examined the returns of New York Stock Exchange (NYSE) small-cap companies compared to the returns of NYSE large-cap companies over the period 1926–1975.[156] What Banz found was that the returns of small-cap companies were *greater* than the returns for

---

152 This chapter is an overview of the relationship between size and return that is limited to analyzing the relative historical performance of "large-cap stocks" and "small-cap stocks," and does not include the much-expanded analyses of the "size effect" as it relates to the development of cost of equity capital found on the D&P/Kroll online Cost of Capital Navigator platform at dpcostofcapital.com. The Cost of Capital Navigator guides the Analyst through the process of estimating the cost of equity capital, a key component of any valuation analysis. The Cost of Capital Navigator includes the critical information and data from the 1999– 2021 CRSP Deciles Size Study and Risk Premium Report Study that were previously published in the *Valuation Handbook – U.S. Guide to Cost of Capital* from 2014 to 2017, and, before that, in the Ibbotson Associates/ Morningstar *Stocks, Bonds, Bills, and Inflation® (SBBI®) Valuation Yearbook* and *Risk Premium Report*, respectively, from 1999 to 2013. The valuation data and information in the Cost of Capital Navigator is the actual "as published" valuation data from those former publications. The 1999–2013 Ibbotson Associates/Morningstar size premia, industry risk premia, and other valuation data that are available within the Cost of Capital Navigator are used with permission from Morningstar, Inc. The Cost of Capital Navigator is web-based, so you can access it from your desktop, laptop, or tablet. To learn more and to purchase the Cost of Capital Navigator, visit dpcostofcapital.com.

153 The focus of the *Stocks, Bonds, Bills, and Inflation® (SBBI®) Yearbook* is to analyze the historical performance data of U.S. asset classes, as represented by the seven Morningstar/Ibbotson Associates "SBBI" series. The seven "SBBI" indices are: (i) SBBI U.S. Large Stocks, (ii) SBBI U.S. Small Stocks, (iii) SBBI U.S. Long-Term Corporate Bonds, (iv) SBBI U.S. Long-Term Government Bonds, (v) SBBI U.S. Intermediate-Term Government Bonds, (vi) SBBI U.S. 30-Day Treasury Bills, and (vii) SBBI U.S. Inflation. For detailed information about the SBBI series, see Chapter 3, "Description of the Basic Series".

154 The CRSP Deciles Size Study and Risk Premium Report Study, both of which provide size premia and other risk premia based upon data licensed from the Center for Research in Security Prices (CRSP) at the University of Chicago Booth School of Business, are fully available in the D&P/Kroll online Cost of Capital Navigator platform at dpcostofcapital.com. CRSP® is a registered trademark and service mark of Center for Research in Security Prices, LLC and has been licensed for use by D&P/Kroll. The D&P/Kroll publications and services are not sponsored, sold or promoted by CRSP®, its affiliates or its parent company. To learn more about CRSP, visit www.crsp.com.

155 A detailed discussion of company size and return (and size premia) based upon CRSP deciles 1-10 and the 10th decile split (10a, 10b, 10w, 10x, 10y, and 10z) is available in the Resources section of the Cost of Capital Navigator at dpcostofcapital.com (subscription required).

156 Rolf W. Banz, "The Relationship between Return and Market Value of Common Stocks," *Journal of Financial Economics* (March 1981): 3–18. This paper is often cited as the first comprehensive study of the size effect.

large-cap companies. Banz's 1981 study is often cited as the first comprehensive study of the size effect. Banz found that there is a significant (negative) relationship between size and historical equity returns as size decreases, returns tend to increase, and vice versa.

Possible Explanations for the Greater Returns of Smaller Companies

Some valuation analysts treat small firms as equivalent to scaled-down large firms. This is likely an erroneous assumption.

There are theoretical reasons for the greater returns of smaller companies (i.e., the "size effect"), which might include: (i) small stocks are less liquid (with higher associated transaction costs), (ii) small stocks are riskier and harder to diversify, (iii) small stocks have higher betas which tend to be underestimated, (iv) investors must do more analysis per dollar invested, (v) investment data is less available.

Valuation analysts also cite more practical reasons that small firms have risk characteristics that differ from those of large firms. For example, large firms may have greater ability to enter the market of the small firm and take market share away. Large companies likely have more resources to "weather the storm" in economic downturns. Large firms can generally spend more cash on R&D, advertising, and typically even have greater ability to hire the "best and brightest." Larger firms may have greater access to capital, broader management depth, and less dependency on just a few customers. A larger number of analysts typically follow large firms relative to small firms so there is probably more information available about large firms. Small firms have fewer resources to fend off competition and redirect themselves after changes in the market occur.[157]

Any one of these differences (not an all-encompassing list) would tend to increase investors' required rate of return to induce them to invest in small companies rather than investing in large companies.

The size effect is not without controversy, nor is this controversy something new. Traditionally, small companies are believed to have greater required rates of return than large companies because small companies are inherently riskier. It is not clear, however, whether this is due to size itself, or to other factors closely related to or correlated with size, and thus the qualification that Banz noted in his 1981 article remains pertinent today:[158,159]

"It is not known whether size [as measured by market capitalization] per se is responsible for the effect or whether size is just a proxy for one or more true unknown factors correlated with size."

---

[157] M. S. Long and J. Zhang, "Growth Options, Unwritten Call Discounts and Valuing Small Firms", *EFA 2004 Maastricht Meetings Paper no. 4057*, March 2004. Available at http://www.ssrn.com/abstract=556203.

[158] Even after controlling for size, research suggests that liquidity is still a systematic factor and a predictor of returns. See: Ibbotson, Roger G., and Daniel Y.-J Kim, "Liquidity as an Investment Style: 2018 Update," available at www.zebracapm.com Updated version of: Ibbotson, Roger G., Chen, Zhiwu, Kim, Daniel Y.-J., and Hu, Wendy Y. "Liquidity as an Investment Style," *Financial Analysts Journal*, May/June 2013.

[159] "Liquidity" is discussed in detail in Chapter 9, "Liquidity Investing."

Aspects of the Company Size Effect

The company size phenomenon is remarkable in several ways. First, the greater risk of small-cap stocks does not, in the context of the capital asset pricing model, fully account for their higher returns over the long term. In the capital asset pricing model (CAPM), only systematic, or beta risk, is rewarded; small-cap stock returns have exceeded those implied by their betas. Second, the calendar annual return differences between small- and large-cap companies are serially correlated. This suggests that past annual returns may be of some value in predicting future annual returns. Such serial correlation, or autocorrelation, is practically unknown in the market for large-cap stocks and in most other equity markets but is evident in the size premium series. Third, the size effect is seasonal. For example, small-cap stocks outperformed large-cap stocks in January in a large majority of the years. Such predictability is surprising and suspicious in light of modern capital market theory. These three aspects of the size effect – long-term returns in excess of systematic risk, serial correlation, and seasonality – will be discussed in the following sections.

The Size Effect: Empirical Evidence

Summary statistics of annual total returns for Large-Cap stocks and Small-Cap stocks are illustrated in Exhibit 7.1 over the 1926–2020 period. The differences in return between large-cap stocks and small-cap stocks is apparent.[160] For example, the annual arithmetic mean return of large-cap stocks was just over 12% over the 1926–2020 period, while the annual arithmetic mean return of small-cap stocks was just over 16%.

Note that this increased return comes at a price: risk (as measured by standard deviation) increases from just under 20% for large-cap stocks to just over 31% for small-cap stocks. The relationship between risk and return is a fundamental principle of finance. History tells us that small companies are riskier than large companies. Investors are compensated for taking on this additional risk by the higher returns provided by small companies.

---

160 Traditionally, researchers have used market value of equity (i.e., market capitalization, or simply "market cap") as a measure of size in conducting historical rate of return studies. However, market cap is not the only measure of size that can be used to predict return, nor is it necessarily the best measure of size to use. In the online D&P/Kroll Cost of Capital Navigator platform, the size effect is examined in relation to *eight* measures of company size (including market cap): (i) market capitalization, (ii) book value of equity, (iii) 5-year average net income, (iv) market value of invested capital (MVIC), (v) total assets, (vi) 5-year average EBITDA, (vii) sales, and (viii) number of employees. The Cost of Capital Navigator guides the Analyst through the process of estimating the cost of equity capital, a key component of any valuation analysis. The Cost of Capital Navigator includes the critical information and data from the 1999–2021 CRSP Deciles Size Study and Risk Premium Report Study, as published in the former *Valuation Handbook – U.S. Guide to Cost of Capital* from 2014 to 2017, and, before that, in the former Ibbotson Associates/ Morningstar *Stocks, Bonds, Bills, and Inflation® (SBBI®) Valuation Yearbook* and *Risk Premium Report*, respectively, from 1999 to 2013. The valuation data and information in the Cost of Capital Navigator is the actual "as published" valuation data from those former publications. The 1999–2013 Ibbotson Associates/Morningstar size premia, industry risk premia, and other valuation data that are available within the Cost of Capital Navigator are used with permission from Morningstar, Inc. CRSP® is a registered trademark and service mark of Center for Research in Security Prices, LLC and has been licensed for use by D&P/Kroll. The D&P/Kroll publications and services are not sponsored, sold or promoted by CRSP®, its affiliates or its parent company. To learn more about CRSP, visit www.crsp.com. To learn more and to purchase the Cost of Capital Navigator, visit dpcostofcapital.com.

Exhibit 7.1: Illustration of Summary Statistics of Large-Cap Stocks and Small-Cap Stocks (%) 1926–2020

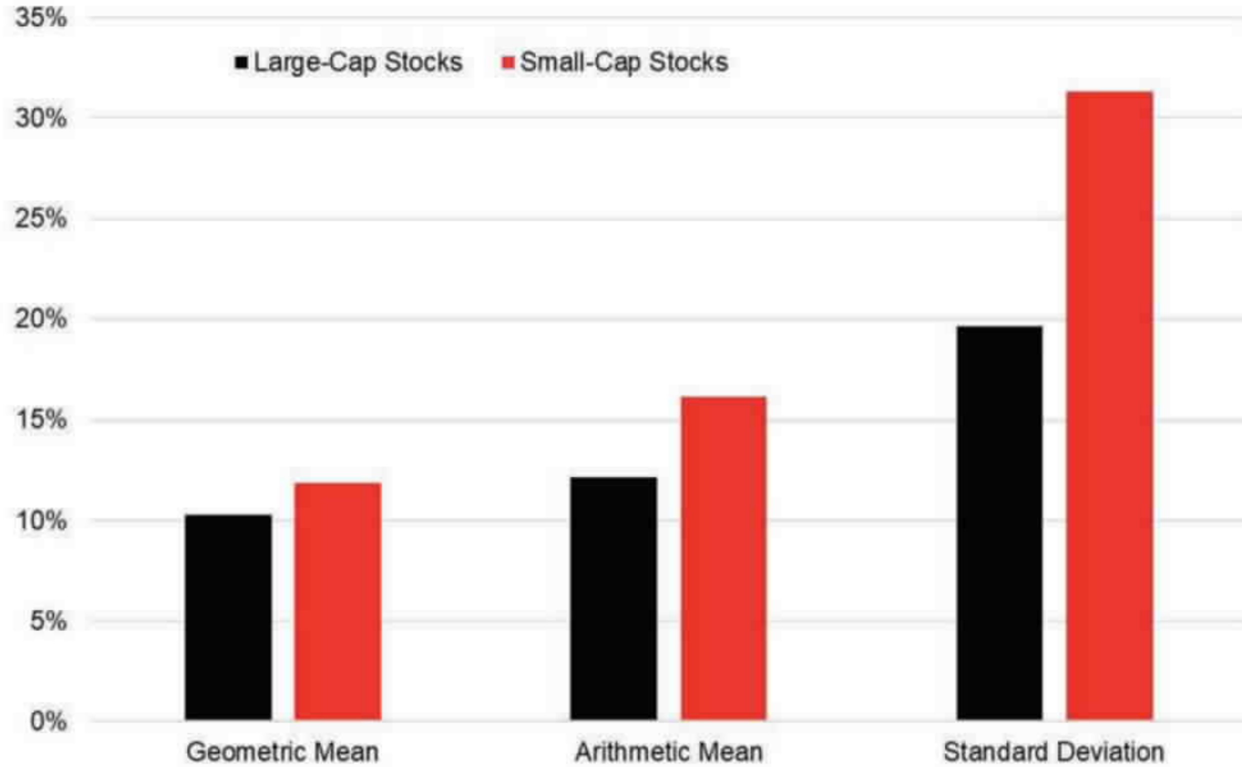

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

The differences in the performance of large-cap stocks and small-cap stocks can have important implications for investors. Exhibit 7.2 is a graphical depiction of the value of $1 invested at the end of 1925 in large-cap and small-cap stocks and held through December 31, 2020 (a total of 95 years).

Exhibit 7.2: The Value of $1 Invested in Large-Cap and Small-Cap Stocks, 1926–2020 Index (Year-end 1925 = $1.00)

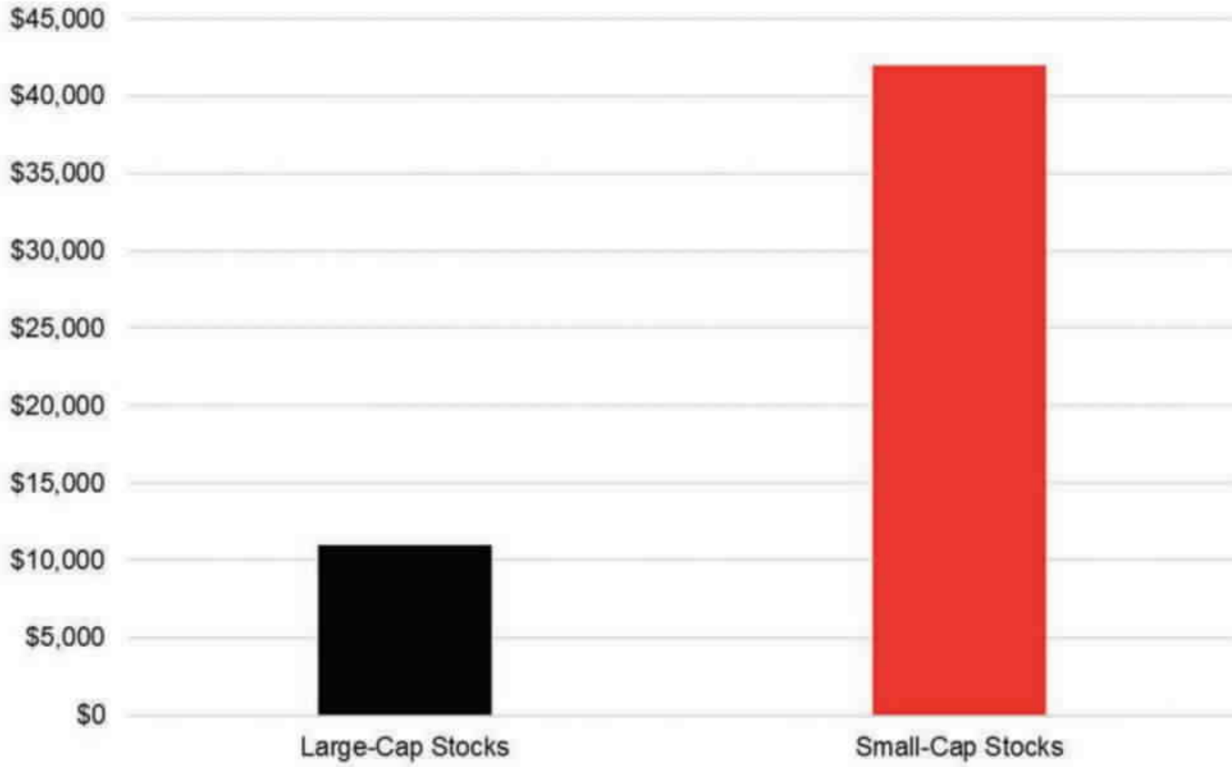

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

The geometric (i.e., compound annual) returns of these series (as illustrated in Exhibit 7.1) can be used to calculate terminal index values as follows: [161,162]

Terminal Index Value = (1 + Geometric Mean Return)^n

Where *n* is the number of periods (in this case, 95 years).

[161] For more information on calculating annual total and income returns, see Chapter 5, "Annual Returns and Indexes".

[162] *Precalculated* annualized monthly returns for each of the six SBBI series plus inflation, for each year over the 1926–2019 time horizon, are presented in table format in the full-version *2020 SBBI® Yearbook* in that book's appendices. For more information, visit dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook.

Do Small-Cap Stocks Always Outperform Large-Cap Stocks?

The increased risk faced by investors in small stocks is quite real. It is important to note, however, that the risk/return profile is over the *long-term*. The long-term expected return for any asset class can be quite different from short-term expected returns. Investors in small-cap stocks should expect losses and periods of *underperformance* relative to large-cap stocks. While this might lead some market observers to speculate that there is no size premium, statistical evidence suggests that periods of smaller stocks' underperformance should be expected. The evidence also suggests that the longer small-cap companies are given to "race" against large-cap companies, the greater the chance that small-cap companies outpace their larger counterparts.

In Exhibit 7.3, a detailed summary of the results of various "races" between small-cap companies and large-cap companies is shown, where the holding periods are limited to *exactly* 1 month, 60 months (5 years), 120 months (10 years), 240 months (20 years), and 360 months (30 years). The entire January 1926–December 2020 period is examined, as well as three more recent start date windows: April 1981–December 2020, January 1990–December 2020, and January 2000–December 2020. All three of the three more recent periods are *after* Banz wrote his March 1981 article that identified the size effect, and so they are labeled "Post Banz."[163]

In Exhibit 7.3 the number of periods examined is shown first, followed by the outperformance percentage of the total periods in parentheses.

[163] Banz, Rolf W. "The Relationship between Return and Market Value of Common Stocks". *Journal of Financial Economics* (March 1981): 3–18. Professor Banz's 1981 article is often cited as the first comprehensive study of the size effect.

Exhibit 7.3: Small-cap Companies' Performance minus Large-cap Companies' Performance Over Periods of Exactly 1, 60, 120, 240, and 360 Months January 1926–December 2020

Holding Period	All Dates Jan 1926–Dec 2020	Post Banz Apr 1981–Dec 2020	Post Banz Jan 1990–Dec 2020	Post Banz Jan 2000–Dec 2020
Exactly 1 month				
Small Stocks Outperform	568 (50%)	226 (47%)	183 (49%)	130 (52%)
Large Stocks Outperform	572 (50%)	251 (53%)	189 (51%)	122 (48%)
Exactly 60 months (5 years)				
Small Stocks Outperform	598 (55%)	184 (44%)	176 (56%)	106 (55%)
Large Stocks Outperform	483 (45%)	234 (56%)	137 (44%)	87 (45%)
Exactly 120 months (10 years)				
Small Stocks Outperform	686 (67%)	199 (56%)	199 (79%)	89 (67%)
Large Stocks Outperform	335 (33%)	159 (44%)	54 (21%)	44 (33%)
Exactly 240 months (20 years)				
Small Stocks Outperform	791 (88%)	193 (81%)	133 (100%)	13 (100%)
Large Stocks Outperform	110 (12%)	45 (19%)	0 (0%)	0 (0%)
Exactly 360 months (30 years)				
Small Stocks Outperform	752 (96%)	113 (96%)	13 (100%)	–
Large Stocks Outperform	29 (4%)	5 (4%)	0 (0%)	–

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the following Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Small-Cap Stocks: IA SBBI® US Small Stock TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

In the top row of Exhibit 7.3 (in which the holding period is restricted to a single month), large-cap companies *barely* outperformed small-cap companies in the January 1926–December 2020 period (572 versus 568 months). In the "Post-Banz" April 1981–December 2020 and January 1990–December 2020 time horizons, large-cap stocks outperformed small-cap stocks 53% and 51% of the time, respectively. In the more recent January 2000–December 2020 time horizon small-cap companies outperformed 52% of the time.

As the holding period is increased and the time that small-cap companies and large-cap companies are given to "race" against each other is lengthened, small-cap stocks tend to increasingly outperform large-cap stocks. For example, over the entire range January 1926–December 2020 (see leftmost column of Exhibit 7.3), as the holding period is increased to 60 months (5-years), to 120 months (10-years), to 240 months (20-years) and finally to 360 months (30-years), small stocks increasingly outperform large stocks (55%, 67%, 88%, and 96% of the time, respectively).

This same pattern of increasing outperformance of small stocks as the holding period is increased can also be seen in the three "Post Banz" periods.

Long-term Returns in Excess of Systematic Risk

The capital asset pricing model, or CAPM, does not fully account for the higher returns of small-cap stocks. The textbook CAPM can be expressed as follows:

$$k_e = R_f + \beta \times (RP_m)$$

Where:

k_e = Cost of equity capital

R_f = Risk-free rate

β = Beta

RP_m = Equity risk premium (also referred as ERP)

According to the CAPM, the expected return on a security should consist of the riskless rate plus an additional return to compensate for the systematic risk of the security. The return in excess of the riskless rate is estimated in the context of the CAPM by multiplying the equity risk premium by beta. The equity risk premium is the return that compensates investors for taking on risk equal to the risk of the market as a whole (systematic risk). Beta measures the extent to which a security or portfolio is exposed to systematic risk.

The beta of each decile indicates the degree to which the decile's return moves with that of the overall market. A beta greater than one indicates that the security or portfolio has greater systematic risk than the market; according to the CAPM equation, investors are compensated for taking on this additional risk.

CAPM is an attempt to predict *future* returns. CAPM can be used to see how well it would have done predicting returns that we *already know* (i.e., historical returns; returns that have already occurred). This is called "back-testing." If what "actually happened" is greater than "what CAPM would have predicted," then CAPM fell short of explaining what actually happened.

Smaller companies tend to have returns that are not fully explained by their higher betas, so return in excess of that predicted by CAPM tends to *increase* as one moves from the largest companies to the smallest companies. This size related phenomenon prompted a revision to the CAPM to include a size premium. A size premium (as used in the CAPM equation) is thus a measure of "what actually happened" minus "what textbook CAPM predicted":

Size Premium = Actual Excess Return – Excess Return Predicted by CAPM

A size premium is a common adjustment that analysts make to the textbook CAPM when developing cost of equity capital estimates.[164,165] This is sometimes referred to as the "modified" CAPM:

$$k_e = R_f + \beta \times (RP_m) + RP_s$$

Where:

k_e = Cost of equity capital

R_f = Risk-free rate

β = Beta

RP_m = Equity risk premium (also referred as ERP)

RP_s = Size Premium

The size effect is not without controversy, nor is this controversy something new. Traditionally, small companies are believed to have greater required rates of return than large companies because small companies are inherently riskier. It is not clear, however, whether this is due to size itself, or to other factors closely related to or correlated with size (e.g., liquidity).[166]

164 The CRSP Deciles Size Study and Risk Premium Report Study, both of which provide size premia and other risk premia based upon data licensed from the Center for Research in Security Prices (CRSP) at the University of Chicago Booth School of Business, are fully available in the D&P/Kroll online Cost of Capital Navigator platform at dpcostofcapital.com (subscription required). CRSP® is a registered trademark and service mark of Center for Research in Security Prices, LLC and has been licensed for use by D&P/Kroll. The D&P/Kroll publications and services are not sponsored, sold or promoted by CRSP®, its affiliates or its parent company. To learn more about CRSP, visit www.crsp.com.

165 A detailed discussion of company size and return (and size premia) based upon CRSP deciles 1-10 and the 10th decile split (10a, 10b, 10w, 10x, 10y, and 10z) is available in the Resources section of the Cost of Capital Navigator at dpcostofcapital.com. (subscription required).

166 For more information, see the Resources section of the Cost of Capital Navigator at dpcostofcapital.com. A comprehensive discussion of the size effect is in the Cost of Capital Navigator "U.S. Cost of Capital Module's" Resources section in Chapter 4, "Basic Building Blocks of the Cost of Equity Capital – Size Premium."

Chapter 8
Growth and Value Investing

Investment style can be defined broadly as an overarching description of groups of stocks or portfolios based on shared characteristics. Probably the first discussion and consideration of style concerned large-company versus small-company investing, and even this distinction was not prominent until the 1960s. Styles of investing are now broken down into more detail and used for performance measurement, asset allocation, and other purposes. Mutual funds and other investment portfolios are often measured against broad growth or value benchmarks. In some cases, investment-manager-specific style benchmarks are constructed to separate pure stock-selection ability from style effects.

Most investors agree on the broad definitions of growth and value, but when it comes to specifics, definitions can vary widely. In general, growth stocks have high relative growth rates in regard to earnings, sales, or return on equity. Growth stocks usually have relatively high price to-earnings and price-to-book ratios. Value stocks will generally have lower price-to-earnings and price-to-book values and often have higher dividend yields. Value stocks are often turnaround opportunities, companies that have had disappointing news, or companies with low growth prospects. Value investors generally believe that a value stock has been unfairly beaten down by the market, leading the stock to sell below its "intrinsic" value. Therefore, they buy the stock with the hope that the market will realize the stock's full value and eventually bid the price up to its fair value.

Fama-French Growth and Value Series

For the analysis in this chapter the Fama-French growth and value "benchmark" portfolios (discontinued) have been replaced by the Fama-French growth and value "research" portfolios.[167]

While *individual* period returns of the Fama-French growth and value research portfolios are significantly different in some cases than individual period returns of the discontinued Fama French growth and value benchmark portfolios, the *summary* statistics of the two are quite similar. Moreover, the *relative performance* of the Large Growth, Large Value, Small Growth, and Small Value remains essentially the same.[168]

---

[167] A July 2012 note on Professor Kenneth R. French's data library site states "...we believe the research factors are more useful than the benchmark factors...". The benchmark factors, which the site continued to publish, were employed by Morningstar in the *SBBI Yearbook* through the 2015 edition, and D&P/Kroll continued to using the benchmark series through the 2019 edition of the *SBBI Yearbook* (D&P/Kroll has published the *SBBI Yearbook* since 2016; for more information about the full-version *2021 SBBI® Yearbook*. visit dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook). In spring of 2019, the benchmark series were discontinued, and the "research" series replaced them replaced them in the *SBBI Yearbook*.

[168] The time horizon in previous versions of the *SBBI Yearbook* over which summary statistics for the Fama-French growth and value benchmark portfolios were calculated was 1928–present. The time horizon in this book over which summary statistics for the Fama-French growth and value research portfolios are calculated is 1927–present.

The following commentary and corresponding data make use of the Fama-French growth and value data series.[169,170]

Fama-French Index Construction Methodology[171]

The Fama-French growth and value portfolios are constructed at the end of each June. The portfolios are the intersections of two portfolios formed on size (market equity) and three portfolios formed on the ratio of book equity to market equity. The size breakpoint for year *t* is the median NYSE market equity at the end of June of year *t*. Book equity to market equity for June of year *t* is the book equity for the last fiscal year end in *t-1* divided by market equity for December of *t-1*. The book equity to market equity breakpoints are the 30th and 70th NYSE percentiles:

	Small Value	Large Value
70th Book Equity/Market Equity Percentile		
	Small Neutral	Large Neutral
30th Book Equity/Market Equity Percentile		
	Small Growth	Large Growth

The portfolios for July of year t to June of *t+1* include all NYSE, AMEX, and NASDAQ stocks for which (i) market equity data for December of *t-1* and June of *t* is available, and (ii) (positive) book equity data for *t-1*.

Historical Returns of the Fama-French Series

Using the Fama-French series ("F-F"), Exhibit 8.1 depicts the growth of $1.00 invested in F-F small-growth, F-F small-value, F-F large-growth, and F-F large-value stocks from the end of 1926 to the end of 2020. All results assume reinvestment of dividends and exclude transaction costs. The top two performers during this time period were small-value and large-value stocks, followed by small-growth and large-growth stocks. Over the period from 1927 to 2020 (94 years), small-value stocks outperformed all other stock series in the graph. One dollar invested in small-value stocks at the end of 1926 grew to over $270,000 by the end of 2020. Alternatively, one dollar

[169] Source of Fama-French growth and value return series data used in this chapter: Monthly Historical Research Returns from the Kenneth R. French data library. Returns from the Kenneth R. French data library. Fama-French growth and value "research" portfolios are revised often. To learn more visit: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. All summary calculations performed by D&P/Kroll.

[170] Eugene F. Fama is the 2013 Nobel laureate in economic sciences, the Robert R. McCormick Distinguished Service Professor of Finance at the University of Chicago Booth School of Business, and an advisory editor of the *Journal of Financial Economics*. Ken French is the Roth Family Distinguished Professor of Finance at the Tuck School of Business at Dartmouth College. Fama and French's paper "The Cross-Section of Expected Stock Returns" was the winner of the 1992 Smith Breeden Prize for the best paper in *The Journal of Finance*. See Eugene Fama and Kenneth French, "The Cross-Section of Expected Stock Returns," *Journal of Finance* (June 1992): 427–486. Also see Eugene F. Fama and Kenneth R. French, "A five factor asset pricing model," *The Journal of Financial Economics* 116 (2015): 1–22.

[171] Source: Kenneth R. French data library at: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

invested in large value, large growth, and small growth stocks at the end of 1926 grew to approximately $35,000, $7,000, and $4,000, respectively, by the end of 2020.[172]

Exhibit 8.1: F-F Small-Value Stocks, F-F Small-Growth Stocks, F-F Large-Value Stocks, and F-F Large Growth Stocks Index (Year-End 1926 = $1.00) 1927–2020

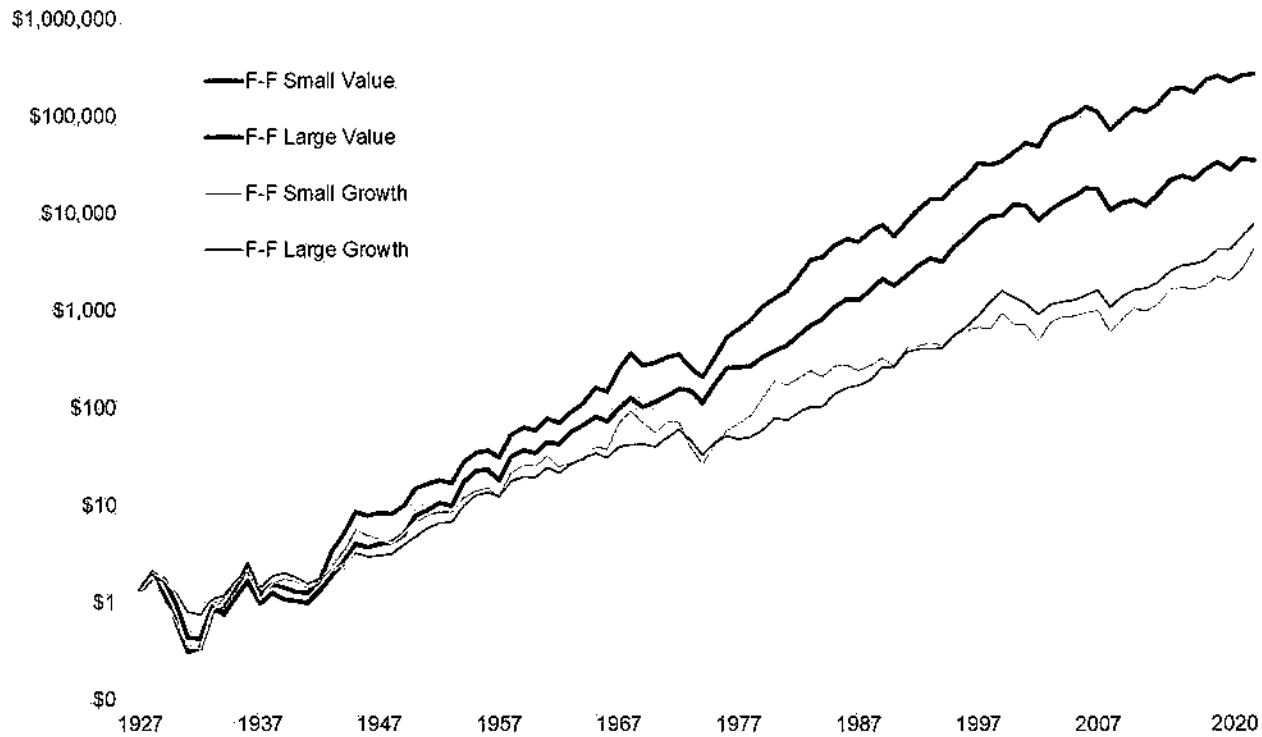

Source of underlying data: Dr. Kenneth R. French's website at: https://mba.tuck.dartmouth.edu/pages/faculty/ ken.french/data_library.html. Series used: Under "6 Portfolios Formed on Size and Book-to Market (2 x 3)": (i) Small Value, (ii) Small Growth, (iii) Big (i.e., "Large") Value, and (iv) Big (i.e., "Large") Growth.

Summary Statistics for the Fama-French Series

Exhibit 8.2 illustrates summary statistics of annual total returns for the Fama-French growth and value series from 1927 to 2020. The summary statistics are geometric mean, arithmetic mean, and standard deviation. Value significantly outperformed growth across the market capitalization spectrum. In the large-cap arena, the extra return of value over growth was at the expense of increased risk, as the standard deviation of large-value was over 26% versus approximately 20.0% for large-growth. In the small-cap series, small value significantly outperformed small-growth, and did so with the same volatility (approximately 32%).[173]

172 For more information on calculating index values, see Chapter 5, "Annual Returns and Indexes". Precalculated index values from 1927 through December 2020 are presented in the full-version *2021 SBBI® Yearbook* for the F-F growth and value series. For more information, visit dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook.

173 For more information on calculating annual returns and other summary statistics, see Chapter 5, "Annual Returns and Indexes". Precalculated summary statistics from 1927 through December 2020 are presented in the full-version *2021 SBBI® Yearbook* for the F-F growth and value series. For more information, visit dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook.

Exhibit 8.2: Illustration of Fama-French Growth and Value Series Summary Statistics of Annual Returns (%) 1927–2020

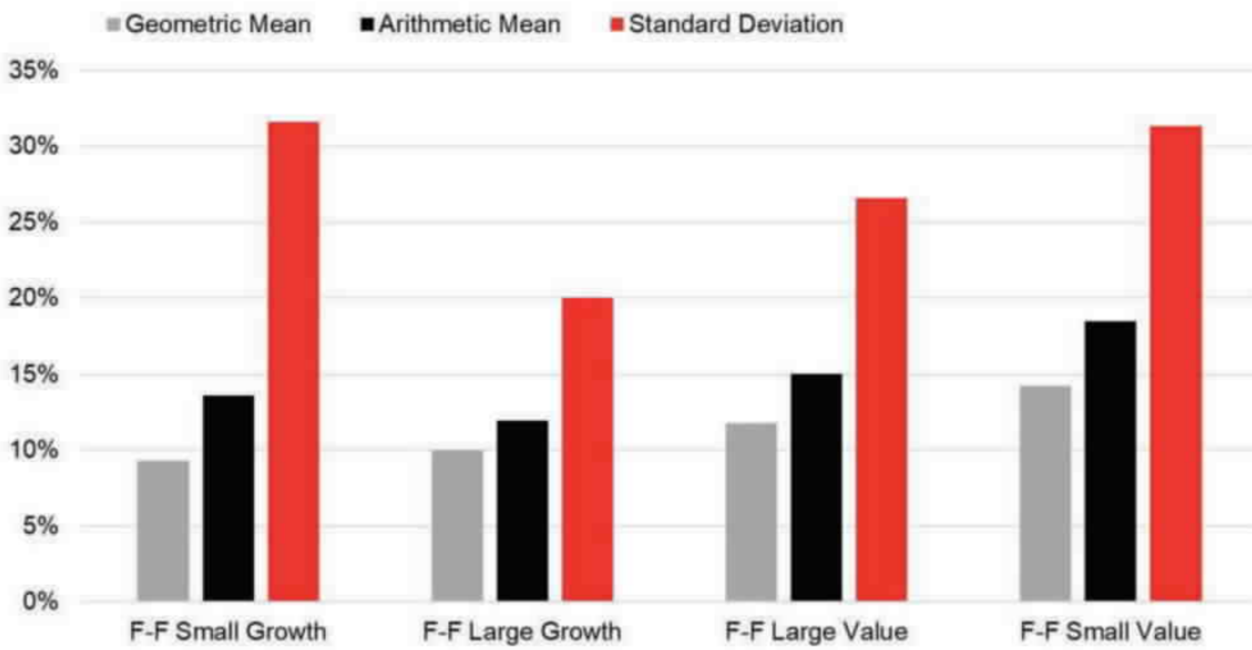

Source of underlying data: Dr. Kenneth R. French's website at: https://mba.tuck.dartmouth.edu/pages/faculty/ ken.french/data_library.html. Series used: Under "6 Portfolios Formed on Size and Book-to-Market (2 x 3)": (i) Small Value, (ii) Small Growth, (iii) Big (i.e., "Large") Value, and (iv) Big (i.e., "Large") Growth. Calculations by D&P/Kroll.

Relative Performance of the Fama-French Growth and Value Series by Decade

Exhibit 8.3 shows the relative performance of the Fama-French growth and value series by decade as measured by compound annual return (best performer at top, worst performer at bottom). Small-value stocks beat small-growth stocks in all decades except the 1930s and the 2010s. It is also interesting to note that small-value stocks were never the worst performing among all four stock series in any decade, with the exception of the 2010s.

In Exhibit 8.3, "Value" (either Large Value or Small Value) was the *best* performer or the *second* best performer 15 times. Alternatively, "Growth" (either Large Growth or Small Growth) was the *best* performer or the *second* best performer 5 times.[174]

---

174 For more information on calculating geometric (i.e., compound) returns over periods, see Chapter 5, "Annual Returns and Indexes". Precalculated compound annual returns by decade from 1927 through December 2019 are presented in the full-version *2020 SBBI® Yearbook* for the F-F growth and value series. For more information, visit dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook.

Exhibit 8.3: The Relative Performance of the Fama-French Growth and Value Series by Decade (Best Performer at Top, Worst Performer at Bottom)

1920s*	1930s	1940s	1950s
F-F Large Value	F-F Small Growth	F-F Small Value	F-F Large Value
F-F Large Growth	F-F Small Value	F-F Large Value	F-F Small Value
F-F Small Value	F-F Large Growth	F-F Small Growth	F-F Small Growth
F-F Small Growth	F-F Large Value	F-F Large Growth	F-F Large Growth

1960s	1970s	1980s
F-F Small Value	F-F Small Value	F-F Small Value
F-F Large Value	F-F Large Value	F-F Large Value
F-F Small Growth	F-F Small Growth	F-F Large Growth
F-F Large Growth	F-F Large Growth	F-F Small Growth

1990s	2000s	2010s
F-F Large Growth	F-F Small Value	F-F Large Growth
F-F Small Value	F-F Large Value	F-F Small Growth
F-F Large Value	F-F Large Growth	F-F Large Value
F-F Small Growth	F-F Small Growth	F-F Small Value

* Based on the period 1927–1929.

Source of underlying data: Dr. Kenneth R. French's website at: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. Series used: Under "6 Portfolios Formed on Size and Book-to- Market (2 x 3)": (i) Small Value, (ii) Small Growth, (iii) Big (i.e., "Large") Value, and (iv) Big (i.e., "Large") Growth. Calculations by D&P/Kroll.

Correlation of Fama-French Series

Exhibit 8.4 presents the annual cross-correlations and serial correlations for the Fama-French growth and value series.

Exhibit 8.4: Fama-French Growth and Value Series Serial and Cross-Correlations of Historical Annual Returns 1927–2020

	F-F Large Growth	F-F Large Value	F-F Small Growth	F-F Small Value	U.S.Treasury Bills	Inflation
F-F Large Growth	1.00					
F-F Large Value	0.78	1.00				
F-F Small Growth	0.81	0.77	1.00			
F-F Small Value	0.70	0.90	0.84	1.00		
U.S. Treasury Bills	-0.04	0.00	-0.11	-0.04	1.00	
Inflation	-0.06	0.05	0.00	0.04	0.42	1.00
Serial Correlation	0.02	-0.07	0.02	0.06	0.91	0.63

Source of underlying data: Dr. Kenneth R. French's website at: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. Series used: Under "6 Portfolios Formed on Size and Book-to- Market (2 x 3)": (i) Small Value, (ii) Small Growth, (iii) Big (i.e., "Large") Value, and (iv) Big (i.e., "Large") Growth. Calculations by D&P/Kroll.

Conclusion

What can explain this value effect? Readers of Benjamin Graham and David L. Dodd's "Security Analysis," first published in 1934, would say that the outperformance of value stocks is due to the market coming to realize the full value of a company's securities that were once undervalued.[175]

The Graham and Dodd approach to security analysis is to do an independent valuation of a company using accounting data and common market multiples, then look at the stock price to see if the stock is under- or overvalued. Several academic studies have shown that the market *overreacts* to bad news and *underreacts* to good news. This would lead us to conclude that there is more room for value stocks (which are more likely to have reported bad news) to improve and outperform growth stocks which already have high expectations built into them.

[175] The sixth edition of this book was published in 2008. See also Cottle, S., Murray, R.F., & Block, F.E. 1988. Graham and Dodd's *Security Analysis*, 5th ed. (New York: McGraw-Hill).

Chapter 9
Liquidity Investing

What Is Liquidity?[176]

Liquidity has many different, but similar meanings. In every case it is related to the ease of movement. Even within the context of financial markets, liquidity has several different meanings. In the banking system, liquidity measures the degree to which loans are made. In the securities markets, liquidity is the ease with which transactions can be made. In valuation, this liquidity impacts value, so that the more liquidity an asset has the more value it has, all other things being equal. The absence of liquidity lowers the value of the asset by the amount of a liquidity discount.

In this chapter, we focus on liquidity as the ease of trading securities in general, especially equities. We focus on liquidity's impact on valuation and in particular its impact on security returns. We will demonstrate that less liquid securities have higher expected returns.

Valuation as Present Value of Cash Flows

In equilibrium, an asset has a value that equals its present value, or the discounted sum of its expected cash flows. These future cash flows are unobservable except for risk-free assets. For stocks, there is great disagreement as to what these expected cash flows might be. This disagreement is the primary reason that stocks are traded. A secondary reason is that they are bought or sold to meet liquidity needs.

The other component of a present value calculation is the discount rate. Similar to the expected cash flows, these discount rates are unobservable. We can usually observe the riskless discount rates from a term structure of riskless bonds which we unravel from U.S. government discount bonds. But there are usually other premiums that we would add to the riskless term structure. The most common one is an equity risk premium which is often modified by a beta in the CAPM framework. We might also add a premium for size and another one for value (or distress). We argue here that another premium should be added for lack of liquidity.

The difference of opinion that investors have about expected cash flows leads to the additional risk of a security. The risk of the security reflects not only the changing economy and company cash-flow expectations, but also the divergence of opinion that changes from moment to moment. This risk reduces the value of a security. Ironically though, this divergence of opinion also leads to most of the trading of a security, thereby making the security more liquid for trades, whether they be active or liquidity traders. The higher liquidity increases the security's value.

[176] This chapter was written by Roger G. Ibbotson, Professor Roger G. Ibbotson, Professor Emeritus of Finance at the Yale School of Management, Chairman of Zebra Capital LLC, and former Chairman and founder of Ibbotson Associates, now part of Morningstar, Inc.

We do not mean to imply that most investors actually make these present-value calculations. Instead, investors rely on simple metrics, such as the price/earnings ratio, or PE ratio, trying to buy stocks with relatively high but unspecified cash flow projections, at relatively low PE ratios. Or they may simply feel that a stock's price is too low or high relative to its estimated value, leading them to buy or sell a security.

The Liquidity Premium

Most conventional present-value calculations ignore the liquidity premium. These calculations usually implicitly assume that securities are perfectly liquid. If they are somewhat liquid, a liquidity discount is often made to the present value, at the end of the calculation. Thus, a liquid stock is priced at the present value of the expected cash flows, discounted by the riskless rate and various other risk premiums, such as a beta-adjusted equity risk premium, a size premium, and a value premium. The final present value is then reduced by some percentage due to its lack of liquidity.

The other way to calculate a present value is to add a liquidity premium into the discount rate. Less-liquid securities would then have their cash flows discounted at higher rates. The benefit of this approach is that this liquidity premium can be thought of as causing a higher discount rate. These discount rates are equivalent, under certain conditions, to the expected return that an investor receives for investing in less-liquid securities.

The liquidity premium is the extra return an investor would demand to hold a security that cannot costlessly be traded. This premium is not exactly a risk premium since it more reflects a transaction cost. We can think of the premium as related to risk, however, because it is the risk of having to buy or sell a security quickly. The less liquid and more hurried the transaction, the higher the cost.

The liquidity premium is potentially interesting to investors who can afford to hold a security over time, instead of continuously trading it. For investors with longer-term horizons, the trading costs become trivial because they happen so infrequently. The liquidity premium is a benefit to the longer-term investor. It means that the less liquid securities will have higher returns and these higher returns are not likely to be affected by trading costs.

It is sometimes argued that part of the expected return that is demanded from real estate, private equity, or venture capital comes from their relative liquidity.[177] In addition to any of their return for other risk characteristics, investors want an extra return for holding an illiquid asset. Thus, investors would want to invest in alternative illiquid assets only if they thought they would receive extra compensation for their lack of liquidity.

The liquidity premium also is substantial within publicly traded securities. There is a difference in the return of the more highly traded securities versus the less traded securities, even though most all public securities can be readily traded. We now examine the relative impact of liquidity across

[177] Ibbotson, Roger, Siegel, Laurence B., and Diermeier, Jeffrey, "The Demand for Capital Market Returns," *Financial Analysts Journal*, January/February 1984.

publicly traded stocks on the New York Stock Exchange (NYSE), the NYSE MKT (formerly the NYSE Amex), and the NASDAQ Stock Market.

Liquidity and Stock Returns

In the U.S. stock market, liquidity has substantial impact on stock returns. We examine the monthly data for the largest 3,500 U.S. stocks by capitalization over the period 1972 through 2020. These stocks are traded on either the NYSE, the NYSE MKT, or the NASDAQ. All are publicly traded and relatively liquid, but of course some are more liquid than others.

We separate the stocks into four quartiles separated from the prior year by the turnover rate. The turnover rate is the number of shares traded during the year divided by the number of shares outstanding for the stock. The stocks with the highest turnover rates are the most liquid, and the stocks with the lowest turnover rates the least liquid. The return, share volume, and capitalization data are from the Center for Research in Security Prices at the University of Chicago Booth School of Business.

Exhibit 9.1 summarizes the results for the four liquidity quartiles. The exhibit illustrates the historical magnitude of the liquidity premium over the period from 1972–2020. Note that there is a substantial difference in the returns of the least-liquid quartile versus the most-liquid quartile, as well as a continual progression of higher returns as we move to less liquid quartiles. The less-liquid stocks are not necessarily more risky. Measured by the standard deviation, risk seems to increase with liquidity.

Exhibit 9.1: Liquidity Quartiles of the NYSE/NYSE MKT/NASDAQ, Annualized Returns (%) 1972–2020

Quartile	Geometric Mean	Arithmetic Mean	Standard Deviation
1-Less Liquid	14.41	16.11	19.39
2	14.06	15.94	20.49
3	12.39	14.58	21.86
4-More Liquid	7.83	11.20	26.77

Calculated by Zebra Capital Management at www.zebracapital.com. This is an update to the research published in Ibbotson, Roger G., and Daniel Y.-J Kim, "Liquidity as an Investment Style: 2018 Update," available at www.zebracapital.com, which itself is an updated version of the original paper: Ibbotson, Roger G., Chen, Zhiwu, Kim, Daniel Y.-J., and Hu, Wendy Y. "Liquidity as an Investment Style," *Financial Analysts Journal*, May/June 2013.

Exhibit 9.2 shows the same four quartiles of liquidity, but here presented as indexes of cumulative wealth. The quartiles consist of equally weighted portfolios with all dividends reinvested. The least-liquid quartile of stocks is at the top of the graph, and $1.00 invested at the end of 1971 grows to $733.03 by the end of 2020. One dollar invested in the second-least-liquid quartile grows to $630.51 over the period. One dollar invested in the third-least-liquid quartile (the second-most-liquid- quartile) grows to $305.42 over the period. One dollar invested at year-end 1971 into the most-liquid quartile grows to only $40.15 over the period. These large differences in terminal wealth reflect investments at different share turnover rates but include most types of companies in each liquidity quartile.

Exhibit 9.2: Wealth Indices of Investments in Low to High Quartiles of Liquidity in NYSE/NYSE MKT/ NASDAQ Stocks, Cumulative Total Returns: Index (Year-End 1971 = $1.00) 1972–2020

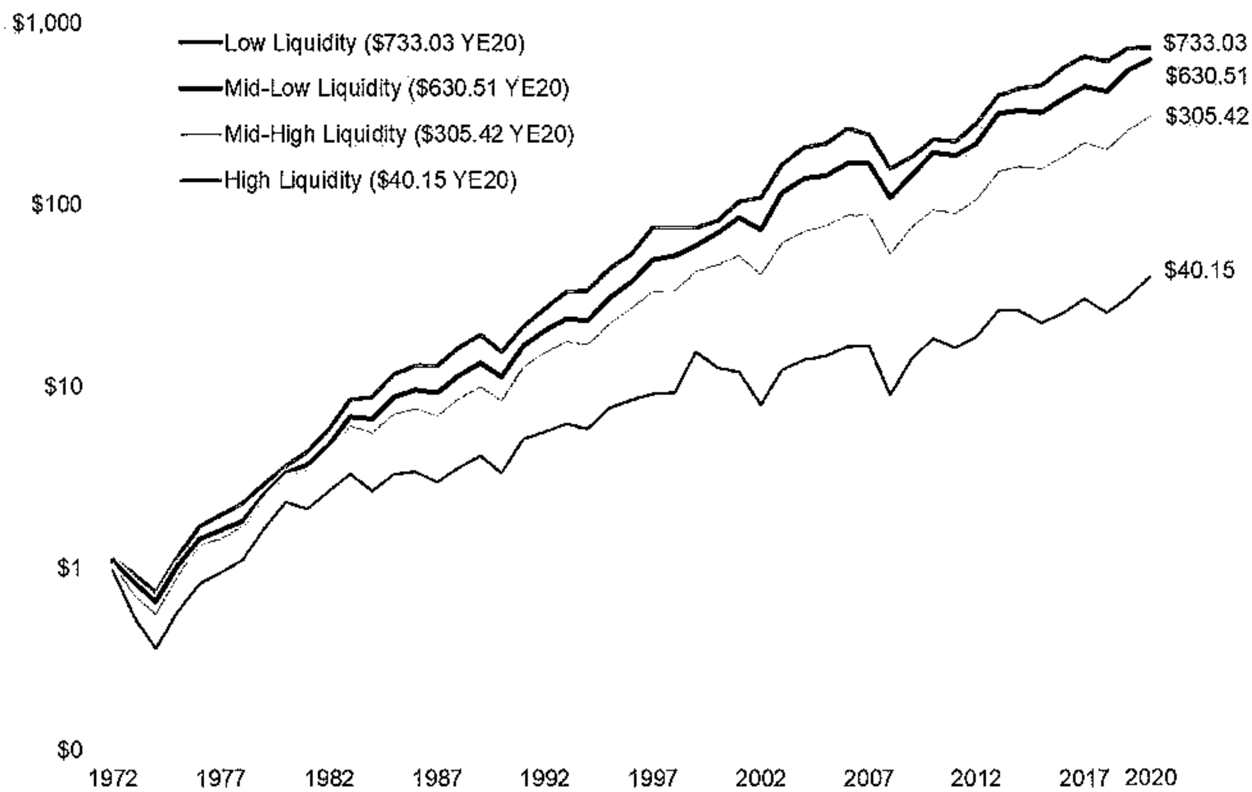

Source: Zebra Capital Management at www.zebracapital.com.

Liquidity as an Investment Style[178]

Similar to small-versus-large or value versus-growth, liquid-versus-illiquid can be viewed as an investment style. Returns are on average higher for small, value, or illiquid stocks. In this way, liquidity can be thought of as another risk factor, with a risk premium. There are some years in which each style outperforms, as well as some years of underperformance, buteach style has a *long-run* positive payoff for investing in it.

Returns on stocks typically are greater than the returns on riskless (or default-free) bonds. This extra expected return is called the equity risk premium. The styles of investing can also add or

[178] Ibbotson, Roger G., and Daniel Y.-J Kim, "Liquidity as an Investment Style: 2018 Update," available at www.zebracapital.com. Updated version of: Ibbotson, Roger G., Chen, Zhiwu, Kim, Daniel Y.-J., and Hu, Wendy Y. "Liquidity as an Investment Style," *Financial Analysts Journal*, May/June 2013.

detract from the investor's return. In fact, styles explain about half of the cross-sectional variation in equity mutual funds, with stock selection, market timing, and fees explaining the other half. Styles seem to explain more of the variation in mutual fund portfolio returns than do industry sectors.[179]

The premiums in the equity market are as follows:

- **Equity Risk Premium:** The excess return of stocks relative to risk-free (default-free) government bonds. This premium can be measured over various bond horizons, and the bonds may themselves contain a horizon premium.

- **Size Premium:** The excess return on small stocks versus the return on larger stocks.

- **Value Premium:** The excess return on value stocks versus growth stocks.

- **Liquidity Premium:** The excess return on less-liquid stocks versus more-liquid stocks.

Liquidity Versus Size

It is natural to think that liquidity and size would be related. The total number of shares of a company that are traded in a given period (say a year) are the number of shares outstanding times the turnover rate. Turnover is a measure of liquidity, adjusted for the number of shares outstanding.

Exhibit 9.3 breaks the universe of stocks into four turnover quartiles and four size-capitalization quartiles, each independently sorted. The numbers in the exhibit are the compound annual (geometric mean) rate of returns for each category. Note that small stocks tend to outperform large stocks in general, but not for the most-liquid stocks. In fact, for the most-liquid stocks shown in column four, the pattern is reversed. The poorest performing category is the highly liquid stocks that are the smallest in size (i.e., that upper-right quartile with a return of 0.33% per year).

The best-performing category is column one which represents the least-liquid stocks. The worst-performing category is column four, the most-liquid stocks. There is a clear pattern of generally decreasing returns as the liquidity of the stocks increase. The best-performing categories are the small, relatively less liquid stocks in the upper left corner of Exhibit 9.3. The Micro-Cap mid-low liquidity and Micro-Cap low liquidity categories performed the best over the 1972–2020 time horizon (15.57% and 15.13%, respectively). In previous SBBI® Yearbooks, the Micro-Cap low liquidity category performed the best, but the pattern is unchanged: small, less-liquid stocks tend to perform better than their more liquid counterparts, across all size categories.

---

[179] Xiong, James X., Roger G. Ibbotson, Thomas M. Idzorek, and Peng Chen., "The Equal Importance of Asset Allocation and Active Management," *Financial Analysts Journal,* March/April 2010.

As shown in the low-minus-high liquidity column (i.e., "Liquidity Effect (%)"), the impact of liquidity is strongest for the smallest companies and weakest for the largest companies. However, the impact of liquidity is strong and consistent across all categories. Liquidity appears to be a much better predictor of returns than is size. Note the mixed results for size shown in the bottom small-minus large row (i.e.,"Size Effect (%)").

Exhibit 9.3: Size and Liquidity Quartile Portfolios, Independently Sorted Each Year Compound Annual Returns (%) 1972–2020

	Low Liquidity	Mid-Low Liquidity	Mid-High Liquidity	High Liquidity	Liquidity Effect (%)
Micro-Cap					
Geometric Mean (%)	15.13	15.57	9.72	0.33	**14.79***
Small-Cap					
Geometric Mean (%)	14.83	14.21	12.08	6.07	**8.76**
Mid-Cap					
Geometric Mean (%)	13.55	13.53	12.96	8.46	**5.08***
Large-Cap					
Geometric Mean (%)	11.47	12.42	11.85	9.32	**2.15**
Size Effect (%)	**3.66**	**3.15**	**-2.13**	**-8.98***	

*Difference due to rounding.

Source: Compound annual returns (%) from 1972–2020. Calculated by Zebra Capital Management at www.zebracapital.com. This is an update to the research published in Ibbotson, Roger G., and Daniel Y.-J Kim, "Liquidity as an Investment Style: 2018 Update," available at www.zebracapital.com, which itself is an updated version of the original paper: Ibbotson, Roger G., Chen, Zhiwu, Kim, Daniel Y.-J., and Hu, Wendy Y. "Liquidity as an Investment Style," *Financial Analysts Journal*, May/June 2013.

Liquidity Versus Value/Growth

As noted in Chapter 8, value tends to outperform growth over time. In this chapter, less-liquid stocks are shown to outperform more liquid stocks. In this section, we examine how liquidity and value/growth interact.

The stocks are ranked by turnover rates and separated into quartiles. Similarly, the stocks are ranked by the earnings-to-price ratios and separated into quartiles. The high-earnings-to-price companies are considered value companies, while the low earnings-to-price companies are growth companies. The inverse, of course, is the PE ratio, with the growth companies having high PE ratios, and the value companies having low PE ratios.

The earnings used are the trailing reported earnings. The earnings data is from *Compustat*, owned by Standard & Poor's. The portfolios are rebalanced once per year with the earnings lagged by two months to reflect delays in compiling the accounting earnings.

Exhibit 9.4 presents the quartile results for the different levels of liquidity and value/growth. Note that both liquidity and value/growth have a strong impact on stock market returns across all categories. The results appear to be additive. There is an excess return for investing in either low-liquidity or value stocks, and the best return of all was earned by investing in the upper-left category: high-value, low liquidity stocks, which have a realized return of 18.33%. The worst category is the lower-right corner, high liquidity growth stocks, which have a return of 2.84%.

Exhibit 9.4: Summary Statistics of Value vs. Growth and Liquidity Quartile Portfolios, Independently Sorted Each Year; Compound Annual Returns (%) 1972–2020

	Low Liquidity	Mid-Low Liquidity	Mid-High Liquidity	High Liquidity	Liquidity Effect (%)
High-Value					
Geometric Mean (%)	17.60	15.99	15.46	9.57	**8.02***
Mid-Value					
Geometric Mean (%)	14.56	14.24	12.75	11.66	**2.90**
Mid-Growth					
Geometric Mean (%)	12.77	12.79	10.99	7.55	**5.22**
High-Growth					
Geometric Mean (%)	10.46	12.73	9.81	3.78	**6.68**
Size Effect (%)	**7.14**	**3.27***	**5.66***	**5.80***	

Source: Compound annual returns (%) from 1972–2020. Calculated by Zebra Capital Management at www.zebracapital.com. This is an update to the research published in Ibbotson, Roger G., and Daniel Y.-J Kim, "Liquidity as an Investment Style: 2018 Update," available at www.zebracapital.com, which itself is an updated version of the original paper: Ibbotson, Roger G., Chen, Zhiwu, Kim, Daniel Y.-J., and Hu, Wendy Y. "Liquidity as an Investment Style," *Financial Analysts Journal*, May/June 2013.

Conclusion

The results confirm that liquidity impacts returns across styles and locations. Investing in less liquid securities generates higher returns. Liquidity seems to be an investment style that is different from size or value. This result seems to hold up in almost any equity market subset and in any location.

The following section is an excerpt from a CFA Institute Research Foundation monograph entitled, Popularity: A Bridge Between Classical and Behavioral Finance *by Roger G. Ibbotson and colleagues Thomas M. Idzorek, CFA, Paul D. Kaplan, CFA, and James X. Xiong, CFA.*[180]

What's Next?

For many years, academics have sought to explain and understand asset prices, with a strong emphasis on market premiums and market anomalies. These premiums and anomalies can be explained by social or behavioral phenomenon in many settings. In a 2014 article, Roger Ibbotson and Tom Idzorek said, "Most of the best-known market premiums and anomalies can be explained by an intuitive and naturally occurring (social or behavioral) phenomenon observed in countless settings: popularity.[181]

Popularity

The existence of various market premiums and anomalies is well established in the finance literature. To date, however, no single agreed-upon explanation for them has emerged. Investment finance is largely divided into two camps, classical and behavioral. Classical finance is based mainly on the idea that investors are risk averse, so market premiums are generally interpreted as risk premiums. In behavioral finance, premiums are considered to be the result of either cognitive errors that investors systematically make or preferences for company or security characteristics that might not be related to risks. We believe that most of the best-known market premiums and anomalies can be explained by an intuitive and naturally occurring (social or behavioral) phenomenon observed in countless settings: popularity.

What Is Popularity?

Popularity is the condition of being admired, sought after, well-known, and/or accepted. A wide range of possible categories – people, food, fashion, music, places to live, types of pet, vacation destinations, television shows, and so on – contain an implicit popularity spectrum or rank. Each of the categories has various criteria for estimating popularity.

For our purposes, the quality of the ranking criteria is not important; what is important is that any given category comprises a natural ordering in which some constituents are more popular than others. Such relative popularity evolves over time. Some aspects of popularity are systematic, or more or less permanent (for example, modern society seems to prefer thin to fat, tall to short). Other aspects of popularity may be transitory or exist only as fads (for example, necktie width, high-waisted jeans, men wearing wigs). Whether the result of systematic trends or idiosyncratic evolution, these rankings are in flux. Some popular items become relatively less popular, and

[180] Available for download at: https://www.cfainstitute.org/en/research/foundation/2018/popularity-bridge-between-classical-and-behavioral-finance, or go to the CFA Institute website at cfainstitute.org and search for "popularity". Copyright 2018, CFA Institute Research Foundation. Reproduced from *Popularity: A Bridge between Classical and Behavioral Finance* with permission from CFA Institute Research Foundation. All rights reserved.

[181] Ibbotson, R.G., Idzorek, T.H. "Dimensions of Popularity," *Journal of Portfolio Management*, Vol. 40 No. 5, (Special 40th Anniversary Issue 2014), P. 68–74.

some of the unpopular items become relatively more popular. While unsustainable, some popular items will temporarily become even more popular. For example, liquidity is permanently popular, but on a relative basis during times of market distress, it is especially sought after. Society places a greater relative value (monetary or otherwise) on the more popular items.

In *Popularity: A Bridge Between Classical and Behavioral Finance*, popularity refers to investor preferences – that is, how much an asset is liked or disliked. Of course, the primary preference for investors is to seek returns. Investors do not know what the returns will be, but they can distinguish one asset from another in terms of their observable characteristics, for which they may have clearly defined preferences. Thus, even with the same set of expected cash flows, investors may have more demand for one asset over another, which gives the preferred asset a higher current price and a lower expected return. An asset could be liked (or disliked) for *rational* or *irrational* reasons.[182]

In this way, popularity spans ideas from both classical and behavioral finance, thus providing a bridge between the two camps.

In classical finance, the primary preference, beyond maximizing expected return, is to take less risk. This fact has given rise to various models that usually assume no other preferences. In the most well-known model, the capital asset pricing model (CAPM), the only "priced" characteristic is exposure to undiversifiable market risk. We consider a broader set of preferences that lead to other priced characteristics, which might include the rational preferences to reduce catastrophic losses, increase liquidity, be tax efficient, and so on. We also consider preferences that might be more in line with what the literature considers "behavioral," such as desiring to hold companies with strong brands, investments with strong past price increases, or companies that have strong ESG (environmental, social, and governance) characteristics.

The popularity framework presented in *Popularity: A Bridge Between Classical and Behavioral Finance,* includes a generalization of a wide range of characteristics in classical finance and behavioral finance that influence how investors value securities. We can classify these characteristics into two broad categories with two subcategories each as follows:

Classical Finance

- **Risks.** In classical finance, risk usually refers to fluctuations in asset values, but risk can be interpreted more broadly as any risks to which a rational investor, who assumes away any real-world frictions in the holding and trading of securities, would be averse. Thus,

[182] Throughout *Popularity: A Bridge Between Classical and Behavioral Finance*, we describe preferences, or the reasons for preferences, as being either rational or irrational. Rational reasons for preferences are those considered in classical finance, broadly defined. The reasons include expected returns, risk, liquidity, taxes, and trading costs. Generally, rational preferences are pecuniary. Irrational reasons for preferences generally are those identified in behavioral finance and result from the various biases and heuristics identified in that literature. Irrational preferences are generally nonpecuniary. Although Ibbotson, Diermeier, and Siegel (1984) acknowledged the possibility of nonpecuniary security characteristics playing a role in asset pricing (such as in the art market), their focus was on pecuniary characteristics that we consider to be subject to rational preferences. Our popularity framework extends their idea to irrational preferences.

risks may be multidimensional, including various types of stock or bond risks, or may arise from catastrophic events.

- **Frictional**. These characteristics are often assumed away in classical finance, but a rational investor would consider them. Examples include taxes, trading costs, and asset divisibility.

Behavioral Finance

- **Psychological**. Investors consider these characteristics because of their psychological impact. For example, buying a company with a small carbon footprint might make an investor feel good.

- **Cognitive**. Investors consider these factors or fail to accurately interpret such factors because of systematic cognitive errors. For example, investors may overvalue the importance of a company's brand when evaluating its stock because they do not realize that the value of the brand is already embedded in the market price of the stock.

Our fourfold classification of security characteristics partially overlaps with the threefold classification in Statman (2017), in which investors are described as holding securities for utilitarian, expressive, and emotional reasons. Utilitarian reasons correspond risk and frictional characteristics, and expressive and emotional reasons correspond to psychological characteristics.

In *Popularity: A Bridge Between Classical and Behavioral Finance*, we focus primarily on the stock market, although we believe the concepts can be applied to fixed-income securities, real estate, and numerous other real assets. Periodically, as necessary, we attempt to distinguish between characteristics of a company and characteristics of the security in question – both of which can have attributes that are more or less popular among investors. Assets are priced not only by their expected cash flows but also by the popularity of the other characteristics associated with the company or security. The less popular stocks have lower prices (relative to the expected discounted value of their cash flows), thus higher expected returns. The more popular stocks have higher prices and, therefore, lower expected returns. Popularity can be related to risk (an unpopular characteristic), and it can also be related to other rational preferences. But popularity can also be related to behavioral concepts. For instance, investors may want to brag about their past winners (or purchase recent winners – for example, in the practice called "window dressing") or hold recognizable securities that are consistent with their social values. Any aspect that can affect the popularity of a stock will affect its demand and thus its price.[183]

Popularity is a bridge between classical finance and behavioral finance because both types of finance rely on preferences. Popularity is an expression of these preferences, whether they are\

[183] By demand, we mean the sum of the demand of all market participants.

rational, irrational, or somewhere in between.[184] Popularity does not make a value judgment but, instead, takes preferences as a given and recognizes that preferences can change over time. *Popularity: A Bridge Between Classical and Behavioral Finance* is presented in an equilibrium framework, so asset prices and expected returns reflect the aggregate impact of investor preferences.

[184] The same preference may be rational for one investor and irrational for another investor. For example, it is rational for a taxable investor to consider tax efficiency and irrational for nontaxable investor to seek out tax efficient investments.

Chapter 10
Using Historical Data in Wealth Forecasting and Portfolio Optimization

When forecasting the return on an asset or a portfolio, investors are (or should be) interested in the entire probability distribution of future outcomes, not just the mean or "point estimate." An example of a point estimate forecast is "large-cap stocks will have a return of 12% in 2021." It is more helpful to know the uncertainty surrounding this point estimate than to know the point estimate itself. One measure of uncertainty is standard deviation. The large cap stock return forecast can be expressed as 12% representing the mean and 20% representing the standard deviation.[185]

If the returns on large-cap stocks are normally distributed, the mean (expected return) and the standard deviation provide enough information to forecast the likelihood of any return. Suppose one wants to ascertain the likelihood that large-cap stocks will have a return of -25% or lower in 2021. Given the above example, a return of -25% is [12 - (-25)] / 20 = 1.9 standard deviations below the mean. The likelihood of an observation of 1.9 or more standard deviations below the mean is 2.9%. This can be looked up in any statistics textbook in the table showing values of the cumulative probability function for a normal distribution. Thus, a likelihood that the stock market will fall by 25% or more in 2020 is 2.9%. This is valuable information, both to the investor who believes that stocks are a sure thing and to the investor who is certain that they will crash tomorrow.

However, historical stock returns are not exactly normally distributed, and a slightly different method needs to be used to make accurate probabilistic forecasts. A description of the model used to forecast the distribution of stock returns appears later in this chapter.

Probabilistic Forecasts

Probabilistic forecasts might seem to be too wide to be useful – the most widely quoted forecasters, after all, sometimes make very specific predictions. However, the forecast of a probability distribution actually reveals much more than the point estimate. The point estimate reflects what statisticians call an "expected value," but the actual return will likely be higher or lower than the point estimate. By knowing the extent to which actual returns are likely to deviate

[185] *Precalculated* summary statistics of annual returns (1926–2020) are presented in table format in the full-version *2021 SBBI® Yearbook* for the following Ibbotson Associates(IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: Large-Cap Stocks (total return, income return, and capital appreciation return), Small-Cap Stocks (total return), Long-term Corporate (i.e., 20-year) Bonds (total return), Long-term (i.e., 20-year) Government Bonds (total return, income return, and capital appreciation return), Intermediate-term (5-year) Government Bonds (total return, income return, and capital appreciation return), (30-day) U.S. Treasury Bills (total return), and Inflation. For more information, visit dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook.

from the point estimate, the investor can assess the risk of every asset, and thus compare investment opportunities in terms of their risks as well as their expected returns. As Harry Markowitz showed nearly a half century ago in his Nobel Prize-winning work on portfolio theory, investors care about avoiding risk as well as seeking return. Probabilistic forecasts enable investors to quantify these concepts.

The Lognormal Distribution

In the lognormal model, the natural logarithms of asset return relatives are assumed to be normally distributed. A return relative is one plus the return. That is, if an asset has a return of 15% in a given period, its return relative is 1.15 (1 + 0.15).

The lognormal distribution is skewed to the right. This means that the expected value, or mean, is greater than the median. Furthermore, if return relatives are lognormally distributed, returns cannot fall below negative 100%. These properties of the lognormal distribution make it a more accurate characterization of the behavior of market returns than does the normal distribution.

In all normal distributions, moreover, the probability of an observation falling one standard deviation below the mean equals the probability of falling one standard deviation above the mean; each has a probability of about 34%. In a lognormal distribution, these probabilities differ and depend on the parameters of the distribution.

Forecasting Wealth Values and Rates of Return

Using the lognormal model, it is fairly simple to form probabilistic forecasts of both compound rates of return and ending period wealth values. Wealth at time *n* (assuming reinvestment of all income and no taxes) is:

$$W_n = W_0(1+r_1)(1+r_2)\ldots(1+r_n)$$

Where:

W_n	=	The wealth value at time *n*
W_0	=	The initial investment at time 0
r_1, r_2, *etc.*	=	The total returns on the portfolio for the rebalancing ending at times 1, 2, and so forth

The compound rate of return or geometric mean return over the same period, r_G, is:

$$r_G = \left(\frac{W_n}{W_0}\right)^{\frac{1}{n}} - 1$$

Where:

r_G = The geometric mean return

W_n = The ending period wealth value at time n

W_0 = The initial wealth value at time 0

n = The inclusive number of periods

By assuming that all $(1+r_n)$ values are lognormally distributed with the same expected value and standard deviation and are all statistically independent of each other, it follows that W_n and $(1+r_G)$ are lognormally distributed. In fact, even if the $(1+r_n)$ values are not themselves lognormally distributed but are independent and identically distributed, W_n and $(1+r_G)$ are approximately lognormal for large enough values of *n*. This "central-limit theorem" means that the lognormal model can be useful in long-term forecasting even if short term returns are not well described by a lognormal distribution.

Calculating Parameters of the Lognormal Model

To use the lognormal model, we must first calculate the expected value and standard deviation of the natural logarithm of the return relative of the portfolio. These parameters, denoted *m* and *s* respectively, can be calculated from the expected return (*m*) and standard deviation (*s*) of the portfolio as follows:

$$m = \ln(1+\mu) - \left(\frac{s^2}{2}\right)$$

$$s = \sqrt{\ln\left[1 + \left(\frac{\sigma}{1+\mu}\right)^2\right]}$$

Where:

ln = The natural logarithm function

To calculate a particular percentile of wealth or return for a given time horizon, the only remaining parameter needed is the z score of the percentile. The z-score of a percentile ranking is that percentile ranking expressed as the number of standard deviations that it is above or below the mean of a normal distribution. For example, the z-score of the 95th percentile is 1.645 because in a normal distribution, the 95th percentile is 1.645 standard deviations above the 50th percentile or median, which is also the mean. Z-scores can be obtained from a table of cumulative values of the standard normal distribution or from software that produces such values.

Given the logarithmic parameters of a portfolio (*m* and *s*), a time horizon (*n*), and the z-score of a percentile (*z*), the percentile in question in terms of cumulative wealth at the end of the time horizon (W_n) is:

$$e^{\left(mn+zs\sqrt{n}\right)}$$

Similarly, the percentile in question in terms of the compound rate of return for the period (r_G) is:

$$e^{\left(m+z\frac{s}{\sqrt{n}}\right)}-1$$

Mean-Variance Optimization

One important application of the probability forecasts of asset returns is mean variance optimization. Optimization is the process of identifying portfolios that have the highest possible expected return for a given level of risk or the lowest possible risk for a given expected return. Such a portfolio is considered “efficient,” and the locus of all efficient portfolios is called the efficient frontier. A simple two-asset efficient frontier constructed from large-cap stocks and U.S. Treasury Bills is shown in Exhibit 10.1. All investors should hold portfolios that are efficient with respect to the assets in their opportunity set.

Exhibit 10.1: Efficient Frontier; Large-cap Stocks and U.S. Treasury Bills 1926–2020

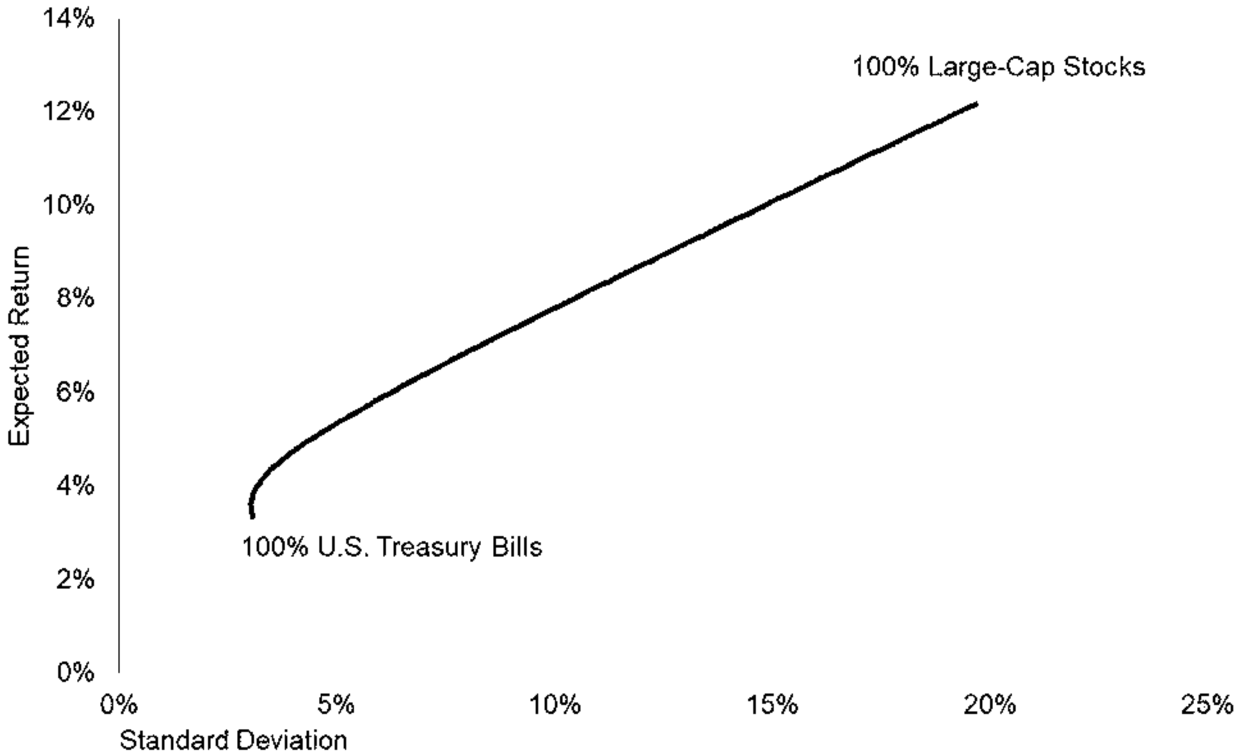

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, and (ii) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

The most widely accepted framework for optimization is Markowitz, or mean-variance, optimization (MVO), which makes the following assumptions: (i) The forecast mean, or expected return, describes the attribute that investors consider to be desirable about an asset; (ii) The risk of the asset is measured by its expected standard deviation of returns; and (iii) The interaction between one asset and another is captured by the expected correlation coefficient of the two assets' returns. MVO thus requires forecasts of the return and standard deviation of each asset and the correlation of each asset with every other asset.[186]

In the 1950s, Harry Markowitz developed both the concept of the efficient frontier and the mathematical means of constructing it (mean-variance optimization).[187]

[186] The standard deviation is the square root of the variance; hence the term "mean-variance" in describing this form of the optimization problem.

[187] Markowitz, H.M. 1959. *Portfolio Selection: Efficient Diversification of Investments* (New York: John Wiley & Sons).

Estimating Returns, Risks, and Correlations

To simulate future probability distributions of asset and portfolio returns, one typically estimates parameters of the historical return data. The parameters that are required to simulate returns on an asset are its mean and standard deviation. To simulate returns on portfolios of assets, one must also estimate the correlation of each asset in the portfolio with every other asset. Thus, the parameters required to conduct a simulation are the same as those required as inputs into a mean-variance optimization.[188] The techniques used to estimate these parameters are described below.

Means, or Expected Returns

The mean return (forecast mean or expected return) on an asset is the probability-weighted average of all possible returns on the asset over a future period. Estimates of expected returns are based on models of asset returns. While many models of asset returns incorporate estimates of gross national product, the money supply, and other macroeconomic variables, the model employed in this chapter does not. This is because we assume (for the present purpose) asset markets are informationally efficient, with all relevant and available information fully incorporated in asset prices. If this assumption holds, investor expectations (forecasts) can be discerned from market-observable data. Such forecasts are not attempts to outguess, or beat, the market. They are attempts to discern the market's expectations, i.e., to read what the market itself is forecasting.

For some assets, expected returns can be estimated using current market data alone. For example, the yield on a riskless bond is an estimate of its expected return. For other assets, current data are not sufficient. Stocks, for example, have no exact analogue to the yield on a bond. In such cases, we use the statistical time series properties of historical data in forming the estimates.

To know which data to use when estimating expected returns, we need to know the rebalancing frequency of the portfolios and the investment horizon. In our example we will assume an annual rebalancing frequency and a 20-year planning horizon. The rebalancing frequency gives the time units in which returns are measured.

With a 20-year horizon, the relevant riskless rate is the yield on a 20-year coupon bond. This riskless rate is the baseline from which the expected return on every other asset class is derived by adding or subtracting risk premiums.

[188] It is also possible to conduct a simulation using entire data sets, that is, without estimating the statistical parameters of the data sets. Typically, in such a nonparametric simulation, the frequency of an event occurring in the simulated history is equal to the frequency of the event occurring in the actual history used to construct the data set.

Large-Cap Stocks

The expected return on large-cap stocks is the riskless rate, plus the expected risk premium of large-cap stocks over bonds that are riskless over the investment horizon.

Bonds and Bills

For default-free bonds with a maturity equal to the planning horizon, the expected return is the yield on the bond. For bonds with other maturities, the expected bond horizon premium should be added to the riskless rate (for longer maturities) or subtracted from the riskless rate (for shorter maturities). Because expected capital gains on a bond are zero, the expected horizon premium is estimated by the historical average difference of the income returns on the bonds.[189]

For U.S. Treasury Bills, the expected return over a given time horizon equals the expected return on a Treasury bond of a similar horizon, less the expected horizon premium of bonds over bills. The long-term horizon premium is estimated by the historical average of the difference of the income return on bonds and the return on bills This forecast typically differs from the current yield on a U.S. Treasury Bill because a portfolio of U.S. Treasury Bills is rolled over (the proceeds of maturing bills are invested in new bills, at yields not yet known) during the time horizon described.

Standard Deviations

Standard deviations can be estimated from historical data as described in Chapter 6. There is no evidence of a major change in the variability of returns on large-cap stocks, so the entire period 1926 to 2020 can be used to estimate the standard deviation of these asset classes. For long-term government bonds and Treasury bills, the period 1970 to 2020 can be used to estimate these inputs. This is because the departure from the Bretton Woods fixed-rate currency exchange agreement in the early 1970s caused a structural shift in the U.S. interest-rate environment. As a result, bond volatility spiked and has remained well above levels experienced before the regime shift, rendering pre-1970 risk comparisons inappropriate.

Correlations

Correlations between the asset classes are estimated from historical data as described in Chapter 6. Correlations between major asset classes change over time. Exhibit 10.2 shows the historical correlation of annual returns on large-cap stocks and long-term government bonds over 20-year rolling periods from 1926–1945 through 2000–2020.

---

[189] The expected capital gain on a par bond is self-evidently zero. For a zero coupon (or other discount) bond, investors expect the price to rise as the bond ages, but the expected portion of this price increase should not be considered a capital gain. It is a form of income return.

Exhibit 10.2: 20-Year Rolling-Period Correlations of Annual Returns of Large-Cap Stocks and Long-term Government Bonds 1926–2020

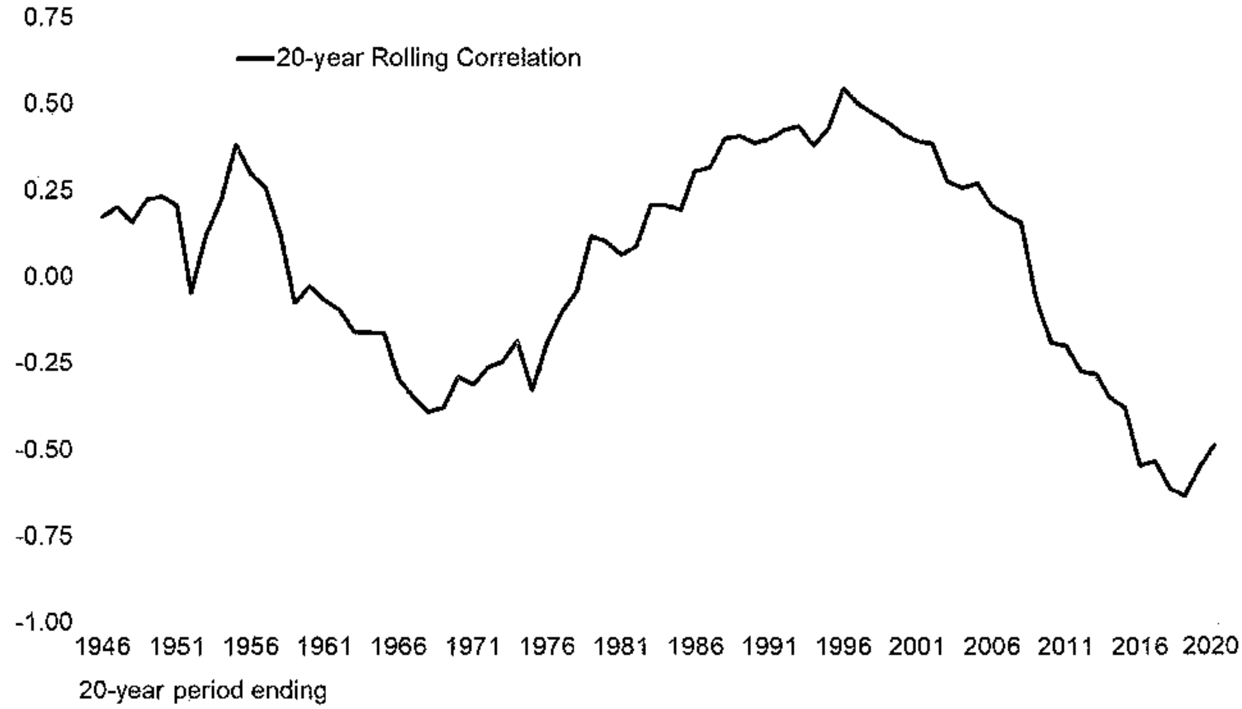

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, and (ii) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Using Inputs to Form Other Portfolios

In Exhibit 10.3, inputs are provided that can be used in forming portfolios.[190]

[190] *Precalculated* "Building Blocks for Expected Return Construction" are presented in table format in the full-version *2020 SBBI® Yearbook* as of December 31, 2020 for the following: (i) **Yields** (Long-term (20-year) U.S. Treasury Coupon Bond Yield; Intermediate-term (5-year) U.S. Treasury Coupon Note Yield; Short-term (30-day) U.S. Treasury Bill Yield), (ii) **Fixed Income Risk Premiums** (Expected default premium; Expected long-term horizon premium; Expected intermediate-term horizon premium, (iii) **Equity Risk Premiums** (Long-horizon expected equity risk premium; Intermediate-horizon expected equity risk premium; Short-horizon expected equity risk premium; Small-cap premium). For more information, visit: dpcostofcapital.com/stocks-bonds-bills-inflation-sbbi-yearbook.

Exhibit 10.3: Optimization Inputs: Year-end 2020 Large-Cap Stocks, Long-term Government Bonds, and U.S. Treasury Bills (%)

			Correlation		
	Expected Return (%)	**Standard Deviation (%)**	**Stocks**	**Bonds**	**Bills**
Stocks	8.6	19.7	1.00		
Bonds	1.4	11.7	0.01	1.00	
Bills	0.1	3.4	-0.02	0.17	1.00

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Long-Term (i.e. 20-year) Government Bonds: IA SBBI® US LT Govt TR USD, and (iii) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Given a complete set of inputs, the expected return and standard deviation of any portfolio (efficient or other) of the asset classes can be calculated. The expected return of a portfolio is the weighted average of the expected returns of the asset classes:

$$r_p = \sum_{i=1}^{n} x_i r_i$$

Where:

r_p = The expected return of the portflio p

n = The number of asset classes

x_i = The portfolio weight of asset class i, scaled such that:

$$\sum_{i=1}^{n} x_i = 1$$

Where:

r_i = The expected return of asset class i

For example, referring to the inputs in Exhibit 10.3, a portfolio comprised of large-cap stocks only would have an expected return of 9.4% and a standard deviation of 19.8%. If the portfolio mix is changed to, say, 60.0% large-cap stocks, 35.0% long-term government bonds, and 5.0% U.S. Treasury Bills, the expected return of this new portfolio mix can be calculated by applying the above formula (again, using the inputs in Exhibit 10.3):

5.7% = (60.0% x 8.6%) + (35.0% x 1.4%) + (5.0% x 0.1%)

The standard deviation of the portfolio depends not only on the standard deviations of the asset classes, but also on all of the correlations. It is given by:

$$\sigma_p = \sqrt{\sum_{i=1}^{n} \sum_{i=1}^{n} x_i x_j \sigma_i \sigma_j p_{ij}}$$

Where:

σ_p	=	The standard deviation of the portfolio
x_i and x_j	=	The portfolio weights of asset classes *i* and *j*
σ_i and σ_j	=	The standard deviations of returns on asset classes *i* and *j*
p_{ij}	=	The correlation between returns on asset classes *i* and *j*
		(note that r_{ij} equals one and that r_{ij} is equal to r_{ji}).

The standard deviation of the new portfolio (60.0% large-cap stocks, 35.0% long-term government bonds, and 5.0% U.S. Treasury Bills) can be calculated using the inputs from Exhibit 10.3 as shown in Exhibit 10.4:

Exhibit 10.4: Calculation of Example Portfolio Comprised of 60.0% Large-cap stocks, 35.0% Long-term Government Bonds, and 5.0% U.S. Treasury Bills

Stocks (asset class 1)	**Bonds (asset class 2)**	**Bills (asset class 3)**
Stocks & Stocks	Stocks & Bonds	Stocks & Bills
$x_1^2\sigma_1^2 p_{1,1}$ =	$x_1 x_2 \sigma_1 \sigma_2 p_{1,2}$ =	$x_1 x_3 \sigma_1 \sigma_3 p_{1,3}$ =
$0.60^2 \times 0.197^2 \times 1.00$ =	$0.60 \times 0.35 \times 0.197 \times 0.117 \times 0.015$ =	$0.60 \times 0.05 \times 0.197 \times 0.034 \times -0.024$ =
=0.013923	=0.000070	=-0.000005
Bonds & Stocks	Bonds & Bonds	Bonds & Bills
$x_1 x_2 \sigma_1 \sigma_2 p_{1,2}$ =	$x_2^2\sigma_2^2 p_{2,2}$ =	$x_2 x_3 \sigma_2 \sigma_3 p_{2,3}$ =
$0.35 \times 0.60 \times 0.117 \times 0.197 \times 0.015$ =	$0.35^2 \times 0.117^2 \times 1.00$ =	$0.35 \times 0.05 \times 0.117 \times 0.034 \times 0.168$ =
=0.000070	=0.001689	=0.000012
Bills & Stocks	Bills & Bonds	Bills & Bills
$x_1 x_3 \sigma_1 \sigma_3 p_{1,3}$ =	$x_2 x_3 \sigma_2 \sigma_3 p_{2,3}$ =	$x_3^2\sigma_3^2 p_{3,3}$ =
$0.05 \times 0.60 \times 0.034 \times 0.197 \times -0.024$ =	$0.05 \times 0.35 \times 0.034 \times 0.117 \times 0.168$ =	$0.05^2 \times 0.034^2 \times 1.00$ =
=-0.000005	=0.000012	=0.000003

By summing these terms and taking the square root of the total, the result is a standard deviation of 12.6%.

$$\sqrt{\begin{array}{l}0.013923+0.000070+-0.000005+\\0.000070+0.001689+0.000012+\\-0.000005+0.000012+0.000003\end{array}}=12.6\%$$

Enhancements to Mean-Variance Optimization

Ibbotson Associates was an early adopter of mean-variance optimization to develop asset class model guidelines and continues to assist the industry in the development of enhancements to the traditional mean-variance approach as well as the state-of-the-art techniques described later in the chapter. Over the last half century, the Markowitz mean-variance optimization (MVO) framework has become the textbook approach for creating these optimal asset allocations, but the approach has several shortcomings.

Shortcomings of Traditional Optimization Techniques

One notable shortcoming is that the output (optimal asset allocation weights) is very sensitive to the inputs (expected returns, standard deviations, and correlations). Input sensitivity often leads to highly concentrated allocations in only a small number of the available asset classes. For example, if a typical optimization starts with an opportunity set of about 10 asset classes, just a few of these asset choices might end up in the resulting optimal allocation with the remaining asset choices not even getting a mention.

Mean-variance optimization is a powerful tool, but it needs to be used with caution. For instance, basing mean-variance optimization inputs on shorter periods can contribute to extreme results. Basing the mean-variance optimization inputs on longer periods, such as those presented elsewhere in this book, can help mitigate the extreme asset allocations mixes. Also, there is usually a more consistent ratio of return to risk amongst the different asset classes when using longer periods.

Placing maximum and minimum allocation constraints on each asset is the most common solution to the problem of highly concentrated asset allocations. For instance, we could specify a minimum allocation of 5% and a maximum allocation of 15% for each of the nine asset choices. This would ensure that each asset gets represented in the final allocation and that no single asset completely dominates in the final allocation mix. Unfortunately, these artificial minimums and maximums are arbitrary and usually end up limiting the ability of the optimizer to properly act on the information contained in the inputs.

Black-Litterman and Resampling Techniques

Two popular enhancements to traditional optimization techniques have emerged in recent years that can help overcome these difficulties. While both of these methods can help develop well diversified asset allocations, they approach the problem in very different ways. The first of these, the Black-Litterman model, attempts to create better inputs. The second, resampled mean variance optimization, attempts to build a better optimizer.

The Black-Litterman model was created by Fischer Black and Robert Litterman in the late 1980s. The Black-Litterman model combines investors' views regarding expected returns and the expected returns predicted by the capital asset pricing model to form a single blended estimate of expected returns. When this new combined estimate is used as an input within a traditional mean-variance optimization framework, it produces well-diversified portfolios that include not only market-based asset allocations but also allocations in assets that received favorable views.

The second approach, resampled mean-variance optimization, grew out of the work of a number of authors, but it is most closely associated with the work of Richard Michaud. While traditional mean-variance optimization treats the capital market assumptions as if they were known with complete certainty, resampled mean-variance optimization recognizes that the capital market assumptions are forecasts and are therefore not known with complete certainty.

Conceptually, resampled mean-variance optimization is a combination of Monte Carlo simulation and the more traditional Markowitz mean-variance optimization approach.[191] The simulation randomly resamples possible returns from a forecasted return distribution or randomly resamples possible returns from a historical distribution. The simulated returns lead to a simulated set of capital market assumptions that are used in a traditional mean-variance optimizer, and the asset allocations are recorded. After combining the asset allocations from the numerous intermediate optimizations, the resulting asset allocations are those that, on average, are predicted to perform best over the range of potential outcomes implied by the capital market assumptions. Research has shown that asset allocations selected from a resampled efficient frontier may outperform those from a traditional efficient frontier.[192]

In addition to the problem of getting results that are highly concentrated in just a few of the assets available, there are two more criticisms of the traditional mean-variance optimization framework.

First, the traditional approach focuses on a subset of the total portfolio. Traditionally, the focus is on finding a mix of asset classes that maximizes the expected return, subject to a risk constraint. However, because the purpose of most asset portfolios is to fund a specified future cash-flow

[191] Monte Carlo simulation is a problem-solving technique utilized to approximate the probability of certain outcomes by performing multiple trial runs, called simulations, using random variables. The probability distribution of the results is calculated and analyzed in order to infer which outcomes are most likely to be produced.

[192] See Markowitz, H. & Usmen, N. 2003. "Resampled Frontiers vs. Diffuse Bayes: An Experiment." *Journal of Investment Management*, Vol. 1, No. 4.

stream – a liability – the true risk for the portfolio is not the standard deviation of the assets or the performance of the assets relative to that of peers, but not being able to fund the future liability. An asset allocation approach that takes the future liability into account is called liability-relative optimization (or surplus optimization). The usual method employed to accomplish this is to constrain the optimizer to hold short an asset class representing the liability.

Second, the traditional mean-variance optimization framework assumes that the returns of the assets in the optimization are normally distributed. As illustrated in Exhibit 2.3, the return distributions of different asset classes do not always follow a standard, symmetrical bell-shaped curve. Some assets have distributions that are skewed to the left or right, while others have distributions that are skinnier or fatter than others. These more complicated characteristics are called skewness and kurtosis, respectively. The next wave of enhancements to the traditional mean-variance optimization are frameworks that incorporate these additional types of abnormalities into the optimization.

Markowitz 1.0

In 1952, Harry Markowitz, invented portfolio optimization. His genius was based on three principles: risk, reward and the correlation of assets in a portfolio. Over the years, technologies advanced, markets crashed, but the portfolio optimization models used by many investors did not evolve to compensate. This is surprising in light of the fact that Markowitz was a pioneer of technological advancement in the field of computational computer science. Furthermore, he did not stand by idly in the area of portfolio modeling but continued to make improvements in his own models and to influence the models of others. Few of these improvements, however, were picked up broadly in practice.

Because Markowitz's first effort was so simple and powerful, it attracted a great number of followers. The greater the following became, the fewer questioners debated its merits. Markowitz's original work is synonymous with modern portfolio theory and has been taught in business schools for generations and, not surprisingly, is still widely used today.

Then came the crash of 2008, and people started to ask questions. The confluence of the economic trauma and the technological advances of recent decades made the post-crash environment the perfect moment to upgrade to a new model built around Markowitz's fundamental principles of risk, reward and correlation. We dub our updated model "Markowitz 2.0." This section is an adaptation of a 2009 article, "The New Efficient Frontier," by Paul D. Kaplan, Morningstar Canada's director of research, and Sam L. Savage, consulting professor at Stanford University.

Markowitz 2.0

The Flaw of Averages

The 1952 mean-variance model of Harry Markowitz was the first systematic attempt to cure what Savage (2009) called the “flaw of averages.” In general, the flaw of averages is a set of systematic errors that occur when people use single numbers (usually averages) to describe uncertain future quantities. For example, if you plan to rob a bank of $10 million and have one chance in 100 of getting away with it, your average take is $100,000. If you described your activity beforehand as “making $100,000” you would be correct on average. But this is a terrible characterization of a bank heist. Yet as Savage discussed, this very “flaw of averages” is made all the time in business practice and helps explain why everything is behind schedule, beyond budget, and below projection. This phenomenon was an accessory to the global financial crisis that culminated in 2008.

Markowitz’s mean-variance model distinguished between different investments that had the same average (expected) return but different risks, measured as variance or its square root (standard deviation). This breakthrough systematic attempt to cure the flaw of averages ultimately garnered a Nobel Prize for its inventor. However, the use of standard deviation and covariance introduces a higher order version of the flaw of averages in that these concepts are themselves a version of averages.

Making a Great Idea Better

By taking advantage of the very latest in economic thought and computer technology, we can, in effect, add more thrust to the original framework of the Markowitz portfolio optimization model. The result is a dramatically more powerful model that is more aligned with 21st century investor concerns, markets, and financial instruments, such as options.

Our discussion here will focus on five practical enhancements to traditional portfolio optimization that can be made with current technology:

1. First, we use a scenario-based approach to allow for fat-tailed distributions. “Fat-tailed” return distributions are not possible within the context of traditional mean-variance optimization where return distributions are assumed to be adequately described by mean and variance.

2. Second, we replace the single-period expected return with the long-term forward-looking geometric mean as this takes into account accumulation of wealth.

3. Third, we substitute conditional value at risk, or CVaR, which focuses on tail risk for standard deviation, which looks at average variation.

4. Fourth, the original Markowitz model used a covariance matrix to model the distribution of returns on asset classes; we replace this with a scenario-based model that can be generated with Monte Carlo simulation and can incorporate any number of distributions.

5. Finally, we exploit new statistical technologies pioneered by Savage in the field of probability management. Savage invented an astonishing new technology called the Distribution String, or DIST, which encapsulates thousands of trials as a single data element or spreadsheet cell, thus eliminating the main disadvantage of the scenario-based approach – the need to store and process large amounts of data.

The Scenario Approach vs. Lognormal Distributions

One of the limitations of the traditional mean-variance optimization framework assumes the distribution of returns of the assets in the optimization can be adequately described simply by mean and variance alone. The most common depiction of this assumption is to draw the distribution of each asset class as a symmetrical bell-shaped curve, but asset class returns do not always fall into normal distributions.

Over the years, various alternatives have been put forth to replace mean-variance optimization with an optimization framework that takes into account the non-normal features of return distributions. Some researchers have proposed using distribution curves that exhibit skewness and kurtosis (i.e., have fat tails) while others have proposed using large numbers of scenarios based on historical data, or Monte Carlo simulation.

The scenario-based approach has two main advantages over a distribution curve approach: (i) It is highly flexible. For example, nonlinear instruments such as options can be modeled in a straightforward manner; and (ii) it is mathematically manageable. For example, portfolio returns under the scenarios are simply weighted averages of asset class returns within the scenarios. In this way, the distribution of a portfolio can be derived from the distributions of the assets classes without working complicated equations that might lack analytical solutions; only straightforward portfolio arithmetic is needed.

In standard scenario analysis there is no precise graphical representation of return distributions. Histograms serve as approximations, such as those shown in Exhibit 2.3. We augment the scenario approach by employing a smoothing technique so that smooth curves represent return distributions. For example, Exhibit 10.5 shows the distribution curve of annual returns of large-cap stocks under our scenario-based approach. Comparing Exhibit 10.5 with the large-cap stock histogram in Exhibit 2.3, we can see that the smooth distribution curve retains the properties of the historical distribution making it more esthetically pleasing and precise. Further, our model can bring all of the power of continuous mathematics to the scenario approach. This was previously enjoyed only by models based on continuous distributions.

In Exhibit 10.5, the solid gray line represents the distribution of annual returns of large-cap stocks when our smoothed scenario-based approach is used, and the red line represents the distribution curve of annual returns of large-cap stocks when traditional mean-variance analysis is used and we assume that returns follow a lognormal distribution.

Exhibit 10.5: Distribution of Annual Returns: Large-cap Stocks (%) Lognormal Distribution vs. Scenario-Based Model

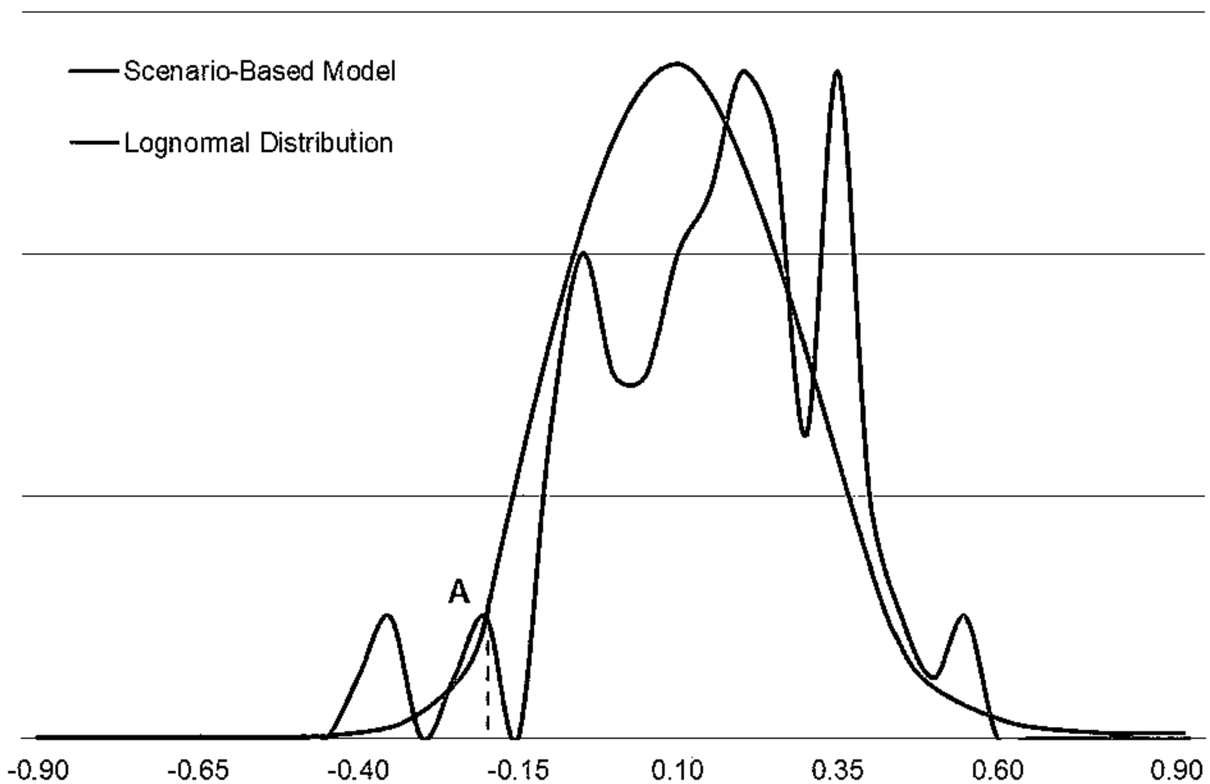

If we extend a vertical line from Point A down to the x-axis, the area to the left (and underneath) each of the curves represents the occurrences of annual returns equal to or less than, in this case, negative 26%. Because these are cumulative distributions, we can calculate the probability that the annual returns of large cap stocks will be less than or equal to negative 26% by dividing the area underneath each of the smaller curves (to the left of Point A) by the total area underneath each of the entire curves.

For example, looking to the scenario-based model, the area to the left of the vertical line under the scenario-based distribution represents 5% of the total area underneath this entire distribution line. This implies that the probability of large cap stocks having a loss of 26% or more is 5%. Correspondingly, the area to the left of the vertical line for the lognormal distribution represents 1.6% of the total area under the entire lognormal distribution line. This implies that the probability of large-cap stocks returning negative 26% or less using the traditional mean-variance model is 1.6%.

As Kaplan et al. (2009) discuss, "tail events" have occurred often throughout the history of capital markets all over the world, but the probabilities associated with them may be systematically underestimated within the context of traditional mean-variance analysis where return distributions are assumed to be lognormal. The scenario-based model proposed by Kaplan and Savage is a real step forward as it better models the nontrivial probabilities associated with tail events.

For a more detailed discussion of tail events and their nontriviality, see Chapter 11, where Kaplan introduces a set of monthly real stock market total returns going back a full 150 years. Using these new returns, we demonstrate that the severity of the financial crisis of 2008 was not unique but was merely the latest chapter in a long history of market meltdowns.

Geometric Mean vs. Single-Period Expected Return

In mean-variance optimization, reward is measured by expected return, which is a forecast of arithmetic mean. However, over long periods, investors are not concerned with simple averages of return rather they are concerned with the accumulation of wealth. We use forecasted long-term geometric mean as the measure of reward because investors who plan on repeatedly reinvesting in the same strategy over an indefinite period would seek the highest rate of growth for the portfolios as measured by geometric mean.[193]

Conditional Value at Risk vs. Standard Deviation

As for risk, much has been written about how investors are not concerned merely with the degree of dispersion of returns (as measured by standard deviation), but rather with how much wealth they could lose. A number of downside risk measures, including value at risk, conditional value at risk, and maximum drawdown, have been proposed to replace standard deviation as the measure of risk in strategic asset allocation. While any one of these could be used, our preference is to use conditional value at risk.

CVaR is related to value at risk. VaR describes the left tail in terms of how much capital can be lost over a given period of time. For example, a 5% VaR answers a question of the form: Having invested $10,000 there is a 5% chance of losing $X or more in 12 months. (The "or more" implications of VaR are sometimes overlooked by investors with serious implications.) Applying this idea to returns, the 5% VaR is the negative of the 5th percentile of the return distribution. CVaR is the expected or average loss of capital should VaR be breached. Therefore CVaR is always greater than VaR.

Scenarios vs. Correlation

In mean-variance analysis, the covariation of the returns of each pair of asset classes is represented by a single number, the correlation coefficient. This is mathematically equivalent to assuming that a simple linear regression model is an adequate description of how the returns on

[193] Ranking investment strategies by forecasted geometric mean is sometimes described as applying the Kelly Criterion, an idea promoted by William Poundstone in his 2005 book, *Fortune's Formula*.

the two asset classes are related. In fact, the R-squared statistic of a simple linear regression model for two series of returns is equal to the square of the correlation coefficient.

However, for many pairs of asset classes, a linear model misses the most important features of the relationship. For example, during normal times, non-U.S. equities are considered to be good diversifiers for U.S. equity investors. But during global crises, all major equity markets move down together.

Furthermore, suppose that the returns on two asset class indices were highly correlated, but instead of including direct exposures to both in the model, one was replaced with an option on itself. Instead of having a linear relationship, we now have a nonlinear relationship which cannot be captured by a correlation coefficient. Fortunately, these sorts of nonlinear relationships between returns on different investments can be handled in a scenario-based model. For example, in scenarios that represent normal times, returns on different equity markets could be modeled as moving somewhat apart from each other, while scenarios that represent global crises could model the markets as moving downward together.

Probability Management Enables Scenario Analysis

Because it may take thousands of scenarios to adequately model return distributions, until recently, a disadvantage of the scenario-based approach has been that it requires large amounts of data to be stored and processed. Even with the advances in computer hardware, the conventional approach of representing scenarios with large tables of explicit numbers remained problematic. The phenomenal speed of computers has given rise to the field of probability management, an extension of data management to probability distributions, rather than numbers. The key component of probability management is the Distribution String, or DIST, that can encapsulate thousands of trials as a single data element. The use of DISTs greatly saves on storage and speeds up processing time so that a Monte Carlo simulation consisting of thousands of trials can be performed on a personal computer in an instant. Monte Carlo simulations that use DISTs are also very adaptable, allowing for almost any return distribution or underlying probability model rather than being contained by parameters. While not all asset management organizations are prepared to create the DISTs needed to drive geometric mean-CVaR optimization, some outside vendors, such as Ibbotson Associates/Morningstar, can fulfill this role. Another facet of probability management is interactive simulation technology, which can run thousands of scenarios through a model before the sound of your finger leaving the <Enter> key reaches your ear. These supersonic models allow much deeper intuition into the sensitivities of portfolios and encourage the user to interactively explore different portfolios, distributional assumptions, and potential black swans. For more information visit: http://www.ProbabilityManagement.org.

Finale: The New Efficient Frontier

Putting it all together, we form an efficient frontier of forecasted geometric mean and conditional value at risk as shown in Exhibit 10.6,[194] incorporating our scenario approach to covariance and new statistical technology. We believe that this efficient frontier is more relevant to investors than the traditional expected return versus standard deviation frontier of MVO because it shows the trade-off between reward and risk that is meaningful to investors, namely, long-term potential growth versus short-term potential loss.

Exhibit 10.6: Geometric Mean – Conditional Value at Risk Efficient Frontier (%)

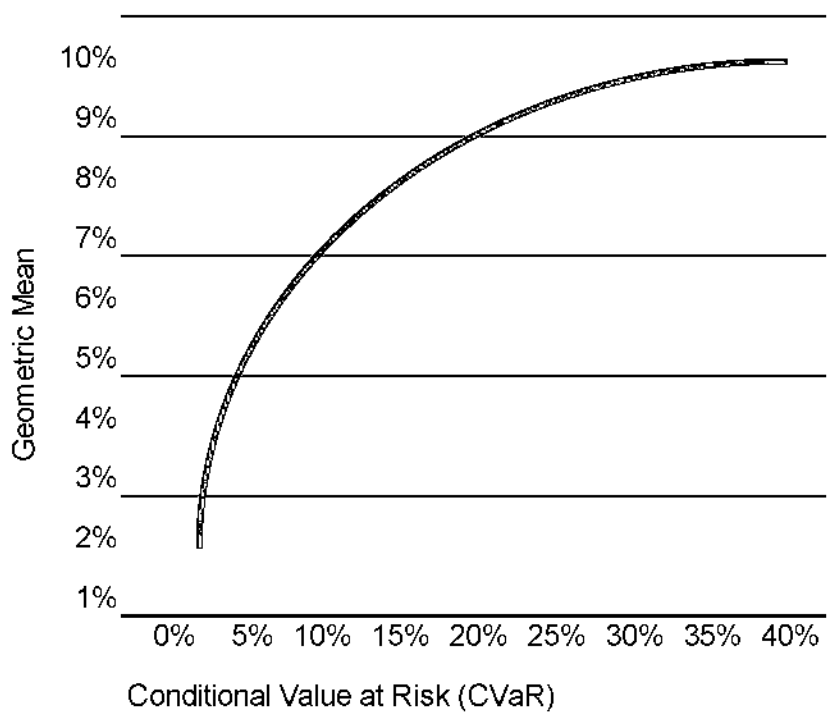

Approaches to Calculating the Equity Risk Premium

Researchers have estimated the expected outperformance of stocks over risk-free bonds – the equity risk premium – using many approaches. Such studies can be categorized into four groups based on the approaches they have taken, using:

- Historical returns between stocks and bonds
- Fundamental information such as earnings, dividends, or overall productivity (supply-side models)

194 Other researchers have also proposed using GM and CVaR as the measures or reward and risk in an efficient frontier. See, for example: Sheikh, A.Z. & Qiao, H. 2009. "Non-Normality of Market Returns: A Framework for Asset Allocation Decision Making". Whitepaper, J.P. Morgan Asset Management.

- Payoffs demanded by equity investors for bearing the additional risk (demand-side models)

- Broad surveys of opinions of financial professionals.

The rest of this chapter will focus on the historical and supply-side methods.

The Historical Equity Risk Premium

The expected equity risk premium (ERP) can be defined as the additional return an investor expects to receive to compensate for the additional risk associated with investing in equities as opposed to investing in riskless assets.

Unfortunately, the expected equity risk premium is unobservable in the market and therefore must be estimated. Typically, this estimation is arrived at using historical data. The historical equity risk premium can be calculated by subtracting the long-term average of the income return on the riskless asset (Treasuries) from the long-term average stock market return (measured over the same period as that of the riskless asset).

In using a historical measure of the equity risk premium, one assumes what has happened in the past is representative of what might be expected in the future. In other words, the assumption one makes when using historical data to measure the expected equity risk premium is the relationship between the returns of the risky asset (equities) and the riskless asset (Treasuries) is stable.

The Stock Market Benchmark

The stock market benchmark chosen should be a broad index that reflects the behavior of the market as a whole. Commonly used indexes include the S&P 500 and the Russell 3000. Although the Dow Jones Industrial Average is a popular index, it would be inappropriate for calculating the equity risk premium because it is too narrow.

We use the total return of our large-cap stock index (currently represented by the S&P 500) as our market benchmark when calculating the equity risk premium.[195] The S&P 500 was selected as the appropriate market benchmark because it is representative of a large sample of companies across a large number of industries. The S&P 500 is also one of the most widely accepted market benchmarks and is a good measure of the equity market as a whole.

Exhibit 10.7 illustrates the equity risk premium calculated using the S&P 500 and the income return on three government bonds of *different* horizons.

[195] The SBBI Large-cap Stocks total return series is essentially the S&P 500 Index.

Exhibit 10.7: Equity Risk Premia Calculated Using the S&P 500 and the Income Return on Three Government Bonds of Different Horizons (%) 1926–2020

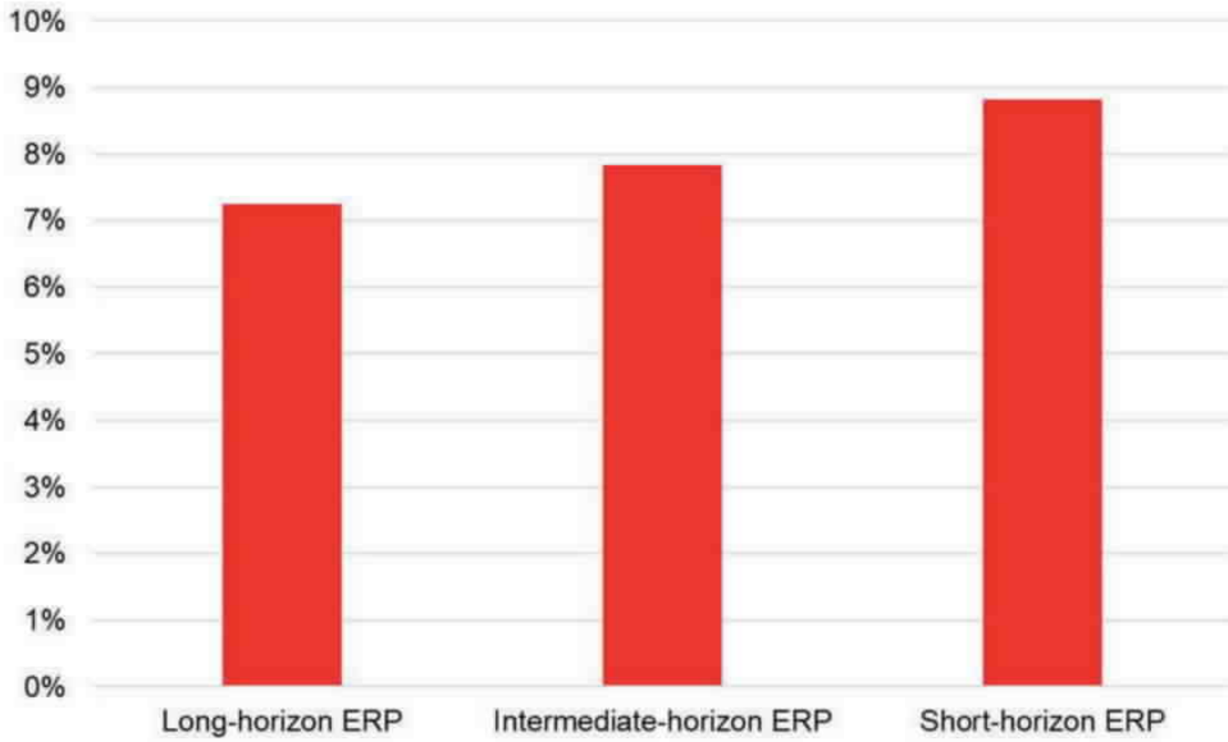

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, (ii) Long-Term (i.e. 20-year) Government Bonds income return series: IA SBBI® US LT Govt IR USD, (v) Intermediate-term (i.e., 5-year) Government Bonds income return series: IA SBBI® US IT Govt IR USD, (vi) U.S. (30-day) Treasury Bills: IA SBBI® US 30 Day TBill TR USD (for U.S. Treasury Bills, the income return and total return are the same). For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Note that the long-horizon ERP is *lower* than the intermediate-horizon ERP, which in turn is *lower* than the short-horizon ERP. This is because the equity risk premium is calculated by subtracting the arithmetic mean of the government bond *income* return from the arithmetic mean of the stock market total return. When calculating the ERPs in these examples:

- The average income return of a *long*-term (20-year) government bond is used when calculating the *long*-horizon ERP. The average annual income return of 20-year government bonds will be *greater* than the average income return of a 5-year government bond.

- The average income return of an *intermediate*-term (5-year) government bond is used when calculating the *intermediate*-horizon ERP. The average annual income return of 5-year government bonds will be *greater* than the average income return of a 30-day Treasury bill.

- The average income return of a *30-day* Treasury bill is used when calculating the *short-*horizon ERP.

Because the ERPs in these examples are calculated as:

ERP = (Avg. Annual Total Return of the S&P 500) – (Avg. Annual Income Return of Risk-Free Security)

It follows that the ERP would *increase* as the value subtracted from the average annual total return of the stock market benchmark *decreases*.

The Market Benchmark and Firm Size

Although not restricted to the 500 largest companies, the S&P 500 is considered a large-cap index. The returns of the S&P 500 are market cap-weighted. The larger companies in the index therefore receive the majority of the weight. If using a large-cap index to calculate the equity risk premium, an adjustment is usually needed to account for the different risk and return characteristics of small stocks, as discussed in Chapter 7.

The Risk-Free Asset

The equity risk premium can be calculated for a variety of time horizons when given the choice of risk-free asset to be used in the calculation. The long-horizon, intermediate-horizon, and short-horizon equity risk premia calculated in Exhibit 10.7 use the income return from (i) a 20-year Treasury bond, (ii) a 5-year Treasury bond, and (iii) a 30 day Treasury bill, respectively.[196]

20-Year vs. 30-Year Treasuries

The U.S. Treasury periodically changes the maturities it issues. For example, in April 1986 the U.S. Treasury stopped issuing 20-year Treasuries, and from October 2001 through January 2006 the U.S. Treasury did not issue 30-year bonds (it resumed issuing 30-year Treasury bonds in February 2006), making the 10-year bond the longest-term Treasury security issued over the October 2001 January 2006 period. Most recently, on January 16, 2020 the U.S. Department of the Treasury announced it plans to issue a 20-year nominal coupon bond in the first half of calendar year 2020, the first time a 20-year maturity will be offered since March 1986.[197,198]

Our methodology for estimating the long-horizon equity risk premium makes use of the income return on a 20-year Treasury bond. While a 30-year bond is theoretically more correct when

196 For U.S. Treasury Bills, the income return and total return are the same.

197 To learn more, visit the U.S. Department of the Treasury website at: https://home.treasury.gov/news/press-releases/sm878.

198 See Kate Davidson, "Treasury to Issue New 20-Year Bond in First Half of 2020", *The Wall Street Journal*, January 16, 2020 at: https://www.wsj.com/articles/treasury-to-issue-new-20-year-bond-in-first-half-of-2020-11579217450

dealing with the long-term nature of business valuation,[199] 30-year Treasury securities have an issuance history that is on-again-off-again. Ibbotson Associates creates a series of returns using bonds on the market with approximately 20 years to maturity because Treasury bonds of this maturity are available over a long history, while Treasury bonds of 30-years are not.

Income Return Another point to keep in mind when calculating the equity risk premium is the income return on the appropriate-horizon Treasury security, rather than the total return, is used in the calculation.

The total return comprises three return components: the income return, the capital appreciation return, and the reinvestment return. The income return is defined as the portion of the total return that results from a periodic cash flow or, in this case, the bond coupon payment. The capital appreciation return results from the price change of a bond over a specific period. Bond prices generally change in reaction to unexpected fluctuations in yields. Reinvestment return is the return on a given month's investment income when reinvested into the same asset class in the subsequent months of the year. The income return is thus used in the estimation of the equity risk premium because it represents the truly riskless portion of the return.

Arithmetic vs. Geometric Mean

The equity risk premium data presented in this book are arithmetic average risk premiums as opposed to geometric average risk premiums. The arithmetic average equity risk premium can be demonstrated to be most appropriate when discounting future cash flows. For use as the expected equity risk premium in either the CAPM or the building-block approach, the arithmetic mean or the simple difference of the arithmetic means of stock market returns and riskless rates is the relevant number.

This is because both the CAPM and the building-block approach are additive models, in which the cost of capital is the sum of its parts. The geometric average is more appropriate for reporting past performance because it represents the compound average return.

Appropriate Historical Period

The equity risk premium can be estimated using any historical time period. For the U.S., market data exist at least as far back as the late 1800s. Therefore, it is possible to estimate the equity risk premium using data that covers roughly the past 125 years.

---

199 An equity risk premium is an input in developing cost of capital estimates (i.e., "expected return", "required return", or "discount rate") for use in a discounted cash flow model. **Note**: The D&P/Kroll "Cost of Capital Navigator" guides financial professionals through the process of estimating the cost of capital, a key component of any valuation analysis. The Cost of Capital Navigator can be used to estimate country-level cost of equity capital globally, for approximately 180 countries, from the perspective of investors based in any one of up to 56 countries. For more information, visit dpcostofcapital.com.

Our equity risk premium covers 1926 to the present. The original data source for the time series comprising the equity risk premium is the Center for Research in Security Prices (CRSP).[200] CRSP chose to begin its analysis of market returns with 1926 for two main reasons. CRSP determined that 1926 was approximately when quality financial data became available. They also made a conscious effort to include the period of extreme market volatility from the late 1920s and early 1930s; 1926 was chosen because it includes one full business cycle of data before the market crash of 1929.

Implicit in using history to forecast the future is the assumption that investors' expectations for future outcomes conform to past results. This method assumes that the price of taking on risk changes only slowly, if at all, over time. This "future equals the past" assumption is most applicable to a random time-series variable. A time-series variable is random if its value in one period is independent of its value in other periods.

Choosing an Appropriate Historical Period

The estimate of the equity risk premium depends on the length of the data series studied. A proper estimate of the equity risk premium requires a data series long enough to give a reliable average without being unduly influenced by very good and very poor short-term returns. When calculated using a long data series, the historical equity risk premium is relatively stable. Furthermore, because an average of the realized equity risk premium is quite volatile when calculated using a short history, using a long series makes it less likely that the analyst can justify any number he or she wants. The magnitude of how shorter periods can affect the result will be explored later in this chapter.

Some analysts estimate the expected equity risk premium using a shorter, more recent period on the basis that recent events are more likely to be repeated in the near future; furthermore, they believe that the 1920s, 1930s, and 1940s contain too many unusual events. This view is suspect because all periods contain unusual events. Some of the most unusual events of the last 100 years took place quite recently, including the inflation of the late 1970s and early 1980s, the October 1987 stock market crash, the collapse of the high-yield bond market, the major contraction and consolidation of the thrift industry, the collapse of the Soviet Union, the development of the European Economic Community, the attacks of Sept. 11, 2001, the global financial crisis of 2008– 2009, and most recently, the market crash in the first quarter of 2020 that was precipitated by the spread of the COVID-19 virus.

It is even more difficult for economists to predict the economic environment of the future. For example, if one were analyzing the stock market in 1987 before the crash, it would be statistically improbable to predict the impending short-term volatility without considering the stock market crash and market volatility of the 1929-1931 period.

---

200 CRSP® is a registered trademark and service mark of Center for Research in Security Prices, LLC and has been licensed for use by D&P/Kroll. The D&P/Kroll publications and services are not sponsored, sold or promoted by CRSP®, its affiliates or its parent company. To learn more about CRSP, visit www.crsp.com.

Without an appreciation of the 1920s and 1930s, no one would believe that such events could happen. The 95-year period starting with 1926 represents what can happen: It includes high and low returns, volatile and quiet markets, war and peace, inflation and deflation, and prosperity and depression. Restricting attention to a shorter historical period underestimates the amount of change that could occur in a long future period. Finally, because historical event-types (not specific events) tend to repeat themselves, long-run capital market return studies can reveal a great deal about the future. Investors probably expect unusual events to occur from time to time, and their return expectations reflect this.

A Look at the Historical Results

It is interesting to look at the realized returns and realized equity risk premium in the context of the above discussion, since a longer historical period provides a more stable estimate of the equity risk premium. The reason is that any unique period will not be weighted heavily in an average covering a longer historical period. It better represents the probability of these unique events occurring over a long period of time.

Exhibit 10.8 helps to clarify this point. Exhibit 10.8 shows the realized equity risk premium for a series of periods through 2020, starting with 1926. In other words, the first value on the graph represents the average realized equity risk premium over the period 1926–2020. The next value on the graph represents the average realized equity risk premium over the period 1927–2020, and so on, with the rightmost value representing the average for a single year, 2020.

Concentrating on the left side of Exhibit 10.8, one notices that the realized equity risk premium when measured over *longer* periods is relatively stable and has a standard deviation of 0.9.

Alternatively, the realized equity premia on the right side of Exhibit 10.8 are measured over *shorter* periods are less stable and have a standard deviation of 3.5.[201]

[201] If the unusually large realized equity risk premia measured over the years 2019–2020 (23.0) and the single year 2020 (17.0) are excluded (the rightmost two bars in Exhibit 10.8), the standard deviation of the realized equity risk premia measured with starting dates 1973–2018 drops to 2.5. This is still more than twice the standard deviation of the realized equity risk premia measured with starting dates 1926–1972 (the left side of Exhibit 10.8) of 0.9.

Exhibit 10.8: Average Long-Horizon Equity Risk Premium Calculated Using Variable Start Dates (1926–2020), and Fixed End Date (2020) (%)

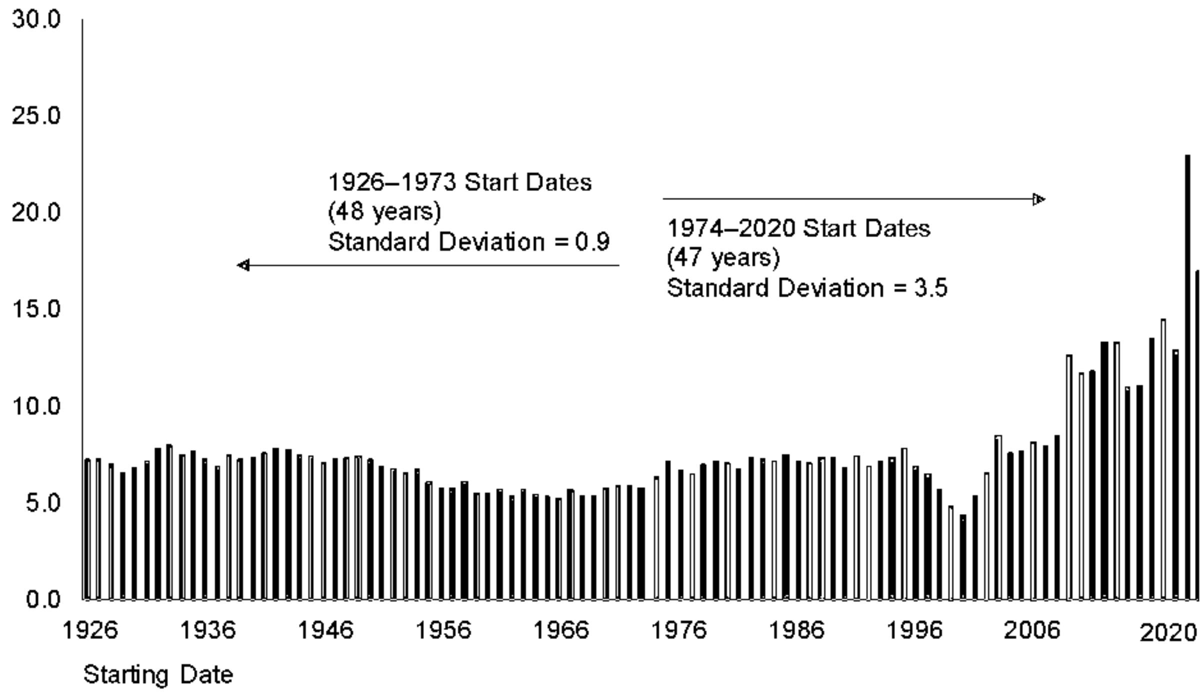

Source of underlying data: Morningstar, Inc. Used with permission. All rights reserved. Calculations by D&P/Kroll. Asset classes represented by the Ibbotson Associates (IA) Stocks, Bonds, Bills, and Inflation® (SBBI®) series, as follows: (i) Large-Cap Stocks: IA SBBI® US Large Stock TR USD Ext, and (ii) Long-Term (i.e. 20-year) Government Bonds income return series: IA SBBI® US LT Govt IR USD. For a detailed description of the SBBI® series, see Chapter 3, "Description of the Basic Series". "Stocks, Bonds, Bills, and Inflation" and "SBBI" are registered trademarks of Morningstar, Inc. All rights reserved. Used with permission.

Some practitioners argue for a shorter historical period, such as 30 years, as a basis for the equity risk premium estimation. The logic for the use of a shorter period is that historical events and economic scenarios present before this time are unlikely to be repeated. However, the impact of adding one additional year of data to a historical average is *lessened* the *greater* the initial period of measurement. As is demonstrated in Exhibit 10.8, shorter-term averages can be affected considerably by one or more unique observations, while longer term averages tend to produce more stable results.

A dramatic example of this is the second rightmost point in Exhibit 10.8, which is the "average" ERP as measured over a single year (2019). In 2019 large-cap stocks (represented by the S&P 500) produced a total return of over 31.49% and the income return of long-term government bonds was 2.55%, implying an "average" ERP of 28.94% (31.49% - 2.55%). Using an estimate of the ERP developed over such a short time horizon is logical only to the extent that one believes that stocks will outperform the risk-free instrument by nearly 29% per year, in perpetuity.

Having said that, the effect of "adding one additional year" when using historical data to estimate the ERP can still lead to counterintuitive conclusions, even when the average is taken over *longer* periods. A very recent example of a result that was “counterintuitive” occurred in the December 2008–2009 Financial Crisis. The historical ERP at the end of 2007 (as calculated over the time period 1926–2007) was over 7%. A year later at the end of 2008, at the height of the financial crisis and risks were likely at an all-time *high*, the historical ERP (as calculated over the time period 1926– 2008) *declined* to less than 7%, implying that risks were actually *lower* than they were a year earlier.

What happened? In 2008 the S&P 500 declined nearly 37%, an unusually large decline for a single year. This single period’s unusually large decline caused the average annual return of the S&P 500 to fall from *over* 12% (as calculated over the 1926–2007 time period) to *less* than 12% (as calculated over the 1926–2008 time period). The historical ERP is calculated as the average annual equity return minus the average annual risk-free rate, so a decline in the average equity return causes a 1 for 1 decline in the ERP, all other things held the same. Such large moves in a single year can produce a “tail wagging the dog” effect.

The Supply-Side Equity Risk Premium

This section is based on the work by Roger G. Ibbotson and Peng Chen, who combined the first and second approaches to arrive at their forecast of the equity risk premium.[202] By proposing a new supply-side methodology, the Ibbotson-Chen study challenges current arguments that future returns on stocks over bonds will be negative or close to zero. The results affirm the relationship between the stock market and the overall economy.

Long-term expected equity returns can be forecasted by the use of supply-side models. The supply of stock market returns is generated by the productivity of the corporations in the real economy. Investors should not expect a much higher or lower return than that produced by the companies in the real economy. Thus, over the long run, equity returns should be close to the long-run supply estimate.

Earnings, dividends, and capital gains are supplied by corporate productivity. Exhibit 10.9 illustrates that earnings and dividends have historically grown in tandem with the overall economy (GDP per capita). However, GDP per capita did not outpace the stock market. This is primarily because the 3-year average P/E ratio increased 2.7 times during the same period. So, assuming the economy will continue to grow, all three should continue to grow as well.

[202] Ibbotson, R.G., & Chen, P. 2003. “Long-Run Stock Returns: Participating in the Real Economy”. *Financial Analysts Journal*, Vol. 59, No. 1, P. 88.

Exhibit 10.9: Capital Gains, GDP Per Capita, Earnings, and Dividends Index (Year-end 1925 = $1.00) 1926–2020

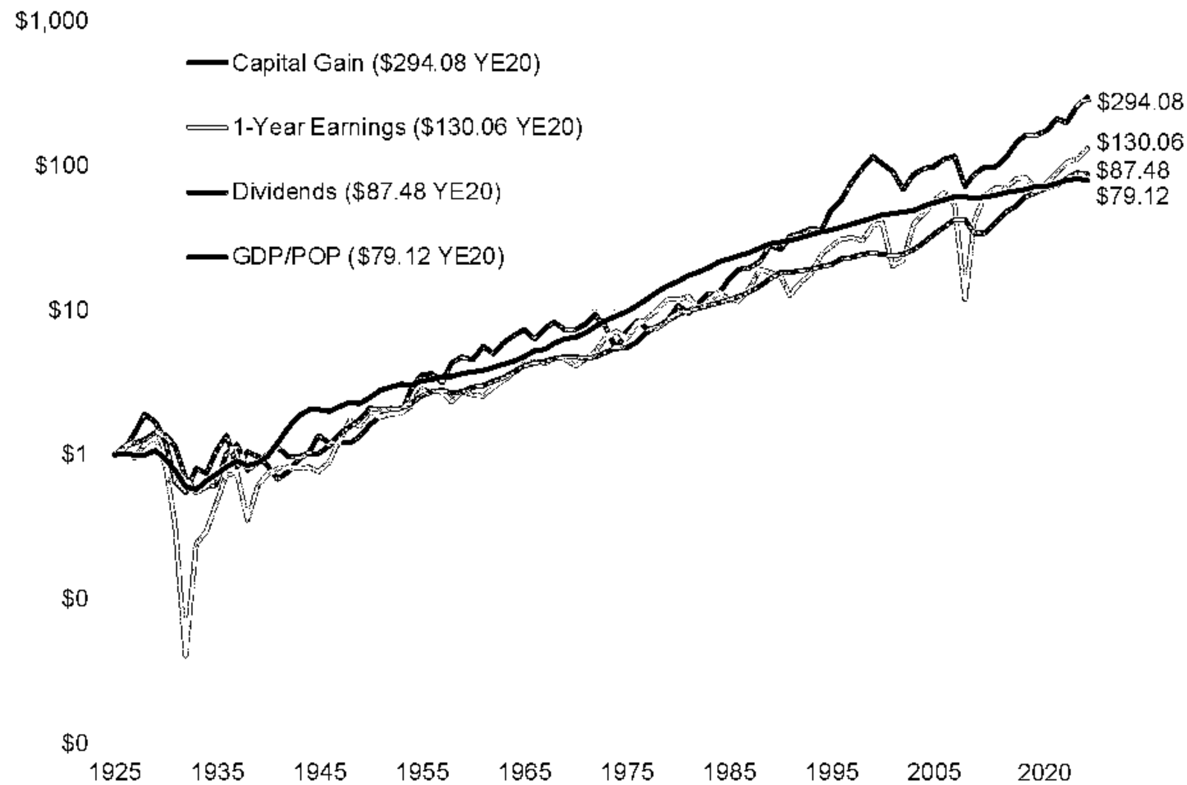

Forward-Looking Earnings Model

Ibbotson and Chen forecast the equity risk premium through a supply-side model using historical data. They used an earnings model as the basis for their supply-side estimate. The earnings model breaks the historical equity return into four pieces, with only three historically being supplied by companies: inflation, income return, and growth in real earnings per share. The growth in the P/E ratio, the fourth piece, reflects investors' changing prediction of future earnings growth. The past supply of corporate growth is forecasted to continue; however, a change in investors' predictions is not. P/E rose dramatically from 1980 through 2001 because people believed that corporate earnings were going to grow faster in the future. This growth in P/E drove a small portion of the rise in equity returns over the same period.

Exhibit 10.10 illustrates the price-to-earnings ratio from 1926 to 2020. The P/E ratio, using one-year average earnings, was 10.23 at the beginning of 1926 and ended the year 2020 at estimated 38.94, an average increase of 1.40% per year. The highest P/E was 136.69 recorded in 1932, while the lowest was 7.08 recorded in 1948. Ibbotson Associates revised the calculation of the P/E ratio from a one-year to three-year average earnings for use in equity forecasting.

Exhibit 10.10: Large-cap Stocks P/E Ratio 1926–2020

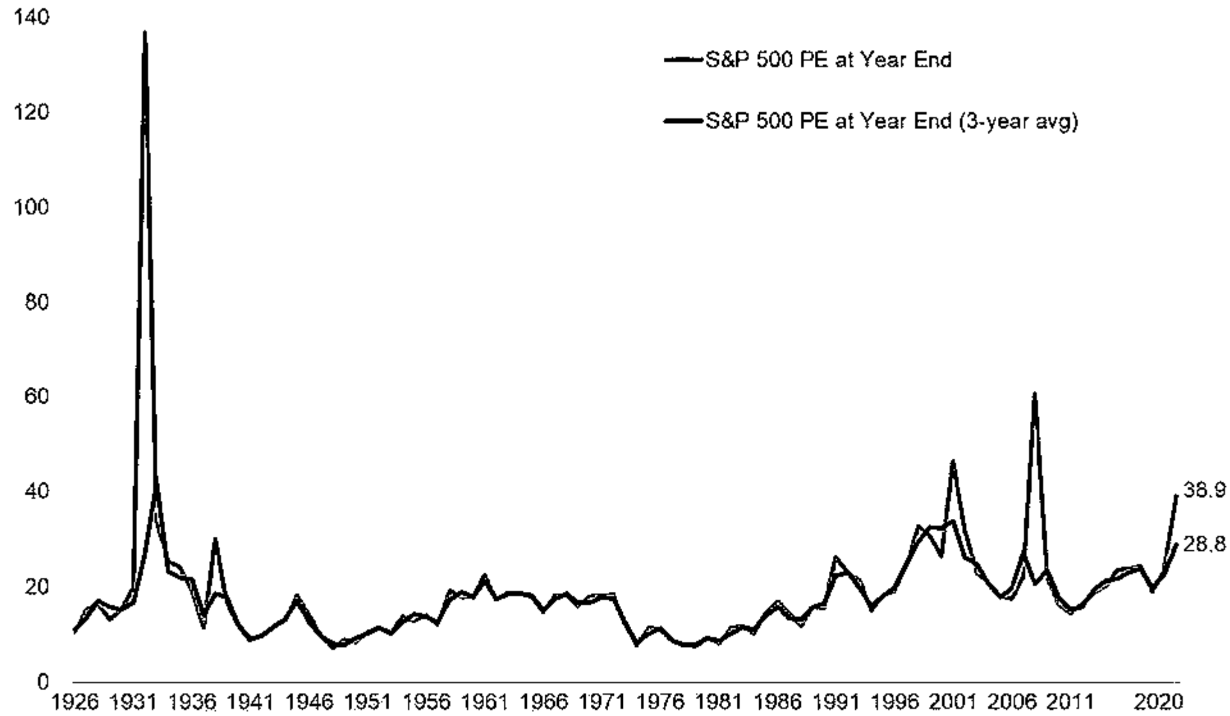

This is because reported earnings are affected not only by the long-term productivity, but also by one-time items that do not necessarily have the same consistent impact year after year. The three-year average is more reflective of the long-term trend than the year-by-year numbers. The P/E ratio calculated using the three-year average of earnings had an increase of 0.96% per year.

The historical P/E growth factor, using three-year earnings, of 0.96% per year is subtracted from the equity forecast because it is not believed that P/E will continue to increase in the future. The market serves as the cue. The current P/E ratio is the market's best guess for the future of corporate earnings and there is no reason to believe, at this time, that the market will change its mind. Using this top-down approach, the geometric supply-side equity risk premium is slightly more than 4% which equates to an arithmetic supply-side equity risk premium of approximately 6%.

Another approach in calculating the premium would be to add up the components that constitute the supply of equity return, excluding the P/E component. Thus, the supply of equity return only includes inflation, the growth in real earnings per share, and income return:

$$SR = \left[(1+CPI)\times(1+g_{REPS}) + Inc + Rinv\right]$$

Where:

SR = The supply of the equity return

CPI = Consumer Price Index (inflation)

g_{REPS} = The growth in real earning per share

Inc = The income return

$Rinv$ = The reinvestment return

The equity risk premium, based on the supply-side earnings model, is calculated on a geometric basis as follows:

$$SERP = \frac{(1+SR)}{(1+CPI)\times(1+RRf)} - 1$$

Where:

$SERP$ = The supply-side equity risk premium

SR = The supply of the equity return

CPI = Consumer Price Index (inflation)

RRf = The real risk-free rate

The geometric estimate can be converted into an arithmetic estimate as follows:[203]

[203] The 1926–present supply-side equity risk premia estimate is calculated by D&P/Kroll for the full-version *2021 SBBI® Yearbook* using (i) the same methodologies and (ii) the same data sources as were used in previous editions of this book, based upon the work by Roger G. Ibbotson and Peng Chen; see: Ibbotson, R.G., & Chen, P. 2003. "Long-Run Stock Returns: Participating in the Real Economy". *Financial Analysts Journal*, Vol. 59, No. 1, P. 88. An update of this work has been published that considers stock buybacks in addition to dividends; see: Philip U. Straehl and Roger G. Ibbotson, "The Long-Run Drivers of Stock Returns: Total Payouts and the Real Economy", *Financial Analysts Journal*, Third Quarter 2017, Volume 73 Number 3. The *Financial Analysts Journal* is a publication of CFA Institute. For more information, visit www.cfainstitute.org.

$$R_A = R_G + \frac{\sigma^2}{2}$$

Where:

R_A = The arithmetic average

R_G = The geometric average

σ = The standard deviation of equity returns

Exhibit 10.11 presents an illustration of the supply-side equity risk premium, on an arithmetic basis, beginning in 1926 and ending in each of the years from 2003 through 2020.[204]

Exhibit 10.11: Supply-Side Equity Risk Premia and Long-Horizon Historical Equity Risk Premia over Time.

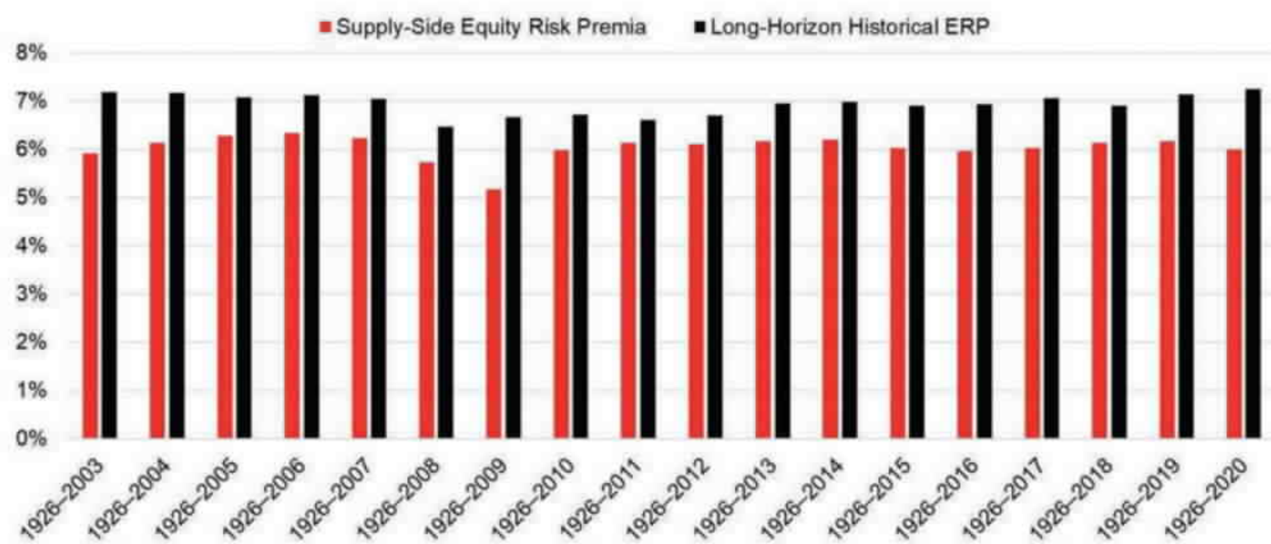

Source of underlying data: Morningstar, Inc. and S&P Global Market Intelligence. All rights reserved. Used with permission. Calculations by D&P/Kroll.

In every year since 2003 the supply-side ERP has been less than the long-term historical ERP. The difference has varied between approximately 0.5% and 1.5% over the course of the 18 observations in Exhibit 10.11.

[204] As published in (i) the *2004–2013 SBBI® Valuation Yearbooks*, (ii) the 2014–2017 *Valuation Handbook – U.S. Guide to Cost of Capital*, and (iii) the Cost of Capital Navigator at dpcostofcapital.com beginning in 2018.

Long-Term Market Predictions

As of December 31, 2020, the supply-side model estimates that stocks will continue to provide significant returns over the long run, averaging more than 9% per year, assuming historical inflation rates. The equity risk premium, based on the top-down supply-side earnings model, is calculated to be just over 4% on a geometric basis and approximately 6% on an arithmetic basis.

Ibbotson and Chen predict future increased earnings growth that will offset lower dividend yields. The fact that earnings will grow as dividend payouts shrink is in line with the Modigliani-Miller theorem which here refers to the irrelevance over whether a firm pays a dividend or reinvests its returns.

The forecasts for the market are in line with both the historical supply measures of public corporations (i.e., earnings) and overall economic productivity (GDP per capita).

Stock Buybacks and Return

Note: This section is updated through December 2018.

In recent decades a new source of stock market supply has emerged as companies increasingly use share buybacks instead of dividends to return cash to shareholders. The impact of buybacks on stock returns has been largely ignored in practice because many practitioners continue to rely on traditional supply models that use dividends as the sole source of corporate payout.

In a 2017 article, Philip U. Straehl and Roger G. Ibbotson developed three total payout models of stock returns showing that US stock returns between 1871 and 2014 can be attributed almost entirely to the supply of both dividends and buybacks.[205,206,207]

Although Straehl and Ibbotson introduced buybacks into the supply-side model, they did not dispute that there are many supply-side approaches that can be taken. Rather they updated and back dated the Ibbotson and Chen 2003[208] paper to cover the period 1871-2014, decomposing historical returns by six different methods each containing an inflation component:

1. Building Blocks: risk-free rate and equity risk premium

2. Capital Gains and Income

---

[205] Philip U. Straehl is head of capital markets and asset allocation at Morningstar Investment Management LLC, Chicago. Roger G. Ibbotson is Professor in the Practice Emeritus of Finance at the Yale School of Management, New Haven, Connecticut, and chairman and chief investment officer at Zebra Capital Management LLC, Stamford, Connecticut.

[206] Philip U. Straehl and Roger G. Ibbotson, "The Long-Run Drivers of Stock Returns: Total Payouts and the Real Economy", *Financial Analysts Journal* (a publication of CFA Institute), Third Quarter 2017, pages 32–52.

[207] This section is a summary of Philip U. Straehl and Roger G. Ibbotson, "The Long-Run Drivers of Stock Returns: Total Payouts and the Real Economy", *Financial Analysts Journal* (a publication of CFA Institute), Third Quarter 2017, pages 32–52. The original article was through 2014; Straehl and Ibbotson updated the commentary herein through 2018.

[208] Ibbotson, Roger G., and Peng Chen. 2003, "Long-Run Stock Returns: Participating in the Real Economy." *Financial Analysts Journal*, vol. 59, no. 1 (January/February).

3. Earnings growth, PE ratio, rate of change, and income

4. Dividends growth, payout ratio rate of change, PE ratio rate of change, and income

5. Book Equity growth, growth in ROE, PE ratio rate of change, and income

6. GDP per capita growth, increase in equity factor share of economy

Straehl and Ibbotson focus in particular on method #4, which concentrates on dividend growth, and method #6 which links the stock market to the overall economy. But they make a major departure from Ibbotson and Chen because they include not only dividends but buybacks into the analysis.

The Rise of Buybacks

A primary objective of Straehl and Ibbotson's study was to shed light on the impact of buybacks on the return generation process. They started by documenting the rise of buybacks as a form of corporate payout relative to dividends.

In 1982, SEC Rate 106-18 provided a safe harbor for firms to conduct share buybacks without the suspicion of share price manipulation. Here we update Exhibit 10.12 from Straehl and Ibbotson to include data through 2018. As can be seen in Exhibit 10.12, there is a major increase in buyback activity starting in the early 1980s. Prior to 1970, buyback activity was so low that it was not included in the study.

Exhibit 10.12: Dividend Yield and Buyback Yield, 1871– 2018

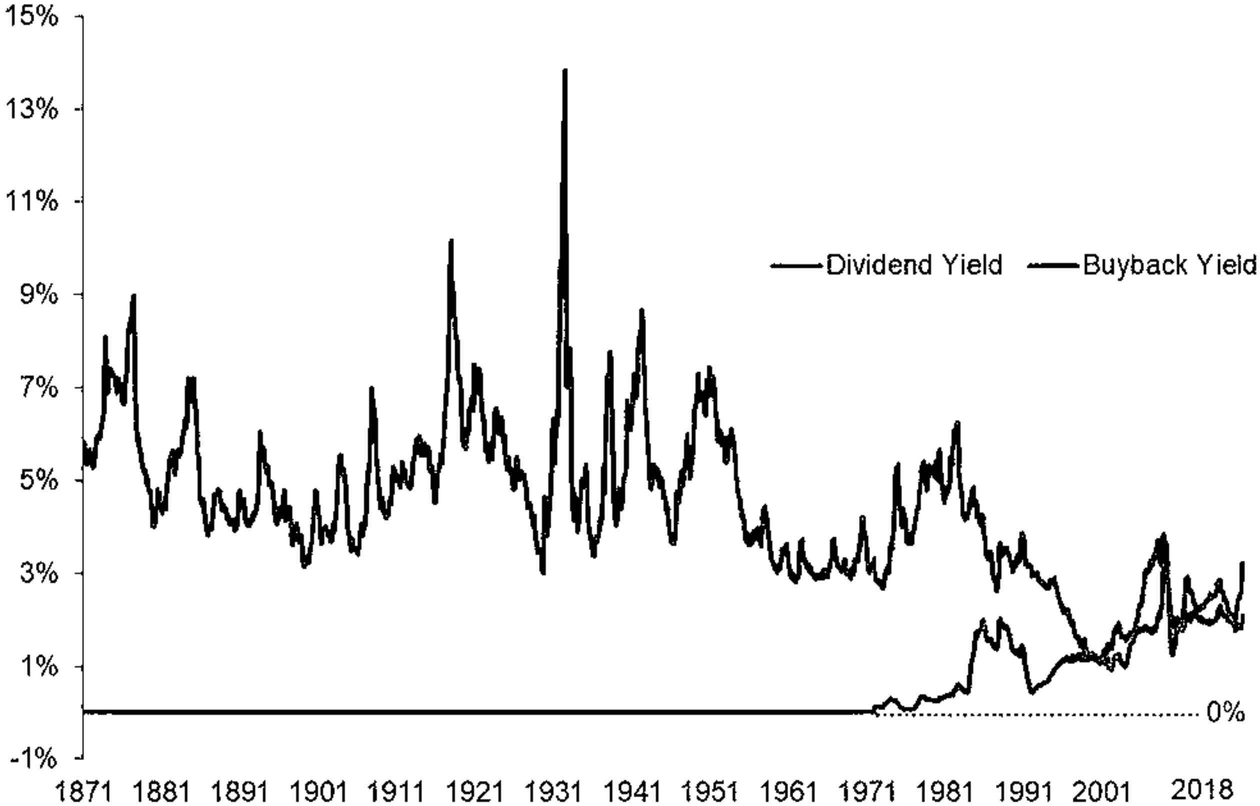

In recent decades, companies increasingly prefer to make payouts in the form of buybacks instead of dividend payments. There are several reasons for this. Most important, the amount of buybacks is completely flexible. A company can aggressively buyback one year and skip the next year without major signaling effects, unlike dividend policy. Also, if a company does not have a good use for its cash, buybacks can increase earnings per share by decreasing the number of shares outstanding. Furthermore, companies can buyback shares when they believe the price is attractive, both potentially boosting the price and benefitting the holding shareholders. Finally, through much of history, the tax treatment of buybacks was more favorable than it was for dividend payouts.

Dividend payout models are typically wrongly applied in the era of buybacks. They often estimate the future returns to be the sum of the current dividend yield plus historical dividend growth. This is wrong for two reasons: the current yield is artificially too low since it only includes one source of income, and the historical growth rate is too low because it ignores the shift in payouts away from dividends.

Thus, the advent of buybacks has created a need for models that can explicitly take into account buybacks. Although buybacks are similar to dividends in that they are a way of paying out cash, buybacks have a different impact on the return generating process than dividends do. For example, the buy and hold investor receives dividends in the form of income, while investors receive buybacks as a price increase because the buy and hold investor's share of the company is increased. Prior studies, including Ibbotson and Chen's 2003 paper, disregarded the fact that return components are sensitive to a company's payout method (i.e., dividends versus buybacks).

Three Total Payout Models of Stock Returns

Miller and Modigliani[209] proved that in a perfect capital market the total return of stocks should be independent of the payout method. The Dividend Per Share Model as typically applied is not independent since higher buybacks make for less historical growth and lower current dividend yields. This is not to say the dividend model is incorrect because it is the future growth that becomes higher as the number of shares diminishes. However, by taking the buybacks explicitly into account, past payout growth can once again be an indicator of future growth, and current yields can reflect the full payouts.

Straehl and Ibbotson present three payout models, all of which include inflation:

1. The Dividend Per Share Model, where the investor gets a dividend yield plus a growth in total payouts, plus the change in payout per share, plus the change in price to total payout. Here, the buy and hold investor gets higher future growth to offset the lower dividend yield.

2. Dividend and Cash Buyback Model, where the investor gets the total yield (dividend plus buyback), plus growth in payout per share adjusted for share decrease, plus change in price to total payout. Here, the buy and hold investor gets the full payouts.

3. Dividend Less Net Issuance Model, where the investor gets the net total yield (dividend plus buyback but diluted by issuance), plus aggregate payout growth, plus rate of change in total payout. Here, the investor gets diluted by issuance but increased ownership by the buybacks.

In all three cases, the historical return is the sum of the components are all equal no matter which of the three methods are used. In the Straehl and Ibbotson historical samples, the total return from 1871-2018 was 9.02%, the total return from 1901-2018 was 9.58%, and the total return from 1970-2018 was 10.20%.

In all three supply-side models the realized return was the same, and most of the real return came from the payouts and the payout growth. However, the nature of the payouts explains what portion of the return comes from the payouts versus the payout growth.

209 Miller, Merton H., and Franco Modigliani 1961, "Dividend Policy, Growth, and the Valuation of Shares" Journal of Business, Vol.34, no.4 (October): 411-433.

Total Payouts and the Real Economy

The stock market should link to the real economy over the long run. More generally, Diermeier, Ibbotson, and Siegel (1984)[210] measured the full scope of financial assets and stressed the importance of capital markets being "macro consistent" with the real economy. In the long run, financial assets cannot continually outgrow the real economy, or financial assets would eventually become the whole economy. On the other hand, financial assets cannot continually underperform, or they would become a smaller and smaller part of the economy. For example, Ibbotson and Chen (2003) showed that earning per share growth for the U.S. stocks were comparable to U.S. GDP per capita growth.

In attempting to link the U.S. stock market to the U.S. real economy, Straehl and Ibbotson focus on growth rates rather than the payouts themselves. In particular, the long-term growth rate of aggregate stock payouts should link to the aggregate GOP growth, and the long-term growth rate of per share payout growth rates should link to the GDP per capital growth rate.

Exhibit 10.13 is an updated chart from Straehl and Ibbotson. It shows that the growth in aggregate stock market (as measured by the S&P Composite Index) roughly matches the aggregate real GDP growth. This link up is better than dividend growth by itself, which would have underestimated payout growth.

Exhibit 10.13: Growth in Aggregate Total Payout vs. GDP Growth 1901–2018

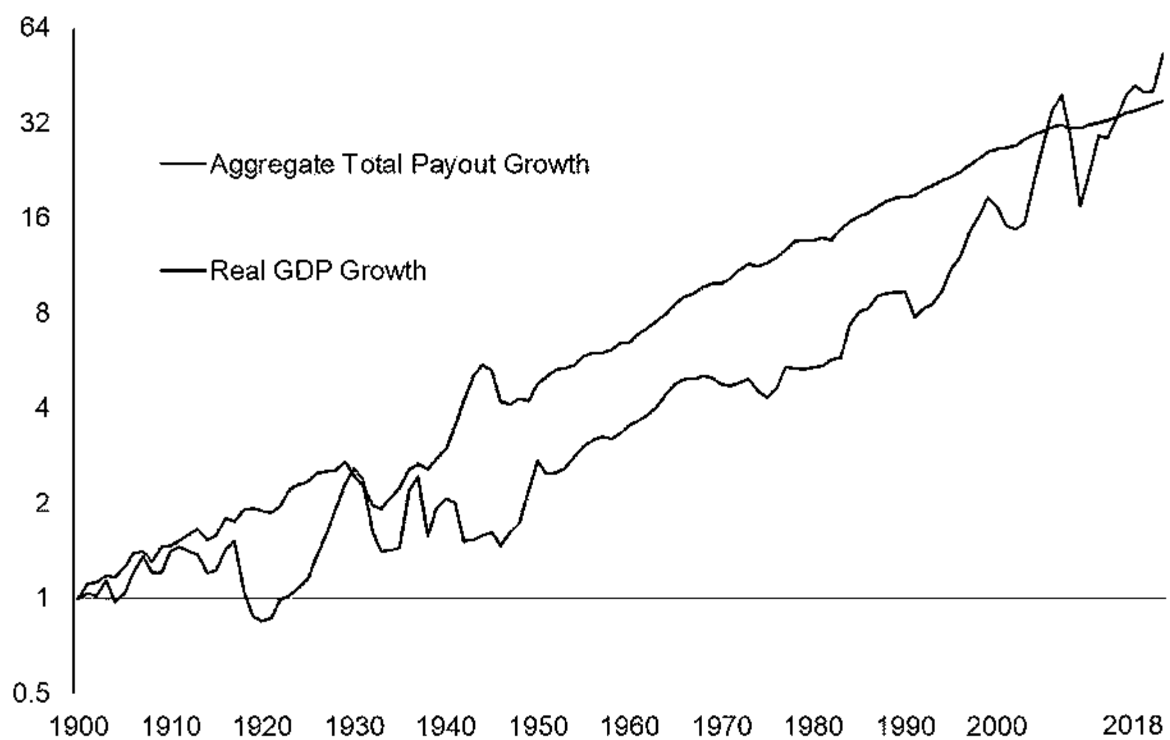

[210] Diermeier, Ibbotson, and Siegel, 1984 "The Supply of Capital Market Returns" Financial Analysts Journal, Vol.40 no.2 (March/April).

Exhibit 10.14 similarly shows real GDP per capital growth is linked to real total payout per share growth. In this case the data goes all the way back to 1871, and again is an updated version of an exhibit in the Straehl and Ibbotson article.

Exhibit 10.14: Growth in Total Payout per Share vs. Growth in GDP per Capita 1872–2018

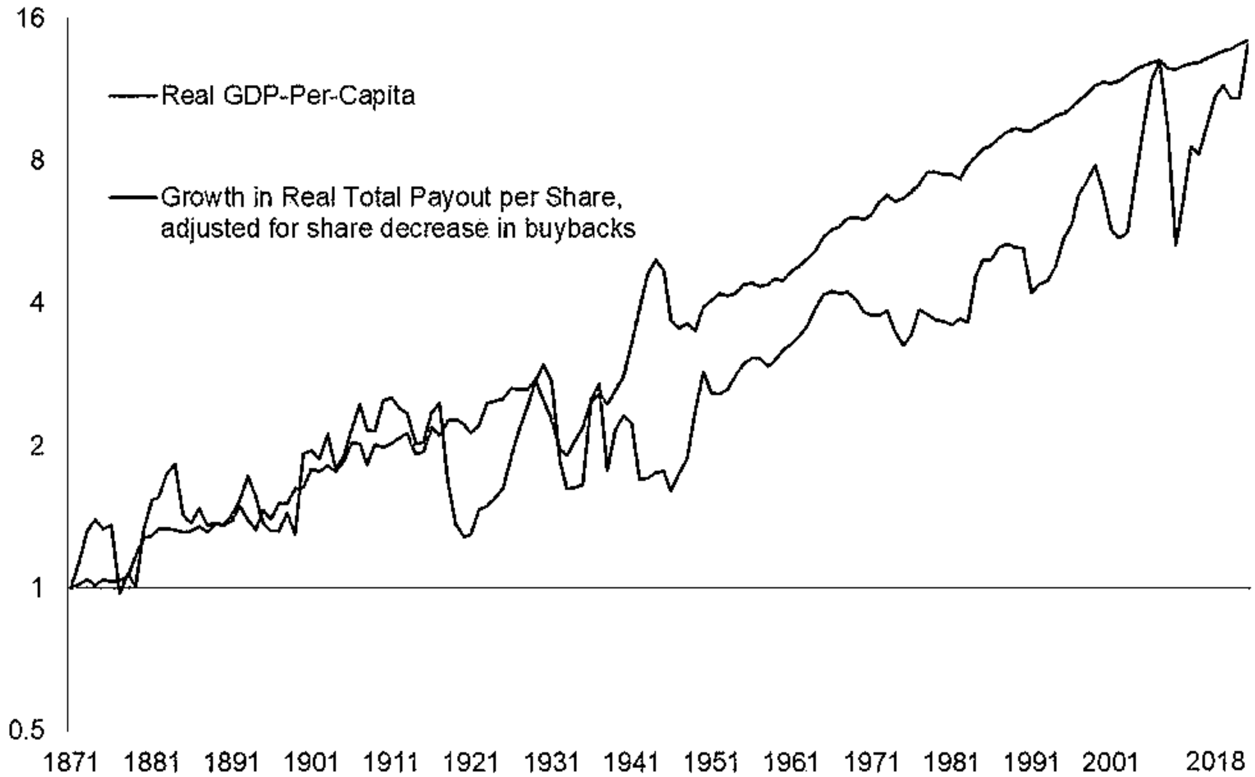

Overall, these results indicate that total payout growth in the stock market is roughly equal to GDP growth over the long run. Of course, this is not necessarily true over shorter intervals, but still total payouts are a useful tool in explaining long run stock returns and their interaction with the overall economy. We now look specifically at forecasting the stock market with supply-side models.

Forecasting Equity Returns

This chapter of the Yearbook is concerned with using historical data to forecast returns which can potentially be used for long-term planning and optimizing portfolios. The long-term historical data that Straehl and Ibbotson provide can be useful in forecasting.

The payout expected real return models presented here have two components: Payout Yield and Real Payout Growth. It is also necessary to add a small interaction term to match up with the historical rates.

Taking the full historical 1871-2018 period that Straehl and Ibbotson studied, we can compare the annual dividend yield growth (1.59%) to the total payout growth (1.79%). We can also compare a recent current dividend yield (1.83%) to a recent payout yield (4.21%). Adding the historical interaction term to both series (0.25% and 0.23% respectively) gives us the two forecasts. The expected real return using the Dividend Yield method is 3.67%. The expected real return using the Total Payout method (which includes buybacks) is 6.23%.

Exhibit 10.15: Long-Run Expected Returns Based On the Current Payout Yield and Historical Growth Over 1871–2018

	Recent Yield	Historical Growth	Interaction	Expected Real Return
Dividend Yield Method	1.83%	1.59%	0.25%	3.67%
Total Payout Method	4.21%	1.79%	0.23%	6.23%

The dividend yield model (the dividend discount model, DDM) is often incorrectly applied, using the current yield with historical growth rates. As can be seen above, both the current payout and the historical growth are too low. The DDM is a theoretically correct model, but a proper interpretation would be to increase the future growth rate caused by today's artificially low yields. The total payout method makes it clear what is going on. Payouts are switching from dividend yields to buybacks, but the overall total payout yields are relatively constant over time. The switch, however, leads to low historical dividend yield growth and low current dividend yields.

Conclusion

The total payout models presented here are useful in forecasting equity returns.

1. Changing payout policies should not change expected returns, as shown by Miller and Modigliani.

2. Payout growth should be linked long-term to the real economy.

3. Using the dividend discount model, although theoretically correct, is often incorrectly implemented with low current yields and low historical dividend growth rates.

4. Total Payout Models give more reasonable supply-side forecasts of returns.

Chapter 11
Stock Market Returns From 1815–2020

Studies on the long-horizon predictability of stock returns, by necessity, require a database of return information that dates as far back as possible. Since the late 1970s, Ibbotson Associates has produced a broad set of historical returns on asset classes dating back to 1926. Researchers interested in the dynamics of the U.S. capital markets prior to 1926 had to rely on indexes of uneven quality. In 2000, Roger G. Ibbotson and William N. Goetzmann, professors of finance, and Liang Peng, then a Ph.D. candidate in finance, all at Yale School of Management, assembled a New York Stock Exchange database of annual returns for the periods prior to 1926. The first part of this chapter covers the sources and construction of this annual return database extending back to 1815.

The second part of this chapter introduces a new set of monthly real stock market total returns developed by Paul Kaplan, now director of research at Morningstar Canada. Kaplan used these new returns to demonstrate that the severity of the financial crisis of 2008 was not unique but was merely the latest chapter in a long history of market meltdowns.

While we firmly believe that a 1926 starting date was approximately when quality financial data came into existence, our hope is that the continuing development of these data sets will allow modern researchers of pre-1926 stock returns, along with future researchers, to test a broad range of hypotheses about the U.S. capital markets as well as open up new areas for more accurate analysis.

1815–1925 Data Series Sources and Collection Methods

Share Price Collection

End-of-month equity prices for companies listed on the New York Stock Exchange were hand collected from three sources published January 1815 to December 1870. For the period 1871 through 1925, end-of-month NYSE stock prices were collected from the major New York newspapers. *The New York Shipping List,* later called *The New York Shipping and Commercial,* served as the "official" source for NYSE share price collection up until the early 1850s. In the mid-1850s, *The New York Shipping List* reported prices for fewer and fewer stocks. This led to the collection of price quotes from *The New York Herald* and *The New York Times.* While neither claimed to be the official list for the NYSE, the number of securities quoted by each far exceeded the number quoted by *The New York Shipping List.*

It is important to note that in instances where no transaction took place in December, the latest bid and ask prices were averaged to obtain a year end price. In total, at least two prices from 664 companies were collected. From a low number of eight firms in 1815, the number of firms in the index reached a high point in May of 1883 with 114 listed firms.