1

RETURNS

This is a book about the conceptual foundations of investing. That does not mean concepts for beating the market. In fact, one of the conceptual foundations that we will return to throughout the book is that there cannot be a trick for beating the market. If there were, and if the trick became well known, who would sell when the trick said buy? The best that can be hoped for is that a strategy for beating the market may work for a while as long as it is not widely known and adopted. Of course, no one would write a book about such a strategy; they would start an investment firm.

That does not mean that understanding the conceptual foundations of investing will not improve an investor's performance. There are a host of investment mistakes that can be avoided by such an understanding. One example involves the trade-off between risk and return. The trade-off seems to imply that if you bear more risk you will have higher long-run average returns. That conclusion is false. It is possible to bear a great deal of risk and get no benefit in terms of higher average return. Understanding the conceptual foundations of finance makes it clear why this is so and, thereby, helps an investor avoid bearing uncompensated risks.

Another choice every investor has to make is between active and passive investing. Making that choice wisely requires understanding the conceptual foundations of investing.

There are numerous other examples we could offer but we are getting ahead of ourselves. Before drawing conclusions, it is essential to lay the proper ground work. In finance and investing everything starts with the concept of returns. Just as the atom is the fundamental unit of analysis in chemistry, the return is the fundamental unit of analysis in investing. The first step in being able to analyze investing properly is becoming comfortable calculating and working with returns. For that reason, our book starts with returns.

The return on an investment is the percentage increase in your wealth associated with holding an investment for a given time period. For example, if you invest $10,000 and earn a 1% return your wealth has increased to $10,100. While this may seem entirely straightforward, much mischief can arise when calculating returns. Because they are the "atoms" of finance, it is critical to understand how returns are calculated and used before turning to more abstract concepts like expected returns or the trade-off between risk and expected return.

One convention we will follow throughout this book is that a "day" will always refer to a trading day. No distinction is drawn, for example, between the trading interval that runs from the close of Friday to the close on Monday as opposed to the close on Monday to the close on Tuesday. Both of these are treated as trading days. The same is true of holidays and three-day weekends. Using this convention, there are typically 252 trading days in a year.

Let's get started with an example. Be prepared to do a little math. There is more to returns than you might expect. The first column of Exhibit 1.1 presents the price of Apple stock for 42 trading days from January 3, 2017 to March 3, 2017. As the exhibit shows, this was a good two months for Apple. The price rose from $116.15 to $139.78.

The third column of the exhibit shows the percentage change in the price of Apple stock on a daily basis. A common mistake is to associate the percentage change in the price of a security with the return. The error is common because on most days it is not a mistake – the return and the percentage price change are the same. But not on every day. That is because Apple pays a dividend and that dividend is part of the return.

There is a problem incorporating the dividend when calculating the return. On what day do you add in the dividend? The obvious answer appears to be on the day it is paid, but that is wrong because markets are forward looking. The correct day is what is called the ex-dividend date (commonly referred to as the "ex-date"), which is the day after the day on which Apple checks its shareholder records to decide who gets the dividend. If you own Apple shares the day before the ex-date, you get the dividend. If you do not buy until the ex-date, you no longer get the dividend. Therefore, the price of the Apple shares drops by the amount of the dividend on the ex-date (holding other factors that may affect the price constant). This means the dividend should be added to the price change on the ex-dividend date when computing returns.

Dividends are not only source of income on securities; bonds typically make payments every six months and mortgages generally pay monthly. All these cash distributions must be taken account of to properly compute returns. This leads to the mathematical

EXHIBIT 1.1 **Apple returns and path of wealth (POW).**

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Date	Apple closing price	Percentage price change (%)	Dividend and ex-date	Apple return (%)	Path of wealth – POW	Average return (%)	POW from average returns
1/3/2017	116.15				100.00		100.00
1/4/2017	116.02	−0.112		−0.112	99.89	0.469	100.47
1/5/2017	116.61	0.509		0.509	100.40	0.469	100.94
1/6/2017	117.91	1.115		1.115	101.52	0.469	101.41
1/9/2017	118.99	0.916		0.916	102.45	0.469	101.89
1/10/2017	119.11	0.101		0.101	102.55	0.469	102.37
1/11/2017	119.75	0.537		0.537	103.10	0.469	102.84
1/12/2017	119.25	−0.418		−0.418	102.67	0.469	103.33
1/13/2017	119.04	−0.176		−0.176	102.49	0.469	103.81
1/17/2017	120.00	0.806		0.806	103.31	0.469	104.30
1/18/2017	119.99	−0.008		−0.008	103.31	0.469	104.79
1/19/2017	119.78	−0.175		−0.175	103.13	0.469	105.28
1/20/2017	120.00	0.184		0.184	103.31	0.469	105.77
1/23/2017	120.08	0.067		0.067	103.38	0.469	106.27
1/24/2017	119.97	−0.092		−0.092	103.29	0.469	106.76
1/25/2017	121.88	1.592		1.592	104.93	0.469	107.26
1/26/2017	121.94	0.049		0.049	104.98	0.469	107.77
1/27/2017	121.95	0.008		0.008	104.99	0.469	108.27
1/30/2017	121.63	−0.262		−0.262	104.72	0.469	108.78
1/31/2017	121.35	−0.230		−0.230	104.48	0.469	109.29
2/1/2017	128.75	6.098		6.098	110.85	0.469	109.80
2/2/2017	128.53	−0.171		−0.171	110.66	0.469	110.32
2/3/2017	129.08	0.428		0.428	111.13	0.469	110.83
2/6/2017	130.29	0.937		0.937	112.17	0.469	111.35
2/7/2017	131.53	0.952		0.952	113.24	0.469	111.87
2/8/2017	132.04	0.388		0.388	113.68	0.469	112.40
2/9/2017	**132.42**	**0.288**	**0.57**	**0.719**	**114.50**	**0.469**	**112.93**
2/10/2017	132.12	−0.227		−0.227	114.24	0.469	113.45
2/13/2017	133.29	0.886		0.886	115.25	0.469	113.99
2/14/2017	135.02	1.298		1.298	116.75	0.469	114.52
2/15/2017	135.51	0.363		0.363	117.17	0.469	115.06
2/16/2017	135.35	−0.118		−0.118	117.03	0.469	115.60
2/17/2017	135.72	0.273		0.273	117.35	0.469	116.14
2/21/2017	136.70	0.722		0.722	118.20	0.469	116.68
2/22/2017	137.11	0.300		0.300	118.55	0.469	117.23
2/23/2017	136.53	−0.423		−0.423	118.05	0.469	117.78
2/24/2017	136.66	0.095		0.095	118.16	0.469	118.33
2/27/2017	136.93	0.198		0.198	118.40	0.469	118.88
2/28/2017	136.99	0.044		0.044	118.45	0.469	119.44

EXHIBIT 1.1 (Continued)

(1) Date	(2) Apple closing price	(3) Percentage price change (%)	(4) Dividend and ex-date	(5) Apple return (%)	(6) Path of wealth – POW	(7) Average return (%)	(8) POW from average returns
3/1/2017	139.79	2.044		2.044	120.87	0.469	120.00
3/2/2017	138.96	−0.594		−0.594	120.15	0.469	120.56
3/3/2017	139.78	0.590		0.590	120.86	0.469	121.13
Arithmetic average return				0.469			
Geometric average return					0.452%		

definition of the return on a security between two dates, t and t − 1.

$$R_t = [(P_t – P_{t-1}) + \text{Cash Payout}_t]/P_{t-1}. \qquad (1.1)$$

If the security in question were a stock, the cash payout would be the dividend and it would be added on the ex-date, but Eq. (1.1) holds for any security.

The fifth column of Exhibit 1.1 presents the sequence of returns on Apple stock. It differs from the percentage price change only in February 9, 2017, which was the ex-date. On that day, which is depicted in bold, the dividend is added to the change in price to compute the return, as shown in Eq. (1.1).

Once you have a series of returns it is possible to calculate one of the most important measures of investment performance, the path of wealth or POW. The POW shows the value of your investment from a given starting point, $100 in Exhibit 1.1. The calculation assumes that any dividends received are reinvested in the security in question – Apple stock in the exhibit. The POW is presented in the sixth column of Exhibit 1.1. It shows that an investor who invested $100 in Apple stock on January 3, 2017 would have an investment worth $120.86 as of the market close on March 3, 2017. It also shows the value of that initial $100 investment for each day in the two-month period.

Investment performance should always be assessed using returns and POWs, not price charts. The problem is that much financial performance data presented in the media are based on price charts, not POWs. This is true not only for individual stocks but also for the best-known indexes. For instance, neither the Dow Jones index nor the S&P 500 index takes account of dividends. Therefore, if you compare the performance, of say, a mutual fund you own with the S&P 500, you have an apples to oranges problem. Mutual fund performance data typically are based on returns, whereas the S&P 500 is a price index that excludes dividends. As a result, the performance of the portfolio of stocks that the S&P 500 is comprised of is significantly better than the price appreciation of the index because many stocks in the index pay dividends. The takeaway is that when comparing two investments you want to be sure to compare POWs. This is not often easy. For example, return data for the S&P 500 and Dow Jones index are not readily available.

It may seem like the dividend issue is a minor annoyance. In Exhibit 1.1, the dividend accounts for a minor part of the total return on Apple stock. But while stock prices move up and down, dividends are never negative. As the investment holding period grows, the impact of dividends becomes more evident. To appreciate the importance of dividends, take a look at Exhibit 1.2, which plots the POW for the U.S. stock market from 1926 to 2017, both including and excluding dividends.

Before interpreting the results, a word on the data. The POWs shown in Exhibit 1.2 are calculated using data from the Center for Research in Securities Prices (CRSP) at the University of Chicago. CRSP provides daily data on the returns for virtually all listed U.S.

EXHIBIT 1.2 U.S. stock market POW with and without dividends: 1926–2017.

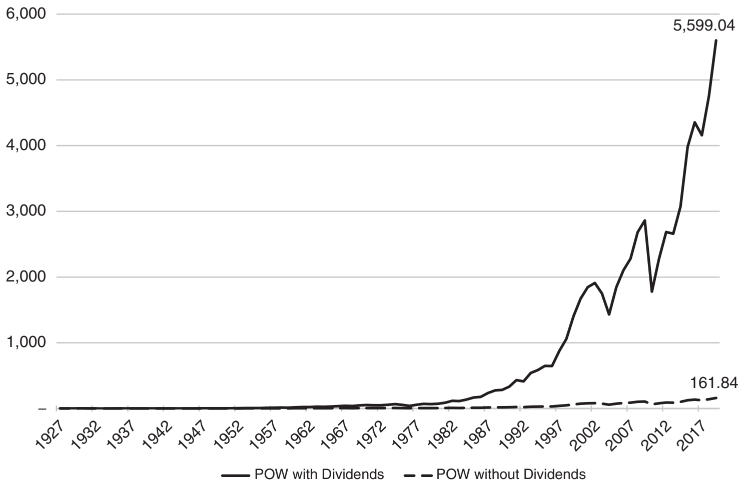

stocks back to 1926. What makes the CRSP data so convenient is that all the hard work of computing returns, such as adding dividends on the ex-date, has been done. CRSP also reports the market return, with and without dividends, for a value-weighted index of all listed stocks. As such, the CRSP index is far more comprehensive than the Dow, and even a good deal more comprehensive than the S&P 500. For that reason, it is used to measure market performance in most academic studies and, unless otherwise noted, when we refer to the market portfolio in this book, we mean the CRSP index.

Turning back to Exhibit 1.2, the results highlight the importance of accounting for dividends when calculating returns. The exhibit shows that an investment of $1 in the CRSP market portfolio in 1926 would be worth $5,599.04 if all dividends were reinvested along the way! If dividends are excluded, the value of the $1 investment grows only to $161.84 in 2017. This demonstrates that it is critical to properly include dividends, or other cash payouts, when computing POWs and not to be misled by price indexes.

Exhibit 1.2 should not be interpreted as saying that stocks that pay dividends offer higher returns than those that do not. The message simply is that if stocks do pay a dividend, it must be taken into account when computing the POW. With regard to comparing stocks that pay dividends versus those that do not, if we hold constant risk and taxes, there is no reason why the average long-run return should be different. Remember that the price of a stock tends to fall by the amount of the dividend on the ex-date. For stocks that do not pay dividends, there is no dividend but there is also no drop, so the return is unaffected. This is another reason to be sure to work with returns and not price changes.

As a further illustration of the utility of POWs, Exhibit 1.3 plots the POW for companies Coke, GE, IBM, and Amazon, along

EXHIBIT 1.3 **POWs for a sample of companies: January 2000–July 2016.**

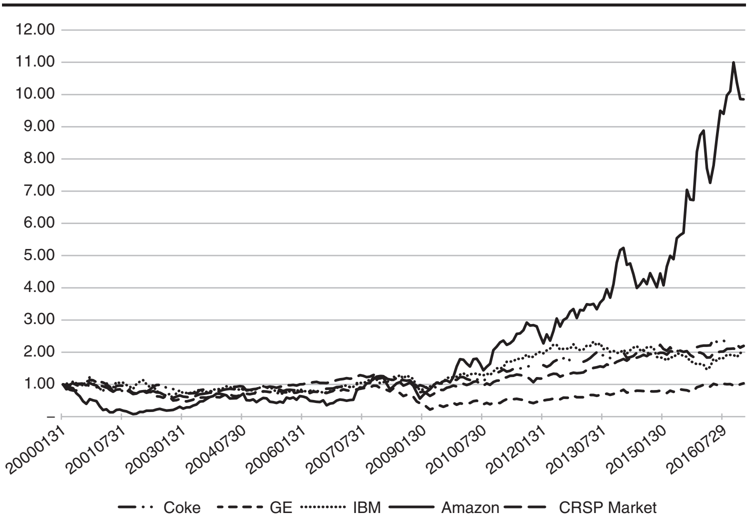

with the CRSP market index. The POWs are calculated using monthly return data from CRSP. One convenient feature of the CRSP data is that it provides monthly returns directly, avoiding the need to build them up from daily data. The exhibit shows that two of the companies, IBM and Coke, basically mirrored the market index while GE significantly underperformed and Amazon markedly outperformed.[1] Calculating POWs in this fashion is the proper way to compare the performance of various securities.

POWs can also be used to compare different measures of the market. Of the three market indexes we have discussed so far, the Dow is a particularly bad measure of the market because it contains only 30 stocks and because it is not calculated based on the market values of the constituent securities. Both the S&P 500 and the CRSP market index are good choices. They are both weighted by the value of the constituent securities. This means an investor could actually buy and hold both of these portfolios and match the index performance.[2] As noted previously, we will generally use the CRSP index in this book because it covers all securities traded on the New York, American, and NASDAQ markets stock exchanges.[3] Given this choice, a natural question to ask is: How much difference does the choice make when assessing market performance? Exhibit 1.4 answers the question.

The exhibit plots the POW using monthly data from 1926 to 2016 for both the S&P 500 and for the CRSP index. The main takeaway is that the two measures are very similar. Although the lines nearly overlap, they are not identical. Over the full period, the

[1]As interesting side note, famed investor Warren Buffett owned shares of IBM, Coke, and GE at times during the period, but never owned any shares of Amazon.
[2]The holding would have to be adjusted slightly to account for stocks leaving or entering the indexes.
[3]Since its acquisition by the New York Stock Exchange in 2009, the American Stock Exchange (AMEX) has been called the "NYSE Amex Equities."

EXHIBIT 1.4 S&P 500 versus CRSP market index: 1926–2016.

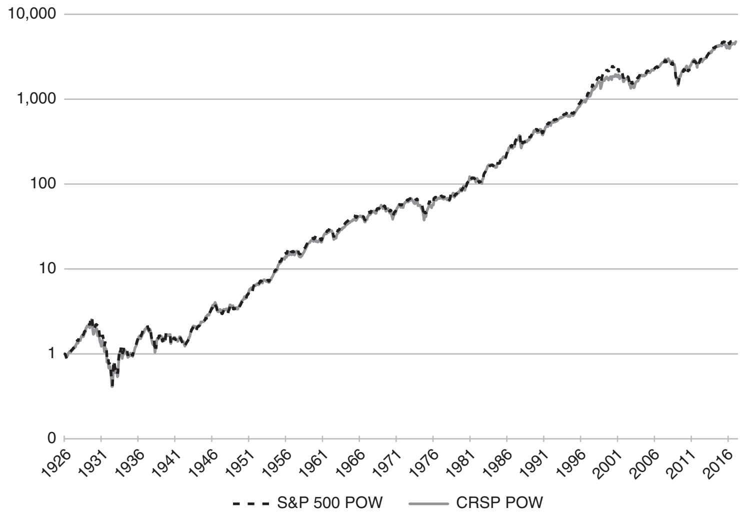

S&P 500 slightly outperforms the CRSP index. Therefore, when someone talks about the market, it is a good idea to ask them what they are talking about.

STOCKS, BONDS, AND BILLS

Calculation of POWs also makes it possible to compare the investment performance of various classes of investments. This is something we make use of throughout this book. As an initial example, Exhibit 1.5 compares what are probably the three most important classes of investments in securities – common stocks, long-term Treasury bonds, and short-term Treasury bills. Bonds and bills are described further in a subsequent chapter. For now, all you need to know is that they are obligations of the U.S. government that promise fixed future payments. The exhibit is

EXHIBIT 1.5 **Stocks, bonds, and bills: 1926–2017.**

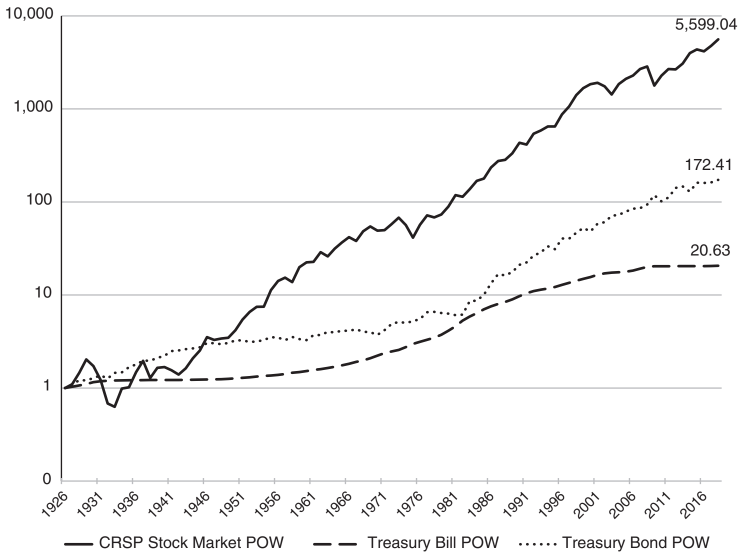

plotted on a logarithmic scale because the performances of the three asset classes are so different. For convenience, the return data underlying the POWs are reported in Exhibit 1.6. The exhibit underscores how what seem like relatively small differences in average returns translate into remarkably large differences in final wealth when compounded over 92 years. Whereas $1 invested in the CRSP index in 1926 grows to $5,599.04 in 2017, the same investment in Treasury bills grows only to $20.63. Treasury bonds are in the middle, with the $1 investment growing to $172.41.

A word to the wise. The vast performance differences between stocks, bonds, and bills is not something that one can automatically expect to continue going forward. What history says about what can be expected going forward is an issue we address in depth later in the book.

EXHIBIT 1.6 Return data for stocks, bonds, bills.

Date	CRSP stock market returns (%)	Treasury bill returns (%)	Treasury bond returns (%)	CRSP stock market POW	Treasury bill POW	Treasury bond POW
				1.00	1.00	1.00
1926	9.85	3.19	9.01	1.10	1.03	1.09
1927	32.87	3.12	11.33	1.46	1.06	1.21
1928	39.14	3.82	−0.52	2.03	1.10	1.21
1929	−15.10	4.74	6.12	1.72	1.16	1.28
1930	−28.90	2.35	6.76	1.23	1.18	1.37
1931	−44.39	1.02	−7.38	0.68	1.20	1.27
1932	−7.94	0.81	14.99	0.63	1.21	1.46
1933	57.41	0.29	1.20	0.99	1.21	1.47
1934	3.18	0.15	13.63	1.02	1.21	1.67
1935	45.45	0.17	6.95	1.48	1.21	1.79
1936	32.32	0.17	9.53	1.96	1.22	1.96
1937	−34.60	0.32	0.43	1.28	1.22	1.97
1938	28.44	0.04	6.78	1.65	1.22	2.10
1939	1.84	0.01	5.62	1.68	1.22	2.22
1940	−7.51	−0.06	12.37	1.55	1.22	2.50
1941	−10.04	0.04	1.48	1.40	1.22	2.53
1942	16.72	0.26	3.22	1.63	1.22	2.62
1943	27.97	0.34	2.08	2.09	1.23	2.67
1944	21.36	0.32	2.81	2.53	1.23	2.75
1945	39.06	0.32	10.73	3.52	1.24	3.04
1946	−6.42	0.35	−0.10	3.29	1.24	3.04
1947	3.29	0.46	−2.63	3.40	1.25	2.96
1948	2.13	0.98	3.40	3.47	1.26	3.06
1949	20.11	1.11	6.45	4.17	1.27	3.26
1950	30.47	1.21	0.06	5.45	1.29	3.26
1951	20.94	1.48	−3.94	6.59	1.31	3.13
1952	13.33	1.64	1.16	7.46	1.33	3.16
1953	0.38	1.78	3.63	7.49	1.35	3.28
1954	50.41	0.86	7.19	11.27	1.36	3.52
1955	25.41	1.56	−0.69	14.13	1.38	3.49
1956	8.58	2.42	−6.27	15.35	1.42	3.27
1957	−10.35	3.13	8.22	13.76	1.46	3.54
1958	44.78	1.42	−5.29	19.92	1.48	3.35
1959	12.65	2.82	−2.51	22.44	1.52	3.27
1960	1.21	2.58	13.32	22.71	1.56	3.71
1961	26.96	2.16	0.19	28.83	1.60	3.71
1962	−9.93	2.72	7.80	25.97	1.64	4.00
1963	21.40	3.15	−0.79	31.53	1.69	3.97
1964	16.35	3.52	4.11	36.68	1.75	4.13
1965	14.06	3.96	−0.27	41.84	1.82	4.12

EXHIBIT 1.6 (Continued)

Date	CRSP stock market returns (%)	Treasury bill returns (%)	Treasury bond returns (%)	CRSP stock market POW	Treasury bill POW	Treasury bond POW
1966	−8.86	4.71	3.96	38.13	1.91	4.29
1967	26.84	4.15	−6.02	48.36	1.99	4.03
1968	12.75	5.29	−1.20	54.53	2.09	3.98
1969	−9.82	6.59	−6.52	49.18	2.23	3.72
1970	1.29	6.38	12.69	49.81	2.37	4.19
1971	15.84	4.32	16.70	57.70	2.47	4.89
1972	17.64	3.89	5.15	67.88	2.57	5.14
1973	−16.92	7.06	−2.49	56.39	2.75	5.02
1974	−26.81	8.08	3.89	41.27	2.97	5.21
1975	37.66	5.82	6.10	56.82	3.15	5.53
1976	26.25	5.16	18.18	71.73	3.31	6.53
1977	−4.84	5.15	0.90	68.26	3.48	6.59
1978	7.33	7.31	−2.93	73.27	3.73	6.40
1979	21.88	10.69	−1.52	89.30	4.13	6.30
1980	32.63	11.52	−3.52	118.44	4.61	6.08
1981	−4.14	14.86	1.16	113.53	5.29	6.15
1982	21.00	10.66	39.74	137.37	5.86	8.60
1983	22.76	8.85	1.28	168.63	6.38	8.71
1984	5.79	9.96	15.81	178.39	7.01	10.08
1985	31.74	7.68	31.96	235.01	7.55	13.30
1986	17.32	6.06	25.79	275.72	8.01	16.73
1987	2.89	5.38	−2.91	283.69	8.44	16.25
1988	17.57	6.32	8.71	333.53	8.97	17.66
1989	29.61	8.22	19.23	432.29	9.71	21.06
1990	−4.27	7.68	6.15	413.85	10.46	22.35
1991	30.65	5.51	18.59	540.71	11.03	26.51
1992	8.22	3.40	7.95	585.14	11.41	28.62
1993	10.75	2.90	16.91	648.04	11.74	33.46
1994	−0.09	3.88	−7.19	647.47	12.19	31.05
1995	35.07	5.53	30.38	874.51	12.87	40.48
1996	21.35	5.14	−0.35	1061.19	13.53	40.34
1997	32.32	5.08	15.46	1404.13	14.22	46.58
1998	19.13	4.78	13.05	1672.80	14.90	52.66
1999	10.38	4.56	−8.66	1846.43	15.57	48.10
2000	3.47	5.76	20.95	1910.59	16.47	58.17
2001	−8.45	3.78	4.09	1749.15	17.09	60.55
2002	−18.22	1.63	17.22	1430.54	17.37	70.98
2003	29.13	1.02	2.45	1847.32	17.55	72.72
2004	13.88	1.20	8.28	2103.67	17.76	78.74
2005	8.45	2.96	7.66	2281.49	18.29	84.77
2006	17.62	4.79	1.14	2683.55	19.16	85.73

EXHIBIT 1.6 (Continued)

Date	CRSP stock market returns (%)	Treasury bill returns (%)	Treasury bond returns (%)	CRSP stock market POW	Treasury bill POW	Treasury bond POW
2007	6.62	4.67	9.74	2861.21	20.06	94.08
2008	−37.83	1.47	25.60	1778.89	20.35	118.16
2009	28.13	0.10	−13.99	2279.23	20.37	101.63
2010	17.78	0.12	9.77	2684.42	20.40	111.56
2011	−0.89	0.04	26.99	2660.65	20.41	141.68
2012	15.51	0.06	3.88	3073.41	20.42	147.18
2013	29.45	0.03	−12.23	3978.64	20.42	129.17
2014	9.45	0.02	24.62	4354.70	20.43	160.97
2015	−4.55	0.01	−0.67	4156.62	20.43	159.89
2016	14.48	0.19	1.38	4758.33	20.47	162.10
2017	17.67	0.79	6.36	5599.04	20.63	172.41
Average	11.69	3.39	6.19			
Volatility	19.48	3.14	9.88			

RETURN MATHEMATICS

Although daily returns are the "atoms" of investment analysis, they are rarely reported as such. It is more common, for instance, to report returns in annual terms. This requires doing the math to convert from one interval to another. The basic formula that does that is

$$W_n = W_0 * (1 + r_1) * (1 + r_2) * (1 + r_3) * \cdots (1 + r_n) \quad (1.2)$$

where W_0 is the initial wealth invested, r_i is the return on day i, n is the total number of trading days the investment is held, and W_n is the final value of the investment.[4] Equation (1.2) is what was used to construct the POW – each trading day W_n is incremented by one day.

[4] Equation (1.2) is cumbersome because the returns compound rather than adding. There is a way around this complexity by using continuously compounded returns rather than standard returns. While continuous returns are often used in academic studies, standard returns remain the norm in almost all practical investment publications. Therefore, we use standard returns in this book. A detailed discussion of continuous returns can be found in any major investment textbook.

Equation (1.2) can also be used to compute returns over periods longer than one day. Suppose that you have a series of daily returns but need a series of weekly returns. Assuming that the week is five trading days long, weekly returns are calculated in two steps. First Eq. (1.2) is used to compute W_5 as

$$W_5 = W_0 * (1 + r_1) * (1 + r_2) * (1 + r_3) * (1 + r_4) * (1 + r_5).$$

Next, the weekly return is defined by the equation

$$W_5 = W_0 * (1 + r_{weekly}).$$

So that,

$$r_{weekly} = w_5/w_0 - 1.$$

There is nothing unique about the weekly return. Monthly or annual returns can be calculated in the same way, though the calculation is a bit cumbersome. A common calculation is converting monthly returns to annual returns using the formula,

$$r_{annual} = (1 + r_{monthly})^{12} - 1. \qquad (1.3)$$

In Eq. (1.3), the two-step procedure has been telescoped into one step.

Converting returns between various intervals can be a pain, because weeks and months do not always have the same number of trading days. Fortunately, sophisticated data sources like CRSP have already done the work. As well as providing daily return data, these sources allow for the downloading of monthly returns or annual returns directly.

A final word of caution: in the financial media conversions of returns from one interval to another are not always done using the compounding formulas described above. It is common, for instance, to multiply a monthly return by 12 to get an annual return. This is an error because it excludes the benefits that accrue from reinvestment of returns earned earlier in the period. Taking account of reinvestment, a return of 1% per month is properly converted to an annual return of 12.68% per year, not 12%.

VOLATILITY

Thus far, we have focused on the level of returns. However, their variability is also important. Return variability is typically called volatility in finance and is calculated as the standard deviation of a sequence of returns.

To give a visual feel for volatility, Exhibit 1.7 plots the annual returns for the CRSP stock market index and Treasury bills over the period from 1926 to 2017. The difference between the two series is immediately obvious. The Treasury bill returns are typically small and often close to zero, the only exception being the high-inflation

EXHIBIT 1.7 **CRSP market versus Treasury bill returns: 1926–2017.**

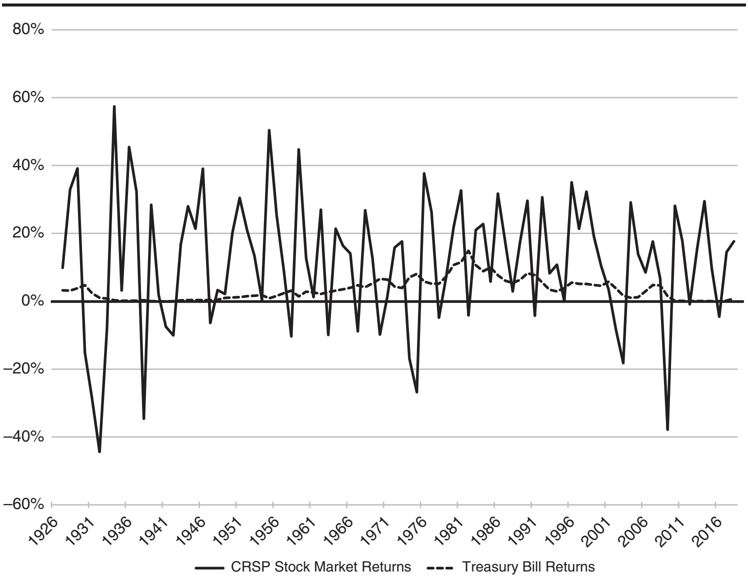

period of the late 1970s and early 1980s (an issue we will discuss in the chapter on inflation). More importantly, from the standpoint of volatility and risk, the Treasury bill returns are never negative. The CRSP market returns are dramatically different. Increases of around 40% occur in several years as do drops of 30%.[5]

Exhibit 1.7 also allows one to get a feeling for what finance professionals call "excess returns." Excess returns are the return on an investment minus the return on short-term Treasury bills. Excess returns can be thought of as a payment for bearing risk. We will discuss risk and return in much greater detail in Chapter 4. For now, notice that because Treasury bill returns are small and positive in every year and stock market returns have huge swings, there is not much difference between the return on stocks and the excess return in any individual year.

The bottom of Exhibit 1.6 shows that the volatility of the market index is 19.5% per year. The average return is 11.7%. This means that a 95% confidence interval constructed around the average return would run all the way from −25.3% to 48.7%! In contrast, the average return on Treasury bills is 3.4% and the volatility is only 3.1%. This suggests that there is a risk–return trade-off with volatility as the measure of risk. The suggestion carries a kernel of truth, but only a kernel. The actual trade-off between risk and return is a good deal more complex. We address the risk–return trade-off in detail in Chapter 4.

AVERAGE RETURNS

Because returns vary day to day, it is often useful to calculate the average return on an investment. For instance, suppose you buy

[5]Note that for the compounding calculations, the combination of a 40% increase and a 30% drop actually results in a 2% decline in the market.

IBM and Apple stock on the same day and would like to compare the average return on the two investments. It turns out that calculating the average is not as straightforward as you would hope. Turning back to Exhibit 1.1, the bottom of the fifth column that shows the Apple average return is 0.469%. More formally, this is known as the arithmetic average, but it is just the standard average with which you are familiar. It is calculated by dividing the sum of the returns by the number of returns.

It turns out that the standard arithmetic average has a peculiar property. To illustrate the problem, the average return is copied to all the cells in column seven. Next, a new POW is computed using the average return. That POW is reported in column eight. The problem is that the new POW shown in column eight is not equal to the original POW in column six. By the end of the period, the value of the investment has grown to \$121.13 instead of \$120.86.

There is nothing special about the Apple example. The ending value of the POW computed using the arithmetic average return will always exceed the actual ending POW unless all the returns are the same. Because of this anomaly, average returns are also computed using a procedure that ensures the ending POW values will be the same. This done by working with Eq. (1.2). The procedure for calculating the average begins by writing down the ending value calculated using actual series of investment returns as shown in Eq. (1.4)

$$W_n = W_0 * (1 + r_1) * (1 + r_2) * \cdots (1 + r_n) = W_0 * (1 + r_{av})^n \tag{1.4}$$

Equation (1.4) is then used to define the geometric average return, r_{av}, as that return that if earned and compounded every period produces the same ending value as the actual series of returns. Solving Eq. (1.4) for r_{av} gives,

$$r_{av} = (W_n/W_0)^{(1/n)} - 1 \qquad (1.5)$$

Applying Eq. (1.5) to the Apple returns gives a geometric average return of 0.452%. The geometric average is less than the arithmetic average, as it has to be because earning the arithmetic average led to a final value that exceeded the actual final value. It turns out that the difference between the geometric average and the arithmetic average depends on the variability of the series of returns. The more variable the returns, the greater this difference will be. But the geometric average is always less than the arithmetic average unless the returns are constant.

The foregoing should serve as a warning. Investors who hear that the average return on an investment was so and so should be careful. If it is not clear how the average was calculated, ask. To illustrate the mischief that can occur recall that the POW computed for the market showed that between 1926 and 2017 an investment of $1 grew to $5,599.04. During that time the arithmetic average return on the market as 11.69%. If the arithmetic average return is compounded for 92 years, it implies that an investment of $1 would have grown to $26,140.55 – almost five times the actual amount. This discrepancy highlights the importance of knowing how averages work.

Finally, returns can be computed for any asset as long as you have periodic data on prices and payouts. The Treasury bill return series used above is an example. But be aware that returns are often reported on assets like real estate for which periodic price data are not available. Under such circumstances, be sure to check what was used in the place of market prices when calculating returns. In the case of real estate, returns are often computed using periodic appraisals. But that means the return data are only as accurate as the appraisals.

USING RETURNS TO TEST INVESTMENT THEORIES

Suppose you read that internet stocks outperform utility stocks because they are more volatile. How do you tell if that is true? Before that question can even be addressed it must be translated into a testable hypothesis – that means it must be expressed in the language of returns. In terms of returns, the statement says that the average return and the average standard deviation of returns are both higher for internet stocks than they are for utility stocks. When stated this way, it becomes a testable hypothesis. Whether the test is meaningful and whether a positive finding implies causality are tricky issues, but that is not the point here. The point is that for statements about investments to have meaningful, testable content, they must be expressed in terms of returns.

As we write this, concern has been expressed in the financial press that if inflation were to accelerate suddenly it would be bad news for the stock market. If that concern is valid, then it should be the case that previous bouts of unexpected inflation were associated with market returns below the long-run average. Now that the concern has been translated into a statement about returns, it can be tested. Once again, we do not want to overplay the importance of the particular suggested test. The relation between inflation and the stock market may be more complicated than that. But however complicated it may be, the relation does not become meaningful and testable until it is stated in terms of returns.

What goes for the two simple examples above is true of the most sophisticated academic theories of asset pricing. They are all stated and tested in terms of returns.

The bottom line is you cannot begin to analyze investments until you are comfortable calculating and working with returns. Fortunately, modern spreadsheet software makes it relatively easy to perform the necessary calculations. In addition, the requisite data to construct sequences of returns and POWs is freely available at both Google Finance and Yahoo Finance. To help you in your efforts, the data used in constructing the exhibits is available at www.wiley.com/go/CornellCFOI (Password: CFI).

RETURNS AND STOCK MARKET HISTORY

The sequence of past returns is the best summary of stock market history. If there is some pattern to the market, it must be discernable in the history of returns. Consider, for instance, technical analysis. Technical analysts construct all sorts of complicated charts that they suggest give insight into where the market is going in the future. Those charts are just a way of summarizing return data, akin to the POWs. If there is a predictable pattern that can be discerned with charts and used to predict future performance, it must be there in the history of returns. This fact has not been lost on generations of graduate students. With continually improving computer technology and ever more complete financial data, they have been combing through return data for every single stock as well as market indexes, down to second-to-second trades, in an attempt to find patterns. In the last 50 years, thousands of papers have been published on the issue and there are undoubtedly thousands more that never got over the publication hurdle.

The basic message is that there are no patterns in stock prices – at least none that can be reliably exploited to earn superior risk-adjusted returns. To be fair there is on ongoing dispute about this issue that involves the question of data mining. We analyze data mining in more detail in Chapter 7. To anticipate that discussion, the issue raised by data mining is that even random series have quirks that look like exploitable patterns. If enough people pore over the same data, they will find those quirks, but the quirks will have no meaning. The dispute is over whether the quirks that have been found are meaningful. Nonetheless, virtually all researchers agree that in the case of stock market data even these quirks are few and far between.

Of course, it is possible that there are relationships that have been overlooked. If you think you have found an exception, the way to test it is using returns. For example, one of the earliest tests involved what is called autocorrelation. Some analysts believed that big positive returns tended to be followed by additional positive returns and vice-versa for negative returns. To a statistician this means that returns are positively autocorrelated. Tests for positive autocorrelation are easy to run and the answer, almost invariably, is that there is no meaningful autocorrelation. What you want to be careful to avoid is ad hoc theorizing based on a few observations. Any alleged pattern worth risking money on is worth testing carefully. And with the return data available today that is not hard. We will revisit to this issue in future chapters.

CONCEPTUAL FOUNDATION 1

The first conceptual foundation of investing is that returns are the fundamental unit of investment analysis. Investment performance should always be measured using a combination of return data and paths of wealth. Before any theory about stock market behavior can be tested, it must first be translated into the language of returns. In later chapters, when we turn to topics such as inflation and the risk–return trade-off, returns will be at the center of the discussion.