PART One

The Relative Pricing of Securities with Fixed Cash Flows

Consumers and businesses are willing to pay more than $1 in the future in exchange for $1 today. A newly independent adult borrows money to buy a car today, agreeing to repay the loan price plus interest over time; a family takes a mortgage to purchase a new home today, assuming the obligation to make principal and interest payments for years; and a business, which believes it can transform $1 of investment into $1.10 or $1.20, chooses to take on debt and pay the prevailing market rate of interest. At the same time, this willingness of potential borrowers to pay interest attracts lenders and investors to make consumer loans, mortgage loans, and business loans. This fundamental fact of financial markets, that receiving $1 today is better than receiving $1 in the future, or, equivalently, that borrowers pay lenders for the use of their funds, is known as the *time value of money*.

Borrowers and lenders meet in fixed income markets to trade funds across time. They do so in very many forms: from one-month U.S. Treasury bills that are almost certain to return principal and interest to the long-term debt of companies that have already filed for bankruptcy; from assets with a simple dependence on rates, like Eurodollar futures, to callable bonds and swaptions; from assets whose value depends only on rates, like interest rate swaps, to mortgage-backed securities or inflation-protected securities; and from fully taxable private-sector debt to partially or fully tax-exempt issues of governments and municipalities.

To cope with the challenge of pricing the vast number of existing and potential fixed income securities, market professionals often focus on a limited set of benchmark securities, for which prices are most consistently and reliably available, and then price all similar assets relative to those benchmarks. Sometimes, as when pricing a UK government bond in terms of other UK government bonds, or when pricing an EUR swap in terms of other EUR swaps, relative prices can be determined rigorously and for the most part accurately by *arbitrage pricing*. This methodology is developed in Chapter 1, where it is also shown that discounting, i.e., calculating present values with discount factors, is really just shorthand for arbitrage pricing.

While discount factors in many ways solve the relative pricing problem, they are not very intuitive for understanding the time value of money that is embedded in market prices. For this purpose, markets rely on spot, forward, and par rates. Chapter 2 introduces these rates and derives the relationships linking them to each other and to discount factors. The trading case study in Chapter 2, inspired by an abnormally shaped EUR forward swap curve, illustrates how fixed income analytics, market technicals (due to institutional factors described in the Overview), and a macroeconomic setting all contribute to a trade idea.

While the interest rates of Chapter 2 provide excellent intuition with respect to the time value of money embedded in market prices, other quantities provide intuition with respect to the returns offered by individual securities. The first half of Chapter 3 defines returns, spreads, and yields. Spreads describe the pricing of particular securities relative to benchmark government bond or swap curves and yields are the widely used, although somtimes misunderstood, internal rates of return on individual securities. The second half of Chapter 3 breaks down a security's return into several component parts. First, how does the security perform if rates and spreads stay the same? Second, how does the security perform if rates change? Third, how does the security perform if spreads change?

Given the central role of benchmarks in Part One, it is worth describing which securities are used as benchmarks and why. Until relatively recently, benchmark curves in U.S. and Japanese markets were derived from the historically most liquid markets, that is, from government bond markets. Recently, however, the benchmark has shifted significantly to swap curves. European markets, on the other hand, have for some time relied predominantly on interest rate swap markets for benchmarks because their swap markets have been, on average across the maturity spectrum, more liquid than government bond markets.

It is not hard to understand why government bond and interest rate swap markets are the preferred choices for use as benchmarks. First, they are the most liquid markets, consistently providing prices at which market participants can execute trades in reasonable size. Second, they incorporate information about interest rates that is common to all fixed income markets.

P1: TIX/b P2: c/d QC: e/f T1: g
JWBT539-c01 JWBT539-Tuckman July 20, 2011 11:47 Printer: Courier Westford

The value of a corporate bond, for example, depends on the interest rate information embedded in the government bond or swap curve in addition to depending on the credit characteristics of the individual corporate issuer.

But what about the choice between government bond and swap curves as benchmarks? Historically, government bonds were the only choice because swaps did not exist until the early 1980s and it took some time for their liquidity to become adequate. But bond markets have a significant disadvantage when used as a benchmark, namely that an individual bond issue is not a commodity in the sense of being a fungible collection of cash flows: bond issues are in fixed supply and have idiosyncratic characteristics. The best-known examples of nonfungibility are *on-the-run* U.S. Treasury bonds that trade at a premium relative to other government bonds because of their superior liquidity and financing characteristics. Put another way, pricing with a curve that is constructed from "similar" bonds, which are not on-the-run bonds, will underestimate the prevailing prices of on-the-runs. By contrast, an interest rate swap is really a commodity, that is, a fungible collection of cash flows. A 10-year, 4% interest rate swap cannot possibly be in short supply because any willing buyer and seller can create a new contract with exactly those terms. In fact, market practice bears out this distinction between bonds and swaps. While bond traders set prices for each and every bond they trade (although they certainly may use heuristics relating various prices to each other or to related futures markets), swap traders strike a curve that is then used to price their entire book of swaps automatically.

In short, global fixed income markets currently use interest rate swaps as benchmarks or base curves and build other curves from spreads or spread curves on top of swap curves. Even in the liquid U.S. Treasury market, strategists assess relative value using spreads of individual Treasury issues against the USD swap curve.

CHAPTER 1

Prices, Discount Factors, and Arbitrage

This chapter begins by introducing the cash flows of fixed-rate, government coupon bonds. It shows that prices of these bonds can be used to extract *discount factors*, which are the market prices of one unit of currency to be received on various dates in the future.

Relying on a principle known as *the law of one price*, discount factors extracted from a particular set of bonds can be used to price other bonds, outside the original set. A more complex but more convincing relative pricing methodology, known as *arbitrage pricing*, turns out to be mathematically identical to pricing with discount factors. Hence, discounting can rightly be used and regarded as shorthand for arbitrage pricing.

The application of this chapter uses the U.S. Treasury coupon bond and Separate Trading of Registered Interest and Principal of Securities (STRIPS) markets to illustrate that bonds are not commodities, meaning that their prices reflect individual characteristics other than their scheduled cash flows. This idiosyncratic component of bond valuation implies that the predictions of the simplest relative pricing methodologies only approximate the complex reality of bond markets.

The chapter concludes with a discussion of day-counts and accrued interest, pricing conventions used throughout fixed income markets and, consequently, throughout this book.

THE CASH FLOWS FROM FIXED-RATE GOVERNMENT COUPON BONDS

The cash flows from fixed-rate, government coupon bonds are defined by *face amount, principal amount*, or *par value*; *coupon rate*; and *maturity date*. For example, in May 2010 the U.S. Treasury sold a bond with a coupon rate of $2\frac{1}{8}\%$ and a maturity date of May 31, 2015. Purchasing \$1 million face amount of these "$2\frac{1}{8}$s of May 31, 2015," entitles the buyer to the schedule of

TABLE 1.1 Cash Flows of the U.S. $2\frac{1}{8}$s of May 31, 2015

Date	Coupon Payment	Principal Payment
11/30/2010	\$10,625	
5/31/2011	\$10,625	
11/30/2011	\$10,625	
5/31/2012	\$10,625	
11/30/2012	\$10,625	
5/31/2013	\$10,625	
11/30/2013	\$10,625	
5/31/2014	\$10,625	
11/30/2014	\$10,625	
5/31/2015	\$10,625	\$1,000,000

payments in Table 1.1. The Treasury promises to make a coupon payment every six months equal to half the note's annual coupon rate of $2\frac{1}{8}\%$ times the face amount, i.e., $\frac{1}{2} \times 2\frac{1}{8}\% \times \$1,000,000$, or \$10,625. Then, on the maturity date of May 31, 2015, in addition to the coupon payment on that date, the Treasury promises to pay the bond's face amount of \$1,000,000. One fact worth mentioning, although too small a detail to receive much attention in this book, is that scheduled payments that do not fall on a business day are made on the following business day. For example, the payments of the $2\frac{1}{8}$s scheduled for Sunday, May 31, 2015, would be made on Monday, June 1, 2015.

For concreteness and continuity of exposition this chapter restricts attention to U.S. Treasury bonds. But the analytics of the chapter apply easily to bonds issued by other countries because the cash flows of all fixed rate government coupon bonds are qualitatively similar. The most significant difference across issues is the frequency of coupon payments, which can be semiannual or annual; government bond issues in France and Germany make annual coupon payments, while those in Italy, Japan, and the UK make semiannual payments.

Returning to the U.S. Treasury market, then, Table 1.2 reports the coupons and maturity dates of selected U.S. Treasury bonds, along with their prices as of the close of business on Friday, May 28, 2010. Almost all U.S. Treasury trades settle $T + 1$, which means that the exchange of bonds for cash happens one business day after the trade date. In this case, the next business day was Tuesday, June 1, 2010.

The prices given in Table 1.2 are *mid-market*, *full* (or *invoice*) prices per 100 face amount. A mid-market price is an average of a lower *bid* price, at which traders stand ready to buy a bond, and a higher *ask* price, at which

TABLE 1.2 Selected U.S. Treasury Bond Prices as of May 28, 2010

Coupon	Maturity	Price
$1\frac{1}{4}$%	11/30/2010	100.550
$4\frac{7}{8}$%	5/31/2011	104.513
$4\frac{1}{2}$%	11/30/2011	105.856
$4\frac{3}{4}$%	5/31/2012	107.966
$3\frac{3}{8}$%	11/30/2012	105.869
$3\frac{1}{2}$%	5/31/2013	106.760
2%	11/30/2013	101.552
$2\frac{1}{4}$%	5/31/2014	101.936
$2\frac{1}{8}$%	11/30/2014	100.834

traders stand ready to sell a bond. A *full* price is the total amount a buyer pays for a bond, which is the sum of the *flat* or *quoted* price of the bond and *accrued interest*. This division of full price will be explained later in this chapter. In any case, to take an example from Table 1.2, purchasing \$100,000 face amount of the $3\frac{1}{2}$s of May 31, 2013, costs a total of \$106,760.

The bonds in Table 1.2 were selected from the broader list of U.S. Treasuries because they all mature and make payments on the same *cycle*, in this case at the end of May and November each year. This means, for example, that all of the bonds make a payment on November 30, 2010, and, therefore, that all their prices incorporate information about the value of a dollar to be received on that date. Similarly, all of the bonds apart from the $1\frac{1}{4}$s of November 30, 2010, which will have already matured, make a payment on May 31, 2011, and their prices incorporate information about the value of a dollar to be received on that date, etc. The next section describes how to extract information about the value of a dollar to be received on each of the payment dates in the May–November cycle from the prices in Table 1.2.

DISCOUNT FACTORS

The *discount factor* for a particular term gives the value today, or the *present value* of one unit of currency to be received at the end of that term. Denote the discount factor for t years by $d(t)$. Then, for example, if $d(.5)$ equals .99925, the present value of \$1 to be received in six months is 99.925 cents. Another security, which pays \$1,050,000 in six months, would have a present value of $.99925 \times \$1,050,000$ or \$1,049,213.

P1: TIX/b P2: c/d QC: e/f T1: g
JWBT539-c01 JWBT539-Tuckman July 20, 2011 11:47 Printer: Courier Westford

Since Treasury bonds promise future cash flows, discount factors can be extracted from Treasury bond prices. In fact, each of the rows of Table 1.2 can be used to write one equation that relates prices to discount factors. The equation from the $1\frac{1}{4}$s of November 30, 2010, is

$$100.550 = \left(100 + \frac{1\frac{1}{4}}{2}\right) d(.5) \tag{1.1}$$

In words, equation (1.1) says that the price of the bond equals the present value of its future cash flows, namely its principal plus coupon payment, all times the discount factor for funds to be received in six months. Solving reveals that $d(.5)$ equals .99925.

By the same reasoning, the equations relating prices to discount factors can be written for the other bonds listed in Table 1.2. The next two of these equations are

$$104.513 = \frac{4\frac{7}{8}}{2} \times d(.5) + \left(100 + \frac{4\frac{7}{8}}{2}\right) d(1) \tag{1.2}$$

$$105.856 = \frac{4\frac{1}{2}}{2} \times d(.5) + \frac{4\frac{1}{2}}{2} \times d(1) + \left(100 + \frac{4\frac{1}{2}}{2}\right) d(1.5) \tag{1.3}$$

Given the solution for $d(.5)$ from equation (1.1), equation (1.2) can be solved for $d(1)$. Then, given the solutions for $d(.5)$ and $d(1)$, equation (1.3) can be solved for $d(1.5)$. Continuing in this fashion through the rows of Table 1.2 generates the discount factors, in six-month intervals, out to four and one-half years, which are reported in Table 1.3. Note how these

TABLE 1.3 Discount Factors from U.S. Treasury Note and Bond Prices as of May 28, 2010

Term	Discount Factor
11/30/2010	.99925
5/31/2011	.99648
11/30/2011	.99135
5/31/2012	.98532
11/30/2012	.97520
5/31/2013	.96414
11/30/2013	.94693
5/31/2014	.93172
11/30/2014	.91584

discount factors, falling with term, reflect the time value of money: the longer a payment of $1 is delayed, the less it is worth today.

THE LAW OF ONE PRICE

Another U.S. Treasury bond issue, one not included in the set of base bonds in Table 1.2, is the $\frac{3}{4}$s of November 30, 2011. How should this bond be priced? A natural answer is to apply the discount factors of Table 1.3 to this bond's cash flows. After all, the base bonds are all U.S. Treasury bonds and the value to investors of receiving $1 from a Treasury on some future date should not depend very much on which particular bond paid that $1. This reasoning is an application of the *law of one price*: absent confounding factors (e.g., liquidity, financing, taxes, credit risk), identical sets of cash flows should sell for the same price.

According to the law of one price, the price of the $\frac{3}{4}$s of November 30, 2011 should be

$$.375 \times .99925 + .375 \times .99648 + 100.375 \times .99135 = 100.255 \qquad (1.4)$$

where each cash flow is multiplied by the discount factor from Table 1.3 that corresponds to that cash flow's payment date. As it turns out, the market price of this bond is 100.190, close to, but not equal to, the prediction of 100.255 in equation (1.4).

Table 1.4 compares the market prices of three bonds as of May 28, 2010, to their present values (PVs), i.e., to their prices as predicted by the law of one price. The differences range from −2.8 cents to +6.5 cents per 100 face value, indicating that the law of one price describes the pricing of these bonds relatively well but not perfectly.

According to the last row of Table 1.4, the $\frac{7}{8}$s of May 31, 2011, trade 2.8 cents *rich* to the base bonds, i.e., its market price is high relative to the discount factors in Table 1.3. In the same sense, the $\frac{3}{4}$s of November 30, 2011, and the $\frac{3}{4}$s of May 31, 2012, trade *cheap*. In fact, were these price discrepancies sufficiently large relative to transaction costs, an arbitrageur might consider trying to profit by selling the rich $\frac{7}{8}$s and

TABLE 1.4 Testing the Law of One Price for Three U.S. Treasury Notes as of May 28, 2010

Bond	$\frac{7}{8}$s 5/31/11	$\frac{3}{4}$s 11/30/11	$\frac{3}{4}$s 5/31/12
PV	100.521	100.255	100.022
Price	100.549	100.190	99.963
PV−Price	−.028	.065	.059

simultaneously buying some combination of the base bonds; by buying either of the cheap bonds and simultaneously selling base bonds; or by selling the rich $\frac{7}{8}$s and buying both of the cheap bonds in the table. Trades of this type, arising from deviations from the law of one price, are the subject of the next section.

ARBITRAGE AND THE LAW OF ONE PRICE

While the law of one price is intuitively reasonable, its justification rests on a stronger foundation. It turns out that a deviation from the law of one price implies the existence of an *arbitrage opportunity*, that is, a trade that generates profits without any chance of losing money.[1] But since arbitrageurs would rush *en masse* to do any such trade, market prices would quickly adjust to rule out any such opportunity. Hence, arbitrage activity can be expected to do away with significant deviations from the law of one price. And it is for this reason that the law of one price usually describes security prices quite well.

To make this argument more concrete, the discussion turns to an arbitrage trade based on the results of Table 1.4, which showed that the $\frac{3}{4}$s of November 30, 2011, are cheap relative to the discount factors in Table 1.3 or, equivalently, to the bonds listed in Table 1.2. The trade is to purchase the $\frac{3}{4}$s of November 30, 2011, and simultaneously sell or *short*[2] a portfolio of bonds from Table 1.2 that replicates the cash flows of the $\frac{3}{4}$s. Table 1.5 describes this *replicating portfolio* and the arbitrage trade.

Columns (2) to (4) of Table 1.5 correspond to the three bonds chosen from Table 1.2 to construct the replicating portfolio: the $1\frac{1}{4}$s of November 30, 2010; the $4\frac{7}{8}$s of May 31, 2011; and the $4\frac{1}{2}$s of November 30, 2011. Row (iii) gives the face amount of each bond in the replicating portfolio, so that this portfolio is long 98.166 face amount of the $4\frac{1}{2}$s, short 1.790 of the $4\frac{7}{8}$s, and short 1.779 of the $1\frac{1}{4}$s. Rows (iv) through (vi) show the cash flows from those face amounts of each bond. For example, 98.166 face amount of the $4\frac{1}{2}$s, which pay a coupon of 2.25% on May 31, 2011, generates a cash flow of $98.166 \times 2.25\%$ or 2.209 on that date. Similarly, -1.779 of the $1\frac{1}{4}$s, which pay coupon and principal totalling $100 + \frac{1.25}{2}$ or

[1]Market participants often use the term arbitrage more broadly to encompass trades that could conceivably lose money, but promise large profits relative to the risks borne.

[2]To short a security means to sell a security one does not own. The mechanics of short selling bonds will be discussed in Chapter 12. For now, assume that a trader shorting a bond receives the price of the bond and is obliged to pay all its coupon and principal cash flows.

TABLE 1.5 The Replicating Portfolio of the $\frac{3}{4}$s of November 30, 2011, with Prices as of May 28, 2010

	(1)	(2)	(3)	(4)	(5)	(6)
(i)	Coupon	$1\frac{1}{4}$s	$4\frac{7}{8}$s	$4\frac{1}{2}$s		$\frac{3}{4}$s
(ii)	Maturity	11/30/10	5/31/11	11/30/11	Portfolio	11/30/11
(iii)	Face Amount	−1.779	−1.790	98.166		100
	Date			**Cash Flows**		
(iv)	11/30/10	−1.790	−.044	2.209	.375	.375
(v)	5/31/11		−1.834	2.209	.375	.375
(vi)	11/30/11			100.375	100.375	100.375
(vii)	Price	100.550	104.513	105.856		100.190
(viii)	Cost	−1.789	−1.871	103.915	100.255	100.190
(ix)	Net Proceeds	.065				

100.625 per 100 face value on November 30, 2010, produces a cash flow of −1.779 × 100.625% or −1.790 on that date. Row (vii) gives the price of each bond per 100 face amount, simply copied from Table 1.2. Row (viii) gives the initial cost of purchasing the indicated face amount of each bond. So, for example, the "cost" of "purchasing" −1.790 face amount of the $4\frac{7}{8}$s is −1.790 × 104.513% or −1.871. Said more naturally, the proceeds from selling 1.790 face amount of the $4\frac{7}{8}$s are 1.871.

Column (5) of Table 1.5 sums columns (2) through (4) to obtain the cash flows and cost of the replicating portfolio. Rows (iv) through (vi) of column (5) confirm that the cash flows of the replicating portfolio do indeed match the cash flows of 100 face amount of the $\frac{3}{4}$s of November 30, 2011, given in the same rows of column (6). Note that most of the work of replicating the $\frac{3}{4}$s of November 30, 2011, is accomplished by the $4\frac{1}{2}$s maturing on the same date. The other two bonds in the replicating portfolio are used for minor adjustments to the cash flows in six months and one year. Appendix A in this chapter shows how to derive the face amounts of the bonds in this or any such replicating portfolio.

With the construction of the replicating portfolio completed, the discussion returns to the arbitrage trade. According to row (viii) of Table 1.5 , an arbitrageur can buy 100 face amount of the $\frac{3}{4}$s of November 30, 2011, for 100.190, sell the replicating portfolio for 100.255, pocket the difference or "net proceeds" of 6.5 cents, shown in row (ix), and not owe anything on any future date. And while a 6.5-cent profit may seem small, the trade can be scaled up: for \$500 million face of the $\frac{3}{4}$s, which would not be an abnormally large position, the riskless profit increases to \$500,000,000 × .065% or \$325,000.

As stated at the start of this section, if a riskless and profitable trade like the one just described were really available, arbitrageurs would rush to do the trade and, in so doing, force prices to relative levels that admit no arbitrage opportunities. More specifically, arbitrageurs would drive the prices of the $\frac{3}{4}$s and of the replicating portfolio together until the two were equal.

The crucial link between arbitrage and the law of one price can now be explained. The total cost of the replicating portfolio, 100.255, given in column (5), row (viii) of Table 1.5, exactly equals the present value of the $\frac{3}{4}$s of November 30, 2011, computed in Table 1.4. In other words, the law of one price methodology of pricing the $\frac{3}{4}$s (i.e., discounting with factors derived from the $1\frac{1}{4}$s, $4\frac{7}{8}$s, and $4\frac{1}{2}$s) comes up with exactly the same value as does the arbitrage pricing methodology (i.e., calculating the value of portfolio of the $1\frac{1}{4}$s, $4\frac{7}{8}$s, and $4\frac{1}{2}$s that replicates the cash flows of the $\frac{3}{4}$s). This is not a coincidence. In fact, Appendix B in this chapter proves that these two pricing methodologies are mathematically identical. Hence, applying the law of one price, i.e., pricing with discount factors, is identical to relying on the activity of arbitrageurs to eliminate relative mispricings, i.e., pricing by arbitrage. Expressed another way, discounting can be justifiably regarded as shorthand for the more complex and persuasive arbitrage pricing methodology.

Despite this discussion, of course, the market price of the $\frac{3}{4}$s was quoted at a level somewhat below the level predicted by the law of one price. This can be attribtured to one or a combination of the following reasons. First, there are transaction costs in doing arbitrage trades which could significantly lower or wipe out any arbitrage profit. In particular, the prices in Table 1.2 are mid-market whereas, in reality, an arbitrageur would have to buy securities at higher ask prices and sell at lower bid prices. Second, bid-ask spreads in the financing markets (see Chapter 12), incurred when shorting securities, might also overwhelm any arbitrage profit. Third, it is only in theory that U.S. Treasury bonds are commodities, i.e., fungible collections of cash flows. In reality, bonds have idiosyncratic differences that are recognized by the market and priced accordingly. And it is this last point that is the subject of the next section.

APPLICATION: STRIPS AND THE IDIOSYNCRATIC PRICING OF U.S. TREASURY NOTES AND BONDS

STRIPS

In contrast to coupon bonds that make payments every six months, *zero coupon* bonds make no payments until maturity. Zero coupon bonds issued by the U.S. Treasury are called STRIPS. For example, \$1,000,000

TABLE 1.6 STRIPS Face Amounts from 1,000,000 Face Amount of the $3\frac{1}{2}$s of May 15, 2020

Date	C-STRIP Face Amount	P-STRIP Face Amount
11/15/10	$17,500	0
5/15/11	$17,500	0
11/15/11	$17,500	0
⋮	⋮	⋮
5/15/19	$17,500	0
11/15/19	$17,500	0
5/15/20	$17,500	$1,000,000

face amount of STRIPS maturing on May 15, 2020, promises only one payment: $1,000,000 on that date. STRIPS are created when a particular coupon bond is delivered to the Treasury in exchange for its coupon and principal components. Table 1.6 illustrates the stripping of $1,000,000 face amount of the $3\frac{1}{2}$s of May 15, 2020, which was issued in May 2010, to create coupon STRIPS maturing on the 20 coupon payment dates and principal STRIPS maturing on the maturity date. Coupon or interest STRIPS are called TINTs, INTs, or C-STRIPS while principal STRIPS are called TPs, Ps, or P-STRIPS. Note that the face amount of C-STRIPS on each date is $1/2 \times 3.5\% \times \$1,000,000$ or $\$17,500$.

The Treasury not only creates STRIPS but retires them as well. For example, upon delivery of the set of STRIPS in Table 1.6 the Treasury would *reconstitute* the $1,000,000 face amount of the $3\frac{1}{2}$s of May 15, 2020. But in this context it is crucial to note that C-STRIPS are fungible while P-STRIPS are not. When reconstituting a bond, any C-STRIPS maturing on a particular date may be applied toward the coupon payment of that bond on that date. By contrast, only P-STRIPS that were stripped from a particular bond may be used to reconstitute the principal payment of that bond.[3] This feature of the STRIPS program implies that P-STRIPS, and not C-STRIPS, inherit the cheapness or richness of the bonds from which they came, an implication that will be demonstrated in the following subsection.

STRIPS prices are essentially discount factors. If the price of the C-STRIPS maturing on May 31, 2015, is 89.494 per 100 face amount, then the implied discount factor to that date is .89494. With this in

[3]Making P-STRIPS fungible would not affect either the total or the timing of cash flows owed by the Treasury, but could change the amounts outstanding of particular securities.

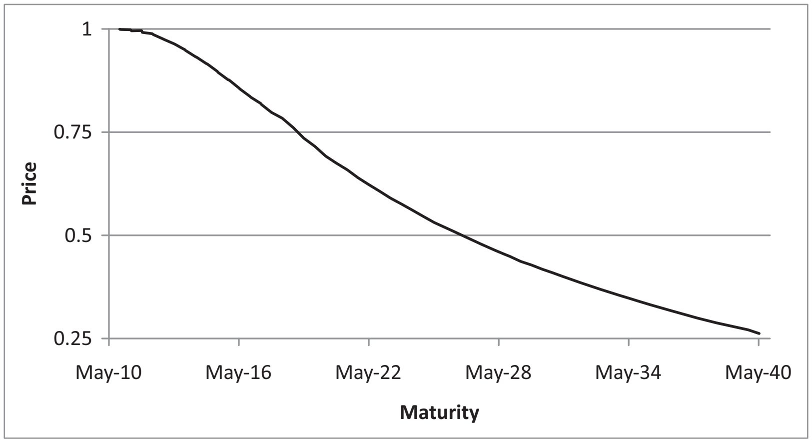

FIGURE 1.1 Discount Factors from C-STRIPS Prices as of May 28, 2010

mind, Figure 1.1 graphs the C-STRIPS prices per unit face amount as of May 28, 2010.

The Idiosyncratic Pricing of U.S. Treasury Notes and Bonds

If U.S. Treasury bonds were commodities, with each regarded solely as a particular collection of cash flows, then the price of each would be well approximated by discounting its cash flows with the C-STRIPS discount factors of Figure 1.1. If however individual bonds have unique characteristics that are reflected in pricing, the law of one price would not be as accurate an approximation. Furthermore, since C-STRIPS are fungible while P-STRIPS are not, any such pricing idiosyncrasies would manifest themselves as differences between the prices of P-STRIPS and C-STRIPS of the same maturity. To this end, Figure 1.2 graphs the differences between the prices of P-STRIPS and C-STRIPS that mature on the same date as of May 28, 2010. So, for example, with the price of P-STRIPS and C-STRIPS, both maturing on May 31, 2015, at 89.865 and 89.494, respectively, Figure 1.2 records the difference for May 31, 2015, as 89.865 – 89.494 or .371. Note that Figure 1.2 shows two sets of P-STRIPS prices, those P-STRIPS originating from Treasury bonds and those originating from Treasury notes.[4]

Inspection of Figure 1.2 shows that there are indeed significant pricing differences between P-STRIPS and C-STRIPS that mature on the same date. This does not necessarily imply the existence of arbitrage opportunities, as discussed at the end of the previous section. However, the results do suggest

[4]The difference between notes and bonds is of historical interest only; see "Fixed Income Markets in the United States, Securities and Other Assets" in the Overview.

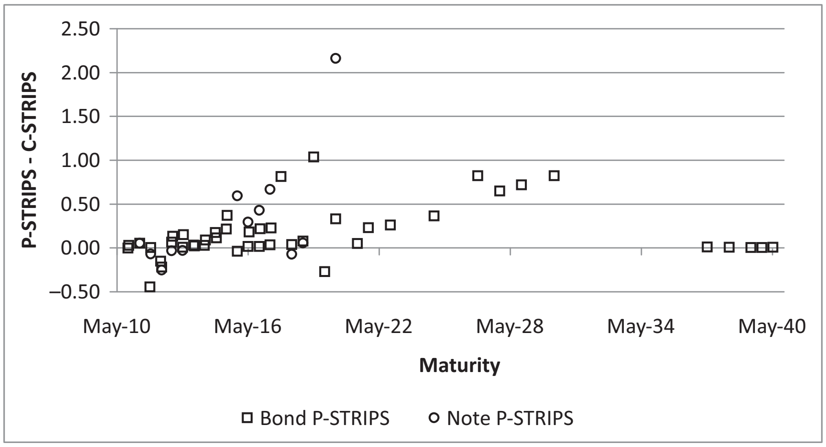

FIGURE 1.2 Differences between the Prices of P-STRIPS and C-STRIPS Maturing on the Same Date per 100 Face Amount as of May 28, 2010

that bonds have idiosyncratic pricing differences and that these differences are inherited by their respective P-STRIPS. Of particular interest, for example, is the largest price difference in the figure, the 2.16 price difference between the P-STRIPS and C-STRIPS maturing on May 15, 2020. These P-STRIPS come from the most recently sold or *on-the-run* 10-year note, an issue which, as will be discussed in Chapter 12, traditionally trades rich to other bonds because of its superior liquidity and financing characteristics. In any case, to determine whether idiosyncratic bond characteristics are indeed inherited by P-STRIPS, Table 1.7 analyzes the pricing of selected U.S. Treasury coupon securities in terms of STRIPS prices. The particular securities selected are those on the mid-month, May-November cycle with 10 or more years to maturity as of May 2010.

Columns (1) to (3) of Table 1.7 give the coupon, maturity, and market price of each bond. Column (4) computes a price for each bond by discounting all of its cash flows using the C-STRIPS prices in Figure 1.1, and column (5) gives the difference between the market price and that computed price. By the simplest application of the law of one price, these computed prices should be a good approximation of market prices. There are, however, some very significant discrepancies. The approximation misses the price of the $3\frac{1}{2}$s of May 15, 2020, the 10-year on-the-run security, by a very large 2.076; the 5s of May 15, 2037, by .924; and the $6\frac{1}{4}$s of 5/15/30 by .708.

Column (6) of Table 1.7 computes the price of each bond by discounting its coupon payments with C-STRIPS prices and its principal payment with the P-STRIPS of that bond. Column (7) gives the difference between the market price and that computed price. To the extent that P-STRIPS prices inherit pricing idiosyncrasies of their respective bonds, these computed prices should be better approximations to market prices than the prices computed

TABLE 1.7 Market Prices Compared with Pricing Using C-STRIPS and with Pricing Using C-STRIPS for Coupon Payments and the Respective P-STRIPS for Principal Payments

(1) Coupon	(2) Maturity	(3) Market Price	(4) C-Pricing	(5) Error	(6) C- and P-Pricing	(7) Error
$3\frac{1}{2}$	5/15/20	101.896	99.820	2.076	101.982	−.086
$8\frac{3}{4}$	5/15/20	146.076	145.738	.338	146.070	.006
$8\frac{1}{8}$	5/15/21	142.438	142.357	.080	142.407	.031
8	11/15/21	141.916	141.750	.167	141.980	−.063
$7\frac{5}{8}$	11/15/22	139.696	139.545	.151	139.805	−.109
$7\frac{1}{2}$	11/15/24	140.971	140.694	.277	141.059	−.087
$6\frac{1}{2}$	11/15/26	131.582	130.894	.687	131.716	−.134
$6\frac{1}{8}$	11/15/27	127.220	126.643	.578	127.291	−.070
$5\frac{1}{4}$	11/15/28	116.118	115.456	.661	116.175	−.058
$6\frac{1}{4}$	5/15/30	130.523	129.815	.708	130.639	−.116
5	5/15/37	113.840	112.916	.924	113.943	−.102
$4\frac{1}{2}$	5/15/38	105.114	104.625	.490	105.214	−.100
$4\frac{1}{4}$	5/1539	100.681	100.425	.256	100.764	−.083
$4\frac{3}{8}$	11/15/39	102.780	102.638	.143	102.905	−.124
$4\frac{3}{8}$	5/15/40	102.999	102.308	.691	102.969	.030

using C-STRIPS prices alone. And, in fact, this is the case. Comparing the absolute values of the two error columns reveals that the approximation in column (6) is better than the approximation in column (4) for every bond in the table.

In conclusion, then, individual Treasury bonds have idiosyncratic characteristics that are reflected in market prices. Furthermore, since P-STRIPS are not fungible across bonds, their prices inherit the idiosyncratic pricing of their respective bond issues.

ACCRUED INTEREST

This section describes the useful market practice of separating the full price of a bond, which is the price paid by a buyer to a seller, into two parts: a quoted or flat price, which is the price that appears on trading screens and is used when negotiating transactions; and *accrued interest*. The full and quoted prices are also known as the *dirty* and *clean* prices, respectively.

Definition

To make the concepts concrete, consider an investor who purchases \$10,000 face amount of the U.S. Treasury $3\frac{5}{8}$s of August 15, 2019, for settlement on June 1, 2010. The bond made a coupon payment of $\frac{1}{2} \times 3\frac{5}{8}\% \times \$10{,}000$ or \$181.25 on February 15, 2010, and will make its next coupon payment of \$181.25 on August 15, 2010. See the time line in Figure 1.3.

Assuming the purchaser holds the bond through the next coupon date, the purchaser will collect the coupon on that date. But it can be argued that the purchaser is not entitled to the full semiannual coupon payment on August 15 because, as of that time, the purchaser will have held the bond for only two and a half months of a six-month coupon period. More precisely, using what is known as the *actual/actual day-count convention*, which will be explained later in this section, and referring again to Figure 1.3, the purchaser should receive only 75 of 181 days of the coupon payment, that is, $\frac{75}{181} \times \$181.25$, or \$75.10. The seller of the bond, whose cash was invested in the bond from February 15 to June 1, should collect the rest of the coupon, i.e., $\frac{106}{181} \times \$181.25$, or \$106.15. A conceivable institutional arrangment is for the seller and purchaser to divide the coupon on the payment date, but this would undersirably require additional arrangements to ensure that this split of the coupon actually takes place. Consequently, market convention dictates instead that the purchaser pay the \$106.15 of accrued interest to the seller on the settlement date and that the purchaser keep the entire coupon of \$181.25 on the coupon payment date.

On May 28, 2010, for delivery on June 1, 2010, the flat or quoted price of the $3\frac{5}{8}$s was 102-26, meaning $102 + \frac{26}{32}$ or 102.8125. The full or invoice price of the bond per 100 face amount is defined as the quoted price plus accrued interest, which, in this case, is $102.8125 + 1.0615$ or 103.8740. For this particular trade, of \$10,000 face amount, the invoice price is \$10,387.40.

At this point, by the way, it becomes clear why discussion earlier in the chapter had to make reference to the fact that prices were full prices. When trading bonds that make coupon payments on May 31, 2010, for settlement on June 1, 2010, purchasers have to pay one day of accrued interest to sellers.

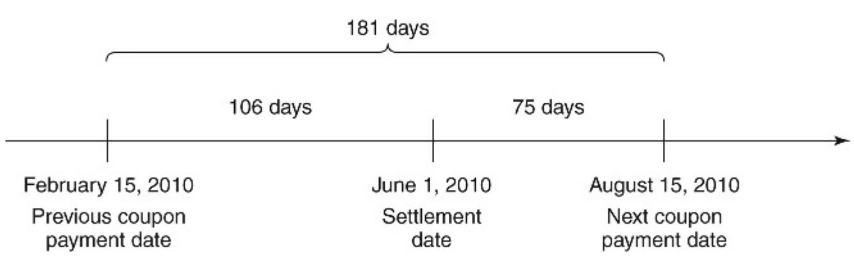

FIGURE 1.3 Example of Accrued Interest Time Line

P1: TIX/b P2: c/d QC: e/f T1: g
JWBT539-c01 JWBT539-Tuckman July 20, 2011 11:47 Printer: Courier Westford

Pricing Implications

The present value of a bond's cash flows should be equated or compared with its full price, that is, with the amount a purchaser actually pays to purchase those cash flows. Conceptually, denoting the flat price by p, accrued interest by AI, the present value of the cash flows by PV, and the full price, as before, by P,

$$P = p + AI = PV \tag{1.5}$$

Equation (1.5) reveals an important point about accrued interest: the particular market convention used in calculating accrued interest does not really matter. Say, for example, that everyone recognizes that the convention in place is too generous to the seller because, instead of being made to wait for a share of the interest until the next coupon date, the seller receives that share at settlement. In that case, by equation (1.5), the flat price would adjust downward to mitigate this advantage. Put another way, the only quantity that matters is the invoice price, which determines the amount of money that changes hands.

Having made this argument, why is the accrued interest convention useful in practice? The answer is told in Figure 1.4, which draws the full and flat prices of the $3\frac{5}{8}$s of August 15, 2019, from February 15, 2010, to September 15, 2010, under several assumptions, with the most important being that 1) the discount function does not change, i.e., $d(t)$ does not change, where t is the number of days from settlement; and 2) the flat price of the bond for settlement on June 1 is 102.8125. In words, then, Figure 1.4 says that the full price changes dramatically over time even when the market is

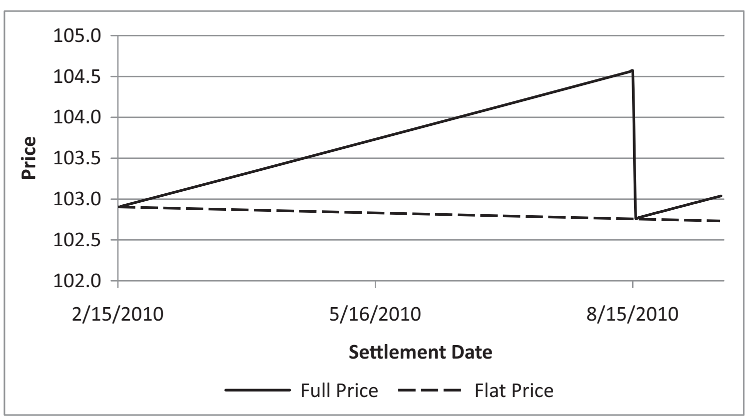

FIGURE 1.4 Full and Flat Prices for the $3\frac{5}{8}$s of August 15, 2019, Over Time with a Constant Discount Function

P1: TIX/b P2: c/d QC: e/f T1: g
JWBT539-c01 JWBT539-Tuckman July 20, 2011 11:47 Printer: Courier Westford

unchanged, including a discontinuous jump on coupon payment dates, while the flat price changes only gradually over time. Therefore, when trading bonds day to day, it is more intuitive to track flat prices and negotiate transactions in those terms.

The shapes of the price functions in Figure 1.4 can be understood as follows. Within a coupon period, the full price of the bond, which is just the present value of its cash flows, increases over time as the bond's payments draw near. But from an instant before the coupon payment date to an instant after, the full price falls by the coupon payment: the coupon is included in the present value of the remaining cash flows at the instant before the payment, but not at the instant after. The time pattern of the flat price, supposing that prevailing interest rates do not change, will be discussed in Chapter 3. Basically, however, the flat price of a bond like the $3\frac{5}{8}$s, which sells for more than its face value, will trend down to its value at maturity, i.e., par.

Day-Count Conventions

Accrued interest equals the coupon times the fraction of the coupon period from the previous coupon payment date to the settlement date. For the $3\frac{5}{8}$s, as for most government bonds, this fraction is calculated by dividing the actual number of days since the previous coupon date by the actual number of days in the coupon period. Hence the term "actual/actual" for this day-count convention.

Other day-count conventions, however, are applied in other markets. Two of the most common are *actual/360* and *30/360*. The actual/360 convention divides the actual number of days between two dates by 360, and is commonly used in money markets, i.e., for short-term, *discount* (i.e., zero coupon) securities, and for the floating legs of interest rate swaps. The *30/360* convention assumes that there are 30 days in a month when calculating the difference between two dates and then divides by 360. Applying this convention, the number of days between June 1 and August 15 is 74 (29 days left in June, 30 days in July, and 15 days in August), as opposed to the 75 days using an actual day count. The 30/360 convention is used most commonly for corporate bonds and for the fixed leg of interest rate swaps.

APPENDIX A: DERIVING REPLICATING PORTFOLIOS

To replicate the $\frac{3}{4}$s of November 30, 2011, Table 1.5 uses the $1\frac{1}{4}$s due November 30, 2010, the $4\frac{7}{8}$s due May 31, 2011, and the $4\frac{1}{2}$s due November 30, 2011. Number these bonds from 1 to 3 and let F^i be the face amount of bond i in the replicating portfolio. Then, the following equations express

P1: TIX/b P2: c/d QC: e/f T1: g
JWBT539-c01 JWBT539-Tuckman July 20, 2011 11:47 Printer: Courier Westford

the requirement that the cash flows of the replicating portfolio equal those of the $\frac{3}{4}$s on each of the three cash flow dates.

For the cash flow on November 30, 2010:

$$\left(100\% + \frac{1\frac{1}{4}\%}{2}\right)F^1 + \left(\frac{4\frac{7}{8}\%}{2}\right)F^2 + \left(\frac{4\frac{1}{2}\%}{2}\right)F^3 = \frac{\frac{3}{4}\%}{2} \tag{1.6}$$

For the cash flow on May 31, 2011:

$$0 \times F^1 + \left(100\% + \frac{4\frac{7}{8}\%}{2}\right)F^2 + \left(\frac{4\frac{1}{2}\%}{2}\right)F^3 = \frac{\frac{3}{4}\%}{2} \tag{1.7}$$

And, for the cash flow on November 30, 2011:

$$0 \times F^1 + 0 \times F^2 + \left(100\% + \frac{4\frac{1}{2}\%}{2}\right)F^3 = 100\% + \frac{\frac{3}{4}\%}{2} \tag{1.8}$$

Solving equations (1.6), (1.7), and (1.8) for F^1, F^2, and F^3 gives the replicating portfolio's face amounts in Table 1.5. Note that since one bond matures on each date, these equations can be solved one-at-a-time instead of simultaneously, i.e., solve (1.8) for F^3, then, using that result, solve (1.7) for F^2 , and then, using both results, solve (1.6) for F^1. In any case, the results are as follows:

$$F^1 = -1.779\% \tag{1.9}$$

$$F^2 = -1.790\% \tag{1.10}$$

$$F^3 = 98.166\% \tag{1.11}$$

Replicating portfolios are easier to describe and manipulate using matrix algebra. To illustrate, equations (1.6) through (1.8) are expressed in matrix form as follows:

$$\begin{pmatrix} 1 + \frac{1.25\%}{2} & \frac{4.875\%}{2} & \frac{4.5\%}{2} \\ 0 & 1 + \frac{4.875\%}{2} & \frac{4.5\%}{2} \\ 0 & 0 & 1 + \frac{4.5\%}{2} \end{pmatrix} \begin{pmatrix} F^1 \\ F^2 \\ F^3 \end{pmatrix} = \begin{pmatrix} \frac{.75\%}{2} \\ \frac{.75\%}{2} \\ 1 + \frac{.75\%}{2} \end{pmatrix} \tag{1.12}$$

Note that each column of the leftmost matrix describes the cash flows of one of the bonds in the replicating portfolio; the elements of the vector to the right of this matrix are the face amounts of each bond for which equation (1.12) has to be solved; and the rightmost vector contains the cash flows of the bond to be replicated. This equation can easily be solved by pre-multiplying each side by the inverse of the leftmost matrix.

In general then, suppose that the bond to be replicated makes payments on T dates. Let C be the $T \times T$ matrix of cash flows, principal plus interest, with the T columns representing the T bonds in the replicating portfolio and the T rows the dates on which those bonds make payments. Let $\overrightarrow{F}$ be the $T \times 1$ vector of face amounts in the replicating portfolio and let $\overrightarrow{c}$ be the vector of cash flows, principal plus interest, of the bond to be replicated. Then, the replication equation is

$$C\overrightarrow{F} = \overrightarrow{c} \tag{1.13}$$

with solution

$$\overrightarrow{F} = C^{-1}\overrightarrow{c} \tag{1.14}$$

The only complication is in ensuring that the matrix C does have an inverse. Essentially, any set of T bonds will do so long as there is at least one bond in the replicating portfolio making a payment on each of the T dates. In this case, the T bonds would be said to *span* the payment dates. So, for example, T bonds all maturing on the last date would work, but T bonds all maturing on the second-to-last date would not work: in the latter case there would be no bond in the replicating portfolio making a payment on date T.

APPENDIX B: THE EQUIVALENCE OF THE DISCOUNTING AND ARBITRAGE PRICING APPROACHES

Proposition: Pricing a bond according to either of the following methods gives the same price:

- Derive a set of discount factors from some set of spanning bonds and price the bond in question using those discount factors.
- Find the replicating portfolio of the bond in question using that same set of spanning bonds and calculate the price of the bond as the price of this portfolio.

P1: TIX/b P2: c/d QC: e/f T1: g
JWBT539-c01 JWBT539-Tuckman July 20, 2011 11:47 Printer: Courier Westford

Proof: Continue using the notation introduced at the end of Appendix A. Also, let $\vec{d}$ be the $T \times 1$ vector of discount factors for each date and let $\vec{P}$ be the vector of prices of each bond in the replicating portfolio, which is the same as the vector of prices of each bond used to compute the discount factors. Generalizing the "Discount Factors" section of this chapter, one can solve for discount factors using the following equation:

$$\vec{d} = \left(C'\right)^{-1} \vec{P} \tag{1.15}$$

where the $'$ denotes the transpose. Then, the price of the bond according to the first method is $c'\vec{d}$. The price according to the second method is $\vec{P}'\vec{F}$ where $\vec{F}$ is as derived in equation (1.14).

Hence, the two methods give the same price if

$$\vec{c}'\vec{d} = \vec{P}'\vec{F} \tag{1.16}$$

Expanding the left-hand side of equation (1.16) with (1.15) and the right-hand side with (1.14),

$$\vec{c}'\left(C'\right)^{-1} \vec{P} = \vec{P}'C^{-1}\vec{c} \tag{1.17}$$

And since both sides of this equation are just numbers, take the transpose of the left-hand side to show that equation (1.17) is true.