TechHackers' QuantNews

February 25, 1999

Welcome to the third issue of Quant*News*.

Convexity and Effective Convexity

The convexity of a bond (or other interest-rate sensitive security) is the relative curvature of the price-yield curve, representing a second order measure of price-yield sensitivity. Combined with modified duration, convexity can be used to approximate the percentage change in price given a percentage change in yield using a second order Taylor series:

$$\frac{\Delta price}{price} \cong -d_m \cdot \Delta yield + \frac{c_x}{2}(\Delta yield)^2,$$

where d_m is the modified duration and c_x is the convexity.

Convexity is thus the second derivative of the price with respect to the yield (scaled and quoted in units of years):

$$c_x = \frac{\frac{\partial^2 price}{\partial yield^2}}{price + ai},$$

where the given *price* is a quoted price, *ai* is the accrued interest, and thus (*price* + *ai*) is the invoice price.

When defining a second derivative, one must decide what other variables are kept constant. Convexity is generally defined with respect to constant cash flows.

For many types of securities, however, the cash flows usually vary with changes in yield. Instruments with floating interest rates are the obvious examples, but not the only ones: for example, (fixed rate) mortgage prepayments often increase when interest rates decline. It is therefore not particularly interesting what the value of c_x is; one would rather see how the price changes when yields move and the corresponding cash flows change as a result. The appropriate measures of price variation are known as effective duration and effective convexity, and are defined with respect to "constant modeling", i.e. the model for determining the cash flows given the yield, rather than the cash flows themselves, are fixed.

So

$$d_m = \frac{-\left.\frac{\partial price}{\partial yield}\right|_{\text{fixed cash flows}}}{price + ai},$$

$$c_x = \frac{\left.\frac{\partial^2 price}{\partial yield^2}\right|_{\text{fixed cash flows}}}{price + ai},$$

and

$$d_{m\,\text{effective}} = \frac{-\left.\frac{\partial price}{\partial yield}\right|_{\text{variable cash flows}}}{price + ai},$$

$$c_{x\ \text{effective}} = \frac{\left.\frac{\partial^2 price}{\partial yield^2}\right|_{\text{variable cash flows}}}{price + ai}.$$

In practice, the effective convexity is usually computed via finite differences. One considers three values: the current price, the price if yields go up by 50 bp, and the price if yields go down by 50 bp. Then

$$c_{x\ \text{effective}} \approx 40000 \cdot \frac{price_{+50} + price_{-50} - 2 \cdot price}{price + ai}.$$

Often the prices are determined via an OAS model[1]. First the spread is determined given the current price and yield curve. Then prices are determined for parallel shifts of the yield curve (e.g. up and down by 50 bp), assuming that the spread stays constant. Lastly, the formula above is applied, as well as a similar formula for effective duration.

Capped Call Complexity

A capped call is a call option modified to have a maximum payoff: in addition to the strike price X, there is a maximum asset price X_{max} such that the maximum value of the option is $X_{\text{max}} - X$.

As usual, a European capped call is only exercisable at maturity, while an American capped call can be exercised at any time before maturity.

If we consider assets such as stocks which do not pay dividends, the distinction between American and European (vanilla) calls is often ignored: as is commonly known, there is no economic advantage to early exercise for such options, so that the two types of options have (essentially) the same price. For capped calls, however, the distinction is extremely significant.

First consider the European capped call. The payoff for this option is

$$\max[\min(S_{\text{maturity}}, X_{\text{max}}) - X, 0]$$

where S_{maturity} is the value of the asset at maturity. A bit of algebra shows that the above equation is the same as

$$\max[S_{\text{maturity}} - X, 0] - \max[S_{\text{maturity}} - X_{\text{max}}, 0],$$

which is the difference between the value of a European call struck at X and one struck at X_{max}.

So what is the value of an American capped call? For the case of assets such as stocks which do not pay dividends, our exercise strategy is simple: when the asset price hits X_{max}, exercise. The value of this strategy is the same as the value of an up-and-out barrier option (with strike X and barrier X_{max}) which pays a rebate of $X_{\text{max}} - X$ if the barrier is hit.

Consider an example: 1 year to maturity, asset price 100, strike 105, cap at 120, volatility 20% and a risk free rate of 5%. Using ***@nalyst*** and ***exotics @nalyst*** we find that the European capped call has the value 4.774, while the American option has the value 6.547. This difference of 1.773 is not just the time value of receiving our payment earlier!

The important difference between the two types of options is that the American option gives us down-side protection when we exercise early. This is also true in the case of vanilla American calls, but there the down-side protection is more than compensated for by the additional up-side potential of continuing to hold the option. For the capped call, once we have reached the cap value, there is no longer any up-side, so the down-side protection wins!

[1] See Quant*News* issue 2

CBOT and Others Reduce Base Coupons on Bond Futures

The CBOT has announced that the base coupon on their Treasury bond and note contracts will be reduced from 8% to 6% effective March 2000. The MATIF has reduced several coupons to 3.5% and other exchanges have reduced their base coupons as well.

MATIF's stated purpose in reducing the base coupon is to focus "on the longer end of target segments." How does a coupon reduction achieve this goal?

The answer is that a coupon reduction *does not* generally achieve this goal, but will do so in the current market where the longer bonds generally have lower coupons. If this is not so, the reverse can occur. For example: consider a 30-year 6.75% bond and a 28-year 5.825% bond, for delivery into both old and new style CBOT[2] futures contracts, with hypothetical futures prices:

	30-yr bond	28-yr bond
8% CF	0.8586	0.7584
8% future price	148	
8% deliv price	127.073	112.243
yield @8% deliv	4.999%	5.007%
6% CF	1.1038	0.9764
6% future price	115	
6% deliv price	126.937	112.286
yield @6% deliv	5.006%	5.005%

The conversion factor CF is essentially the price of the specified bond when the yield is equal to the base coupon value. The delivery price is the product of the futures price and the conversion factor (assuming accrued interest is zero); also computed are the bond yields assuming the delivery price.

All other things being equal (always an issue, but anyway), the bond which has the lowest yield at the delivery price will be cheapest to deliver. We see here that the reduction of the coupon changes the cheapest to deliver bond from the 6.75% 30-year to the 5.825% 28-year.

Thus, if we assume that the true goal is to encourage delivery of longer-maturity bonds, it would seem that the best way to achieve this would be to change the delivery parameters so that shorter bonds are disallowed. The coupon change may also lead to (some of) the same effect, but certainly not uniformly. Why then the emphasis on coupon changes? They have a greater psychological impact, "re-aligning" the futures contract with the cash market.

[2] We use the CBOT contract rather than the MATIF contract for illustrative purposes based on the irrational preferences of your editor.

Are You Y6K Compliant?

The year 2000 problem has gotten a lot of press lately. The short description of the problem is that treating years as two digits, rather than four, causes problems if you need to deal with a span of more than 100 years. This is, however, only the most common cause; other bugs besides the obvious can cause "date errors".

Consider an algorithm for converting a date serial number to a month, day and year. Such an algorithm, combined with its inverse, allows one to build a "date arithmetic" library. The well-known compendium *Numerical Recipes in C* provides such an algorithm—but don't try using it with Visual C++ with optimization enabled, if you care about dates in the 60th century.

Admittedly, few of us care about events 4000 years in the future. And hopefully Microsoft will have fixed this bug before 3999 years have elapsed—but you never know…

Prepayments in Subsets of Loan Cohorts

Do certain groups of loans prepay faster than others? That question is one of many answered by the *Loan History System*, a new software product offered by TechHackers.

Assume that we have a model for the mean Single Monthly Mortalities: $\langle \mathrm{SMM} \rangle (y^i; n)$, where $\{y^i\}$ are the model parameters and n is the month since origination. The data is fit for $\{y^i\}$ by assuming that the real system is subject to random fluctuations about this mean value, so that at any month the real value will be: $\langle \mathrm{SMM} \rangle + \delta_{\mathrm{SMM}}$, where δ_{SMM} is randomly generated from a probability distribution centered at zero. Assuming that the distributions for the different SMMs are mutually independent, we can compute the cumulative average pre-payment rate over some range of months $[n_1, n_2]$:

$$\langle \mathrm{CPR} \rangle_{n_1,n_2} = 1 - \prod_{m=n_1}^{n_2} \left(1 - \langle \mathrm{SMM} \rangle (y^i; m)\right),$$

with fluctuations determined from a distribution of width $\sigma_{\mathrm{CPR}(n_1,n_2)}$, which can be computed from the distributions of the SMMs.

We can compute the expected ratio of the cumulative pre-payment rate of one cohort of loans $\langle \mathrm{CPR_{co1}} \rangle_{n_1,n_2}$ with another cohort $\langle \mathrm{CPR_{co2}} \rangle_{n_1,n_2}$ over the same time interval:

$$E(n_1, n_2) = \langle \mathrm{CPR_{co1}} \rangle_{n_1,n_2} / \langle \mathrm{CPR_{co2}} \rangle_{n_1,n_2}.$$

The distribution of the fluctuations has width

$$\sigma_{E(n_1,n_2)} = E(n_1, n_2) \times \sqrt{\left(\frac{\sigma_{\mathrm{CPR_{co1}}(n_1,n_2)}}{\langle \mathrm{CPR_{co1}} \rangle_{n_1,n_2}} \right)^2 + \left(\frac{\sigma_{\mathrm{CPR_{co2}}(n_1,n_2)}}{\langle \mathrm{CPR_{co2}} \rangle_{n_1,n_2}} \right)^2}.$$

If $|E(n_1, n_2) - 1|$ is large compared to $\sigma_{E(n_1,n_2)}$ the enhancement is statistically significant; if it is small then the result is most likely due to random sampling.

TechHackers Moves to Support Linux

In case you haven't heard, Linux is a robust, pre-emptive multitasking operating system that runs on several types of hardware, including Intel-based PCs. It's also freely distributable under the Gnu Public License. Compared to Unix on proprietary hardware it has a significant price advantage (the hardware is cheaper and the OS is free) while it provides better performance, reliability, customizability, and scalability than its commercial competitors, such as Windows NT.

TechHackers already uses Linux internally. We will soon be moving web, mail, and file servers over to Linux. Our QuantTools analytical libraries simply need recompiling for Linux, and porting the ***@nalyst*** spreadsheet add-ins to Applix on Linux can be done easily as well. "Linux gives us better performance at lower cost—it's much easier to support than NT—and I'm sure our customers will begin to use it in much greater numbers as they realize its benefits" says Bill Tonkin, TechHackers' Managing Director for Software Products, a Linux user since 1996.

You can find out more about Linux at www.linux.org.

In Conclusion…

We hope you've enjoyed issue 3 of Quant*News*. Comments, submissions, and other requests should be sent to your editor at sjanowsky@thi.com.