

CHAPTER 17

Essentials of Structured Product Engineering

1. Introduction

Structured products consist of packaging basic assets such as stocks, bonds, and currencies together with some derivatives. The final product obtained this way will, depending on the product, (1) have an *enhanced return* or *improved credit quality*, (2) *lower costs* of asset-liability management for corporates, (3) build in the *views* held by the clients, (4) often be *principal protected*.¹

Households do not like to build their own cars, computers, or refrigerators themselves. They prefer to buy them from the *producers* who manufacture and assemble them. Every complex product has its own specialists, and it is more cost effective to buy products manufactured by these specialists. The same is true for *financial products*. Investors, corporates, and institutions need solutions for problems that they face in their lives. The packaging solutions for investors' and institutions' needs are called *structured* products. "Manufacturers," i.e., the "structurers," put these together and sell them to clients. Clients consist of investment funds, pension funds, insurance companies, and individuals. Clients may have *views* on the near or medium term behavior of equity prices, interest rates, or commodities. Structured products can be designed so that such clients can take positions according to their views in a convenient way.

Industrial goods such as cell phones and cars are constantly updated and improved. Again, the same is true for structured products. The views, the needs, or simply the risk appetite of clients change and the structurer needs constantly to provide new structured products that fit these new views. This chapter discusses the way financial engineering can be used to service retail clients' particular needs.

¹ According to this, the investor (or the corporate) would not lose the principal in case the expectations turn out to be wrong. We deal with the so-called Constant Proportion Portfolio Insurance (CPPI) and its more recent version, Dynamic Proportion Portfolio Insurance (DPPI) in Chapter 20.

Financial engineering provides ways to construct any payoff structure desired by an investor. However, often these payoffs involve complex option positions, and clients may not have the knowledge, or simply the means, to handle such risks. Market practitioners can do this better. For example, many structured products offer principal protection or credit enhancements to investors. Normally, institutions that may not be allowed to invest in such positions due to regulatory reasons will be eligible to hold the structured product itself once principal protection is added to it. Providing custom made products for clients due to differing views, risk appetite, or regulatory conditions is one way to interpret structured products and in general they are regarded this way.

However, in this book our main interest is to study financial phenomena from the manufacturer's point of view. This view provides a *second interpretation* of structured products. Investment banks deal with clients, corporates, and with each other. These activities require holding inventories, sourcing and outsourcing exposures, and maintaining books. However, due to market conditions, the instruments that banks are keeping on their books may sometimes become too costly or too risky, or sometimes better alternatives emerge. The natural thing to do is to sell these exposures to "others." Structured products may be one convenient way of doing this. Consider the following example. A bank would like to *buy volatility* at a reasonable price, but suppose there are not enough sellers of such volatility in the interbank market. Then a structured product can be designed so that the bank can buy volatility at a reasonable price from the retail investor.

In this interpretation, the structured product is regarded from the manufacturer's angle and looks like a tool in inventory or balance sheet management. A structured product is either an indirect way to sell some existing risks to a client, or is an indirect way to buy some desired risks from the retail client. Given that bank balance sheets and books contain a great deal of interest rate and credit risk related exposures, it is natural that a significant portion of the recent activity in structured products relate to managing such exposures.

In this chapter we consider two major classes of structured products. The first group is the new *equity*, commodity and FX-based structured products and the second is *Libor-based fixed income* products. The latter are designed so as to benefit from expected future movements in the yield curve. We will argue that the general logic behind structured products is the same, regardless of whether they are Libor-based or equity-linked. Hence we try to provide a unified approach to structured products. In a later chapter we will consider the *third* important class of structured products based on the occurrence of an *event*. This event may be a mortgage prepayment, or, more importantly, a credit default. These will be discussed through *structured credit* products. Because credit is considered separately in a different chapter, during the discussion that follows it is best to assume that there is no credit risk.

2. Purposes of Structured Products

Structured products may have at least *four* specific objectives.

The first objective could be *yield enhancement*—to offer the client a higher return than what is normally available. This of course implies that the client will be taking additional risks, or foregoing some gains in other circumstances. For example, the client gets an enhanced return if a stock price increases up to 12%. However, any additional gains would be forgone and the return would be *capped* at 12%. The value of this cap is used in offering an enhanced yield.

The second could be *credit enhancement*. In this case the client will buy a predetermined set of debt securities at a lower default risk than warranted by their rating. For example, a client invests in a portfolio of 100 bonds with average rating BBB. At the same time, the client buys insurance on the first default in the portfolio. The cost of first debtor defaulting will be met by another party. This increases the credit quality of the portfolio to, say, BBB+.

The third objective could be to provide a desired payoff profile to the client according to the client's *views*. For example, the client may think that the yield curve will become steeper. The structurer will offer an instrument that gains value if this expectation is realized.

Finally, a fourth objective may be facilitating *asset/liability management needs* of the client. For example, a corporate treasurer thinks that cost of funds would increase in the future and may want to get a payer interest rate swap. The structurer will provide a modified swap structure that will protect against this eventuality at a smaller cost. In the following we discuss these generalities using different sectors in financial markets.

2.1. Equity Structured Products

First we take a quick glance at the history of equity structured products. This provides a perspective on the most common methodologies used in this sector. The first examples of structured products appeared in the late 1970s. One example was the *stop-loss* strategies. According to these, the risky asset holdings would automatically be liquidated if the prices fell through a target tolerance level. These were precursors of the Constant Proportion Portfolio Insurance (CPPI) techniques to be seen later in Chapter 20. They can also be regarded as precursors of barrier options.

Then, during the late 1980s, market practitioners started to move to principal protected products. Here the original approach was offering “zero coupon bond plus a call” structures. For example, with 5-year treasury rates at r_t , and with an initial investment of $N = 100$, the product would invest

$$\frac{N}{(1 + r_t)^5} \quad (1)$$

into a discount bond with a 5-year maturity. The rest of the principal would be invested in a properly chosen call or put option. This simple product is shown for a one-year maturity in Figure 17-1.

This was followed in the early 1990s with structures that essentially *complicated* the long option position. Some products started to “cap” the upside. The structure would consist of a discount bond, a long call with strike K_L and a *short* call with strike K_U , with $K_L < K_U$. This way, the premium obtained from selling the second call would be used to increase the participation rate, since more could be invested in the long option. This is shown in Figure 17-2, again for a one-year maturity. Other products started using Asian options. The gains of the index to be paid to the investor would be calculated as an *average* of the gains during the life of the contract.

Late 1990s started seeing *correlation* products. A *worst of* structure would pay at maturity, for example 170% of the initial investment plus the return of a worst performing asset in a basket of, say 10 stocks or commodities. Note that this performance could be negative, thus the investor could receive less than 170% return. However, such products were also principal protected and the investor would still recover the invested 100 in the worst case.

In the *best of* case, the investor would receive the return of the best performing stock or commodity given a basket of stocks or commodities. The *observation period* could be over the entire maturity, or could be annual. In the latter case the product would lock in the annual gains of the best-performing stock, which can be different every year. Mid-2000s brought several new versions of these equity-linked structured instruments which we discuss in more detail below, but first we consider the main tools underlying the products.

2.2. The Tools

Equity structured products are manufactured using a relatively small set of *tools* that we will review in this section. We will concentrate on the main concepts and instruments: basically three

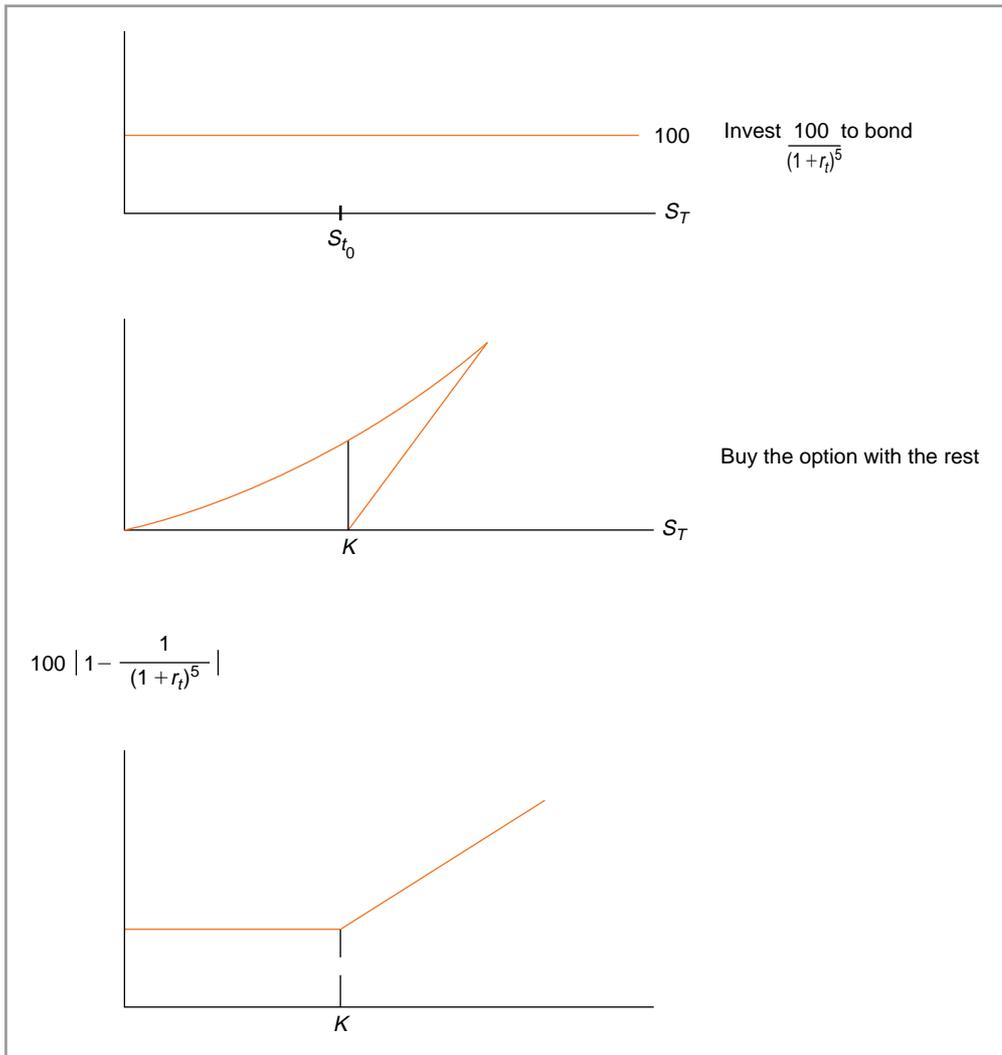


FIGURE 17-1

main types of instruments and a major conceptual issue that will recur in dealing with equity structured products.

First there are *vanilla call* or *put* options. These were discussed in Chapters 8 and 10 and are not handled here. The second tool is *touch* or *digital* options, discussed in a later chapter, but we'll provide a brief summary below.

Touch or digital options are essentially used to provide payoffs (of cash or an asset) if some levels are crossed. Most equity structured products incorporate such *levels*. The third tool is new; it is the so-called *rainbow* options. These are options written on the maximum or minimum of a *basket* of stocks. They are useful since almost all equity structured products involve payoffs that depend on more than one stock. The fourth tool is the *cliquet*. These options are important prototypes and are used in buying and selling *forward starting* options. Note that an equity structured product would naturally span over several years. Often the investor is offered returns of an index during a future year, but the initial index value during these future years would not be known. Hence, such options would have forward-setting strikes and would depend on *forward*

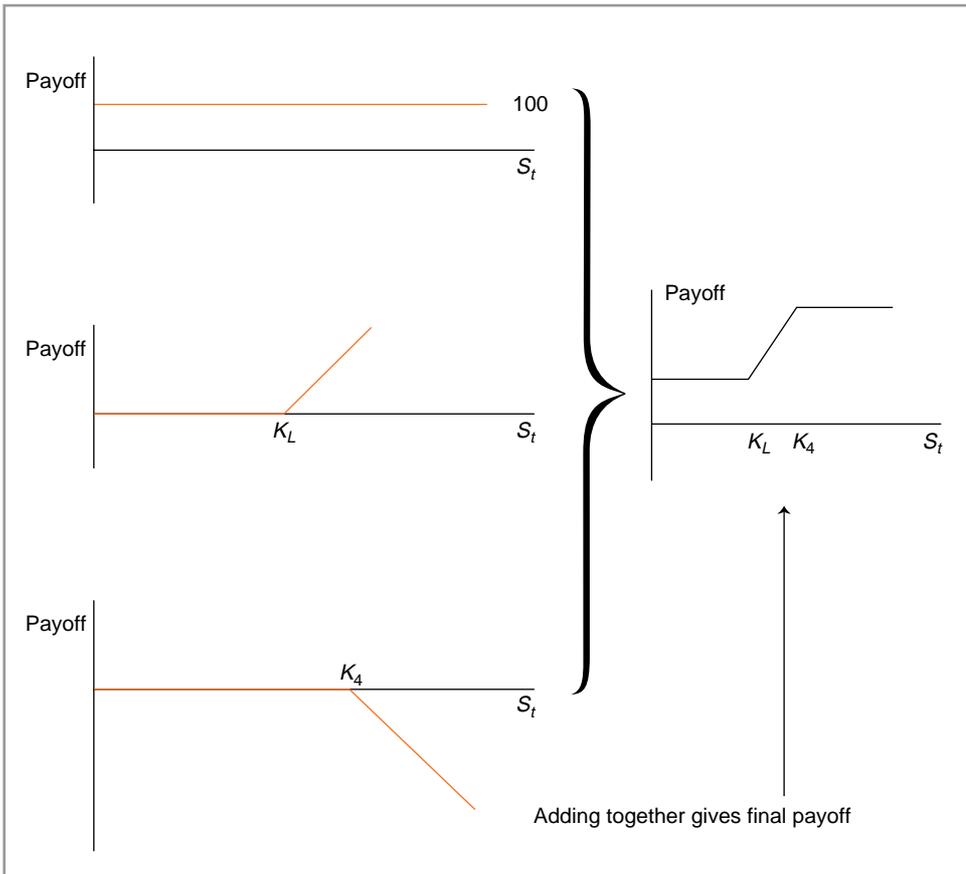


FIGURE 17-2

volatility. Forward volatility plays a crucial role in pricing and hedging structured products, both in equity and in fixed-income sectors.²

2.2.1. Touch and Digital Options

Touch options are similar to the digital options introduced in Chapter 10. European digital options have payoffs that are step functions. If, at the maturity date, a long digital option ends in-the-money, the option holder will receive a predetermined amount of cash, or, alternatively, a predetermined asset. As discussed in Chapter 10, under the standard Black-Scholes assumptions the digital option value will be given by the risk-adjusted probability that the option will end up in-the-money. In particular, suppose the digital is written on an underlying S_t and is of European style with expiration T and strike K . The payoff is $\$R$ and risk-free rates are constant at r , as shown in Figure 17-3. Then the digital call price will be given by

$$C_t = e^{-r(T-t)} \tilde{P}(S_T > K) R \tag{2}$$

² The equity structured products are often principal protected. A discussion of CPPI type portfolio insurance which is relevant here will be considered later. We do not include the CPPI techniques in this chapter.

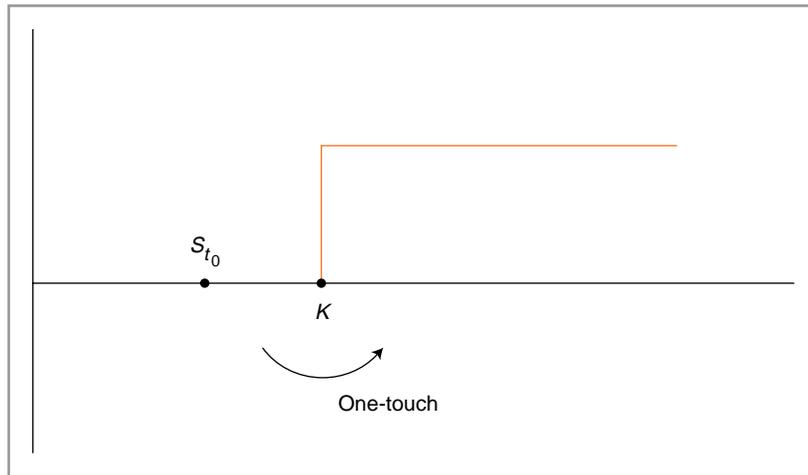


FIGURE 17-3

where \tilde{P} denotes the proper risk-adjusted probability. Digital options are standard components of structured equity products and will be used below.

A *one-touch* option is a slightly modified version of the vanilla digital. A one-touch call is shown in Figure 17-3. The underlying with original price $S_{t_0} < K$ will give the payoff \$1 if (1) at expiration time $T(K < S_T)$ and (2) if the level K is breached only *once*.

A previous chapter discussed a *double-no-touch* (DNT) option which is often used to structure *wedding cake* structures for FX markets.³

The more complicated tools are the *rainbow options*, and the concept of *forward volatility*. We will discuss them in turn before we start discussing recent equity structured products.

2.2.2. Rainbow Options

The term *rainbow options* is reserved for options whose payoffs depend on the trajectories of *more* than one asset price. Obviously, they are very relevant for equity products that have a basket of stocks as the underlying. The major class of such options are those that pay the *worst-of* or *best-of* the n underlying assets. Suppose $n = 2$; two examples are

$$\text{Min} [S_T^1 - K^1, S_T^2 - K^2] \quad (3)$$

where the option pays the smaller of the two price changes on two stocks, and

$$\text{Max} [0, S_T^1 - K^1, S_T^2 - K^2] \quad (4)$$

where the payoff is the larger one and it is floored at zero. Needless to say the number of underlying assets n can be larger than 2, although calibration and numerical burdens make a very large n impractical.

2.2.3. Cliquets

Cliquet options are frequently used in engineering equity and FX-structured products. They are also quite useful in understanding the deeper complexities of structured products.

³ A wedding cake is a portfolio of DNT options with different bases.

A *cliquet* is a series of pre-purchased options with forward setting strikes. The first option's strike price is known but the following options have unknown strike prices. The strike price of future options will be set according to where the underlying closes at the end of each future subperiod. The easiest case is at-the-money options. At the beginning of each observation period the strike price will be the price observed for S_{t_i} .⁴ The number of reset periods is determined by the buyer in advance. The payout on each option is generally paid at the end of each reset period.

EXAMPLE:

A five-year cliquet call on the S&P with annual resets is shown in Figure 17-4. Essentially the cliquet is a basket of five annual at-the-money spot calls.

The initial strike is set at, say, 1,419, the observed value of the underlying at the purchase date. If at the end of the first year, the S&P closes at 1,450, the first call matures in-the-money and the payout is paid to the buyer. Next, the call strike for the second year is reset at 1,450, and so on.

To see the significance of a five-year cliquet, consider two alternatives. In the first case one buys a one-year at-the-money call, then continues to buy new at-the-money calls at the beginning of future years four times. In the second case, one buys a five-year cliquet. The difference between these is that the cost of the cliquet will be known in advance, while the premium of the future calls will be unknown at t_0 . Thus a structurer will know at t_0 what the costs of the structured product will be only if he uses a cliquet.

Consider a five-year maturity again. The chance that the market will close lower for five consecutive years is, in general, lower than the probability that the market will be down after

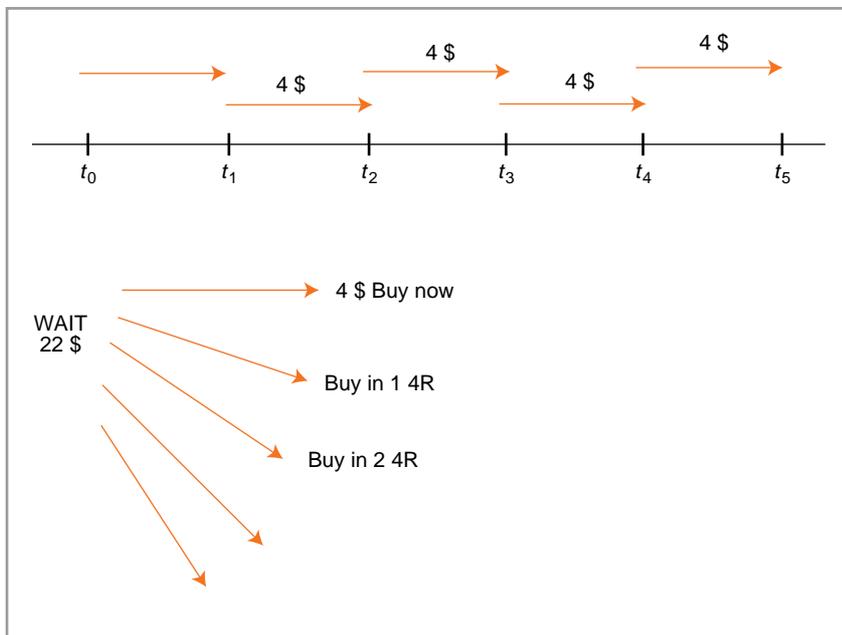


FIGURE 17-4

⁴ Clearly, one can also buy a cliquet where the future strikes set $k\%$ out-of-the-money.

five years. If the market is down after five years, chances are it will close higher in (at least) one of these five years. It is thus clear that a cliquet call will be more expensive than a vanilla at-the-money call with the same final maturity.

The important point is that cliquet needs to be priced using the implied *forward volatility* surface. Once this is done the cliquet premium will equal the present value of the premiums for the future options.

2.3. Forward Volatility

Forward volatility is an important *concept* in structured product pricing and hedging. This is a complicated technical topic and can only be dealt with briefly here. Consider a vanilla European call written at time t_0 . The call expires at T , $t_0 < T$ and has a strike price K . To calculate the value of this call we find an implied volatility and plug this into the Black-Scholes formula. This is called the Black-Scholes implied volatility.

Now consider a vanilla call that will *start at a later date* at t_1 , $t_0 < t_1$. Yet, we have to price the option at time t_0 . The expiration is at t_2 . More important, the strike price of the option denoted by K_{t_1} is unknown at t_0 and is given by

$$K_{t_1} = \alpha S_{t_1} \quad (5)$$

where $0 < \alpha \leq 1$ is a parameter. It represents the moneyness of the forward starting call and hence is an important determinant of the option's cost. The forward call will be an ATM option at t_1 if $\alpha = 1$. Assuming deterministic short rates r , we can write the forward start option value at t_0 as

$$C(S_{t_0}, K_{t_1}, \sigma(t_0, t_1, t_2)) = e^{r(t_2-t_0)} E_{t_0}^{\tilde{P}}[(S_{t_2} - \alpha S_{t_1})^+] \quad (6)$$

where $C(\cdot)$ denotes the Black-Scholes formula, and where the $\sigma(t_0, t_1, t_2)$ is the *forward* Black-Scholes volatility. The volatility is calculated at t_0 and applies to the period $[t_1, t_2]$. We can replace the (unknown) K_{t_1} , using equation (5) and see that the cliquet option price would depend only on the current S_{t_0} and on forward volatility.

Thus the pricing issue reduces to calculating the value of the forward volatility given liquid vanilla option markets on the underlying S_t . This task turns out to be quite complex once we go beyond very simple characterizations of the instantaneous volatility for the underlying process. We consider two special cases that represent the main ideas involved in this section. For a comprehensive treatment we recommend that the reader consult Gatheral (2006).

EXAMPLE: Deterministic Instantaneous Volatility

Suppose the volatility parameter that drives the S_t process is time dependent, but is deterministic in the sense that the only factor that drives the instantaneous volatility σ_t is the time t . In other words we have the risk-neutral dynamics,

$$dS_t = rS_t + \sigma_t S_t dW_t \quad (7)$$

Then the implied Black-Scholes volatility for the period $[t_0, T_1]$ is defined as

$$\sigma_{BS}^{T_1} = \sqrt{\frac{1}{T_1 - t_0} \int_{t_0}^{T_1} \sigma_t^2 dt} \quad (8)$$

In other words, $\sigma_{BS}^{T_1}$ is the average volatility during period $[t_0, T_1]$. Note that under these conditions the variance of the S_t during this period will be

$$\left(\sigma_{BS}^{T_1}\right)^2 (T_1 - t_0) \quad (9)$$

Now consider a longer time period defined as $[t_0, T_2]$ with $T_1 < T_2$ and the corresponding implied volatility

$$\sigma_{BS}^{T_2} = \sqrt{\frac{1}{T_2 - t_0} \int_{t_0}^{T_2} \sigma_t^2 dt} \quad (10)$$

We can then define the forward Black-Scholes variance as

$$\left(\sigma_{BS}^{T_2}\right)^2 (T_2 - t_0) - \left(\sigma_{BS}^{T_1}\right)^2 (T_1 - t_0) \quad (11)$$

Plug in the integrals and take the square root to get the forward implied Black-Scholes volatility from time T_1 to time T_2 , $\sigma_{BS}^f(T_1, T_2)$

$$\sigma_{BS}^f(T_1, T_2) = \sqrt{\frac{1}{T_2 - T_1} \int_{T_1}^{T_2} \sigma_t^2 dt} \quad (12)$$

The important point of this example is the following: In case the volatility changes deterministically as a function of time t , the forward Black-Scholes volatility is simply the forward volatility. Hence it can be calculated in a straightforward way given a (deterministic) volatility surface. What intuition suggests is correct in this case. We now see a more realistic case with *stochastic* volatility where this straightforward relation between forward Black-Scholes volatility and forward volatility disappears.

EXAMPLE: Stochastic Volatility

Suppose the S_t obeys

$$dS_t = rS_t + \sigma I_t S_t dW_t \quad (13)$$

Where the I_t is a zero-one process given by:

$$I_t = \begin{cases} .30 & \text{With probability } .5 \\ .1 & \text{With probability } .5 \end{cases} \quad (14)$$

Thus, we have a stochastic volatility that fluctuates randomly (and independently of S_t) between high and low volatility periods. Then, the average variances for the periods $[t_0, T_1]$ and $[t_0, T_2]$ will be given respectively as

$$\begin{aligned} \left(\sigma^{T_1}\right)^2 (T_1 - t_0) &= \int_{t_0}^{T_1} E\left[(\sigma(I_t) dW_t)^2\right] \\ &= (T_1 - t_0) (.3)^2 \end{aligned} \quad (15)$$

which implies that forward volatility will be .2.

Yet, the forward implied Black-Scholes volatility will not equal .2.

According to this, whenever instantaneous volatility is stochastic, calculating the Black-Scholes forward volatility will not be straightforward. Essentially, we would need to model this stochastic volatility and then, using Monte Carlo, price the vanilla options. From there we would back out the implied Black-Scholes forward volatility. The following section deals with our first example of equity structured products where forward volatility plays an important role.

2.4. Prototypes

The examples of major equity structured products below are selected so that we can show the major methods used in this sector. Obviously, these examples cannot be comprehensive.

We first begin with a structure that imbeds a cliquet. The idea here is to benefit from fluctuations in forward equity prices. *Forward volatility* becomes the main issue. Next we move to structures that contain rainbow options. Here the issue is to benefit from the maxima or minima of stocks in a basket. The structures will have exposure to *correlation* between these stocks and the investor will be long or short correlation. Third, we consider Napoleon type products where the main issue becomes hedging the forward volatility movements. With these structures the volatility exposure will be convex and there will be a *volatility gamma*. If these dynamic hedging costs involving volatility purchases and sales are not taken into account at the time of initiation, the structure will be mispriced. Such dynamic hedging costs involving volatility exposures is another important dimension in equity structured products.

2.4.1. Case I: A Structure with Built-In Cliquet

Cliquets are convenient instruments to structure products. Let s_t be an underlying like stock indices or commodities or FX. Let g_{t_i} be the annual rate of change in this underlying calculated at the end of year.

$$g_{t_i} = \frac{s_{t_i} - s_{t_{i-1}}}{s_{t_{i-1}}} \quad (16)$$

where $t_i, i = 1, 2, \dots, n$ are settlement dates. There is no loss of generality in assuming that t_i is denoted in years.

Suppose you want to promise a client the following: Buying a 5-year note, the client will receive the future annual returns $\lambda g_{t_i} N$ at the end of every year t_i . The $0 < \lambda$ is a parameter to be determined by the structurer. The annual returns are *floored* at zero. In other words, the annual payoffs will be

$$P_{t_i} = \text{Max}[\lambda g_{t_i} N, 0] \quad (17)$$

The λ is called the *participation rate*.

It turns out that this structure is less straightforward than appears at the outset. Note that the structurer is promising *unknown* annual returns, with a known coefficient λ at time t_0 . In fact, this is a cliquet made of one vanilla option and four forward starting options. The forward starting options depend on forward volatility. The pricing should be done at the initial point t_0 after calculating the forward volatility for the intervals $t_i - t_{i-1}$. Figure 17-5 shows how one can use cliquets to structure this product. Essentially the structurer will take the principal N , deposit part of it in a 5-year Treasury note, and with the remainder buy a 5-year cliquet. We would like to discuss this in detail.

First let us incorporate the simple principal protection feature. Suppose 5-year risk-free interest rates are denoted by $r\%$. Then the value at time t_0 of a 5-year default-free Treasury bond will be given by

$$PV_{t_0} = \frac{100}{(1+r)^5} \quad (18)$$

Clearly this is less than 100. Then define the *cushion* Cu_{t_0}

$$Cu_{t_0} = 100 - PV_{t_0} \quad (19)$$

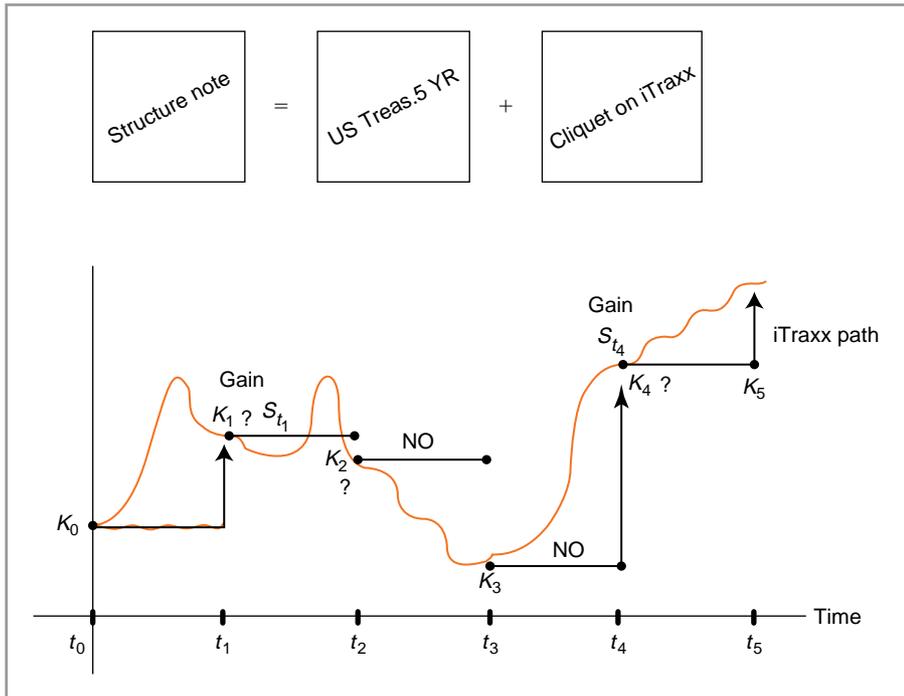


FIGURE 17-5

Note that

$$0 < C u_{t_0} \tag{20}$$

and that these funds can be used to buy options. However, note that we cannot buy *any* option; instead we buy a *cliquet* since λ times the *unknown* annual returns are promised to the investor. The issue is how to price the options on these unknown forward returns at time t_0 . To do this, forward volatility needs to be calibrated and substituted in the option pricing formula which, in general, will be Black-Scholes.

With this product, if the annual returns are positive the investor will receive λ times these returns. If the returns are negative, then the investor receives nothing. Note that even in a market where the long run trend is downward, *some* years the investor may end up getting a positive return.

2.4.2. Case II: Structures with Mountain Options

Structures with payoffs depending on the maximum and minimum of a basket of stocks are generally denoted as *mountain options*. There are several examples. We consider a simple case for each important category.

Altiplano Consider a basket of stocks with prices $\{S_t^1, \dots, S_t^n\}$. A level K is set. For example, 70% of the initial price. The simplest version of an Altiplano structure entitles the investor to a “large” coupon if none of the S_t^i hits the level K during a given time period $[t_i, t_{i-1}]$. Otherwise, the investor will receive lower coupons as more and more stocks hit the barrier. Typically, once 3–4 stocks hit the barrier the coupon becomes zero. The following is an example.

EXAMPLE: *An Altiplano*

Currency: Eur; Capital guarantee: 100%; Issue price: 100.

Issue date: 01-01-2008; Maturity date: 01-01-2013

Underlying basket: {Pepsico, JP Morgan Chase, General Motors, Time Warner, Seven-Eleven}

Annual coupons:

Coupon = 15% if no stocks settle below 70% of its reference price on coupon payment dates.

Coupon = 7% if one stock settles below the 70% limit.

Coupon = 0.5% if more than one stock settles below the limit.

Figure 17-6 shows how we can engineer such a product. Essentially, the investor has purchased a zero coupon bond and then sold five digital puts. The coupons are a function of the premia for the digitals. Clearly this product can offer higher coupons if the components of the reference portfolio have higher volatility.

This product has an important property that may not be visible at the outset. In fact, the Altiplano investor will be *long equity correlation*, whereas the issuer will be *short*. This property is similar to the pricing of CDO equity tranches and will be discussed in detail later. Here we consider two extreme cases.

Suppose we have a basket of k stocks $S_t^i, i = 1, 2, \dots, k$. For simplicity suppose all volatilities are equal to σ . For all stocks under consideration we define the annual probability of *not* crossing the level KS_{t_0} ,

$$P(S_t^i, t \in [t_0, t_1] > KS_{t_0}^i) = (1 - p^i) \quad (21)$$

for all i and $t \in [t_0, T]$. Here the $(1 - p^i)$ measure the probability that the i th stock never falls below the level KS_{t_0} . For simplicity let all p^i be the same at p . Then if the $S_t^i, i = 1, 2, \dots$ are independent, we can calculate the probability of receiving the high coupon at the end of the first year as $(1 - p)^k$. Note that as k increases, this probability goes down.

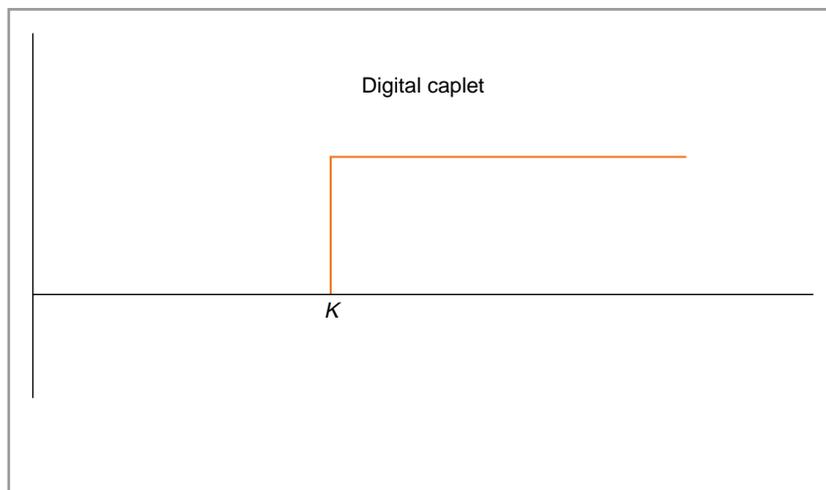


FIGURE 17-6

Now go to the other extreme case and assume that the correlation between S_t^i becomes *one*. This means that all stocks are the same. The probability of receiving a high coupon becomes simply $1 - p$. This is the case since all of these stocks act identically; if *one* does not cross the limit, *none* will. Since $0 < p < 1$ with $k > 1$ we have

$$(1 - p)^k < (1 - p) \quad (22)$$

Thus, the investor in this product will *benefit* if correlation increases, since the investor's probability of receiving higher coupons will increase.

Himalaya The *Himalaya* is a call on the average performance of the best stocks within the basket. In one version, throughout the life of the option, there are preset observation dates, say t_1, t_2, \dots, t_n , at which the best performer within the basket is sequentially removed and the realized return of the removed stock is recorded. The payoff at maturity is then the sum of all best returns over the life of the product.

EXAMPLE: A Himalaya

Currency: Eur; Issue price: 100.

Issue date: 01-01-2007; Maturity date: 01-01-2012.

Underlying basket: 20 stocks possibly from the United States and Europe.

Redemption at Maturity:

If the basket rose, the investor receives the maximum of the basket of remaining securities observed on one of the evaluation dates.

If the basket declined, the investor receives the return of the basket of remaining securities observed on the last evaluation date.

In this case the return is related to the maximum or minimum of a certain basket over some evaluation periods. Clearly, this requires writing rainbow options, including them in a structure, and then selling them to investors.

2.4.3. Case III: The Napoleon and Vega Hedging Costs

A *Napoleon* is a capital-guaranteed structured product which gives the investor the opportunity to earn a high fixed coupon each year, say, $c_{t_0} = 10\%$, *plus* the worst monthly performance in an underlying basket of k underlying stocks S_t^i . If k is large there will be a high probability that the worst performance is negative. In this case the actual return could potentially be much less than the coupon c_{t_0} .

The importance of Napoleon for us is the implication of dynamic hedging that needs to accompany such products. The key issue is that Napoleon-type products cannot be hedged statically and require dynamic hedging. But the main point is that the dynamic hedging under question here is different than the one in plain vanilla options. In plain vanilla options the practitioner buys and sells the underlying S_t^i to hedge the directional movements in the option price. This dynamic hedging results in *gamma* gains (losses). What is being hedged in Napoleon-type products is the volatility exposure. The practitioner has to buy and sell volatility dynamically.

These products have exposure to the so-called *volatility gamma*. The structurer needs to buy option volatility when volatility *increases* and sell it when volatility *decreases*. This is similar to

the *gamma* gains of a vanilla option discussed in Chapter 8, except that now it is being applied to the *volatility* itself rather than the underlying price; hence the term *volatility gamma*. By buying volatility when vol is expensive and selling it when it is cheap, the structurer will suffer hedging costs. The expected value of these costs need to be factored in the initial selling price, otherwise the product will be mispriced.

EXAMPLE: Napoleon Hedging Costs

Suppose there is a basket of 10 stocks $\{S_t^1, \dots, S_t^{10}\}$ whose prices are monitored monthly. The investor is paid a return of 10% plus the worst monthly return among these stocks.

Suppose now volatility is very high with monthly moves of, say, 50%. Then a 1 percentage point change in volatility does not matter much to the seller since, chances are, one of the stocks will have a negative monthly return which will lower the coupon paid. Thus the seller has relatively little volatility exposure during high volatility periods.

If, on the other hand, volatility is very low, say, 9%, then the situation changes. A 1% move in the volatility will matter, leading to a high volatility exposure.

This implies that with low volatility the seller is long volatility, and with high volatility, the volatility exposure tends to zero. Hence the structurer needs to sell volatility when volatility decreases and buy it back when volatility increases in order to neutralize the vol of the position.

This is an important example that shows the need to carefully calculate future hedging costs. If volatility is *volatile*, Napoleon-type structured products will have volatility *gamma* costs to hedging costs that need to be incorporated in the initial price.

2.5. Similar FX Structures

It turns out that cliquets, mountain options, Napoleons or other structured equity instruments can all be applied to FX or commodity sectors by considering baskets of currencies or commodities instead of stocks. Because of this close similarity we will not discuss FX and commodity structures in detail. Wystub (2006) is a very good source for this.

3. Structured Fixed-Income Products

Structured fixed-income products follow principles that are similar to the ones based on equity or commodity prices. But, the analysis of the principles of fixed-income is significantly more complex for several reasons.

First, the main driving force behind the fixed-income structured products is the *yield curve*, which is a k -dimensional stochastic process. Equity or commodity indices are scalar-valued stochastic processes, and elementary structured products based on them are easy to price and hedge. Equity (commodity) products that are based on *baskets* would have a k -dimensional underlying, yet the arbitrage conditions associated with this vector would still be simpler. Second, the basis of fixed-income products is the Libor-reference system, which leads to the Libor market model or swap models. In equity even when we deal with a vector process, there is no need to use similar models. Third, fixed income markets are bigger than the equity and commodity markets combined. The very broad nature of fixed-income products' maturities and credits can make some maturities in fixed income much less *liquid*. Finally, the fixed-income

structured products do have long maturities whereas in equity or commodity-linked derivatives they are relatively short dated.

3.1. Yield Curve Strategies

It is clear that most fixed-income structured products will deal with yield curve strategies. There aren't too many yield curve movements.

1. The yield curve may shift *parallel* to itself up or down—called the *level effect*.
2. The yield curve *slope* may change. This could be due to monetary policy changes, or due to changes in inflationary expectations. The curve can steepen if the Central Bank lowers short-term rates, or flatten if the Central Bank raises short-term rates. This is called the *slope effects*.
3. The “belly” of the curve may go up and down. This is in general interpreted as a *convexity* effect and is related to changes in interest rate volatility.

The next point is that many of these yield curve movements are at least partially *predictable*. After all, Central Banks often announce their future policies clearly to inform the markets. Structurers can use this information to put together CMS-linked products that benefit from the expected yield curve movements. One can also add callability to enhance the yield further.⁵

The reading below is one example of how the structurers look at yield curve strategies.

EXAMPLE:

The popularity of CMS-linked structured notes derives from end users wanting to take advantage of the inverse sterling yield curve, which seems to have stabilized in the long end.

A typical structured note might be EUR5-50 million, with a 20-year maturity. It could pay a coupon of 8% for the first five years, and then an annual coupon based on the 10-year sterling swap rate, capped at 8% for the remainder of the note. It would be noncallable. The 10-year sterling swap rate was about 6.77% last week.

The long end of the sterling yield curve has likely stopped dropping because U.K. life insurers, who have been hedging guaranteed rate annuity products sold in the 1980s, have stopped scrambling for long-dated gilts. They have done so either because they no longer require further hedging, or because they have found more economical ways of doing so. If the long end fails to fall further, investors are more secure about receiving a long-term rate in a CMS, a trader said.

The sterling yield curve, which is flat for about three years, and then inverts, makes these products attractive for investors who believe the curve will disinvert at some point in the future, according to traders. (IFR, January 31, 2000.)

Hence it is clear that fixed-income structured products are heavy in terms of their involvement in Libor, swaption, and call/floor volatilities and their dynamics. Essentially, to handle them the structurer needs to have, at the least, a very good command of the forward Libor and swap models.

⁵ There is some possibility that investors are more interested in yield enhancement and are willing to tolerate some *duration uncertainty*. Accordingly, if a product is called before maturity, investors may not be too disappointed. In fact, for many structured products, many retail investors prefer that the product is called, and that they receive the first year high coupon.

3.2. The Tools

Some of the tools involved in designing and risk managing structured products were discussed earlier. Digital and rainbow options and forward volatility were among these. Fixed income structure products use additional tools: Two familiar tools are modified versions of *Cap/Floors* and *Swaptions* and a third major tool is *CMS swaps*. We review these briefly in this section.

A *digital caplet* is similar to a vanilla caplet. It makes a payment if the reference Libor rate exceeds a cap level. The difference is the payoff. While the vanilla caplet payoff may vary according to how much the Libor exceeds the level, the digital caplet would make a constant payment no matter what the excess is, given that the Libor rate is greater than the cap level.

A *Bermudan swaption* can be defined as an option on a swap rate s_t . The option can be exercised only at some specific dates t_1, t_2, \dots . When the option is exercised, the option buyer has the right to get in a payer (receiver) swap at a predetermined swap rate κ . The option seller has the obligation of taking the other side of the deal. Clearly with this product the option buyer receives swaps of different maturity as the exercise date changes.

CMS swaps are fundamental elements of fixed-income structured products; hence, we review them separately.

3.3. CMS Swaps

A *CMS swap* is similar to a plain vanilla swap except for the definition of the floating rate. They were discussed in Chapter 5. There is a fixed payer or receiver, but the floating payments would *no longer* be Libor-referenced. Libor is a short-term rate with tenors of 1, 2, 3, 6, 9, or 12 months. It can only capture views concerning increasing or decreasing *short-term* rates. In a CMS swap, the floating rate will be *another vanilla swap rate*. This swap rate could have a maturity of 2 years, 3 years, or even 30 years. This way, instruments that benefit from increasing or decreasing long-term rates can also be put together. A 10-year CMS with a maturity of 2 years is shown in Figure 17-7.

A special property of CMS swaps should be repeated at this point. Note that at every reset date, the contract requires obtaining, say, a 10-year swap rate from some formal fixing process. This 10-year swap rate is normally valid for the next 10 years. Yet, in a CMS swap that settles semiannually, this rate will be used for the next 6 months *only*. At the next reset date the new fixing will be used. Thus, the floating rate that we are using is not the “natural rate” for the payment period. In other words, denoting the 10-year floating swap rate by $S_{t_i}^{10}$ we have:

$$\frac{(1 + s_{t_i}^{10} \delta)}{(1 + L_{t_i} \delta)} \neq 1 \quad (23)$$

even though both rates are “floating.” As long as the yield curve is upward sloping, the ratio will in fact be greater than one. But, in the case of a *vanilla* swap, each floating rate L_{t_i} is the natural rate for the payment period and we have:

$$\frac{(1 + L_{t_i} \delta)}{(1 + L_{t_i} \delta)} = 1 \quad (24)$$

This is true regardless of whether we have observed the L_{t_i} or not. For this reason, the CMS swaps require a *convexity adjustment*. This means, heuristically speaking, that the future unknown floating rates cannot simply be replaced by their forward equivalents. For example, if in 3 years we receive a 10-year floating swap rate s_t , during pricing we cannot replace this by the

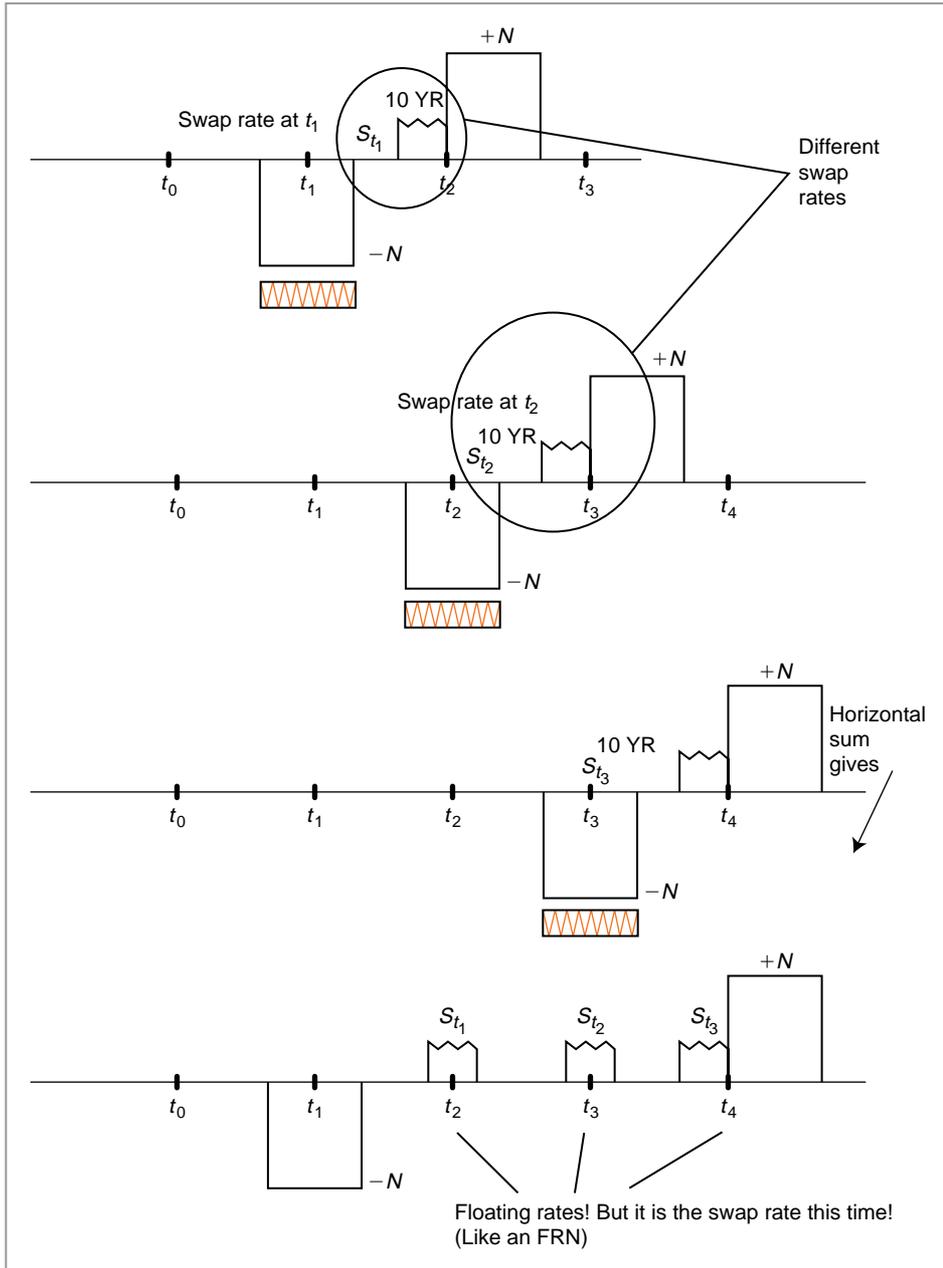


FIGURE 17-7

corresponding forward swap rate s_t^f . Instead we replace it with a forward swap rate adjusted for convexity. See also Figure 17-8.

3.4. Yield Enhancement in Fixed Income Products

Suppose an investor desires an enhanced return, or a corporation wants a hedging solution at a lower cost. The general principle behind putting together such structured products is similar to

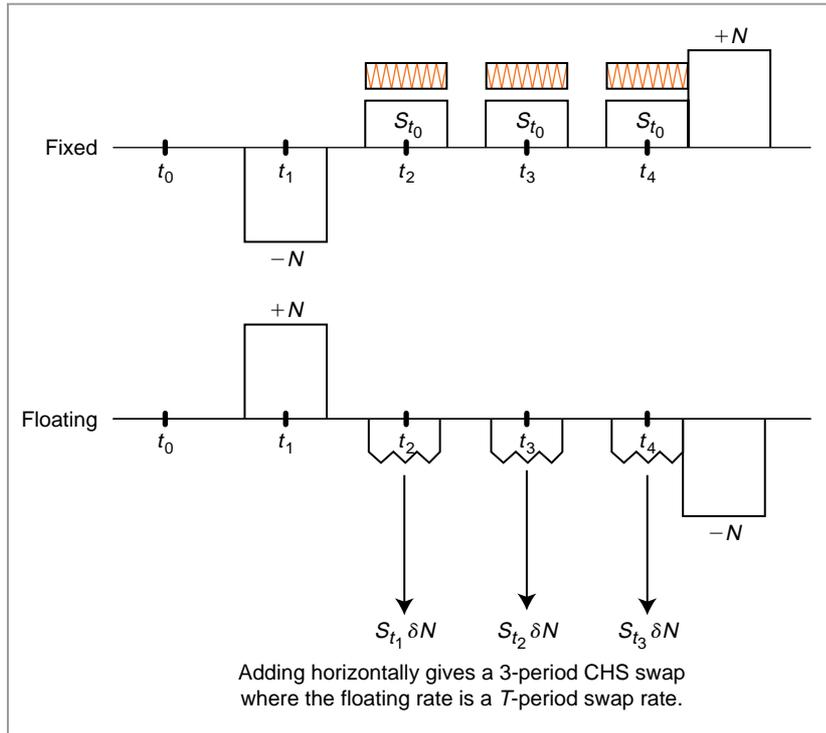


FIGURE 17-8

equity products and is illustrated in the following contractual equation:

Buy a standard asset	+	Sell one or more options	=	An asset with enhanced return	(25)
----------------------------	---	--------------------------------	---	--	------

As in equity structured products, in order to offer a return higher than the one offered by straight bonds, make the client sell one or more options. In fact, as long as the client properly understands the risks and is willing to bear them, the more expensive and more numerous the options are, the higher will be the return. If the client is a corporation and is looking for a cheaper hedge, selling an option would again lower the associated costs.

In structured fixed income products there are at least two standard ways one can enhance yields.

3.4.1. Method 1: Sell Cap Volatility

The first method to enhance yields is conceived so as to make the client sell cap/floor volatility. Remember that a caplet was an insurance written on a particular Libor rate that made a payment if the observed Libor rate went above (below) the cap level (floor level). Then, one can consider daily fixings of Libor and make a digital caplet-type payment when a day's observation stays within a range, say [0,7%]. If the observed Libor exceeds that rate for that day, no interest is received.

This way, the client is *selling* digital caplet volatility and he or she will receive an enhanced yield for bearing this risk. Such products are called Range Accrual Notes (RAN). The client will

earn interest for the proportion of the day's Libor observations that remain within the range. This feature would be suitable for a client who does not expect Libor rates to fluctuate significantly during the maturity period.

Let $F_t^{t_i}$ be the time- t six-month *forward* rate associated with the Libor rate L_{t_i} . The associated settlement of the spot Libor is done, in-arrears, at time t_{i+1} and the day-count adjustment parameter is δ as usual.

Let the index $j = 1, 2, \dots$ denote *days*. A typical caplet starts on day $t_i + (j - 1)$ and has an expiration one day later at $t_i + j$. Each caplet's payoff will depend on the selected reference rate that is followed daily. Often this would be the Libor rate at time $t_i + j$, L_{t_i+j} . Depending on this daily observation the caplets will expire in- or out-of-the-money. In other words, the seller collects *daily* fixings on the Libor rate and sees if the rate stayed within the *range* that day. If it does, there will be a digital payoff for that day (i.e., interest accrues); otherwise, no interest is earned for that particular day.

On the other hand, the actual *amount paid* will depend on *another* predetermined Libor rate. The rate applied to the payoff will be $L_{t_i} + \text{spread}$, settled at t_{i+1} .⁶ According to this the return of the structured product return is a function of the payoffs of m digital options, where m is the number of calendar days in the payment period. Hence, the issue of *whether interest is earned or not* and the *payoff* depend on different Libor rates for each settlement period.

Symbolically, assuming that interest paid is constant at R , the j th day's payoff of the caplets can be written as

$$\text{Pay}_{t_i+j} = \begin{cases} L_{t_i} \frac{1}{360} N & \text{If } L_{t_i+j} \geq L^{\max} \\ 0 & \text{If } L_{t_i+j} \leq L^{\max} \end{cases} \quad (26)$$

where L^{\max} is the upper limit of "range," N is the notional amount, and Pay_{t_i+j} is the daily payoff that depends on the j th observed Libor rate L_{t_i+j} . The investor will be *short* this caplet. How would this enhance the return?

Suppose there are m days during the interest payment period;⁷ then the client is selling m digital caplets. For observation days these caplets are written on that day's Libor rate that we denoted by L_{t_i+j} . For weekends, the previous observation day's Libor is used. So these digital caplets can be regarded as an m -period digital *cap*, made of caplets with daily premiums c_{t_i+j} , if settled at the end of that day. The investor receives the daily premiums and pays off that day's payoff at every $t_i + j$, instead of collecting all the cap premiums at the contract conception t_0 . The total value of these digital caplets at the payment time t_{i+1} will be given by

$$C_{t_0} = \sum_{j=1}^m B(t_0, t_i + j) c_{t_i+j} \quad (27)$$

where c_{t_i+j} are the caplet premiums for the option that starts at time $t_i + j$. Clearly these quantities are known at time t_0 .⁸ Note that this quantity is measured in time t_{i+j} dollars. Then the enhanced yield of the RAN settled at time t_{i+j} will be given by

$$L_{t_i} + c_{t_i+j} \quad (28)$$

⁶ Or alternatively, it could simply be a fixed rate, say R .

⁷ For example, on a 30/360 day basis, and semiannual payment periods, we will have $m = 180$.

⁸ Hence a more complicated notation could be useful. We can let the caplet premia denoted by $\{c(t_0, t_i + (j - 1), t_i + j)\}$ which means that this is the premium calculated at time t_0 , for a caplet that starts at $t_i + (j - 1)$ and expires at $t_i + j$.

At the time of inception t_0 of the note, the relevant Libors will be $\{L_{t_i}\}$ and these will be “equivalent” to s_{t_0} , the swap rate observed at the time of inception. So the enhanced yield can in fact be expressed by the constant R_{t_0} :

$$R_{t_0} = s_{t_0} + c_{t_i+j} \quad (29)$$

If at time t_0 the structurer observes the (1) swap rate s_{t_0} , (2) the forward volatilities of each digital cap c_j , and (3) the discount factors $B(t_0, t_{i+1})$, then the R_{t_0} can be calculated. Thus the investor will receive the $R_{t_0}N$ and will pay the payoffs of daily caplets that expire in-the-money. Naturally, all this assumes a correct calculation of the digital cap premium c_j . Here there are some small technical complications. However, before we get to these we look at an example.

EXAMPLE: EURIBOR Accrual Note

Issuer: Bank ABC; Maturity: 5 years; Issue price: 100.

*Coupon: $\frac{n}{m} * (7\%)$*

Payment dates: Semiannual, in arrears on an ACT/360 basis

Reference rate: 6M Euribor fixed according to Reuters page EURIBOR01.

Range: $0\% - L^{\max}$

n: Number of calendar days in the interest period on which the reference rate fixed at or below L^{\max} . For days which are not observation days, the preceding observation of the reference rate is applied.

m: Number of calendar days in the interest period.

Interest period: From (and including) issue date (first period)/previous coupon payment date (all other periods) to (and excluding) the following coupon payment date.

Observation day: Every business day of the interest period.

Note that the underlying reference rate that we use to determine the payoff of the digital caplets are 6-month Libor rates, which are not the “natural rates” for the 1-day payoffs. Hence, the pricing of these digital caplets would require that a convexity adjustment is applied to the Libor rates, similar to CMS swaps.

3.4.2. Method 2: Sell Swaption Volatility

Making a straight bond callable is the second way of enhancing yields. This will result in the investor being short *swaption volatility*.

The difference is important. In the first case one is writing a series of options on a *single* cash flow, namely the caplet payoff. But in the case of callable bonds, the investor will write options on all the cash flows *simultaneously* and will receive his principal 100 if the bond is called. Thus swaption involves payoffs with baskets of cash flows. These cash flows will depend on different Libor rates. When the swaption is Bermudan, this is similar to selling several options (although dependent on each other) at the same time. Hence the Bermudan swaption will be more expensive and there will be more yield enhancement.

An investor that buys a callable Libor exotic has sold the issuer the right (but not the obligation) to redeem the notes at 100% of the face value at any given *call date*. A note that is callable just once (European) will have a lower yield than a comparable note with multiple calls (Bermudan). The question whether a callable note will be called or not depends on the initially

assumed dynamics of the forward Libor rates versus the behavior of these forward rates and their volatilities in the future.

4. Some Prototypes

In this section we discuss some typical fixed income structured products and their engineering in detail using the tools previously introduced. We consider three typical structured products that are representative.

4.1. The Components

In order to engineer fixed income structured products, the market practitioner will need a small number of components. These are

1. The relevance *discount curve* $B(t_0, t_i)$ in a certain currency. This will be used to discount future “expected” cash flows.⁹
2. A relevant *forward curve* in the same currency. This could be a forward Libor curve, or a forward swap curve. Obviously, this can be obtained from the discount curve using relations such as¹⁰

$$(1 + F(t_0, t_i, t_k) \delta) = \frac{B(t_0, t_i)}{B(t_0, t_k)} \quad i \leq k \quad (30)$$

This is equivalent to needing a market for vanilla swaps, i.e., a tradeable swap curve.

3. A market for CMS swaps, since the structurer may want to receive or pay a *floating* rate that can be any point of the yield curve. A fixed CMS swap rate will be paid against this.¹¹
4. A market for (Bermudan) swaptions if the structure is callable.
5. A market for caps/floors if the structure is of range accrual type.

We now show how some prototypes for fixed income structured products can be manufactured using these components. The prototypes we discuss are *exotics* in the sense that the structurer cannot buy the note from some wholesale market and then sell it. The structurer has to *manufacture* the note. In other words, they are exotics because one side of the market does not exist and the structurer has to know how to price and hedge the product in-house.

4.2. CMS-Linked Structures

There are (at least) two kinds of CMS-based products. Some link the coupon to a CMS *rate*. This would be similar to a floating rate note, but the floating rate would be a long-term rate this time. The second kind will be linked to a CMS *spread*. An investor buying CMS spread-linked structures will not be affected by parallel shifts in the yield curve. Rather, the buyer will be affected by the *slope* of the yield curve. Depending on whether the curve flattens or steepens, the buyer of the spread notes would benefit.

⁹ The expectation is with respect to some working probability measure.

¹⁰ To review this go back to the arbitrage argument used to obtain the FRA rates. In this case the FRA rate is defined more generally.

¹¹ Remember that the CMS swap rate is quoted as a spread to the vanilla swap rate, or the relevant Libor rates:

$$s_{t_0}^{cms} = s_{t_0} + \text{spread}$$

A note linked to a CMS rate enables investors to benefit from shifts in the long or short end of the yield curve over prolonged periods of time. The note pays a fixed coupon which goes up as the short end of the yield curve shifts upward over the lifetime of the product.¹² If made *callable*, the investor also receives enhanced returns. This is one of the most common types of structured fixed income products. A straightforward example is shown below.

EXAMPLE: *A CMS-Linked Note*

Issuer: Bank ABC

Tenor: 5 yrs

Principal: Guaranteed at maturity at 100%

Coupon:

Year 1: 7.00%

Year 2: CMS10 + 2.5%

Year 3: CMS10 + 2.5%

Year 4: CMS10 + 2.5%

Year 5: CMS10 + 2.5%

Call: Callable on each coupon date

Suppose an institutional investor expects 10-year rates to rise successively during the next 5 years. This would be equivalent to five consecutive shifts in the long end of the yield curve. Then, the note provides a way to take exposure to this risk. We now discuss how CMS-linked products can be engineered using standard tools in fixed income markets.

4.3. Engineering a CMS-Linked Note

In this section we engineer a straightforward CMS-linked note. Suppose we have “a view” concerning the yield curve. For example, we expect the long end of the yield curve to shift up gradually during the next 5 years. Essentially we would like to put together a portfolio of elementary assets that generates the promised risk-return characteristics of the relevant CMS-linked note with maturity $T = 5$ years. The example discussed above provides a good framework. For simplicity we assume annual interest payment dates. How do we engineer the CMS-linked note shown above?

1. First we select a CMS rate. For example, let the CMS rate be the floating rate of an m -period swap observed at every t_i (in-advance) and paid at time t_{i+1} (in arrears). There is no harm in thinking that $m = 10$ years.
2. Next we manufacture the high, known, first-year coupon c_{t_0} . There is no harm in thinking that $c_{t_0} = 7\%$ as in the above case. The question is, of course, how to manufacture such a high coupon. This value will be determined during pricing.

¹² If these shifts are predetermined and are written in the contract at contract initiation, then the instrument becomes a step-up note.

3. Next, offer a “floating” coupon for the following 4 years of the form

$$c_{t_i} = cms_{t_i}^m + \alpha_{t_0}^{t_i} \quad (31)$$

where the $cms_{t_i}^m$ is the CMS rate of period t_i .

4. The $\alpha_{t_0}^{t_i}$ will be known constants at time t_0 and will have to be determined during the pricing. Below we consider two different formulations for this term.
5. Make the note callable at call dates t_1, t_2, t_3 , and t_4 .

First, some general observations. The first-year coupon needs to be constant, since even with floating rate instruments the first coupon is always known. The investor is taking a position on “floating” rates, but the first floating rate will be observed at the purchase date. Second, the note is made callable. The structurer is making the investor sell an option, so that the returns can be enhanced by this option’s premium. Third, the option is of Bermudan style, so it can be exercised four times at any of the four future dates. This option would naturally be more expensive than a vanilla swaption and hence the investor can be better compensated.

4.3.1. A Contractual Equation

Now we can put together a replicating portfolio, i.e., obtain a contractual equation for this note.

First, start with the floating CMS coupons $cms_{t_i}^m$. How can the structurer pay such coupons to the investor?¹³ Here the answer is straightforward. The structurer gets into a 5-year *receiver* CMS swap at time t_0 . In this swap the first-year coupon is fixed, but at every $t_i, i = 1, 2, 3, 4$ the going CMS rate $cms_{t_i}^m$ will be received by the structurer and that period’s Libor L_{t_i} will be paid. The structurer will pass the $cms_{t_i}^m$ to the investor. So part of the coupon has now been constructed. This situation is shown in Figure 17-9.

What would the $\alpha_{t_0}^{t_i}$ represent then? Calculate the premium of a 4-year Bermudan swaption—call it C_{t_0} . This swaption is on the *CMS rate* and can be exercised four times during the period $[t_0, T]$ annually at each $t_i, i = 1, 2, 3, 4$. Allocate this premium to the four future years by choosing the $\alpha_{t_0}^{t_i}$ such that,

$$C_{t_0} = B(t_0, t_1) \alpha_{t_0}^{t_1} + B(t_0, t_2) \alpha_{t_0}^{t_2} + B(t_0, t_3) \alpha_{t_0}^{t_3} + B(t_0, t_4) \alpha_{t_0}^{t_4} \quad (32)$$

This gives the $\alpha_{t_0}^{t_i}$.

Finally, note that the structurer will be making Libor referenced payments to the CMS market maker. These payments will come from the original principal $N = 100$. The structurer will place the principal into a deposit account and receive floating Libor.

The replication is complete. The structurer buys a receiver a 5-year CMS swap on the 10-year CMS rate and sells a 5-year Bermudan swaption on the CMS rate that can be exercised annually. The original N received as principal is held in a deposit account.

The cash flows of this portfolio are shown in Figure 17-10. These cash flows are identical to the ones promised by the note. The structurer is essentially buying the portfolio, repackaging it, and then selling it to the client as a structured CMS-linked note. We can summarize such CMS-linked structures using contractual equations. For this particular CMS-linked note we have,

¹³ In other words, how would such payments be hedged?

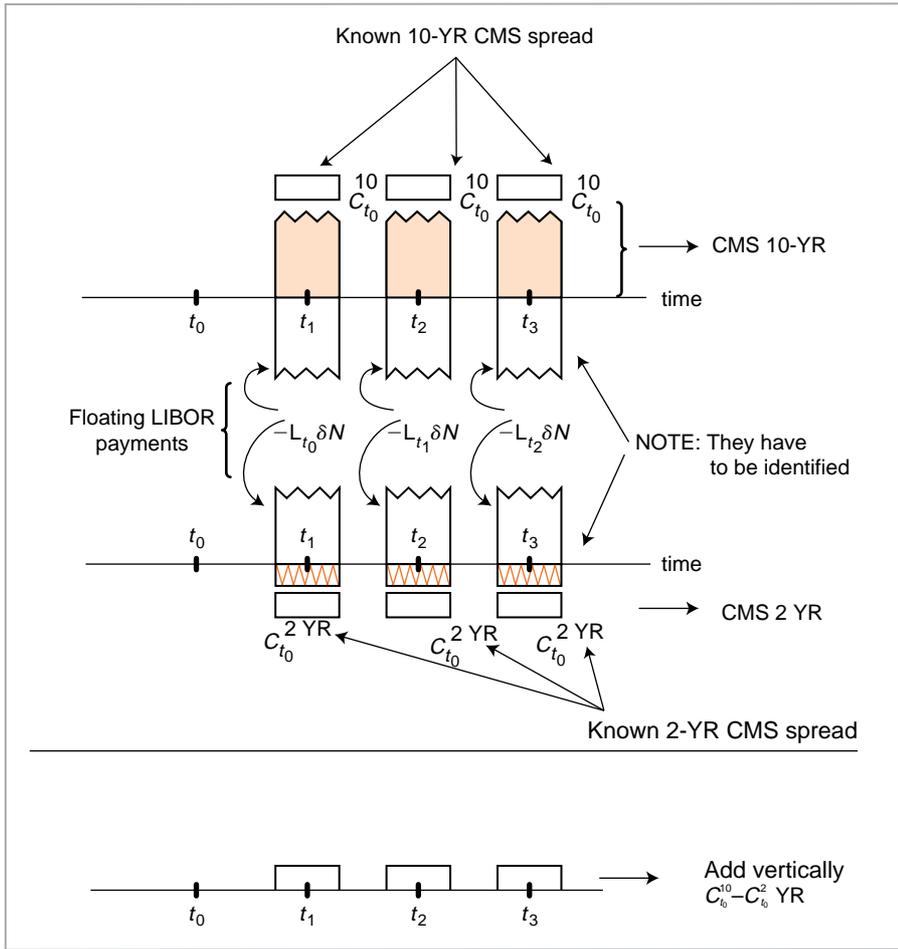


FIGURE 17-9

Callable CMS-linked note	=	Libor deposit	+	Receive 10-year CMS	+	Short Bermudan swaption	(33)
--------------------------------	---	------------------	---	------------------------	---	-------------------------------	------

Note that the replicating portfolio presents further opportunities to the structurer. The structurer may be in need to sell swaption volatility to other clients. Through this CMS-linked note the structurer is buying *swaption volatility* from retail clients. Hence, the note may be a good way of generating a needed supply of swaption volatility at an attractive price. The structurer will naturally sell the swaption at a higher (offer) price than the price of the swaption implicitly bought from the retail client.

A similar comment is valid for the CMS swap. The structurer may in fact be a CMS market maker and may be receiving fixed CMS rates and paying floating rates in a different deal. By marketing the CMS-linked note to the investor the structurer is paying a fixed CMS rate. This is like receiving the asked CMS rate and then paying the bid CMS rate.

In other words, the instruments that need to be purchased from the market may in fact already be in the books of the structurer. The structured product is then a good way of taking these risks, repackaging them, and then selling them, to retail clients.

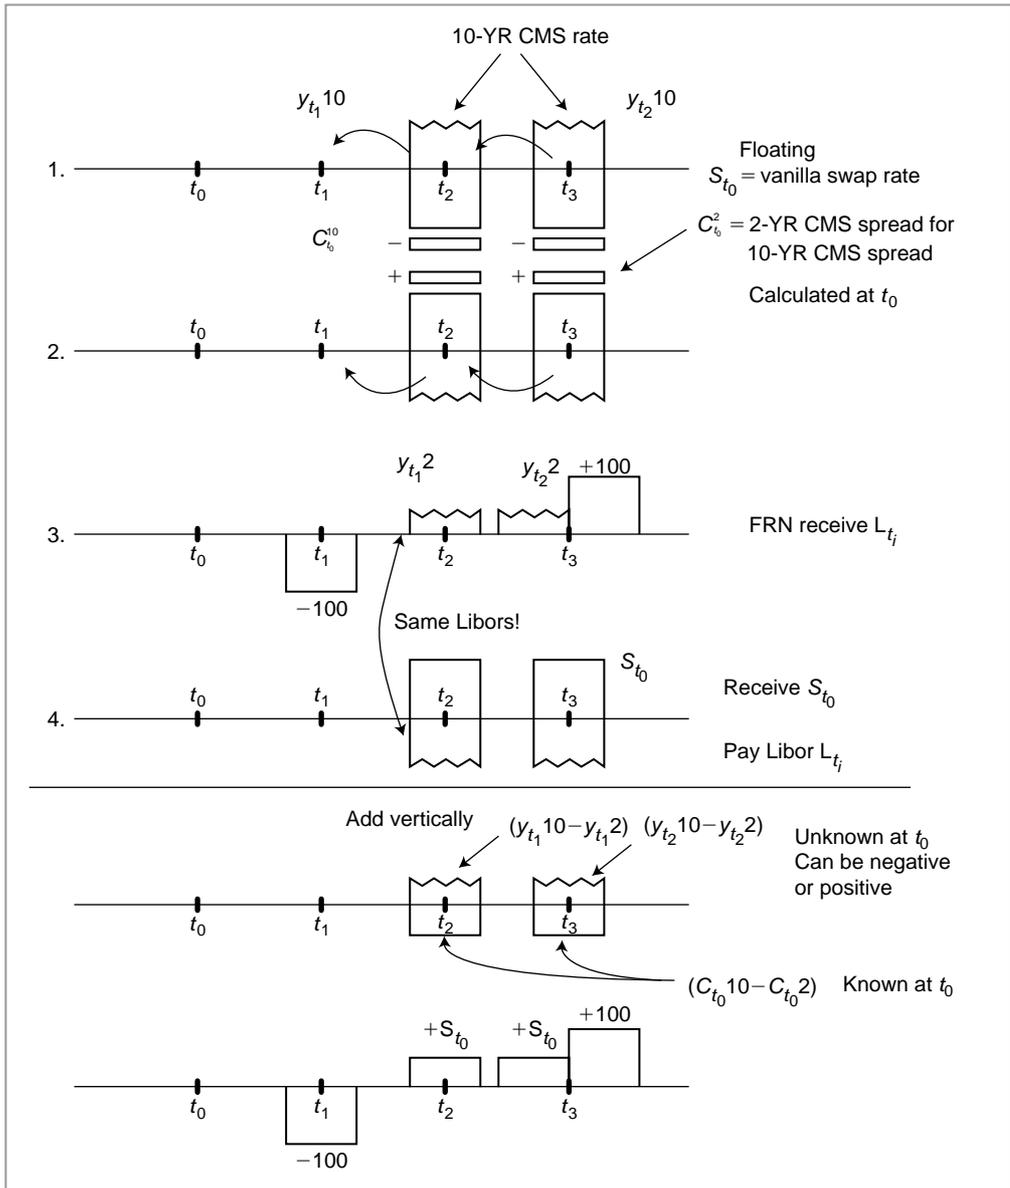


FIGURE 17-10

4.4. Engineering a CMS Spread Note

Suppose an investor expects that the yield curve will *steepen*. In that case, a *CMS-spread structure* is appropriate. First we consider the product itself.

This Libor exotic has three additional properties. First, the instrument is more complicated because the floating annual coupon will depend on *more* than one CMS rate, hence the “spread.” Second, this spread will be offered to the retail investor after multiplying it with a *participation rate*. The participation rate has the potential of significantly enhancing the yields if the expectation turns out to be correct and if the product is not called. The spread in question can become negative. To prevent investors from paying negative coupons, such spread-related coupons are

often *floored at zero*. Third, because the product is written on more than one CMS rate, the value of the structure will explicitly depend on the correlation between these rates.

The example below is typical of CMS spread notes.

EXAMPLE: *A CMS Spread Note*

Issuer: Bank ABC; Tenor: 5 years

Principal: Guaranteed at maturity at 100%

Coupons:

Year 1: 8.00%

Year 2–5: $17 \times (CMS^{30} - CMS^{10})$, max 22%, min 0%

Call: Callable on each coupon date, t_i .

Why would an investor be interested in such a note? Suppose an institutional investor expects that the yield curve will steepen further. This can happen in two ways *at least*.

First, if at constant inflation and hence at constant long rates, the short end interest rates decline. A loosening of monetary policy by the Central Bank due to moderately weakening real economy may be one example. Second, short rates and Central Bank monetary policy may stay the same, but due either to inflationary pressures or strengthening economic activity the long rates may go up. These two cases represent two possible *views* and the spread note will provide one way to take an exposure toward such an event.

This product will offer higher rates if the yield curve keeps steepening gradually, as the coupon is dependent on the differential between the rates. One example is shown in Figure 17-10. Below we synthetically create a CMS spread note starting with more elementary instruments.

4.5. The Engineering

The CMS spread product has (at least) two novel financial engineering features: the product will illustrate the utilization of the *participation rate* and the way one has to *floor* the spreads. In addition, on the pricing side, we will see that correlation becomes an *explicit* additional risk.¹⁴

Now we engineer the CMS spread note mentioned above. The first year coupon is already discussed; it comes from the first year Libor rate which is known, plus part of the option premium sold by the investor. The real novelty of the structure is in the coupons for years $i = 1, 2, 3, 4$. In fact, the coupons are of the form,

$$c_{t_i} = \max [\lambda (cms_{t_i}^m - cms_{t_i}^h) + \alpha_{t_i}, 0] \quad (34)$$

where the $cms_{t_i}^m$, $cms_{t_i}^h$ are two floating CMS rates observed at time t_i , with CMS maturities of m and h years, respectively. As shown in the example, there is no harm in thinking that $m = 10$ and $h = 2$ years, respectively. According to this, the coupon gets bigger or smaller depending on the difference between the 10-year and 2-year swap rates at times t_i in the future. At times t_i we are taking snapshots of the swap curve and paying the client a coupon proportional to the *slope* of the curve. In this particular case the client would get progressively higher coupons if the swap curve becomes steeper and steeper during the following 5 years. The note will benefit from progressive steepening *if* it is not called.

¹⁴ We say explicit, because implicitly one can claim that all CMS products have to deal with Libor rate correlations across the curve.

In order to engineer the coupons, first ignore the floor, and let the α represent a swaption premium allocated to the 5 settlement dates, as before. Now consider engineering the component,

$$\lambda (cms_{t_i}^m - cms_{t_i}^h) \quad (35)$$

These coupons can be replicated using the following position: Pay 2-year CMS rate and receive 10-year CMS rate for five years during the t_1, t_2, t_3, t_4 . This can be accomplished by getting in *two* CMS swaps. The structurer buys λ units of the 10-year CMS and sells λ units of the 2-year CMS. The cash flows are as in Figure 17-11.

The figure incorporates the way CMS market quotes the CMS. It turns out that the market does this as a spread to the plain vanilla swap rate. Thus,

$$cms_{t_i}^{10} = s_{t_0, t_i}^5 + sp_{t_i}^{10,5} \quad (36)$$

And,

$$cms_{t_i}^2 = s_{t_0, t_i}^5 + sp_{t_0}^{2,5} \quad (37)$$

where the s_{t_0, t_i}^5 is the 5-year vanilla (forward) swap rate known at time t_0 for the case swap beginning at time, t_i . The $sp_{t_0}^{10,5}$ is the 10-year CMS *spread* for an instrument of 5-year maturity. The structurer will receive this. The $sp_{t_0}^{2,5}$ on the other hand is the 2-year CMS *spread* for an instrument of 5-year maturity. The structurer will pay this. Note that at time t_0 , both spreads are known for all t_i .

Adding together the components we have:

$$\begin{aligned} c_{t_i} &= \max [\lambda (cms_{t_i}^{10} - cms_{t_i}^2) + \alpha_{t_i}, 0] \\ &= \max \left[\lambda \left(\left(s_{t_0, t_i}^5 + sp_{t_0}^{10,5} \right) - \left(s_{t_0, t_i}^5 + sp_{t_0}^{2,5} \right) \right) + \alpha_{t_i}, 0 \right] \\ &= \max \left[\lambda \left(sp_{t_0}^{10,5} - sp_{t_0}^{2,5} \right) + \alpha_{t_0}, 0 \right] \end{aligned} \quad (38)$$

This is the coupon. Note that at this point of the engineering the floating rates have dropped, and the only unknown on the right-hand side is the participation rate λ . Next we show how λ can be determined.

First remember that the principal N is received from the investor. This is placed in a deposit account that pays the Libor rates L_{t_i} in the future. At time t_0 one can get in a 5-year swap and convert these floating cash flows into a strip of *known* swap rate cash flows at the rate $s_{t_0}^5$. This means at every t_i the structurer will receive the known quantity $s_{t_0}^5$. Consider the following:

$$\lambda \left(sp_{t_0}^{10,5} - sp_{t_0}^{2,5} \right) = s_{t_0}^5 \quad (39)$$

This is an equation where all quantities are known at t_0 except λ . Solve for λ and insert this number in the original coupon rate,

$$c_{t_i} = \max [\lambda (cms_{t_i}^m - cms_{t_i}^h) + \alpha_{t_i}, 0] \quad (40)$$

In this expression the unknowns are the CMS rates and this is the risk the client is assuming. To the structurer, however, the spread $cms_{t_i}^m - cms_{t_i}^h$ does not represent a risk since it can be hedged at a cost of $sp_{t_0}^{10,5} - sp_{t_0}^{2,5}$. The example below shows this simple calculation.

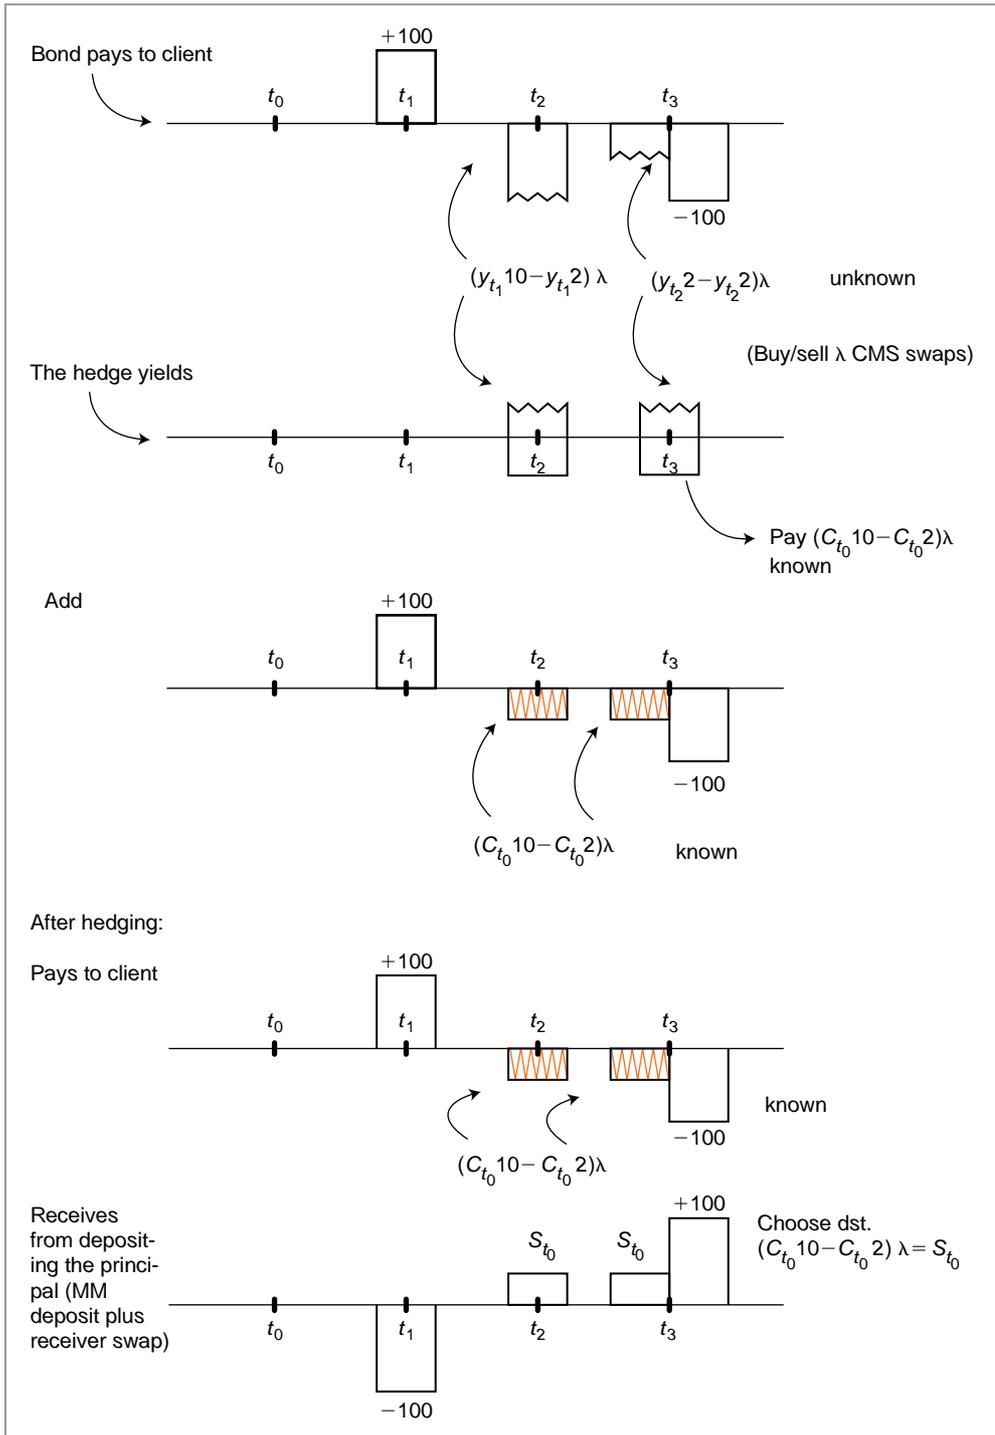


FIGURE 17-11

EXAMPLE:

Suppose the 2-year and 10-year CMS swaps trade at spreads of

$$sp^{10,5} = 50 \text{ bps} \quad (41)$$

$$sp^{2,5} = 20 \text{ bps} \quad (42)$$

The difference is 30 bps. Suppose also that the 5-year swap rate is 4.5%. Then the λ will be given by

$$\frac{450}{30} = 15 \quad (43)$$

This explains the high participation rates. Even if the curve steepens by a small 30 bp, the investor can receive a coupon over 10%: the α_{t_i} plus the 450 bp.

The structurer has determined all the unknown parameters. Essentially the structurer will buy and sell two CMS swaps with different reference rates, buy floors and sell swaptions on CMS rates to manufacture a synthetic of the 5-year CMS note. Then the structurer will repackage these as a structured note and sell it to clients.

4.5.1. A Contractual Equation

This characterization is shown in the contractual equation below.

$$\begin{array}{ccccccc}
 \boxed{\text{Callable CMS-spread note}} & = & \boxed{\text{6-month Libor deposit}} & + & \boxed{\text{Receiver swap rate}} & + & \boxed{\text{Receive } \lambda \times 10\text{-year CMS}} \\
 & & & & & & \\
 + \boxed{\text{Pay } \lambda \times 2\text{-year CMS}} & + & \boxed{\text{Short Bermudan swaption}} & + & \boxed{\text{Long CMS spread floor}} & & \\
 & & & & & & (44)
 \end{array}$$

As in the previous case, the synthetic structure can open several possibilities for the structurer. The structurer can buy swaption volatility, sell cap/floor volatility at advantageous rates from the retail client, and market them at better rates in the interbank market or to other clients. Again, as before, the structurer may in fact have some of the components of the synthetic on his books and the CMS spread note would be a good way of passing them along to other customers and removing them from the balance sheet.

Or, the structurer can take the exposure itself. The structurer can buy/sell all the components in the right-hand side of the contractual equation except the swaption. Note that, then, if the expectation turns out to be correct, the synthetic structure will have a positive value. But, the note is callable (see Figures 17-12 and 17-13). The structurer will call the note and close the position on the hedge side with a good profit. It is partly for this reason that when the callable notes are likely to pay very high coupons they are in fact called. The investors basically receive the high initial coupon.

4.6. Some Other Structures

A special case of structured fixed income products is called *Target Redemption Note* (TARN). This security provides a sum of coupons until a *target level* is reached. The note then terminates

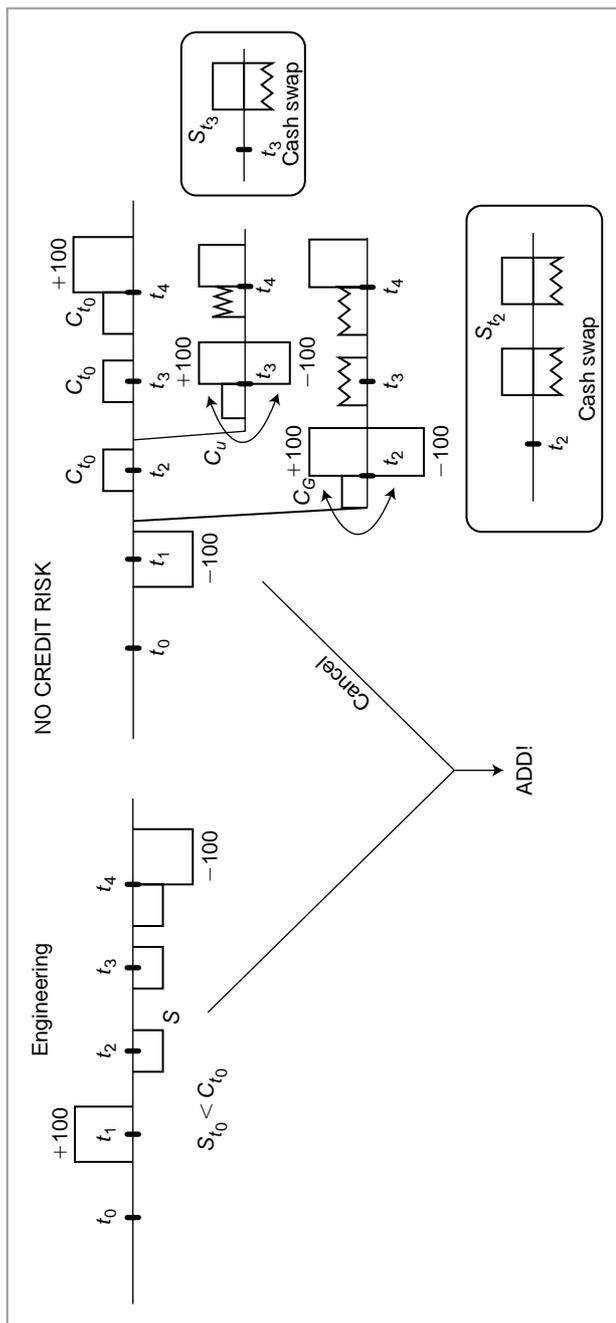


FIGURE 17-12

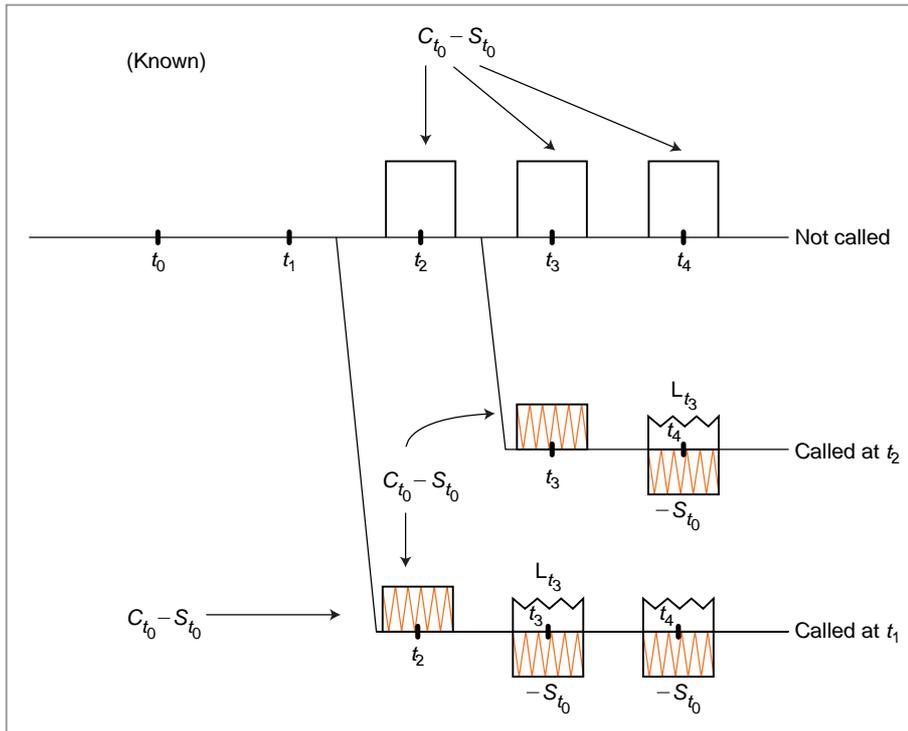


FIGURE 17-13

early. TARNs may be quite popular with investors when interest rates are low, or more correctly when they are heading lower. The additional risk in TARNs is the uncertainty of termination date. Although this is like a callable note, there is a difference. The termination condition is *explicitly stated* in a TARN and can easily be priced using the Libor market model. On the other hand, callable products contain embedded Bermudan swaptions. It is much more difficult to determine when the option will be exercised (i.e., called.) Some investors may prefer the more transparent way the TARNs are redeemed early. Like the others, the instrument is path dependent.

Another example is an *inverse floater*. As typical in such structured products, the first coupon is fixed. The subsequent coupons are set so as to depend on Libor inversely. When Libor rates decline, the coupon automatically increases. Often such coupons accumulate. If the accumulated coupons reach the target level, the note will be redeemed early. The client is paid the par value.

We can also give examples of structured products that are useful for asset-liability management. Consider a *trigger swap* that fixes borrowing costs for T years at a level lower than the current comparable market rates. In this sense, it is an asset-liability management tool. We discuss a simple variant. It is easy to complicate this simple prototype. Products such as trigger swaps belong in the category of fixed income structured products although they are not marketed to investors. The clients are corporations and the product is useful in managing assets and liabilities. Still, the main idea is the same. The corporation has a *view* on the yield curve movements, or simply desires to lower hedging costs, which is the equivalent of yield enhancement.

5. Conclusions

We close this chapter with a comment on modeling structured products. Which model to use and how to calibrate the chosen models is clearly a crucial component of structured product

trading. The structurer cannot buy these products in the wholesale market. The products need to be manufactured in-house. This requires extensive pricing and hedging efforts that will often depend on the model one is working with.

In equity structured products, versions of the stochastic volatility model are found to be quite effective and are widely used. For fixed income versions of the forward Libor model that incorporate some volatility, smile needs to be used.

Suggested Reading

Wystub (2007) and Geman (2005) are two excellent sources that the reader may consult on the FX and commodity sectors. For Libor exotics consider Piterbag (2004). For the new equity structured products Gatheral (2006) is required reading. Overhaus (2005) and Quesette (2007) are two interesting articles that summarize the technical background.

Exercises

1. Consider the swap and Libor curves available in Reuters or Bloomberg.

- (a) Obtain the 3-month discount and forward curves
 (b) Obtain the 2-year forward curve
 (c) Find the components for the following note:
 maturity: 3 years
 callable : each coupon payment date
 payments: annual
 coupons:
 Year 1: R_1
 Year 2: $\alpha_1 \times (2\text{-year CMS}) + \text{previous coupon}$
 Year 3: $\alpha_2 \times (2\text{-year CMS}) + \text{previous coupon}$

Determine the unknowns R_1, α_1, α_2 .

- (d) Find the components for the following note:
 maturity: 3 years
 callable: each coupon payment date
 payments: annual
 coupons:
 Year 1: R_1
 Year 2: $\alpha \times [(3\text{-year CMS}) - (2\text{-year CMS})] + \beta_1$
 Year 3: $\alpha \times [(3\text{-year CMS}) - (2\text{-year CMS})] + \beta_2$

Determine the unknowns R_1, α, β_i .

- (e) In the latter case, when would $\beta_1 = \beta_2$?

2. What follows is the description of a rather complex swap structured by a bank. The structure is sold for the purpose of liability management and involves an exotic option (digital cap) and a CMS component.

At time 0 the bank and the client agree to exchange cash flows semiannually for 5 years according to the following rules:

- The bank pays semiannually, 6m-Libor on the notional amount. This is a vanilla swap.
- The client pays a coupon c_t according to the following formula:

$$c_t = c_{1t} - c_{2t} \quad (45)$$

where

$$c_{1t} = \begin{cases} \text{Libor} + 47 \text{ bp} \% & \text{if } \text{Libor} < 4.85\% \\ 5.23\% & \text{if } 4.85\% < \text{Libor} < 6.13\% \\ 2.98\% & \text{if } \text{Libor} > 6.13\% \end{cases}$$

$$c_{2t} = \{ \text{cms}(30Y) - (\text{cms}(10Y) + 198 \text{ bp}) \}$$

for the first two years. And

$$c_{2t} = 8 \{ \text{cms}(30Y) - (\text{cms}(10Y) + 198 \text{ bp}) \}$$

for the last three years.

- (a) Unbundle this Libor exotic into vanilla products the best you can.
- (b) Why would an investor demand this product? What would be his or her expectations?
3. Show how you would engineer the following *Snowball Note*.
- Issuer: ABC bank
 Notional: \$10 mio
 Tenor: 10 years; Principal: Guaranteed at maturity
 Coupon: Yr 1; Q1: 9.00%
 Q2: Previous Coupon + CMS10 4.65%
 Q3: Previous Coupon + CMS10 4.85%
 Q4: Previous Coupon + CMS10 5.25%
 Yr 2 Q1: Previous Coupon + CMS10 5.45%
 Yr 2 Q2-Q4: Previous Coupon + CMS10 5.65%
 Yr 3: Previous Coupon + CMS10 5.75%
 Yr 4-10: Previous Coupon
 Coupon subject to a minimum of 0%
 Call: Callable on each coupon
- (a) What is the view of the investor?
- (b) What are the risks?
- (c) Now forget about the call provision and calculate the coupons paid under the two following realizations of Libor rates:
 Realization 1 = 5.0, 6.0, 6.5, 7.0, 8.0, 9.0, 10.0
 Realization 2 = Libor stays at 3.5
- (d) How can you characterize these coupons using a swap? What type of swap is this?
- (e) Suppose you have 8 annual FRAs quoted to you. How can you price these coupon payments?
- (f) How do you generate the first year coupon?
- (g) Are the coupons floored at 0?
- (h) Write a contractual equation representing this instrument.
4. Show how you would engineer the following *CMS spread note*.
- Issuer: ABC
 Notional: \$10mio
 Tenor: 10 years
 Principal: Guaranteed at maturity
 Coupon:
 Yr 1: 11.50%
 Yr 2-10: $16 \times (\text{CMS30} - \text{CMS10})$, max of 30%, min 0%
 Call: Callable on each coupon date by the issuer
- (a) What is the view of the investor?
- (b) What are the risks?