SOCIETY OF ACTUARIES Investment Section

Risks & Rewards

ISSUE 58 AUGUST 2011

Actuaries
Risk is Opportunity.®

Modeling And Hedging Dynamic Lapse In Equity-Linked Insurance: A Basic Framework

By Christian-Marc Panneton and Mathieu Boudreault

Variable annuities (VA), segregated funds and participating policies are classes of equity-linked life-insurance contracts with benefits that are tied to the return of an index, such as the S&P500, or an actively managed portfolio. The most important risk underlying these policies is the payment of a guaranteed minimum benefit upon death or survival. This is, in essence, a long-term put option. Most attention has been devoted to the pricing of this embedded put option, but other elements in the contract also deserve attention. One of these is the possibility to lapse or surrender the contract before maturity.

The policyholder may lapse an equity-linked insurance (ELI) contract for several reasons but one may classify these into two categories: dynamic and non-dynamic lapse. Non-dynamic surrenders are driven by reasons that are not related to the moneyness of the embedded guarantee. For example, one policyholder may lack liquidity and choose to give up the contract to recover its value. Because this is easily diversifiable in a portfolio, non-dynamic lapse can be treated as mortality in the sense that it can be priced and reserved using multiple decrement tables for example. Dynamic lapse, on the contrary, will be driven by the value of the guarantee, i.e., the policyholder may choose to surrender a contract at the moment when the remaining premiums are too costly in relation to the value of the guarantee. It is easy to deduce that lapse will go up when the guarantee gets out-of-the-money, but by how much? Thus, pricing and reserving for dynamic lapse can be difficult.

The purpose of this paper is thus to illustrate, using a simple approach, that in general, dynamic lapse can be priced and hedged using financial engineering arguments. We will also discuss more mathematically rigorous approaches to dynamic lapse toward the end of the paper.

CONTINUED ON PAGE 4

Risks & Rewards

ISSUE 58 AUGUST 2011

Published by the Investment Section of the Society of Actuaries

Joseph Koltisko, Newsletter Editor
(Chief Editor of this issue)
New York Life Insurance Co
51 Madison Avenue
rm 1113
New York, NY 10010
ph: 212.576.5625 f: 212-447-4251

Nino A. Boezio, Newsletter Editor
(Chief Editor of next issue)
TD Bank Financial Group
161 Bay St. 3rd Floor
Toronto, ON M5J 2T2
ph: 416.308.3036 f: 416.983.9701

CHAIRPERSON'S CORNER

Lately I have been wondering if we will ever get back to a period of relative boredom in the investment markets. I can recall years at a time of predictable positive returns and manageable volatility. At the time it seemed like an easy way to make a living, but that changed. One could pick several dates that might be the beginning of the financial crisis, but it has been over three years since Bear Stearns fell and the seas got rougher after that. Today, as I look at the headlines, there are still significant unknowns in the macro-economic environment that make a strong grasp of investments as important as ever to our clients, employers and constituents.

The Investment Section council frequently debates how we can best provide value to our members. As we have done this during the course of the year, there has been an abundance of ideas. This letter will provide an update on some of our work. While our projects change, a constant has been an enthusiasm for investments which helps define and promote our common goals.

Looking at the range of topics at the Investment Symposium that was recently held in New York, I can't help but be impressed by the number of significant challenges faced by our employers and clients caused by economic uncertainty and turmoil in the investment markets. I am equally impressed, though, with the great effort that put together a program on so many important issues and the quality of the speakers. Thanks to the organizing committee who did such a great job. The council is especially grateful to Bogdan Ianev who has been our chair for the past two years. I enjoyed the meeting not only for the content, but also the community of actuaries who came to participate.

Another accomplishment from earlier this year was publishing, in collaboration with the ERM Section, a second set of essays called "Risk Management: Part Two—Systematic Risk, Financial Reform, and Moving Forward from the Financial Crisis." This contained many thought-provoking and original essays.

Looking to the future we have been working on many activities that will be of value. A partial listing of them includes:

- We have sponsored sessions and arranged speakers for the SOA Annual Meeting, Valuation Actuary Symposium, and Life and Annuity Symposium.
- In collaboration with the International Section, we have agreed to send a speaker to China to speak at a variable annuity seminar in August.
- We agreed to provide financial support to the Actuarial Research Conference to be held in August, and a section council member is on the planning committee for the conference. We are providing volunteer support for another SOA research effort on liability-driven investing.

Edwin Martin

- The Redington Prize committee has been formed and we will select a prize winner this year, with the prize awarded at the SOA Annual Meeting.

- We have a webinar in the works that will provide another way to get your content from us. We continue to look for webinar topics and have a volunteer group organized around continuing education activities, if you have ideas or time to contribute please let us know.

As I look back on my tenure on the Investment Section Council, I have been lucky to be involved with a group that has displayed such good effort and ability across the board. I would like to thank the other Investment Section Council members and the many other volunteers who make the section activities a reality. Volunteers are the fuel that powers the investment section machinery. I would remind everyone reading this that there are many ways to join the party—from something as simple as writing an article or posting on our LinkedIn group, to helping to plan a meeting or webinar—every bit counts. On a personal note, it has been a great experience. I hope that a few of you will decide that it is something you want to do.

One thing I haven't mentioned—Risks and Rewards! The newsletter continues to be a major outlet for communication with Investment Section members, highlighting our ideas, successes and areas of major interest. Enjoy another issue full of timely and interesting articles.

> "… UNDER RISK NEUTRAL VALUATION, NO ADJUSTMENT MUST BE MADE TO THE LAPSE PROBABILITY OR TO THE MONEYNESS RATIO."

MODELING FRAMEWORK

Assume that in an ELI, a policyholder has the right to lapse its contract at a single pre-determined moment before the maturity of the policy T. Moreover, suppose that if the policyholder does surrender, it will do so when the guarantee reaches a given level of moneyness ζ, which is the ratio of the value of the underlying asset over the guaranteed amount. For the sake of the presentation, we ignore mortality risk to focus on the dynamics of the investment guarantee and the lapse risk, and we suppose no dividends are paid on the asset. Letting S_t be the value of the underlying asset at time t and X the strike price (the guaranteed amount), then the cost of the guarantee, including the dynamic lapse opportunity, is simply written as:

- $X - S_T$, if there is no prior lapse, i.e., $\left(\frac{S_\tau}{X} \le \zeta\right)$, and $X \ge S_T$, i.e., the guarantee matures in-the-money;
- 0, if lapse occurs at τ i.e., $\left(\frac{S_\tau}{X} \le \zeta\right)$, or the guarantee finishes out-of-the money, i.e., $X < S_T$.

PRICING

This payoff is a simple case of a reset put option, which has been tackled by Gray and Whaley (1999). Thus, assuming that the underlying asset (and a risk-free asset) is tradable and Black-Scholes' assumptions hold, one can invoke no-arbitrage arguments and use a dynamic replicating portfolio or risk-neutral valuation to price such a contract. Following Gray and Whaley (1999), one obtains:

$$P_t = e^{-r(T-t)} X N_2\left(-a_2, -b_2, \sqrt{\frac{\tau - t}{T - t}}\right) - S_0 N_2\left(-a_1, -b_1, \sqrt{\frac{\tau - t}{T - t}}\right), \qquad t < \tau,$$

where

$$a_1 = \frac{\ln\frac{S_t}{\zeta X} + (r + \tfrac{1}{2}\sigma^2)(\tau - t)}{\sigma\sqrt{\tau - t}}, \qquad a_2 = a_1 - \sigma\sqrt{\tau - t}$$

$$b_1 = \frac{\ln\frac{S_t}{X} + (r + \tfrac{1}{2}\sigma^2)(T - t)}{\sigma\sqrt{T - t}}, \qquad b_2 = b_1 - \sigma\sqrt{T - t}$$

and r is the (continuously compounded) risk-free rate. Moreover, $N_2(x_1, x_2, \rho)$ is the joint cumulative distribution function (c.d.f.) of a standard bivariate normal distribution[1] with correlation ρ. When $\tau < t < T$ and no lapse occurred, the price of this option is given by the Black-Scholes formula for a plain vanilla European put option.

In summary, both the investment guarantee and the lapse features of an ELI can be priced using no-arbitrage arguments. The practical implications are that under risk-neutral valuation, no adjustment must be made to the lapse probability or to ζ in order to obtain the correct price of the option. This may appear counterintuitive to some readers because under the risk-neutral probability measure, the stock grows at the risk-free rate which would imply a lower lapse probability. To correct this, some readers might be tempted to adjust the moneyness ratio ζ to match intuition. However, this is the wrong way to proceed because the risk-neutral lapse probability is only valid for pricing purposes. Absence of arbitrage pricing principles state that one may replicate the payoff of the contract (lapse and guarantee) with stock and bonds, or equivalently, use risk-neutral expectations (expected return of r and risk-neutral lapse probabilities). Using other pricing assumptions may create arbitrage opportunities[2].

HEDGING

In the previous lines, we have derived the no-arbitrage price of an investment guarantee combined with a lapse possibility at τ. We also need to dynamically hedge these benefits. Given that the hedging portfolio is continuously rebalanced, one needs to hold delta shares of the stock at each instant. Taking the partial derivative of the price with respect to S_t (which is delta), one gets:

$$\frac{\partial P_t}{\partial S_t} = -N_2(-a_1, -b_1, \rho) - \frac{e^{-r(T-t)}X}{\sqrt{2\pi}S_t}\left(\frac{e^{-½a_2^2}}{\sigma\sqrt{\tau - t}}N\left(\frac{a_2\rho - b_2}{\sqrt{1-\rho^2}}\right) + \frac{e^{-½b_2^2}}{\sigma\sqrt{T-t}}N\left(\frac{b_2\rho - a_2}{\sqrt{1-\rho^2}}\right)\right) + \frac{1}{\sqrt{2\pi}}\left(\frac{e^{-½a_1^2}}{\sigma\sqrt{\tau - t}}N\left(\frac{a_1\rho - b_1}{\sqrt{1-\rho^2}}\right) + \frac{e^{-½b_1^2}}{\sigma\sqrt{T-t}}N\left(\frac{b_1\rho - a_1}{\sqrt{1-\rho^2}}\right)\right), \quad t < \tau$$

where $\rho = \sqrt{\frac{\tau - t}{T-t}}$ and N is the c.d.f. of a standard normal distribution (the proof is available upon request to the authors). Once again, after τ, if no lapse occurred, the guarantee is hedged like a European put option. Since rebalancing cannot occur continuously, one may need to use other Greeks to improve the hedge of the guarantee and lapse options. This can be done by taking appropriate partial derivatives of the price function and this will not be illustrated here.

In summary, both the guarantee and the lapse features of an ELI can be hedged using a dynamic strategy that is exactly equivalent to delta-hedging more basic financial options.

ILLUSTRATIONS

Dynamics of a basic contract

In this section, we illustrate the dynamics of a basic ELI contract that features both an investment guarantee and a lapse option. We demonstrate how the value of the contract changes with respect to many variables like the time to maturity, the value of the underlying asset, etc. We also show the efficiency of the replicating strategy.

Assume that a 10-year contract is issued, with a lapse occurring at time five when the moneyness ratio exceeds 120 percent. No mortality will be considered. The risk-free rate is set at 4 percent and the volatility of the underlying asset is 25 percent. For an initial investment of \$100 and a guarantee fixed at \$100, the (no-arbitrage) price of the guarantee with a single lapse possibility is \$11.09. Without the lapse feature of the contract, the investment guarantee is an at-the-money put option, which would be priced at \$12.19.

Figure 1 shows a random path of the stock price and the resulting option value. In this scenario, the stock price is below 120 percent (the exercise price at duration 5), so no lapse occurs. However, the guarantee finishes in-the-money.

Figure 1 Simulated path of the underlying asset along with the price of the guarantee with the lapse feature

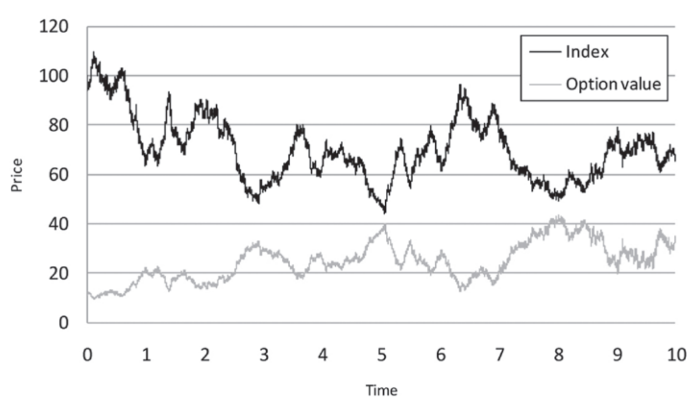

CONTINUED ON PAGE 6

> IN SUMMARY, BOTH THE GUARANTEE AND THE LAPSE FEATURES OF AN EQUITY-LINKED INSURANCE CONTRACT CAN BE HEDGED USING A DYNAMIC STRATEGY THAT IS EXACTLY EQUIVALENT TO DELTA-HEDGING MORE BASIC FINANCIAL OPTIONS.

According to theory, the replicating portfolio requires continuous rebalancing to be exact, which is, in practice, impossible for various reasons. We will illustrate the effectiveness of a delta-hedged portfolio using 10 trades per day for every trading day (250 trading days every year). Figure 2 shows that under these conditions, the tracking error of the hedge portfolio is very small. Note that Figure 2 has a greater emphasis on negative values because replication errors are mostly negative under this scenario.

Figure 2 Tracking error between the replicating portfolio and the true option value

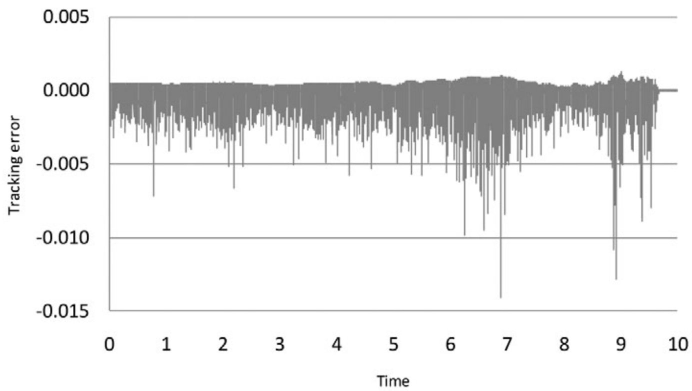

Figure 3 Changes in contract value with respect to the time to maturity, for various contract features

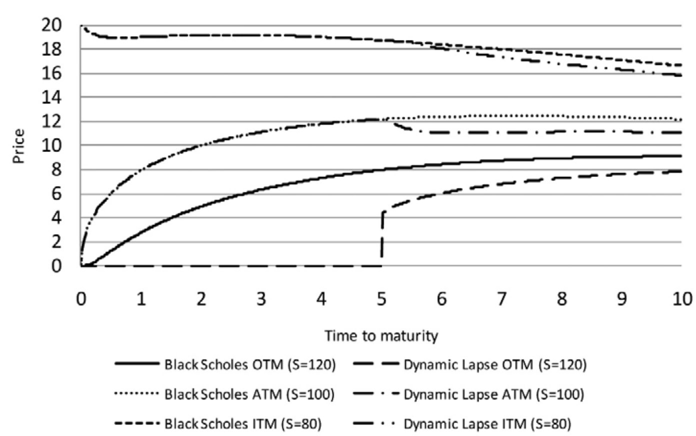

In Figure 3, we show how the contract value changes with time to maturity. This has been done for various moneyness ratios, and for contracts that feature the lapse option or not. We observe that the price of the contract with the lapse option is lower than the corresponding Black-Scholes option right after the lapse opportunity moment at five years. From the five-year point, the contract is either a put option if no lapse occurred or has a value of zero if lapse happened.

DYNAMICS OF A MORE COMPLEX CONTRACT

The ideas developed earlier can be used in a more realistic context. Rather than allowing for only one lapse opportunity, we will assume three lapse opportunities. To handle such options, one must condition the result on the value of the underlying asset at each lapse opportunity.

To illustrate the dynamics of such a contract, we assume the previous policy has a maturity of 12 years and the policyholder may lapse at years three, six or nine. For dynamic lapse, we suppose 15 percent of the contracts will lapse at the aforementioned periods when the moneyness ratio exceeds 120 percent. Non-dynamic lapses occur at an annual rate of 4 percent. Finally, a reset feature is added at duration 5, that is, if the stock price is greater than the guaranteed value at time 5, the new guaranteed amount becomes S_5.

The contract is more complex and so is the price formula. Because it is cumbersome, it will not be presented here. Its complexity comes from the fact that 16 possible scenarios have to be taken into account.[3] The approach taken to derive it is similar, i.e., we need to condition on whether lapse occurs or not at each pre-specified period, and whether a reset has been applied or not. The pricing formula has been validated with 100 million simulations, and the two are precise to the third decimal, which is about \$10.54.

Figure 4 (page 7) compares the value of the aforementioned ELI contract (put option) with and without various features. They are described as follows:

- 1st contract (with non-dynamic lapses): a 12-year plain vanilla European put option with 4 percent annual non-dynamic lapses;
- 2nd contract (with both types of lapses): same as previous, but lapse opportunities are added at times three, six and nine. Note that 15 percent of the contract holders will use this feature if the moneyness ratio is greater than 120 percent;
- 3rd contract (with a single reset and non-dynamic lapses): a 12-year European put option with 4 percent annual non-optimal lapses and a reset at five years;
- 4th contract (with a single reset and both types of lapses): all features are included.

Comparing the second contract with the first, dynamic lapse has been added. We observe in Figure 4 that lapses reduce the contract value, and this is amplified by the fact that the payoff is zero upon a lapse. More importantly, the decrease is more important when the price of the underlying asset is greater and this is because it has a greater (risk-neutral) lapse probability.

When we compare the first contract with the third contract, only a reset feature is added. The effect of the reset is generally to increase the value of the option because of a possible strike increase. A smile is observed because for small asset values, the option is similar to a put, meaning its payoff decreases when S increases. However, when the underlying asset is high enough, the likelihood of reset increases more rapidly, which explains the smile.

Finally, when we compare the fourth contract and the third, we can isolate the effects of dynamic lapse when a reset feature is present. Once again, the effect of dynamic lapse is to reduce the value of the contract, and the effect is more dramatic towards higher values of the stock price, to reflect a greater likelihood of lapse. We implicitly assume the policyholder applies its lapse decision at time three without regard to the reset that may kick in at year five, which can be sub-optimal but this is consistent with actuarial practice.

Figure 4 Comparison of the contract value under numerous lapse opportunities and contract features

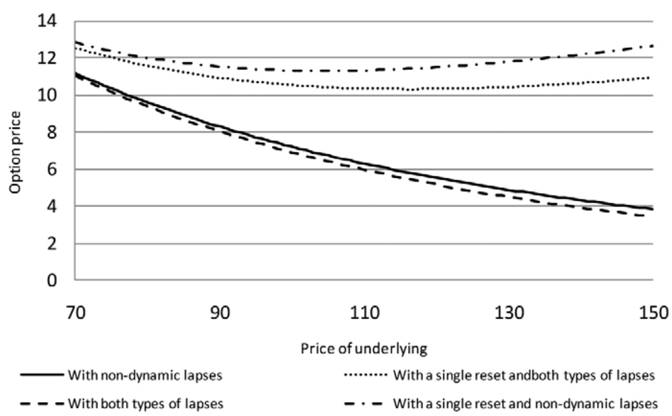

FURTHER READING

Dynamic lapse generally stems from the fact that premiums are paid as a penalty on the credited return. When the guarantee is out-of-the-money, the policyholder may decide that its guarantee is not worth enough for the premiums he or she has to pay in the future. The policyholder faces an important decision: when is the optimal moment to lapse an ELI policy? Or, in other words, what is the optimal level of moneyness for which a policyholder will lapse the policy? This is mainly an American-style type of payoff, and these questions should be answered with mathematical models that determine the optimal exercise moment. To the best of our knowledge, in actuarial science, at least two researchers have devoted significant attention to this issue: Moshe Milevsky and Anna Rita Bacinello.

In the absence of market frictions, Milevsky et al. (2001, 2002), argue that policyholders would lapse their contract and enter a new one every time the underlying asset attains a new maximum. The effect is similar to continually resetting the guarantee. The mechanism that will prevent such abuses is

CONTINUED ON **PAGE 8**

deferred surrender charges. Indeed, they have a similar purpose as transaction costs in financial markets to prevent arbitrage. The authors thus treat premiums, surrender charges and the optimal lapse decision of the policyholder as a whole problem. They obtain closed-form solutions for the optimal moneyness level and the range of deferred surrender charges that imply the existence of this policy.

The approach taken by Bacinello (2003a,b, 2005, 2008) is also similar, i.e., lapses or surrenders are treated as the optimal exercise of an American-style option. It is applied to participating policies using multinomial trees (Bacinello (2003a,b, 2005)) and simulation (i.e., least squares Monte Carlo, see Bacinello (2008)). As one of the applications, she has analyzed the effects of various variables on the premiums like the level of participation, the risk-free rate, the mean return of the stock, etc.

Compared with the actuarial literature, our approach differs in mainly two ways. First, lapses may only occur at predetermined times, which is the standard assumption in the industry, whereas in the previous papers, the policyholder may lapse at any moment. Second, the optimal decision problem is not treated in our approach, as we use a deterministic lapse criterion. This is also the standard industry practice since data on the optimality of policyholder behavior is not available. Thus, our lapses are dynamic because they depend on the level of moneyness, but this is not necessarily an *optimal* moneyness level, in the sense of American-style options. Optimality is the most conservative assumption for an insurance company since the policyholder will maximize its policy value.

We conclude this section with some remarks. In finance, surrenders or lapses are somewhat similar to abandonment options in corporate finance (see real options in Hull (2011) for example). Finally, multiple resets of strike prices within call and put options are widely discussed in the finance literature (see for example Liao and Wang (2003) and references therein).

CONCLUSION

Using a fairly simple framework, we have written the payoff of an equity-linked life insurance policy that has a single dynamic lapse opportunity as a special case of the reset put option. Because the underlying asset is traded, the cost of the guarantee can be seen as a derivative issued on the asset, which can be replicated with stocks and bonds. Consequently, replicating portfolios and risk-neutral valuation are equivalent pricing approaches that will preserve the absence of arbitrage in such a market.

We have also illustrated the effects of multiple lapse opportunities and a reset option in a more complex contract to illustrate its impact on the price of the policy. Because the latter is still mathematically tractable, it can be used as a control variate to improve the quality of simulations of more complex policies (see variance reduction techniques in Glasserman (2003) for example).

We reviewed the academic literature and compared our framework to more technical approaches. The message from our paper and the literature is clear: no matter how and when the policyholder lapses, it should always be possible to replicate the cash flows from lapses and surrenders. This is especially true when policyholder behavior is modeled (say, through a moneyness ratio) and the underlying asset is traded. Thus, replicating portfolios or risk-neutral valuation should be used equivalently. Using real-world lapse probabilities will introduce arbitrage, unless the appropriate discount factor is used.

REFERENCES

- Bacinello, A.R. (2003a), "Fair Valuation of a Guaranteed Life Insurance Participating Contract Embedding a Surrender Option", *Journal of Risk and Insurance* **70**, 461-487.

- Bacinello, A.R. (2003b), "Pricing Guaranteed Life Insurance Participating Policies with Annual Premiums and Surrender Option", *North American Actuarial Journal* **7**, 1-17.

- Bacinello, A.R. (2005), "Endogenous Model of Surrender Conditions in Equity-Linked Life Insurance", *Insurance: Mathematics and Economics* **37**, 270-296.

- Bacinello, A.R. (2008) "A Full Monte Carlo Approach to the Valuation of the Surrender Option Embedded in Life Insurance Contracts", Chapter in the book by Perna & Sibillo (2008), "Mathematical and statistical methods in insurance and finance", Springer.

- Drezner, Z. (1978), "Computation of the Bivariate Normal Integral", *Mathematics of Computation* **32**, 277–279.

- Glasserman, P. (2003), "Monte Carlo Methods in Financial Engineering", Springer.

- Gray, S.F. and R.E. Whaley (1999), "Reset Put Options: Valuation, Risk Characteristics and an Application", *Australian Journal of Management* **24**, 1-20.

- Hull, J.C. (2011), "Options, Futures and Other Derivatives", 8th edition, Prentice-Hall.

- Liao, S.L. and C.W. Wang (2003), "The Valuation of Reset Options with Multiple Strike Resets and Reset Dates", *Journal of Futures Markets* **23**, 87-107.

- Milevsky, M.A. and T.S. Salisbury (2001) "The Real Option to Lapse a Variable Annuity: Can Surrender Charges Complete the Market", Working Paper, York University.

- Milevsky, M.A. and T.S. Salisbury (2002) "The Real Option to Lapse and the Valuation of Death-Protected Investments", Working Paper, York University.

CONTINUED ON **PAGE 10**

SUMMARY OF DERIVATIONS

To derive the formula for the price of this option, one needs to condition on whether or not the contract is lapsed at τ, i.e.,

$$P_t = e^{-r(T-t)} E\left[X - S_T \mid \frac{S_\tau}{X} \leq \zeta, X \geq S_T\right] Pr\left[\frac{S_\tau}{X} \leq \zeta, X \geq S_T\right], \qquad t < \tau$$

with expectations and probabilities taken under the risk-neutral probability measure. Rearranging the terms, we have:

$$P_t = e^{-r(T-t)} X Pr[S_\tau \leq \zeta X, X \geq S_T] - e^{-r(T-t)} E[S_T \mid S_\tau \leq \zeta X, X \geq S_T] Pr[S_\tau \leq \zeta X, X \geq S_T].$$

However, the expectations and probabilities simplify to:

$$e^{-r(T-t)} X Pr[S_\tau \leq \zeta X, X \geq S_T] = e^{-r(T-t)} X N_2\left(-a_2, -b_2, \sqrt{\tau/T}\right)$$

$$e^{-r(T-t)} E[S_T \mid S_\tau \leq \zeta X, X \geq S_T] Pr[S_\tau \leq \zeta X, X \geq S_T] = S_t N_2\left(-a_1, -b_1, \sqrt{\tau/T}\right)$$

which yield the desired formula.

Finally, the derivations of the stock portion of the replicating portfolio can be obtained by request to the authors as it is tedious and cumbersome, but not very complicated. However, one needs to remember that the conditional distribution of *X* given *Y*, when (*X*,*Y*) follows a bivariate normal distribution, is normally distributed.

END NOTES

1. The function N_2 is not available in Excel but can be easily coded. John Hull's Technical Note #5 provides a straightforward algorithm to compute the latter function, which is based upon Drezner (1978). The website is http://www.rotman.utoronto.ca/~hull/TechnicalNotes/TechnicalNote5.pdf. Alan Genz also provides papers on the numerical computation of the cumulative distribution function of a bivariate normal distribution.
2. This does not prevent the company from using experience data in determining the policyholder behavior. For example, in a slight modification of the model, one may define a moneyness ratio that changes deterministically with time (accumulation/discounting, for example). In this case, this would be similar to changing the strike price in an option and once again, absence of arbitrage arguments would be used to price this contract. Under stochastic interest rates, the problem gets more complex because of the presence of a second source of risk. Fortunately, the same principles hold, i.e., it would be possible to trade available assets to completely hedge the payoff of the contract. However, these modifications clearly go beyond the main purpose of this paper.
3. In this case, the computation of the c.d.f. of a four-dimensional multivariate normal distribution is necessary. Alan Genz's Web page has numerous algorithms to compute the c.d.f. of multivariate normal distributions. R, MATLAB and SAS also have packages that perform such computations.

Christian-Marc Panneton, M.Sc., FSA, FCIA, CFA, CERA, is a financial risk manager at Industrial Alliance, Québec, Québec. He can be reached at christian-marc.panneton@inalco.com

Mathieu Boudreault, Ph.D., FSA, is an assistant professor in the Department of Mathematics at the Université du Québec à Montréal. He can be reached at boudreault.mathieu@uqam.ca.

ARBITRAGE AND STOCK OPTION PRICING: A FRESH LOOK AT THE BINOMIAL MODEL

By Dick Joss

The avoidance of an arbitrage potential has played a very big role in the theory of stock option pricing. In fact, the presumed ability of speculators to move option prices because of an arbitrage opportunity was the key conclusion reached in the 1973 classic paper on option pricing written by Dr. Myron Scholes and Dr. Fischer Black. But the famous Black-Scholes option pricing formula has not matched actual option quotations well in recent years, and perhaps it is time to revisit their well-known anti-arbitrage argument.

THE BINOMIAL MODEL

The Black-Scholes anti-arbitrage argument is based on the concept that a speculator would be able to take a hedged equity position in the underlying security that would match the payoff under the option contract. If the cost of the hedged equity position did not match the option price, the speculator would have the opportunity to make free money. It is common to illustrate this issue in finance textbooks by using a simple binomial model. For example, one illustration shows that a stock currently sells for $100, that it might either go up to a value of $200 or down to a value of $50, and that these are the only possibilities. The illustration then proceeds to calculate the call option price using only a risk-free rate of return and the strike level.

Continuing the illustration, assuming a risk-free rate of 8 percent and a strike level of $125, the price of a one-year call option must be $26.85. It is easily demonstrated that at this option price the investor winds up in exactly the same position whether he buys an option or takes a hedged position in the security. If the investor buys the option and the stock increases in value to $200, the investor will have $75; whereas if the stock decreases in value to $50, the investor will have $0. In lieu of buying the option with his $26.85, the investor could have borrowed (at 8 percent interest) an additional $23.15, thus buying a 0.5 shares of stock worth $50. If the stock goes up in value, the investor's stock will be worth $100, and after paying off his loan, the investor is left with $75. If the half share of stock goes down in value to $25, the investor sells the stock to pay off his loan to be left with $0. These are the same potential payoffs that would have occurred had the investor owned the call option.

Furthermore, if the option price in the above illustration were anything other than $26.85, it is possible to create a scenario where the investor enters the market with no money, but is able to leave the market with a net gain. For example, if the option were priced at $30, the investor could write two options producing $60 in cash, borrow an additional $40 at 8 percent interest, and use the $100 to purchase one share of stock. If the stock value decreases to $50, the investor sells the share and repays the loan ($43.20) to net a profit of $6.80 with no cash outlay. If the stock value increases to $200, the investor uses $150 to pay off on the two options and $43.20 to repay the loan, again yielding a net profit of $6.80. In other words, no matter how the stock performs, the investor gets a guaranteed return. This illustration certainly dramatizes the important role that this type of arbitrage avoidance might play in option pricing.

However, there is a significant problem with the above illustration, that being that it is completely dependent upon a known fact about future market behavior. For example, it was the known fact that the only possible future values for the $100 share of stock were a gain to $200 or a loss to $50, which enabled the investor to identify the mispriced $30 option and realize his $6.80 gain.

But the future of the market is not known. And now consider what happens to our investor in the simplified market if the stock actually rises to $300, not $200. The investor who formerly thought he was guaranteed a free $6.80 is now faced with the prospect of paying off $175 on each of the two options, plus a loan repayment of $43.20 for a total liability of $393.20. His only asset is the one share of stock worth $300. Suddenly his $6.80 guaranteed return has turned into a $93.20 loss. The investor would have also been disappointed had the stock dropped in value to $30 instead of $50. In this situation the investor loses $13.20 on his loan repayment. Clearly the arbitrage potential outlined above has significant risks.

TWO DIFFERENT HEDGING STRATEGIES

What happened was that the investor assumed the volatility characteristic of the stock could be described by the $200/$50 scenario, when the investor should have assumed that it was described by the $300/$30 scenario. In this second scenario, the price of the call option which allows the hedged position to exactly match the option payoff is $46.81. The investor thought the $30 option was overpriced, when it was in fact under priced. This led the investor into adopting the wrong hedging strategy, thus locking in his losses instead of his gains.

For option contracts that appear to be overpriced, the correct hedging strategy is to write the options, borrow from the bank, and buy shares of the underlying security. For option contracts that appear to be under priced, it is the exact opposite strategy that is the correct hedging approach. The investor needs to buy the option contract, sell the underlying security short, and invest the difference at the risk-free rate. Picking the wrong strategy will lock in losses, not gains. Furthermore, it is impossible to know in advance which strategy will be correct, as it is impossible to know in advance the volatility of the stock for the duration of the given option contract.

A VERY TELLING ADDITIONAL ILLUSTRATION

Before leaving the simple binomial model illustration, it is interesting to look at one more, and somewhat unusual, modification. In this case, assume that the two possible changes for the stock are a rise to $200 or a rise to $110. Given that the strike price is still $125, there are the same potential returns for the option investor as in the first illustration: either a gain to $75 or a loss to $0. In this situation it is still possible to determine a hedged equity arbitrage avoiding price of the stock option. In this case the price of the option then becomes -$1.55. While it might seem unusual to have a negative option price, this is the unique price which allows for the perfectly hedged position.

The hedged position consists of writing a contract for -$1.55 borrowing $84.89, and using the difference to purchase .8334 of a share of stock. If the stock goes up to $200 the investment will be worth $166.68, which after payment of the loan leaves exactly $75.00. If the stock goes up to $110, the .8334 of a share will be worth $91.67. This is the amount necessary to pay off the $84.89 loan, leaving a wealth of $0. Furthermore, if the price of the option were anything other than -$1.55, the hedged equity arbitrage potential in this case exists exactly as it did above. Clearly, just using the hedged equity replicating value as the option price cannot possibly be right. The correct price of the option contract in this illustration cannot be a negative number.

THE BIG BINOMIAL PROBLEM

Clearly, the problem with this illustration is that all possible returns on the stock exceed the risk-free rate, meaning that the arbitrage has nothing to do with the option price, and is based entirely on the supposed arbitrage inherent in the underlying security. But this is always the case with the binomial model. To see this, begin with a monomial model, where there is only one known fact about the market (the risk-free rate) and one assumed possible return for the stock. Clearly, with this model an investor would always know which hedging strategy to take in order to make an arbitrage gain. If the assumed return exceeded the risk-free rate, the investor would know to borrow at the risk-free rate and purchase shares of stock. If the assumed return were less than the risk-free rate, the investor would know to sell shares of the stock short and invest at the risk-free rate.

But this model hardly constitutes an arbitrage opportunity. One merely needs to broaden the assumed return to two possibilities, one bigger than the risk-free rate, and one less than the risk-free rate, and the arbitrage potential disappears. The investor does not know which of the two opposite hedging strategies to use.

CONTINUED ON **PAGE 14**

Now consider the case where two facts about the market are known (the risk-free rate and the details of a specific stock option contract). In this situation if only two possible rates of return are assumed, then it is again always possible to identify the correct hedging strategy. This is true no matter how outrageously priced the known option contract is—even if it is -$1.55. This is a situation that is no different than the monomial model shown above. All of the supposed arbitrage potential derives from any difference between the assumed rates of return for the underlying stock and the risk-free rate of return.

But the problem is that having only two possible rates of return for an assumption is far too limiting. By expanding the list to three or four possible returns, the investor is again lost as to which of two exactly opposite hedging strategies to use. Combining the first two binomial illustrations into one "binomial volatility" model where the stock price may either follow the $200/$50 scenario or the $300/$30 scenario clearly shows the impossible choice facing a speculator who hopes to make an arbitrage gain. The speculator does not know which of the two exactly opposite hedging strategies to adopt.

IMPACT ON OPTION PRICING

Modern mathematical presentations of the Black-Scholes formula go to great lengths using sophisticated mathematical techniques such as stochastic calculus, Ito's Lemma, and Girsanov's Theorem to demonstrate that the Black-Scholes value for an option is the unique value for an option which permits the option's potential payoff to be replicated by a hedged equity position using shares of the underlying security. This is not unlike the elementary illustration shown above where the unique value of $26.85 yielded a hedged equity position which exactly matched possible option payoffs.

However, to conclude that this unique Black-Scholes hedged equity replicating value is the option price requires an additional argument. This second argument notes that if the price of the option exceeded the Black-Scholes value, there exists a hedging strategy which would yield the investor a guaranteed return, and that if the price of the option were less than the Black-Scholes value a different hedging strategy could be used to also provide the investor with a guaranteed return.

But for the investor to realize his gain from a mispriced option, the investor must actually create the right hedged position, and change it continually assuming that there are no transaction costs. Clearly, this type of "arbitrage" contains a significant element of risk. The investor who adopts the wrong strategy winds up locking-in losses, not gains.

Hedged equity arbitrage requires an investor:

a. To make certain assumptions about the future market behavior in order to decide whether a particular option is overpriced or under priced so that the correct hedging strategy can be adopted.

b. To continually monitor the hedged position and make appropriate changes at no transaction cost.

c. To be subject to the risk that if the changes aren't made continually or if market behavior is different than what was assumed, the supposed guaranteed gain could be turned into a loss.

All of these points lead to the question: Is hedged equity arbitrage really arbitrage? The practical realities associated with the hedged equity approach are too significant to overcome. This last point is important because real arbitrage opportunities, such as put-call parity, only deal with the difference in put and call prices, not the actual prices. This means that additional real market forces, such as the expected return on the underlying security or the investor's option premium (additional risk assumed by investing in option contracts), can come into play in determining market-driven option prices.

> "IT IS TIME TO REVISIT THE ENTIRE THEORY OF RISK-NEUTRAL OPTION PRICING."

SUMMARY

The Black-Scholes formula and the general theory of risk-neutral option pricing have been the dominant focus of academic option pricing for the past 30 years. However, these theories rely on the ability of a speculator to choose between two exactly opposite hedging strategies with such sufficient accuracy that it constitutes an arbitrage potential. Clearly, when one considers the practical realities of the market place, this supposed arbitrage potential does not exist.

It is mathematically impossible to provide a guaranteed hedged equity arbitrage return using only a single option contract, either a put option or a call option, over a single finite time period. This is because the investment process to realize such a gain requires the investor to decide whether the particular option is either overpriced or under priced, and apply the appropriate hedging strategy. But it is impossible to tell whether any given option was either overpriced or under priced until the actual market plays out.

If actual market volatility is higher than that which was assumed by the investor seeking the hedged equity arbitrage gain, then the option might turn out to have been under priced not overpriced. If actual market volatility is less than that which was assumed by the investor seeking the hedged equity arbitrage gain, then again the investor could suffer a loss instead of a gain. Even if perfect assumptions were selected, actual market behavior could deviate from those assumptions for the duration of the option contract, thus producing results which are the exact opposite of what was anticipated. No matter what hedging strategy is adopted by the investor, there is always the possibility of market results which yield a loss.

In short, hedged equity trading against a single option contract is not an arbitrage opportunity. It is time to revisit the entire theory of risk-neutral option pricing. Perhaps new formulas that reflect additional market forces, such as the expected rate of return on the underlying security, will serve as springboards for a whole new theory of option pricing. Given the recent turmoil in the financial services industry, the market place seems to be clamoring for new ideas.

Dick Joss, FSA, is retired. He can be contacted at rrjoss@comcast.net.

NON-TRADITIONAL ACTUARY: JAMES GANNON

By Risk & Rewards

Traditionally, actuaries have tended to follow one of two career paths. Either they worked for insurance companies helping to design insurance and annuity products that met the needs of policy holders, or they worked for consulting firms which helped companies design and administer their own employee benefit plans.

But increasingly, actuaries are finding some "non-traditional" sources of employment. To help shine a spotlight on some of these careers, R&R plans to include a featured interview with an actuary working in a non-traditional area in each of the next several issues. The interview for this issue is with James Gannon who works with the Russell Investment Group in Seattle, Wash. The Russell Investment Group is a Washington, USA corporation, which operates through subsidiaries worldwide, including Russell Investments, and is a subsidiary of The Northwestern Mutual Life Insurance Company.

James was recently elected to a three-year term on the Investment Section Council. The interview was conducted in January, 2011.

1) HOW DID YOU FIRST LEARN ABOUT THE ACTUARIAL PROFESSION?

When I was a sophomore, a college professor and academic advisor suggested that I take the first actuarial exam (Exam 100: Calculus and Linear Methods) and even offered to pay my exam fee if I would give him a ride to the exam. I passed the exam and then started to do some research on the field. I also later found that I had a cousin who was an actuary, so I was able to periodically call him to ask questions.

2) WHERE AND WHEN DID YOU GO TO COLLEGE; WHAT DEGREE OR DEGREES DID YOU EARN?

I graduated from Lehigh University in 1996 with a B.S. in Statistics. After college, I was ready to enter the workforce and really never considered graduate school because I was tired of studying and classrooms. So, of course, I entered a profession that had me studying and taking exams for the next decade after graduation.

3) WHAT WAS YOUR FIRST JOB OUT OF COLLEGE?

I started working with the benefits consulting firm Kwasha Lipton in Fort Lee, N.J. I was hired as an actuarial analyst and mostly prepared actuarial liability valuations for pension plans. I remained there for three and a half years.

4) WERE THERE ANY OTHER JOBS BETWEEN YOUR FIRST JOB AND YOUR CURRENT POSITION?

After Kwasha Lipton, I took a job at Towers Perrin in New York City. Again I was involved in the actuarial liability for pension plans, but was able to branch out into other projects within the actuarial pension field, including non-discrimination testing, early retirement windows, budgeting, etc.

5) WHAT ATTRACTED YOU TO YOUR CURRENT POSITION IN THE INVESTMENT MANAGEMENT FIELD, AND WHEN DID YOU START WORKING WITH YOUR CURRENT EMPLOYER?

I was contacted by a colleague at Russell Investments who was looking to fill a position in their Consulting group. I was attracted by the opportunity to learn an entirely new field (investment consulting) where I would still be able to use my knowledge of pension plans as a backdrop. I started working at Russell in January 2007.

6) WHAT ARE SOME OF THE MORE FREQUENT TASKS WITH YOUR CURRENT POSITION?

My current responsibilities are two-fold. First, I have management responsibilities, as I oversee the analyst staff for the Investment Consulting practice at Russell, and am responsible for determining their client assignments and assisting with career development. Second, my practice responsibility includes managing the process by which Russell provides and delivers asset allocation (or asset liability) studies for both our consulting and investment management clients. This second responsibility means that I have to maintain our processes according to the current regulatory environment for clients (Pension Protection Act) and the best practices and emerging trends in investment opportunities.

7) WHAT ARE SOME OF THE TASKS ASSOCIATED WITH YOUR CURRENT POSITION THAT YOU FIND THE MOST ENJOYABLE?

Listening to clients discuss their current issues. My team meets with many Russell clients during the year. These meetings allow us to better serve all of our clients by identifying and addressing broad industry trends. For instance, this process created the idea for a paper I wrote with my colleague Bob Collie in April 2009 on a form of dynamic asset allocation called Liability-Responsive Asset Allocation (LRAA), which is a process by which pension plans reduce the risk in their asset allocation as funded status increases. This concept has been very influential and well-received in the marketplace.

8) ARE THERE SOME TASKS ASSOCIATED WITH YOUR CURRENT POSITION THAT YOU FIND LESS APPEALING, AND IF SO, WHAT ARE THEY?

I travel frequently throughout North America visiting clients. I think that the face-to-face client meeting still has many advantages over a phone call or video conference. But sometimes I can do without the airport delays, odd connections, rental car mix-ups, and unhealthy eating that comes with traveling. Also, traveling, even with today's technology, can still cause one to be a little less productive.

9) ARE THERE ANY CONCERNS THAT YOU SEE WORKING AS AN ACTUARY FOR A FIRM THAT HAS NOT BEEN CONSIDERED TO BE A TYPICAL ACTUARIAL EMPLOYER?

The biggest concern I had was how I would stay connected to the actuarial profession. I believe it to be a great profession which can bring valuable insights and solutions to many different fields, and the Society of Actuaries (SOA) is an organization that is very supportive of its membership. It was easy to maintain this contact when I was in an office with 40 other actuaries, but now I must be a little more proactive in maintaining my connections. I usually do this by attending conferences (Enrolled Actuaries Meeting or Conference of Consulting Actuaries) or SOA-sponsored webcasts, through participation in the SOA's Investment Section Council, and by visiting the SOA website.

10) WHY DID YOUR FIRM CONSIDER ADDING ACTUARIES TO ITS STAFF?

A large segment of Russell Investments' clients are defined benefit pension plans, and the leaders of Russell's consulting practice recognized nearly 30 years ago that actuaries have a deep knowledge of the workings of pension plans. They have employed a core group of actuaries ever since. This allows us to assist our clients and work with their other service providers in helping them with their investment issues, especially on topics like liability-driven investing and reducing their risk in their pension plan asset allocation.

11) ARE THERE SOME CREATIVE SOLUTIONS TO PROBLEMS THAT YOU HAVE HANDLED IN YOUR CURRENT POSITION THAT YOU CAN SHARE? IF SO, WHAT ARE THEY?

I think the most creative solution was introducing Liability-Responsive Asset Allocation to our clients. The most common way to invest for pension plans was to set a static asset allocation and then review it every three to five years with an asset liability study. The Liability-Response Asset Allocation process we developed allows plans to take advantage of market movements and to invest according to the current situation and health of their pension plans. Pension boards can operate within a more dynamic invest-

CONTINUED ON **PAGE 18**

> "GIVEN THE PROPER TRAINING IN FUNDAMENTAL SKILLS SUCH AS **PROBLEM-SOLVING, ANALYTICAL ANALYSIS AND ATTENTION TO DETAIL,** ACTUARIES WILL EXCEL ACROSS A RANGE OF INDUSTRIES … "

ment framework, which allows them to consistently look for opportunities to align their current investment allocation with their current corporate situation and goals as well as the current market environment.

12) DO YOU HAVE ANY "WORDS OF WISDOM" THAT YOU MIGHT OFFER TO ACTUARIES WHO MIGHT BE CONSIDERING A CAREER OUTSIDE THE TRADITIONAL INSURANCE COMPANY OR CONSULTING FIRM?

I would encourage them to look for opportunities outside of traditional fields if they are interested. I have long thought that actuaries are a natural fit for many corporate and consulting positions beyond those in the areas of insurance and pensions. Given the proper training in fundamental skills such as problem-solving, analytical analysis and attention to detail, actuaries will excel across a range of industries involving risk management, credit analysis and problem-solving. The firms adding actuaries to their staff—or in leadership roles—will benefit as well.

SOLVENCY II, FDIC AND RATINGS-FREE CAPITAL RATIOS

By Cicero I. Limberea

This paper analyzes the current Solvency II credit spread stress test framework and proposes a ratings-free capital reserving framework as required by the FDIC.

Solvency II is used by some U.S. insurers as a guideline for the stress tests, economic capital and VaR calculations. For insurers that do underwriting in Europe, the framework is mandatory. Solvency II provides a distinct framework for credit spreads, apart from default risk which is calculated separately and summed to the credit spread exposure to come out to the overall credit risk band.

In practice, summing the exposures generates very high credit VaR exposures.

The credit spreads exposure part of market risk is defined in the Solvency II framework[1] as

$$Mkt_{sp}^{struct} = \sum_i MV_i \bullet n(dur_i) \bullet G(rating_i)$$

where

$G(rating_i)$ = a function of the rating class of the credit risk exposure which is calibrated to deliver a shock consistent with VaR of 99.5 percent. The shocks are ratings-dependent and differ across financial instruments (i.e., the bonds brackets are different from the MBS/ABS brackets)

$n(dur_i)$ = a function of the duration of the credit exposure

Since the first order Taylor approximation[2] of the price-rate function approximates $\Delta P/P=-D\Delta y$, the framework assumes that the credit spread moves cointegrate within interest rates moves.

However, this view was challenged at the height of the credit crisis as the Federal Reserve's interest-rate cuts did not immediately lower borrowing costs for many companies. In the first part of 2008, investors demanded even more yield to buy investment-grade U.S. corporate bonds. According to Merrill[3] data, high-risk, high-yield securities rose a quarter-point over Treasury in the first part of 2008. "The increase in credit spreads has sort of worked against our policy. The fact that the spreads went up so dramatically really resulted in an effective tightening of financial conditions that our cuts were partly meant to address,'' San Francisco Fed President Janet Yellen told *Bloomberg News* in March, 2008.

This increase in credit spreads reversed in 2009 as the spreads tightened, but it is unclear whether it was the Fed policy that became effective only after a lag, or the strong quarterly earnings releases lowered corporate default expectations, or a combination of the two.

The lag view is shared by Morris, Neal and Rolph[4] in "Credit spreads and interest rates: a cointegration approach," a research study published by the Kansas City Fed in 1998. Using earlier empirical evidence, the authors statistically prove that the short-term cointegration of credit spreads is different from the long term and that in the short term the rate tightening has a credit spread widening effect, consistent with data during the 2007–2009 credit crisis. In addition, they proved inconsistencies in the models of Leland and Toft,[5] Longstaff and Schwartz,[6] and Merton.[7]

CONTINUED ON **PAGE 20**

This lag view would make the duration approach in Solvency II relevant only for the long term. However, the loss of market value measured in the above formula is prescribed with the available spurious short-term data, which would have displayed wild swings during the credit crisis and thus has little ability to detect a long-run positive relation between spreads and rates.

In recognition of these data measurement difficulties, and thus being forced to adopt a Jarrow, Lando and Turnbull[8] approach with lack of correlation between credit spreads and interest rates (which emphasizes the default factor), we model the credit-spread stress-test free of any interest rate influence, as follows:

PV of CDS spread = PV of expected default loss
Expected Default Loss = LGD*Probability of Default

Where LGD stands for expected loss given default and equates to Protection Notional *(1-Estimated Recovery Rate)

we find that a 99.5 percent severity VaR default approach would include the clean, interest-free credit spread moves effect.

Such a default test is provided by CEIOPS in their Consultation Paper # 28[9] as:

$$SCR = \min(\Sigma LGD_i;\ q \times V)$$

where

LGD_i = Loss-given-default for type 1 exposure of counterparty i

q = Quantile factor
V = Variance of the loss distribution of the type 1 exposures

The sum is taken over all independent counterparties in the same ratings bands to make this consistent with Duffee and Singleton[10] (1996) who model individual credit classes separately as they assume that there is a flight to quality in times of stress. Thus the variance of the loss distribution differences between rating bands can become high.

Assuming a lognormal normalized bond prices distribution, at the 99.5 percent confidence level, q is 3.[11] In the CEIOPS formula the Estimated Recovery Rates are approximated via a floor calibrated as three times the sum of variances per each ratings band.

Such an approach would suffice as an overall credit risk test assuming sufficient counterparty diversification which can be achieved via a separate correlation test. For a diversified corporate bond and structured security portfolio such an approach produced realistic results, without the function of the duration of the credit exposure which is an interest related test and seems unnecessary for credit spreads.

Currently, the U.S. Federal Deposit Insurance Corporation (FDIC) seeks alternatives to credit ratings in bank capital guidelines,[12] based on its observations from the recent credit crunch years, and has asked the public for comments.

In a ratings-free environment, since the loss distributions for instruments with high default levels vs. low levels can become high in times of stress, taking three times the overall variance of the loss distribution per the CEIOPS consultative paper, can be quite high.

Thus, the multiple is likely, in this case, to be instrument specific and default-level specific, a higher multiple for higher default levels and a lower multiple for lower default levels.

The variance of loss exposures could possibly be replaced with the variance of market prices, since the VIX index is a liquid measure of uncertainty and expresses traders' opin-

> "TAKING THREE TIMES THE OVERALL VARIANCE OF THE LOSS DISTRIBUTION PER THE CEIOPS CONSULTATIVE PAPER, CAN BE QUITE HIGH."

ions of estimated future default levels, as it reflects the daily variance of the path of option prices.

However, as seen in the recent credit crisis, illiquidity can demand a large discount, which can be unrelated to historical default levels.

PV (Estimated future bond flows) + Defaulted past flows*Recovery rate = FMV = Price + μ
μ is the discount

In such cases, linking the pessimism or out-of-moneyness of the markets through μ to the capital levels is unrealistic, as such the capital levels should be a function of the FMV where the recovery rates are calibrated exactly based on CDS prices. μ however could be minimized in days with high trading volumes, so the ratings-free capital levels are likely to look at volumes at least as a validating factor.

The case of structured securities with the respective credit support from inferior tranches is a special case. In order to calculate the pure default level per tranche, the subordination support has to be backed out of the price:

PV (Estimated future bond flows) + (Defaulted past flows-Defaulted past flows paid by subordinate tranches) *Recovery rate = FMV

Where Defaulted past flows paid by subordinate tranches is calculated and known by the trustee.

Defaulted past flows paid by subordinate tranches has to be added again to the tranches providing subordination as it reflects the credit quality of these tranches, notwithstanding that the price of such tranches will reflect the expected subordination to the senior tranches and as such it will not be used in the capital calculation.

The capital reserving methodology for these tranches should be based off an expected default level calculation rather than an observable price approach.

BIBLIOGRAPHY:

Bloomberg Financial News database 2006-2010

Committee of European Insurance and Occupational Pensions Supervisors - Consultative Paper # 28 .

Duffie, Darrell and Kenneth J. Singleton (2003). Credit Risk: Pricing, Measurement, and Management. Princeton University Press.

Leland, Hayne E. and Klaus Bjerre Toft, "*Optimal Capital Structure, Endogenous Bankruptcy, and the Term Structure of Credit Spreads,*" Journal of Finance, Vol. 51, No. 3, (July 1996), pp. 987-1019.

Robert Alan Jarrow, David Lando, Stuart Turnbull, A Markov Model for the Term Structure of Credit Risk Spreads, *Review of Financial Studies,* Vol. 10, No. 2, (Summer 1997), pp. 481-523.

F. A. Longstaff "Interest Rate Volatility and the Term Structure: A Two-Factor General Equilibrium Model" (with Eduardo S. Schwartz), *The Journal of Finance* 47, 1259-1282, 1992.

Robert C. Merton: Theory of rational option pricing (1971).

Neal, Robert, Rolph, Douglas S. and Morris, Charles, "Interest Rates and Credit Spread Dynamics" (December 2000).

Taylor, Brook, "Methodus Incrementorum Directa et Inversa" [Direct and Reverse Methods of Incrementation] (London, 1715), pages 21-23 (Proposition VII, Theorem 3, Corollary 2).

END NOTES

[1] http://www.ceiops.eu/index.php?option=content&task=view&id=576, pp11/29.

[2] Taylor, Brook, Methodus Incrementorum Directa et Inversa [Direct and Reverse Methods of Incrementation] (London, 1715), pages 21-23 (Proposition VII, Theorem 3, Corollary 2).

[3] http://www.bloomberg.com/apps/news?pid=newsarchive&sid=aUBDfKosxdlg

[4] Neal, Robert, Rolph, Douglas S. and Morris, Charles, Interest Rates and Credit Spread Dynamics (December 2000).

[5] Leland, Hayne E. and Klaus Bjerre Toft, "*Optimal Capital Structure, Endogenous Bankruptcy, and the Term Structure of Credit Spreads*", Journal of Finance, Vol. 51, No. 3, (July 1996), pp. 987-1019.

[6] F. A. Longstaff "Interest Rate Volatility and the Term Structure: A Two-Factor General Equilibrium Model" (with Eduardo S. Schwartz), *The Journal of Finance* 47, 1259-1282, 1992.

[7] Robert C. Merton: *Theory of rational option pricing* (1971).

[8] http://www.ceiops.eu/index.php?option=content&task=view&id=576

[9] Robert Alan Jarrow, David Lando, Stuart Turnbull, A Markov Model for the Term Structure of Credit Risk Spreads, *Review of Financial Studies*, Vol. 10, No. 2, (Summer 1997), pp. 481-523

[10] Duffie, Darrell and Kenneth J. Singleton (2003). Credit Risk: *Pricing, Measurement, and Management*. Princeton University Press

[11] http://www.ceiops.eu/index.php?option=content&task=view&id=576, pp11/29

[12] http://www.bloomberg.com/news/2010-08-10/fdic-seeks-comment-on-alternatives-to-credit-ratings-in-capital-guidelines.html

Cicero I. Limberea, CPA is a director at PwC Washington, D.C. He has an MIA in Economic Policy Management from Columbia University in New York.

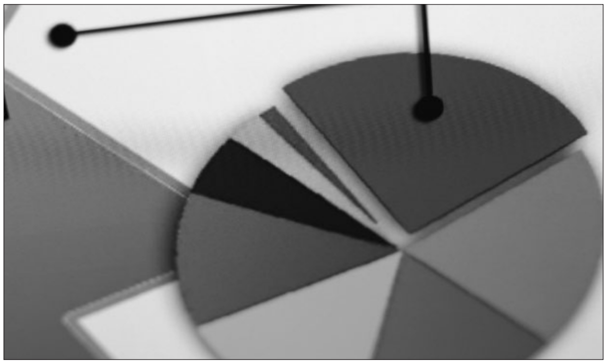

EFFECTIVE RATES OF INTEREST: BEWARE OF INCONSISTENCIES

By Stevens A. Carey

Current introductory textbooks on the mathematics of finance inconsistently define an "effective rate of interest." Many, if not most, of these definitions are general and cover not only compound interest but also simple interest. These general definitions yield the same results only for interest accumulation functions, such as continuous compound interest at a constant rate, which satisfy any one of three equivalent regularity conditions known as Markov accumulation, the consistency principle and the transitivity of the corresponding time value relation. Simple interest, as it is generally known, does not satisfy these conditions. Outside the context of compound interest, care must be taken to ensure that there is a common understanding of an effective rate of interest.

MUCH ADO ABOUT NOTHING?

What is an effective rate of interest? Isn't the definition relatively standard? As the word effective suggests, the effective rate for a particular time period seems to be nothing more than the rate that reflects the actual economic effect of the interest during that period, namely:

> *the actual percentage increase by which a unit investment would grow during the applicable period*.

Doesn't that put an end to the inquiry?

The answer is no. Admittedly, the italicized definition above may be both intuitive and accurate when operating within the confines of compound interest (if appropriate assumptions are made, and the period in question is one or more consecutive compounding periods or there is continuous compounding at a constant rate). However, in the general case (which could include, for example, simple interest), this definition may be ambiguous. Moreover, the textbooks that do provide a general definition of an effective rate are not consistent. This article will explain the ambiguity and inconsistency.

Also, the italicized definition above assumes that the unit investment remains fully invested without withdrawal or additional principal investment. Yet in practice, the term effective rate is used to describe many types of investments that may involve multiple withdrawals or additional principal investments or both (e.g., the annual effective rate of a loan that involves monthly payments). This article will also discuss the assumptions (often unspoken) that underlie such use of the term effective rate.

ALTERNATIVE GENERAL DEFINITIONS IN CURRENT TEXTBOOKS

The general definitions in modern textbooks differ from the italicized definition at the beginning of this article because they limit their definitions to unit investments *made as of a certain time*. There is a good reason for these limitations: they eliminate the ambiguity in the general italicized definition. Specifically, these general definitions seem to boil down to the following alternatives (assuming no withdrawals or additional principal):

Time 0 Investment Definition
the actual percentage increase over the applicable period of a unit investment *made as of time 0.*

New Investment Definition
the actual percentage increase over the applicable period of a unit investment *made as of the commencement of the applicable period.*

Clearly, these definitions yield the same result if the period in question commences as of *time 0*. But many textbooks discuss effective rates for periods that commence after *time 0*.

TIME OF ENTRY

The only meaningful distinction between these two definitions is the *time of entry* of the unit investment upon which each definition is based: in one case, the investment is made as of *time 0* (a *time 0* investment); and in the other case, the investment is made as of the commencement of the period in question (a *new* investment). In the context of compound interest (assum-

CONTINUED ON PAGE 24

ing a single rate of interest), this may be a distinction without a difference as long as there is continuous compounding at a constant rate or the relevant period starts at the beginning of a compounding period. Under such circumstances, an investment amount (including the then amount of a *time 0* investment) at the commencement of the period may always be viewed as though it were a *new* principal investment made as of the commencement of the period (and, assuming proportionality, may be assumed to be a unit investment for purposes of determining proportionate growth). However, if, for example, interest accumulates as it would in a simple interest account (where each deposit earns simple interest at the same rate from the time it is made), then there could be a meaningful difference.

> Example. Assume the following facts: (1) there is a bank account that does not permit withdrawals until the account is closed, although deposits may be made at any time; (2) the bank account credits simple interest at the rate of 100 percent per annum for each deposit from the time it is made; and (3) the bank account will be closed three years from the date it is opened. Given these facts, what would the effective rate be for this account during the second year?

This example illustrates why the italicized statement of effective rate at the beginning of this article may not be well defined: the percentage increase in a unit investment for the second year could vary drastically, depending on when the unit investments were made. Each of the two alternative general textbook definitions avoid this ambiguity by assuming a singular time of entry. However, because they assume different times of entry, they yield different results: under the *Time 0* Investment Definition, the effective rate would be 50 percent ($1 deposited at *time 0* would grow from $2 to $3 during year two, representing a 50 percent increase); and under the *New* Investment Definition, the effective rate would be 100 percent ($1 deposited at the beginning of year two would grow from $1 to $2 during year two, representing a 100 percent increase). Indeed, if the parties contract for a constant rate of simple interest (as in the example above), then given a unit period (e.g., a year), the effective rate for sequential unit periods would remain constant under the *New* Investment Definition, but would decline under the *Time 0* Investment Definition.

GENERALIZED ACCUMULATION FUNCTIONS

Before further analyzing this inconsistency, it will be helpful to have a way to describe how an investment grows during a given period depending on its *time of entry*. Unfortunately, an accumulation function of one variable is limited to *time 0* investments and does not (by itself) indicate how an investment grows if it is made at any time other than *time 0*. To remedy this shortcoming, it is possible to define an accumulation function of two variables, $A(t_1,t_2)$, indicating the value at time t_2 of a unit investment made at time t_1. With this generalized accumulation function, the typical accumulation function, $a(t)$, is merely a special case for *time 0* investments: $a(t) = A(0,t)$.

Restatement of Definitions. Now, the effective rate definitions stated in words earlier may be restated more precisely as follows:

Effective Rate for period (t_1,t_2)	
New Investment Definition	*Time 0* Investment Definition
$A(t_1,t_2) - 1$	$[a(t_2)/a(t_1)] - 1$

Simple Interest Example. Revisiting the previous example, with reference to an accumulation function of two variables:

> Example. Assume that interest accrues at 100 percent simple annual interest rate in accordance with the following accumulation function: $A(t_1,t_2) = 1 + (t_2 - t_1)$. Under this accumulation function, which represents the common understanding of simple interest (although the rate is artificially high to make the calculations easier):
>
> - the effective rate for the *n*th year under the *New* Investment Definition would be $1 + (n - [n - 1]) - 1 = 1 = 100\%$; and

> "MARKOV ACCUMULATION: THE GROWTH FACTOR OF A UNIT INVESTMENT, OVER ANY UNIT INTERVAL IS THE SAME AS THAT OF A TIME ZERO INVESTMENT."

- the effective rate for the nth year under the *Time 0* Investment Definition would be $[(1 + n)/(1 + [n - 1])] - 1 = [(1 + n)/n] - 1 = (1 + n - n)/n = 1/n$.

It follows, in particular, that the effective rate for the second year of a simple interest unit investment would be 50 percent under the *Time 0* Investment Definition and 100 percent under the *New* Investment Definition.

How can this inconsistency between the two definitions be addressed?

WHEN TIME OF ENTRY DOESN'T MATTER

Fortunately, for certain types of well behaved interest, the *time of entry* is irrelevant (for purposes of defining the effective rate). In other words, for certain accumulation functions (assuming again proportionality and no withdrawals or additional principal), interest accumulation does not vary by reason of the *times of entry* of two equal investment amounts as of a particular date. For example, for such an accumulation function, an investment that has grown to $100 as of a particular time would thereafter grow in the same manner as a new $100 investment made at such time (assuming proportionality and no withdrawals or additional principal). This occurs when the accumulation function has the following property:

- **Markov Accumulation:** What has been called Markov accumulation (which reflects the notion that the growth factor of a unit investment, as to which there are no withdrawals or additional principal, over any interval (t_1,t_2) is the same whether it is a *new* investment, i.e., an investment made as of time t_1, or a *time 0* investment):

$$A(t_1,t_2) = a(t_2)/a(t_1)$$

It is easy to show that Markov accumulation is equivalent to each of the following (assuming in each case proportionality and no withdrawals or additional principal):

- **transitivity:** The "transitivity" property of the corresponding time value relation (which reflects the notion that if one dated cash flow is equivalent to each of two other dated cash flows, then those two other dated cash flows must also be equivalent):

$$(a,t_0) \equiv (b,t_1) \text{ and } (b,t_1) \equiv (c,t_2) \Rightarrow (a,t_0) \equiv (c,t_2)$$

In other words, for any cash flows a, b and c, and times $t_0 \le t_1 \le t_2$, $aA(t_0,t_1) = b$ and $bA(t_1,t_2) = c$ implies $aA(t_0,t_2) = c$.

- **consistency principle:** What has been called the "consistency principle" (which reflects the notion that the growth of an investment is not affected if it is withdrawn and immediately reinvested in the same investment):

$$A(t_0,t_1)A(t_1,t_2) = A(t_0,t_2)$$

In other words, if a unit investment made as of time t_0 were withdrawn as of time t_1, when it would have grown to $A(t_0,t_1)$, and then immediately reinvested until time t_2, when the reinvested amount would have grown to $A(t_0,t_1)A(t_1,t_2)$, the result would be the same as if the unit investment remained invested for the entire period, $A(t_0,t_2)$.

The equation above may be rewritten as follows:

$$A(t_1,t_2) = A(t_0,t_2)/A(t_0,t_1)$$

Viewed in this way, the consistency principle looks like a generalized version of Markov accumulation (which is the special case where $t_0 = 0$): for any interval, (t_1,t_2), the growth factor for a *new* investment, $A(t_1,t_2)$, is the same as the growth factor for an investment made as of time t_0, $A(t_0,t_2)/A(t_0,t_1)$, regardless of when it is made (i.e., it is independent of t_0).

DETERMINING THE EFFECTIVE RATE

Now that the definition of effective rate has been addressed, it may be useful to examine effective rate calculations in practice. These calculations may be another source of confusion because

CONTINUED ON PAGE 26

EFFECTIVE RATES OF INTEREST | FROM PAGE 23

"THE EFFECTIVE RATE FOR A PARTICULAR PERIOD DOES NOT INDICATE HOW THE EFFECTIVE RATES EVOLVES DURING THAT PERIOD."

the definitions discussed earlier are limited to unit investments that remain fully invested. Yet it is common for finance professionals to refer to the effective rate of many different investments (e.g., loans) that may not remain fully invested. This apparent discrepancy is explained below.

Accumulation Function. As a preliminary matter, information about the agreed upon interest is required in order to determine the effective rate. An accumulation function would be an ideal source of information, but in practice, one is rarely given an accumulation function. (Even if one knows the relevant accumulation function, it has already been shown that effective rates may be different depending on what definition is used, unless there is Markov accumulation. If there is Markov accumulation, then the effective rate for the interval (t_1,t_2) is simply $A(t_1,t_2) - 1$). When the accumulation function is not expressly stated, effective rate calculations seem largely to be based on assumptions as to what the underlying accumulation function is. Once an accumulation function is established, then it may be easy to determine the effective rate for the period in question.

Determining Effective Rate for a Period Based on Effective Rate for Another Period. Often an effective rate for a particular period is known and the effective rate for another period (usually a larger period that includes the given period or a smaller period within the given period) is to be determined. Before considering such problems, observe that the effective rate for a particular period does not indicate how the effective rate evolves during that period. Infinitely many accumulation functions may reach the same result.

> Example. Consider a semi-annual period commencing at *time 0*, and assume the effective rate for this period is 10 percent. Among the infinitely many accumulation functions for which the effective rate is 10 percent for this period (all that is required is that a(1) = 1.1) are the following continuous accumulation functions, in each case assuming a 10 percent semi-annual interest rate (and a time unit equal to ½ of a year):

(1) *(Everyday) Simple Interest*: $A_S(t_1,t_2) = 1 + .10(t_2 - t_1)$
(2) *Modified (Markov) Simple Interest:*
$A_{MS}(t_1,t_2) = [1 + .10t_2]/[1 + .10t_1]$
(3) *Continuous Compound Interest:*
$A_{MC}(t_1,t_2) = (1 + .10)^{t_2 - t_1}$

Straightforward Compounding Calculations—Increasing the Unit Period. Perhaps the most common effective rate calculation involves a nominal annual rate and a compounding period (equal to some unit fraction of a year) from which one can determine the effective rate for a year commencing at the beginning of a compounding period. By convention (in the context of compound interest), the nominal rate establishes a constant proportionate effective rate for the relevant compounding period, so this calculation is tantamount to using an effective rate for a compounding period (that is a unit fraction of a year) to establish the effective rate for the year.

> Example. 20 percent per annum, compounded semi-annually, implies an effective annual rate equal to 21 percent: $(1 + .20/2)^2 = (1.10)^2 = 1.21$.

This is an easy calculation because the accumulation function values for the times in question are implicit: $A(t_1,t_2) = (1 + .10)^{t_2 - t_1}$, where t_1 and t_2 are non-negative integral numbers of semi-annual periods and $t_1 \leq t_2$. It is easy to see that this function, although discontinuous, is Markov: $A(t_1,t_2) = (1 + .10)^{t_2 - t_1} = (1 + .10)^{t_2}/(1 + .10)^{t_1} = a(t_2)/a(t_1)$. Consequently, there is no ambiguity.

Decreasing the Unit Period. The reverse calculation is not as straightforward.

> Example. Given a 10 percent semi-annual effective rate, what is the effective rate for a quarter?

This question is problematic because the underlying accumulation function may not be known for non-integral time values, and there are many different accumulation functions that match the given results for integral time values. One is the accumulation function for continuous compound interest at a rate of 10 percent per semi-annual period, namely, $A(t_1,t_2) = (1 + .10)^{t_2 - t_1}$, for all positive non-zero values of t_1 and t_2, which is Markov. To simplify things, many textbooks assume that this accumulation function applies. But the information provided may also be consistent with accumulation functions that are not Markov, such as $A(t_1,t_2) = (1 + .10)^{/t_2 - t_1/} (1 + .10 [(t_2 - t_1) - /t_2 - t_1/])$, where $/x/$ = the largest integer which is not more than x, and t_1 and t_2 are numbers of semi-annual periods. This alternative accumulation function indicates that for any particular deposit, there is simple interest between the semi-annual compounding. Because this alternative accumulation function is not Markov, there may be confusion. For example, assuming this alternative accumulation function were to apply (and a quarter was one-quarter of a year and a semi-annual period was one-half of a year), the effective rate for the first quarter would be 5 percent, but the effective rate for the second quarter might be 5 percent or approximately 4.76 percent, depending on which of the two definitions described earlier were used.

Constant Effective Rate for Unit Period of Particular Duration. As illustrated in the prior paragraph, a constant effective rate for a unit period of a particular duration (e.g., 10 percent per semi-annual period) does not necessarily establish the underlying accumulation function or imply a constant effective rate for unit periods of all other durations (e.g., a quarter).

Cash Flows. Sometimes no rate is specified at all. There may be only a (finite) sequence of cash flows during the period in question. Does it make sense to talk about the effective rate that applies to a series of cash flows? Not without further assumptions. Unless there are only two cash flows (a payment by the investor at the beginning, and a payment to the investor at the end, of the period in question), which fit neatly within the fully invested unit investment definition (assuming proportionality), there may be the same problem: there may be more than one accumulation function that is consistent with the cash flows. As a result, there could in theory be multiple answers:

> Example. Consider an investment account in which the sole activity is an initial deposit of $100, followed by a withdrawal after the first one-half year of $10 of interest and a withdrawal at the end of the year of the then total $110 balance of outstanding principal and interest, thereby generating the following semi-annual cash flows: -100, 10, 110. What is the underlying accumulation function? Unfortunately, there is no single answer. For example, assuming a time unit equal to one-half year, both the following accumulation functions are consistent with these cash flows: $A(t_1,t_2) = 1 + (.1)(t_2 - t_1)$, which represents a 10 percent simple semi-annual interest rate; or $A(t_1,t_2) = (1.1)^{(t_2 - t_1)}$, which represents a 10 percent semi-annual rate, compounded semi-annually. But for the year in question, the effective annual rate is 20 percent for the first accumulation function and 21 percent for the second accumulation function.

With sufficient assumptions, however, it may be possible to establish a single underlying accumulation function which, in turn, may establish a unique associated effective rate. (This article will not discuss multiple IRRs; a discussion of that subject may be found in most finance textbooks.) Unfortunately, these assumptions are rarely articulated. Instead, a number of books simply state that the effective rate for a series of cash flows is the IRR. However, at least one book explains that one uses the IRR to back into an imputed interest rate that is called the effective interest rate. This imputed rate is based on assumptions that lead to a compound interest rate (so that the use of effective rate in connection with cash flows is limited and would not include, for example, simple interest).

IRR CALCULATION

Given a sequence of cash flows, the following steps are typi-

CONTINUED ON **PAGE 28**

cally taken to determine the IRR (which include assumptions that establish an accumulation function which, in turn, leads to an effective rate in the manner defined at the outset):

- The first step is to create equal consecutive cash flow periods beginning with the cash flow at *time 0* (which is zero if there is no such cash flow) and ending with the last nonzero cash flow. This is easily done by taking the greatest common divisor of the original cash flow period durations and adding zero cash flows as necessary.

- The next step is to assume that each such regular cash flow period is a compounding period and that there is a constant effective rate for each compounding period, so there is a sequence of regular cash flows, CF_0, CF_1, CF_2 ... CF_n, where CF_k is the cash flow occurring immediately after k compounding periods for $k = 0, 1, 2 ..., n$ (for some number n).

- The final step is to calculate what is sometimes called the *yield rate* or *internal rate of return* (IRR) for the compounding period by solving the following equation for i (assuming $i > -1$):

$$\sum_{k=0}^{n} CF_k (1+i)^{-k} = 0.$$

Equivalently (assuming $i > -1$), one may solve the following equation:

$$\sum_{k=0}^{n} CF_k (1+i)^{n-k} - 0.$$

This IRR calculation assumes a compound interest rate, i, thereby keeping the calculations in familiar territory. Moreover, i is assumed to be a constant interest rate per compounding period, and therefore the underlying accumulation function $A(t_1,t_2)$ is $(1 + i)^{t_2 - t_1}$ where each of t_1 and t_2 is an integer, representing a corresponding number of compounding periods, and $0 \le t_1 \le t_2 \le n$. Because this accumulation function is Markov (so that the growth of a *time 0* unit investment during any period between two compounding dates is representative of the growth during that period of any unit investment regardless of when made, assuming no withdrawals or additional principal), it is possible to express the accumulation function as a function of one variable: $a(t) = (1 + i)^t$, where t denotes an integral number of compounding periods between 0 and n. Look familiar? Of course, this is the fundamental formula of compound interest (for unit investments and an integral number of time periods).

LOAN EFFECTIVE RATES

One of the more common uses of IRRs is to determine the effective periodic interest rate in a loan transaction. As stated by one author:

> "In the context of a loan regarded as an investment, the internal rate of return is the rate of interest for which the loan amount is equal to the present value of the loan payments. In other words, the internal rate of return on a loan transaction is simply the interest rate at which the loan is made."

But sometimes the periodic effective loan rate (e.g., a monthly rate) is already known and the effective rate for the year is to be determined. Again, this might seem confusing because the definitions of effective rate in both the older and current textbooks listed above are generally based on the growth of amounts that remain *fully invested* during the applicable period (without withdrawal or additional principal) and here the loan balance may not stay fully invested for the year. This problem is typically addressed by making the same assumption that is made in the IRR calculation, namely that the payment (cash flow) periods are compounding periods: "using an effective [annual] rate of interest with a compounding frequency that matches the payment frequency." In this way, the effective rate determination is a straightforward calculation: it is based on the same underlying compound interest accumulation function used in connection with the calculation of the IRR (or effective periodic [e.g., monthly] interest rate) of a loan when given the periodic (e.g., monthly) payments. Thus, although a 12 percent fully amortizing loan with equal monthly payments would have

> IF INTEREST ACCUMULATION IS NOT MARKOV, THEN THE DEFINITION ABOVE MAY NOT MAKE SENSE WITHOUT FURTHER CONTEXT.

the same monthly payments for numerous different accumulation functions (including everyday 12 percent annual simple interest, 12 percent annual interest compounded monthly and 12 percent annual interest compounded annually), it is generally assumed, for purposes of determining the effective annual rate, that there is monthly compounding. As stated in a popular real estate finance textbook:

> "For example, a 12% loan with monthly payments actually applies a simple interest rate of 1% due at the end of each month…this implies an effective annual rate (EAR) of $(1.01)^{12} - 1 = 12.68\%$, compounding the simple monthly rate at the monthly frequency."

COMPUTATIONAL ASSUMPTIONS – COMPOUND INTEREST

To review, although the definitions of effective rate (in most introductory textbooks on the mathematics of finance reviewed by the author) assume that there are no withdrawals or additional principal investments, the term effective rate is often used in practice to describe an investment that may involve one or more withdrawals or additional principal investments or both (e.g., the effective annual interest rate of a loan with monthly payments). In practice, the determination of the effective rate seems to be based on assumptions about the underlying accumulation function. The assumptions tend to be more extensive than merely assuming Markov accumulation. To simplify computations, it is frequently also assumed that there is a constant effective rate for periods of any particular duration (i.e., for any duration h, the effective rate for a period of duration h would be the same). For a Markov accumulation function, this means $A(t_1, t_1 + h) = A(t_2, t_2 + h)$ for all h, t_1 and t_2. While the everyday simple interest accumulation function has this property, it is not Markov so it would therefore be excluded. In fact, it is easy to show that any continuous accumulation function that is Markov and yields a constant effective rate over any unit period must be of the form $a(t) = (1 + i)^t$. Thus, for practical purposes, these IRR assumptions limit the use of effective rates for cash flows to compound interest.

CONCLUSION

The effective rate for a particular time period may be defined both generally and intuitively when dealing with interest that meets a certain fundamental, regularity requirement known alternatively as Markov accumulation, the consistency principle or the property of transitivity. Under such circumstances, it is nothing more than:

> *the actual percentage increase by which a unit investment would grow during the applicable period (assuming no withdrawals or additional principal).*

If the relevant accumulation function is also continuous and the effective rate, as defined above, is constant (in the sense that for any duration h, the effective rate for all periods of duration h is the same), then interest must be compounded continuously and the accumulation function must be of the form $a(t) = (1 + i)^t$. Perhaps for this reason many older textbooks, and even some current textbooks, limit the discussion of effective rates to compound interest.

But if interest accumulation is not Markov (recognizing, for example, that simple interest, as commonly understood, is not Markov), then the italicized definition above may not make sense without further context; to avoid inconsistencies in this general setting, the definition of effective rate for a particular period may be (and sometimes is) limited to the percentage increase over that period of either *new* investments or *time 0* investments (assuming proportionality and no withdrawals or additional principal). For example, if interest accumulates with compounding at discrete times (e.g., at the end of each calendar year, quarter or month), and simple interest between those compounding times, then the effective rate for any period that does not commence as of a compounding time may vary depending on the definition used.

Whatever approach is taken, the underlying assumptions should be clear to avoid confusion.

STEVENS A. CAREY is a transactional partner with Pircher, Nichols & Meeks, a real estate law firm with offices in Los Angeles and Chicago. This article is a condensed version of an article previously published by the author: Carey, "Effective Rates of Interest", The Real Estate Finance Journal (Winter 2001) © 2010 Thompson/West. Please refer to the prior article for more detail (including acknowledgments, citations, quotes of effective rate definitions from various textbooks, and proofs).

LEGERDEMATH: ANATOMY OF A BANKING TRICK

By Omer Rosen

In my previous article[1], "Legerdemath," I made brief mention of Treasury-rate locks:

Most brazenly, we taught clients phony math that involved settling Treasury-rate locks by referencing Treasury yields rather than prices.

A number of readers expressed a doubt that using a settlement method based on Treasury prices was appropriate. What follows is as good an explanation of Treasury-rate lock settlements as 2,000 words will allow. I have simplified some of the bond math and concepts and will end with an analogy that I hope will elucidate what the math did not. However, as this post hardly qualifies as an easy read, feel free to ask questions in the comments section. Confession: I fudged the word count a few sentences ago to increase the likelihood of you reading on.

Forget for a moment, everything you have heard or think you know about Treasury bonds. Taken in isolation, the purchase of a Treasury bond is nothing more than the purchase of a fixed set of future cash flows. If you find the term cash flows confusing, think instead of the following: buy a bond today, receive predetermined amounts of money on predetermined dates in the future.

In this column I will be referencing a 10-year Treasury bond paying a coupon of 5.00 percent, with a notional amount of $100. For convenience, I will christen this bond Bondie. Sans jargon, the fixed set of cash flows received when purchasing Bondie would be $2.50 every six months for 10 years and an additional $100 at the end of the tenth year.

There are two basic ways to describe the value of this fixed set of cash flows, either by price or by yield. Price answers a simple question: How much would it cost you to purchase this fixed set of cash flows? This price will change over time, in much the same way that the price of a stock changes over time. Yield expresses the return earned by purchasing these cash flows at a certain price.

If you had to pay $100 in order to receive the fixed set of cash flows I described above, then your yield would be 5.00 percent. If you had to pay more to purchase these same cash flows, say $105, then the return you would be earning (the yield) would be lower than 5.00 percent—it would be 4.3772 percent. Intuitively this should make sense—the more you have to pay for a given set of cash flows the lower your return will be. Or, more simply, when prices go up, yields come down. Conversely, if you had to pay only $95 for these same cash flows, the yield earned would be higher than 5.00 percent—it would be 5.6617 percent.

Algebraically speaking, price and yield are linked by an equation where all the other variables are known. Therefore, if you know the yield of a given bond you can calculate the price of that bond and vice versa. In plain terms, saying you are willing to pay $100 for Bondie is the same as saying you are willing to buy Bondie at a yield of 5.00 percent (i.e., at a price that will allow you to earn a return of 5.00 percent). It is similar to how one can describe the speed of a car either by the number of miles per hour it is traveling or by the time it takes it to travel one mile—if you know one you can solve for the other, and if one goes up the other comes down.

To belabor the point, if a car is traveling around a one-mile track at an average speed of one mph, then it is easy to solve for the time needed to complete a single lap: 60 minutes. Either "one mph" or "a 60-minute mile" provides you access to the same knowledge about the speed of the car during that lap. And, if the car's speed were to increase, the time it would take to complete another lap would decrease (At two mph a mile would only take 30 minutes). The same inverse relationship holds true between prices and yields.

Now back to Treasury-rate locks. When a company puts on a Treasury-rate lock, it is doing nothing more than taking a short position in a Treasury bond. A short position is a bet that will pay off for the company if Treasury prices go down and go

CONTINUED ON **PAGE 32**

against the company if prices go up. Why would they do this? That is a subject for another column and I ask that you accept as an article of faith that sometimes this bet, rather than being a gamble, reduces risk and uncertainty for a company.

The short position can be viewed as an agreement under which the client will sell the bank Treasury bonds at a certain price on a set date in the future. This price is determined based on current market conditions. For example, let us say, that based on what current market conditions dictate, the client agrees to sell Bondie to the bank at $95 one month hence. A month passes and Bondie is now trading at $100. The client will have to go into the market, buy Bondie at the current price of $100, and then sell it at a loss of $5 to the bank at the previously agreed upon price of $95. For expediency's sake, the client just pays the bank the $5 it has lost and the bank takes care of all the buying and selling behind the scenes. The calculation of $5 in the above manner—subtraction—is an example of the price-settlement method of Treasury-rate locks.

However, when it comes to bonds, corporate clients do not think in terms of price; they think in terms of yield because yield is expressed in the language of interest rates, the same language companies are familiar with from business concepts such as rates of return and borrowing costs. In theory, this should add only a simple step to the settlement process. The company locks in a sale of Bondie at the same level as before, $95, but rather than quoting them that price, the bank quotes them the corresponding yield of 5.6617 percent. We can refer to this yield as the locked-in yield.

A month passes and the Treasury rate lock is settled. Rather than telling the client that Bondie is now trading at $100, the bank tells them that the yield is now 5.00 percent, having fallen by 0.6617 percent. But 0.6617 percent is not a dollar value that can be paid out as a settlement. To calculate the settlement, both yields, 5.6617 percent and 5.00 percent, need to first be converted back to their respective corresponding prices, $95 and $100. Taking the difference between the two prices results in the same settlement value we calculated before: $5.

But the client is never shown how to settle based on prices. Instead they are introduced to a nonsensical and more complicated method called yield settlement. The sole purpose of this settlement method is to trick the client into allowing the bank extra profit.

Whereas price settlement asks the question, "By how much did Treasury prices change?" yield settlement asks, "By how much did Treasury yields change?" As mentioned in the previous paragraph, the yield decreased by 0.6617 percent. But how does one convert 0.6617 percent into a dollar value that can be paid out?

First, a unit conversion is necessary. For clarity and convenience, finance makes use of a unit called a basis point. Each basis point is equal to 0.01 percent. Using this new unit, the above decrease of 0.6617 percent can be expressed as 66.17 basis points. Of course, this solves nothing, only modifying our most recent question slightly: now we ask, how much is each of the 66.17 basis points worth in dollar terms?

At this point the client is introduced to a concept called DV01 (Dollar Value of One Basis Point). DV01 is defined as the change in price of a bond for a one basis-point change in yield. For example, if the yield on a bond changes from 5.00 percent to 5.01 percent or from 5.00 percent to 4.99 percent, by how much would the corresponding price of that bond change? This change in price is the DV01. If yields shifted by 66.17 basis points, DV01 will answer the question of how much each of these basis points is worth.

The starting point for this calculation is the yield at the time of settlement. In our example, the yield at the time of settlement is 5.00 percent. At this yield, the corresponding price of Bondie is $100. If the yield were to rise by one basis point to 5.01 percent, the corresponding price of the bond would fall to $99.922091, a decrease of 7.7909 cents. If instead the yield were to decrease by one basis point to 4.99 percent, the corresponding price would rise to $100.077983, an increase of 7.7983 cents. By convention, the average of these two changes in bond prices is taken to be the DV01. So, at a yield of 5.00 percent, the DV01 would be 7.7946 cents per one basis-point move ((7.7983 + 7.7909) ÷ 2). If the yield changes by one basis point, price is said to move by 7.7946 cents. Or, in more plain terms, each basis point has been assigned a value of 7.7946 cents.

The DV01 is then multiplied by the difference between the current yield and the locked-in yield. In our example the difference between 5.00 percent and 5.6617 percent is 66.17 basis points. From the previous paragraph we know that each of these 66.17 basis points is worth 7.7946 cents. Multiplying 66.17 by 7.7946 we arrive at a settlement value of $5.1577. This is the yield-settlement method of Treasury-rate locks.

Apart from being confusing, the yield-settlement method has resulted in a settlement value that is greater than the $5 calculated using the price-settlement methodology. For a good-sized rate lock, say $500 million dollars worth of 10-year Treasuries, the client would pay the bank an extra $788,500 (500 million x (5.1577 – 5.00) ÷ 100) when settling using the yield-based methodology. This "extra" is profit for the bank.

I ask that you stop reading here for a moment. I have stated from the beginning that yield settlement is incorrect. However, when reading the explanation of yield settlement, did you find yourself agreeing with the logic? At what point, if any, did you spot the flaw? And can you guess what happens if prices had gone the other way? If prices had gone down instead of up, say to $90, the bank would have owed the client money. However, yield settlement would have allowed the bank to earn a profit by paying the client less than it actually owed them. No matter what happens to prices, yield settlement allows the bank to earn extra profit.

Now picture yourself as a client receiving a tutorial on Treasury-rate locks. You are being instructed by a banker on a matter that seems procedural, in a manner that seems advisory and helpful, without any warning that something might be amiss. You are led through the yield-based settlement process and taught how the DV01 is calculated. If you have access to a Bloomberg terminal you are shown where the DV01 can be found on the relevant Treasury bond's profile page. Perhaps presentation materials are sent over detailing the mechanics of rate locks and different possible outcomes depending on various possible market movements. And all this is part of a larger interaction, a relationship even, during which the banker is nothing but genuinely friendly and informative. Furthermore, there is a good chance that someone from a different part of the bank, someone who has advised you before, was the one that introduced the two of you in the first place. Would you question your banker?

Clients, among them some of the largest corporations in the world, never did. Confident in the tools provided them and blinded by specious logic, the client never even thinks to question the underlying methodology. And, especially since the client is never made aware of price settlement, the methodology does sound logical: Check to see by how many basis points Treasury yields moved. Calculate the dollar value of each basis point. Multiply the two and arrive at a settlement value.

However, this methodology is an approximation that always works out in the bank's favor. Why? Because each of the 66.17 basis points has erroneously been assigned the same value of 7.7946 cents. The DV01 calculated at a certain yield is only valid for a one basis-point move away from that yield. Therefore, while the first basis-point shift away from 5.00 percent is indeed worth 7.7946 cents, successive ones are not. Put another way, DV01 at 5.00 percent is different than DV01 at 5.01 percent is different than DV01 at 5.02 percent is different than DV01 at every other yield. And so the value of the basis-point change from 5.00 percent to 5.01 percent is different than the value of the basis-point change from 5.01 percent to 5.02 percent is different than the value of all successive basis-point changes. In fact, even the original DV01 is inaccurate because it was taken to be an average of two different movements. Multiplying the 66.17 basis-point change by a

CONTINUED ON **PAGE 34**

single DV01 ignores all this and assumes that the relationship between changes in yield and changes in price is constant—that each one basis-point move results in a fixed change in price no matter what the yield. Yield settlement takes the graphical representation of the relationship between prices and yields—a curve—and flattens it into a straight line.

Admittedly, all this can be a bit confusing. After all, if price and yield are both valid ways of expressing the value of a bond, shouldn't you also be able to measure the change in value of a bond by looking at either the change in its price or the change in its yield? The math says no. Resorting to hyperbole, teaching the client yield-based settlement is akin to selling them on time travel.

Return for a moment to the example of a car driving along a one-mile track (a conceptual, though not mathematical, equivalent to rate lock settlements). In this analogy, "mph" will play the role of "yield" and "travel time" will play the role of "price." Assume the car is traveling at a speed of one mph. If the car speeds up to two mph, the time required to travel a mile decreases from 60 minutes to only 30 minutes—a 30-minute decrease in travel time. This 30-minute decrease plays the role of "DV01."

Now assume that the car is traveling at a speed of 120 mph. If again the car's speed increases by one mph, here to 121 miles per hour, does the time needed to travel a mile again decrease by 30 minutes? Since a mile only takes 30 seconds to complete at a speed of 120 miles per hour, short of a DeLorean and some lightning, reducing the completion time by 30 minutes would be impossible. The actual reduction in travel time—the "DV01"—would be only a fraction of a second at this high speed. "DV01" is not a constant in this analogy either.

To extend the analogy, calculating a rate lock settlement would be akin to calculating the difference in travel times for each of two laps. If lap one were completed at a speed of 120 mph and lap two at a speed of 1 mph, how would you calculate the difference in travel time between the first and the second lap? Would you take the difference between 120 mph and one mph and multiply that difference by the 30-minute "DV01" calculated above? Doing so would imply an impossibly high difference between the two lap times: 3,570 minutes ((120 – 1) x 30). This calculation is the parallel of the yield-settlement method.

For makes and models without a flux capacitor, you would simply look at the difference between the times the car took to complete each lap. If a stopwatch is not handy, the following quick math provides the answer: a 120-mph lap takes 30 seconds to complete and a one-mph lap takes 60 minutes to complete. The difference in travel time between the two laps is therefore 59.5 minutes. This calculation is the parallel of the price-settlement method. As you can see, the 3,570 minutes calculated using the other method is far off the mark.

In price/yield relationships the same problem exists—that problem being the realities of math. Yet banks I encountered almost always instructed clients to use the yield-based settlement method. And so a product that is meant to return the difference between two Treasury prices, a matter of elementary subtraction, is perverted for profit.

If yields change by very little, this profit does not amount to much. Fortunately, depending on one's point of view, banks have other tricks for profiting from rate locks and do not rely solely on yield-based settlement. In fact, miseducating clients with yield-based settlement is almost an afterthought, just a bonus that pays off with large movements in yield. Because as yields move by more and more basis points two things happen: First, there are more basis points to infect with an erroneously constant DV01. Second, the constant DV01 becomes an even worse approximation for the proper DV01 of each basis point.

In behavior that might be considered yet more sinister, sometimes banks had to implicitly agree with one another to use yield settlement. This transpired if a client decided to divvy up a single rate-lock transaction, with each bank getting a piece of the deal and each bank knowing that settlement of the rate lock would have to be a coordinated affair.

"THIS METHODOLOGY IS AN APPROXIMATION THAT ALWAYS WORKS IN THE BANK'S FAVOR."

All this mathiness is hidden in plain sight. Some examples of yield settlement can be found online. Or you can just ask a company that put on a rate lock to dig up some trade confirmations and see what settlement methodology was used. There are hundreds, if not thousands, of such documents in corporate offices around the country, each one part of an unwarranted transfer of millions of dollars from clients to banks.

END NOTES

[1] http://bostonreview.net/BR36.1/rosen.php

Omer Rosen, a former derivatives banker, is working on an autobiographical novel about love, 9/11, and the financial collapse. He recently started blogging for The Huffington Post and can be reached at writingomer@gmail.com. twitter: @omerrosen website: www.legerdemath.com

THE PERILS OF IGNORING CURRENCIES

By Adnan Akant

Figure 1. Aggregate Mix of Canadian Fund Assets

Mutual Fund Total Assets: C$ 635.7 billion*

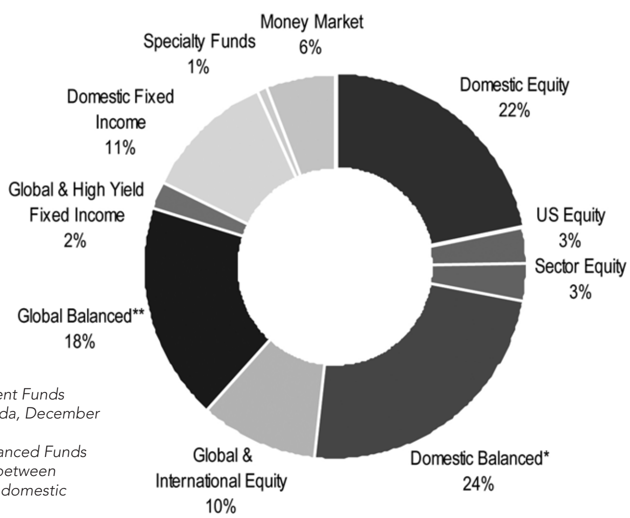

Source: Investment Funds Institute of Canada, December 31, 2010
- *Domestic Balanced Funds mix includes between*
- *70% and 90% domestic assets*

*** Global Balanced Funds can include over 70% Canadian asset*

Pension Plan Total Assets C$ 827.5 billion*

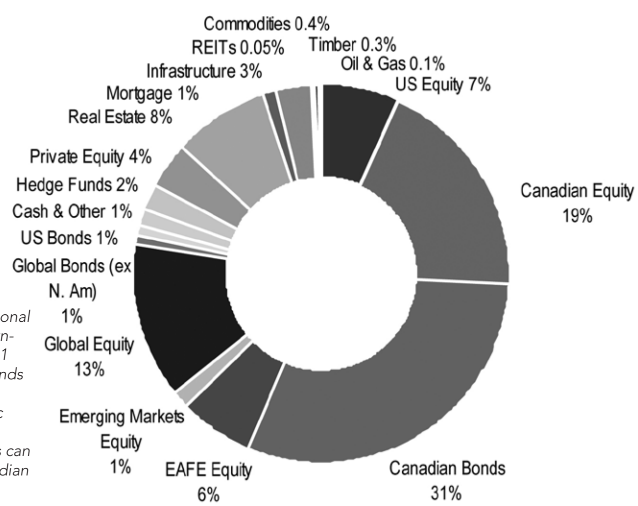

Source: Canadian Institutional Investment Network, download as of February 4, 2011
- *Domestic Balanced Funds mix include between*
- *70% and 90% domestic assets*

*** Global Balanced Funds can include over 70% Canadian assets*

WHY IS CURRENCY MANAGEMENT IMPORTANT?

With the globalization of financial markets, cross-border diversification by investors in all parts of the world continues to expand. Investment opportunities have increased in the last decade with the rise of emerging markets, as well as the growing interaction of developed market economies with emerging economies and with each other through trade and investment. Canadian investors have taken advantage of these international investment opportunities as can be seen in Figure 1. Foreign holdings comprise more than 30 percent of Canadian mutual fund assets and more than 20 percent of Canadian pension plan assets as of the beginning of 2011, a significant exposure.

With this exposure to international investments comes currency risk. Unlike the expected return implied in the risk premium for assets such as equities and bonds, investors cannot expect a return from purely passive holdings of currencies. Equities pay dividends and corporate profits will typically grow with the economy, such that a passively-held market portfolio of stocks can be expected to generate positive returns over longer horizons. Similarly, bonds have fixed coupons and a well diversified portfolio of bonds can also be expected to provide positive returns over time. A similar strategy of holding a portfolio of foreign currencies has no analogous pay off: the expected return of passively-held portfolios of foreign currency is generally zero. Not only do currencies have no expected passive return, they also generate a large amount of risk due to volatility. It could be said that currencies are a form of "returnless risk" when held passively.

Another peculiar finding is that currency returns have very low correlation to returns on bonds and equities in general. In the case of Canada, there appears to be a measurable positive correlation between the Canadian dollar and world equities, especially in the recent decade. This correlation has been explained by the fact that the Canadian economy benefits from higher commodity prices which are themselves correlated to world growth, and hence world equity markets. Over longer periods, however, this correlation has not always held, and may again disappear as the commodity links to the Canadian dollar may weaken in the future.

Given the volatility of currency, unhedged international bond and equity returns can easily be dominated by adverse currency moves. For example, in 2010 European equities returned a positive 11.8 percent in Euro terms, but showed a negative 0.7 percent return in Canadian dollar terms due to the weakness of the Euro.

In spite of the large impact of currency on international asset returns, equity managers typically do not focus on currency management, while international bond managers have shown mixed results from undertaking active currency management. The reason for this is that currency management is a specialist discipline that requires separate focus. Dedicated currency managers typically have the best track records, and many have added significant value from active currency management as will be seen later in this presentation.

THE HISTORY OF EXCHANGE RATES IN CANADA

The modern era of floating exchange rates started with the breakdown of the Bretton Woods system in the early 1970s, and the suspension of the convertibility of the U.S. dollar to gold in August 1971. By the mid '70s most major currencies became floating and the present system came into existence. Figure 2 shows the 36-year evolution of exchange rates for the Canadian dollar from 1975 to now. The blue, green, and red lines depict the nominal exchange rates for the Canadian dollar measured in terms of the U.S. dollar, the Euro and the Japanese yen, starting from a normalized index value of 100, or par in 1975.

It can be seen that the value of the Canadian dollar when measured in U.S. dollars and Euros is very similar today to what it was in 1975. Even though there have been fluctuations of up to 25 percent above and below starting levels, the values today are close to the values in 1975 in the case of the U.S. dollar and the (historically reconstructed) Euro vs. the Canadian dollar. This is also true in the case of the real-effective exchange rate (REER) value of the Canadian dollar shown in orange. The REER is a measure that takes into account a trade-weighted basket of foreign currencies for Canada, adjusted for relative prices due to inflation. (The measure of REER we utilize in this graph is that calculated by the Bank of England). The fact that the REER value for the Canadian dollar, and even the nominal values versus the U.S. dollar and Euro over a 36-year history remain near starting values in 1975, illustrates the point discussed earlier about nearly zero expected returns from passively holding foreign currencies.

In the case of the Japanese yen, the index value today for the Canadian dollar has fallen to about a quarter of its initial level in 1975, unlike the levels near par today for the U.S. dollar, the Euro and the REER. There are two reasons for this seeming anomaly. First, inflation rates have been very different for Japan compared to inflation rates in the United States, Europe and Canada over the past 36 years. Whereas the United States, Europe, and Canada have generally had similar rates of inflation over this period, Japan has had much lower levels of inflation, including periods of outright deflation over the past 20 years. This has been reflected in Japanese interest rates as well as the inflation rate. Holdings in the yen, after adjusting for Japanese interest rates, would have produced poor total returns in Canadian dollar terms over this period. The sharp

Figure 2. Historical Exchange Rates for Canada

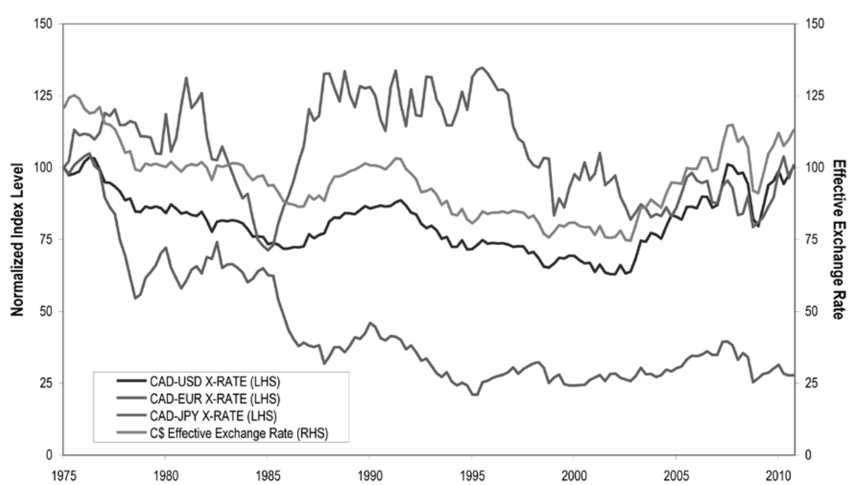

CONTINUED ON **PAGE 38**

appreciation in the yen would not have been sufficient to offset the higher interest rate earned in Canada during these 36 years, a testament to the workings of the classic "carry" strategy. Secondly, the Japanese yen today is at an extremely overvalued level and may yet correct somewhat in the future.

In summary, Figure 2 shows that the exchange rate values for the Canadian dollar vs. the U.S. dollar, the Euro and a trade-weighted basket of foreign currencies adjusted for inflation have remained near the same level after 36 years, though large fluctuations of up to 25 percent in either direction occurred during this time, illustrating the high risk inherent in currency volatility.

THE DIFFICULTY OF MODELING CURRENCY VALUES FROM FUNDAMENTALS

In the first decade of the modern floating-rate regime that started in the early '70s, academic researchers attempted to formulate models that explain and predict the movement of currencies. Though much effort went into developing such models, results were disappointing. By the early '80s a powerful negative finding was reported by Meese and Rogoff in their seminal paper "Empirical Exchange Rate Models of the Seventies."[1] In this paper the authors argued that a random walk model performed as well as any estimated model at horizons up to three years, even when models utilize actual realized values of future explanatory variables such as GDP growth, trade balances, interest rates, inflation, etc. … In other words fixed fundamental models of exchange rates could not explain or predict currency movements in the medium term, even when given the precise evolution of future economic data, which in reality is of course not available when running such models.

As years passed, many other researchers attempted to find models of economic data that would explain currency movements. In 2001, a special conference on empirical exchange rate models was held at the University of Wisconsin for the 20th anniversary of the original work by Meese and Rogoff. A subsequent paper by Rogoff titled "The failure of empirical exchange rate models: no longer new, but still true," published by Blackwell Publishing in May 2002, summarizes the continued failure of exchange rate models.[2] To this day researchers have not been able to explain currencies in the medium term using fundamental economic data.

Figure 3. PPP Currency Valuation vs. EUR/USD Exchange

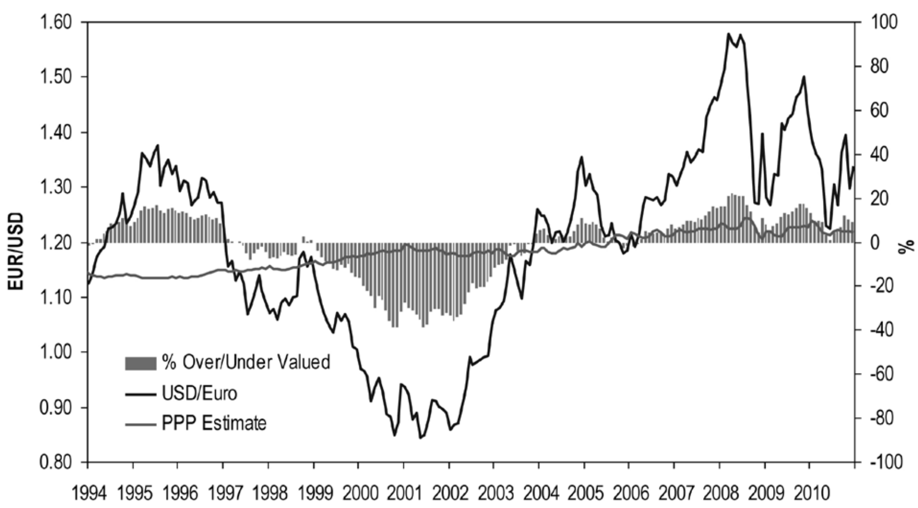

Source: Bloomberg, as of December 31, 2010
Note: Base year = 1980

Another way to see the difficulty of valuing currencies is by observing nominal exchange rates versus their purchasing power parity (PPP) valuations. PPP is a theory of long-term equilibrium exchange rates based on the relative price level of two countries. The concept is based on the notion that identical goods will have similar prices in different markets. There are obviously many limitations to this PPP principle since not all goods are traded and income levels differ greatly in different countries. Nevertheless PPP can provide an anchor for gauging whether a currency is cheap or expensive.

In Figure 3 below, we track the nominal and PPP values of the Euro versus the U.S. dollar since 1994 (note that the values for the Euro were reconstructed for the period prior to the introduction of the Euro in 1999). It can be seen from the red bars that the actual nominal value of the Euro ranged from a low of about 35 percent below PPP in 2001 to a high of +20 percent in 2008.

The Euro's fluctuations around its PPP level last for years and go to extreme levels. For example, the current overvaluation of the Euro has lasted nearly eight years since 2003 and reached an extreme of 20 percent in 2008. Prior to that, the Euro was undervalued for six years reaching a low point of 35 percent below PPP in 2001.

Extreme fluctuations from PPP, which take years to correct, suggest that the notion of an equilibrium valuation based on PPP is not a very useful tool for managing currency risk. This failure, in fact, likely lies at the root of the Meese-Rogoff negative finding, as discussed in the May 2002 paper by Rogoff mentioned earlier.

Having examined the modern history of foreign exchange rates and the difficulty in modeling currency movements from fundamental considerations, we now turn to the practical issues facing an international investor.

OBJECTIVES OF CURRENCY OVERLAY

There are two major objectives that can be viewed as distinct in a currency overlay program. The first objective is the reduction of currency risk stemming from exposure to international investments. The second objective is to enhance returns from active currency management. The objectives of currency overlay are illustrated in Figure 4.

RISK REDUCTION

The typical approach utilized for risk reduction is partial hedging of foreign assets to the base currency. Assume a passive asset allocation is calculated using modern portfolio theory (MPT) while allowing a currency hedge ratio for foreign assets as a parameter in the optimization. Such an optimization, illustrated in Figure 5, would typically produce an efficient frontier that dominates the efficient frontier calculated with no hedging of foreign assets.

Figure 4. Objectives of Currency Overlay

1)	**Risk Reduction:**	Passive partial hedging of a foreign currency benchmark to reduce absolute return volatility of foreign assets
		And / Or
2)	**Return Enhancement:**	Exploiting alpha opportunity in "currency-as-an-asset-class"

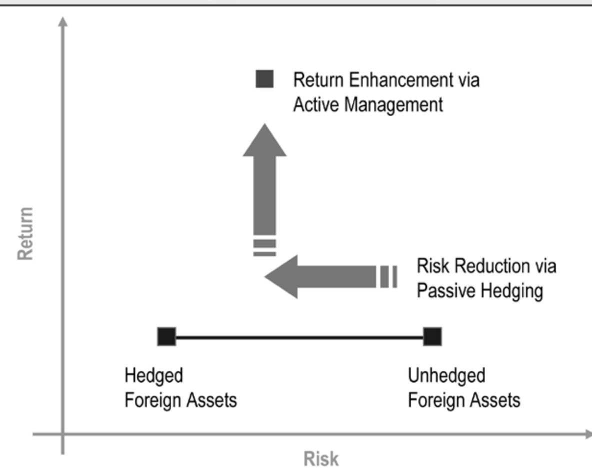

Figure 5. Risk Reduction Via Passive Hedging: An Idealized Efficient Frontier

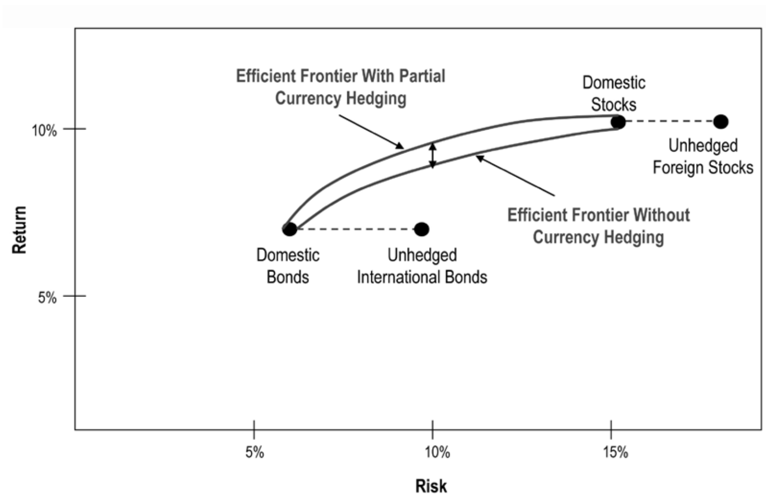

CONTINUED ON **PAGE 40**

The idealized example illustrated in Figure 5 shows that domestic and foreign assets have similar expected returns but that foreign assets have significantly higher volatility. At each point on the superior efficient frontier, which allows partial hedging, the optimal portfolio consists of a unique mix of domestic and foreign stocks and bonds with an appropriate partial hedge ratio.

In general, the optimal hedge ratio rises the higher the amount of foreign assets in the portfolio, at any given portfolio risk level. In a study made by Citigroup for a 20-year period, the optimal hedge ratio from a U.S. investor perspective when foreign equities reach a 20 percent allocation, was nearly 50 percent, as shown in Figure 6.

Even though this result is for a specific 20-year period for a U.S. investor, broadly similar findings can be expected for other historical periods and base currencies, including for the Canadian dollar.

Figure 6. Currency Hedging Becomes More Critical as International Allocations Grow

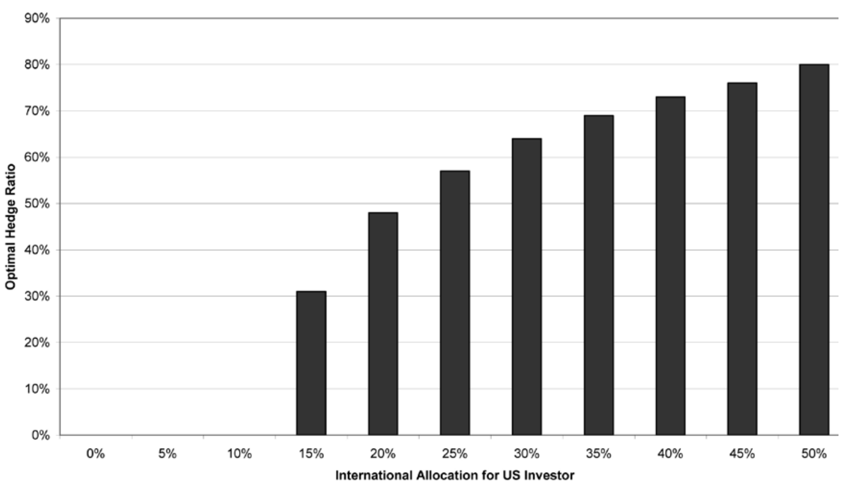

For the years 1978–1998
Note: We use the MSCI EAFE + Canada Index as the Unhedged International Portfolio. The Domestic Portfolio is composed of 605 MS U.S. Index and 40% U.S. Bond index. The spot and forward data was obtained from MSCI.
Source: Citigroup

RETURN ENHANCEMENT VIA ACTIVE STRATEGIES

Currency managers as a group have generally added value through active management. A recent study by Pojarliev and Levich of the Barclay's Currency Traders' Index data over the 17-year period from January 1990 through December 2006 comes to some interesting conclusions. The study finds that currency managers as a group (106 funds as of 2006) had excess returns of 0.25 percent per month (or 3 percent annualized) with a standard deviation of 3.04 percent per month (or 10.5 percent annualized). This provides evidence that currency managers did add value, though the information ratio for the group at 0.28 was not overly impressive (yet far more attractive than the typically negative results found for large samples of active equity managers).

Pojarliev and Levich go beyond simply calculating total excess return by decomposing this excess return into an alpha and a sum of beta factors plus a residual error term. The result of their study is summarized in Table 1 and Equation 1.

	Currency Returns	Risk-Free Returns	Excess Returns	Carry Factor	Trend Factor	Value Factor	Volatility Factor*
January 1990 – December 2006 (N = 204)							
Mean monthly return (%)	0.62	0.37	0.25	0.15	0.21	0.03	-0.01
Standard deviation (%)	3.06	0.16	3.04	0.78	1.86	0.38	1.55

** The "return' for the volatility factor is , in fact, the first difference of the implied volatility rather than a return.*
Source Barclay's Currency Traders Index ("BCTI" - 106 funds in 2006)

The underlying thesis of the authors is that there are four simple rule-based naïve strategies, namely carry, trend, value and volatility, that can be thought of as representing beta factors which earn a risk premium in currency markets. Investors can follow these well-defined styles and earn a return just as stock investors can generally expect to earn a return over time from buying and holding market portfolios of equities (e.g., the S&P 500 Index). When total excess returns are looked at in this fashion, the alpha is the excess return which is above and beyond the return earned by following the naïve strategies of carry, trend, value and volatility. In this much stricter analysis of currency manager skill, Pojarliev and Levich find that the overall group of currency managers actually has a negative alpha, with the entire mean excess return more than fully explained by the simple strategies of carry and trend (note that value and volatility beta factors show almost no contribution to mean excess return either). It would appear from this analysis that a currency investor would be well advised to simply follow carry and trend as naïve strategies rather than selecting currency managers to add value.

In a follow-up paper, however, Pojarliev and Levich show that it makes sense to allocate to currency managers by selecting a group based on any of four different methods of selection. In each case the selection process is based on an analysis of a three-year "in-sample" period from April 26, 2005 to March 26, 2008. The evaluation of the various currency manager selections is then made based on a 27-month subsequent "out-of-sample" period of April 2, 2008 to June 30, 2010.

The four methods of selecting currency manager portfolios are based on the beta factor model discussed in Table 1. The currency portfolio results shown in Table 2 focus on four specific ways to select currency managers using "in-sample" data, and provide the excess return and information ratio achieved in each case in the "out-of-sample" period.

Equation 1. Alpha and Beta Factors in Currency Excess Returns

$$R_t = \alpha + \sum_i \beta_i F_{i,t} + \varepsilon_t$$

R is the excess return generated by the currency manager, defined as the total return (R_t*) less the periodic risk-free rate ($R_{F,t}$)

α is a measure of active manager skill,

F is a beta factor, that requires a systematic risk premium in the market (e.g. carry or trend factors)

ε_t is an error residual (i.e. random noise)

Table 2. Out-of -Sample relative performance to the MSCI World Index by Allocatoion

	2% Allocation to Currency Managers* (April 2, 2008 to June 30, 2010)		
	Excess Return	**Tracking Error**	**Information Ratio**
Portfolio 1: Equity + Total Return FX	123 bps	76 bps	1.62
Portfolio 2: Equity + Beta Chasing FX	57 bps	74 bps	0.77
Portfolio 3: Equity + Alpha Hunting FX	182 bps	66 bps	2.78
Portfolio 4: Equity + Alpha Generating FX	257 bps	92 bps	2.80

** A 2% cash allocation is levered up by 10 to achieve a 16.95% (or 20/118) allocation to currency and 83.05% allocation to equity (1–0.1695)*

Source: Pojarliev, M., and Levich, R.M. 2010. "Are All Currency Managers Equal?" NYU finance working paper. Portfolios 1–4 determined based on the in-sample period of April 26, 2005 to March 26, 2008

Portfolio 1. The Total-Return Portfolio: the currency allocation is invested in an equal-weighted exposure to the top three managers with the highest total return generated during the "in-sample" period. The "out-of-sample" excess return produced is a solid 123 basis points with a high information ratio of 1.62.

CONTINUED ON **PAGE 42**

Portfolio 2. The Beta-Chasing Portfolio: the currency allocation is invested in an equal-weighted exposure to the top three "beta-grazers" during the "in-sample" period; i.e., the managers with the highest estimated R-squared for the beta factor model of excess returns (Equation 1). The "beta grazers" excess returns have the best fit to the factor model, and these managers are therefore most likely chasing these beta factors as their investment approach. In the "out-of-sample" period the "beta grazers" show an excess return of 57 basis points with a respectable information ratio of 0.77.

Portfolio 3. The Alpha-Hunter Portfolio: the currency allocation is invested in an equal-weighted exposure to the top three "alpha hunters" during the "in-sample" period, i.e., the managers with the lowest estimated R-squared for the beta factor model of excess returns. The "alpha hunters" excess returns have the worst fit to the factor model, and these managers are therefore the least likely to be following active strategies based on the beta factors: they are managers likely to be focused mostly on generating alpha as defined by the factor model equation. In the "out-of-sample" period these "alpha hunters" show a strong excess return of 182 bps with a superb information ratio of 2.78.

Portfolio 4. The Alpha-Generator Portfolio: the currency allocation is invested in an equal-weighted exposure to the top three "alpha generators" during the "in-sample" period, i.e., the managers with the highest point estimate for the alpha of the beta factor model of excess returns. The "alpha generator" managers show the most skill in the "in-sample" period in generating alpha as defined by the beta factor model of excess return. In the "out-of-sample" period these "alpha generators" show a very strong excess return of 257 bps with again a superb information ratio of 2.80.

The study demonstrates that even though the overall group of currency managers does not provide impressive risk-adjusted excess return in the earlier paper, the four methods used in the second paper show success in selecting solid portfolios of active currency managers, each with different investment approaches for adding value. The persistence of investment styles from the "in-sample" to the "out-of-sample" period is notable and provides some evidence of sustainability of excess return generation and skill from one period to the next. Furthermore, the managers with demonstrated alpha skills defined by the beta factor model of excess returns, show much higher value-added compared to the managers chasing beta factors.

WHAT MAKES CURRENCY ALPHA POSSIBLE?

We have examined the difficulty academics have in explaining the movement and valuation of currencies based on fundamental economic factors. We have also examined the empirical evidence that many active currency managers have, nevertheless, shown the ability to add solid value. The question then is: what makes currency alpha possible?

First, we should review a few facts about currency markets. The recent period of free-floating exchange rates evolved in the early 1970s, after the breakdown of the Bretton Woods system of fixed exchange rates. Since that period, the currency market has proven to be the largest financial market by far, where daily volume has reached over $4 trillion with an annual growth rate of about 10 percent in 2010. There are literally many thousands of participants large and small, spread around the world trading currency 24 hours a day, nearly every day of the year. The liquidity of the currency market is also unsurpassed among financial markets: transaction costs are minimal and major currencies can be traded in very large size. Given the size and liquidity of the currency market many casual observers expect a degree of efficiency which precludes the possibility of alpha generation by active managers. As we have seen earlier in this presentation, however, the evidence shows that many active currency managers consistently add value.

The existence of profit potential in active management may be the result of a few distinct characteristics of currency markets. In spite of the enormous volume and liquidity of currency markets, a very large portion of the daily flow is not driven by profit seekers in currency. Transactional flows for trade in goods and services, corporate activity for mergers and acquisitions or Foreign Direct Investment, purchases or sales of international stocks or bonds, hedging activity for exports or imports, official currency flows for reserve management or intervention, flows generated by tourists traveling abroad or guest workers remitting savings to their home country, etc. … are not primarily driven by the motive of profiting directly from currency moves.

The participants who are attempting to profit from currency moves directly are primarily currency dealers such as international banks and active currency managers such as hedge funds, private speculators, and specialized currency overlay managers. The first group, the bank dealers, tends to focus on very short-term moves with the largest activity taking place intraday. The second group, which includes the true active currency managers who seek to profit from moves over multi-week or longer periods, is indeed a small fraction of the entire flow and of the participants in the currency markets. The data is hard to ascertain accurately but we estimate this flow and group to account for less than 10 percent of the overall currency market. The fact that investors seeking to profit from currency movements over periods exceeding a single day constitute a minor portion of the market may partly explain why this group can exploit significant opportunities and shows some success in adding value.

Currency markets serve as a significant "release valve" for global imbalances and play a large role in macro-economic equilibrium adjustments. In most countries it is far easier from a political, as well as an economic perspective, to let the currency adjust than to move official interest rates, shift fiscal policy, or deploy other tools to relieve financial and economic stresses. For this reason, currencies can move significantly, providing profit opportunities for those who are able to anticipate capital flows stemming from global macro imbalances. This paradigm is certainly true for countries that have sovereign monetary policy and free capital flows, for example the United States or the United Kingdom. For countries that choose to peg their currencies at a fixed rate, the situation is obviously different. With fixed exchange rates, a country has to choose between a sovereign monetary policy and free capital flows. The famous "tri-lemma" of foreign exchange, illustrated in Figure 7, states that a country can only have two of the following three: sovereign monetary policy, free capital flows and pegged exchange rates.

Most major countries have a sovereign monetary policy and free capital flows and thus free-floating exchange rates (with the Eurozone viewed as a single country for this purpose). An

Figure 7. The "Tri-lemma" of Currency Markets

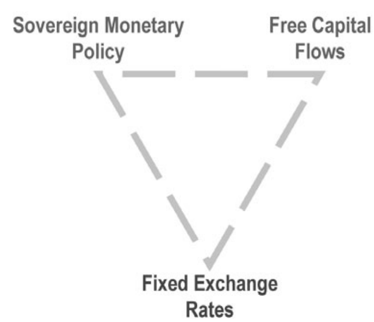

example of pegged exchange rates and sovereign monetary policy is China (U.S. dollar peg) where free capital flows are therefore not possible. An example where pegged exchange rates and free capital flows coexist is Hong Kong (U.S. dollar peg), in which case a sovereign monetary policy is therefore not feasible: Hong Kong interest rates must closely match U.S. rates.

Given that macro imbalances around the world often occur, and that currencies are a primary method by which such imbalances are reduced, active currency management makes a great deal of sense, provided imbalances can be identified, and the timing of capital flows that act to alleviate such imbalances can be anticipated. In fact, many active strategies seek to anticipate policy changes and capital flows that generate adjustments

CONTINUED ON **PAGE 44**

in currencies. Capital flows themselves typically respond to expected rates of return, which are driven not only by sovereign economic policy but also by a number of other factors. In what follows, we discuss FFTW's approach to active currency management as a practical example of how a currency manager seeks to organize its process for generating return in currencies.

FFTW'S ACTIVE CURRENCY PROCESS

In the case of FFTW, the currency process, which is illustrated in Figure 8, is organized into two separate styles to capture both an alpha and a beta component corresponding to the framework of Equation 1 discussed earlier.

The beta component is based on systematic models that follow three approaches: trend, carry and yield-delta. The trend and carry approaches are classic and well understood. The yield-delta approach is based on changes in interest rate spreads, rather than the level of rate spreads which drive carry (hence the term yield-delta). We have found that such yield spread momentum is a good basis for a systematic model that is robust and has strong negative correlation to the carry approach. These three beta factors taken together show solid excess returns with desirable diversification benefits over time. For example, during the recent financial crisis of 2008, the overall beta component had good returns with trend and yield-delta performing very strongly while the carry approach showed significant underperformance.

Figure 8. Currency Alpha Process: FFTW Example

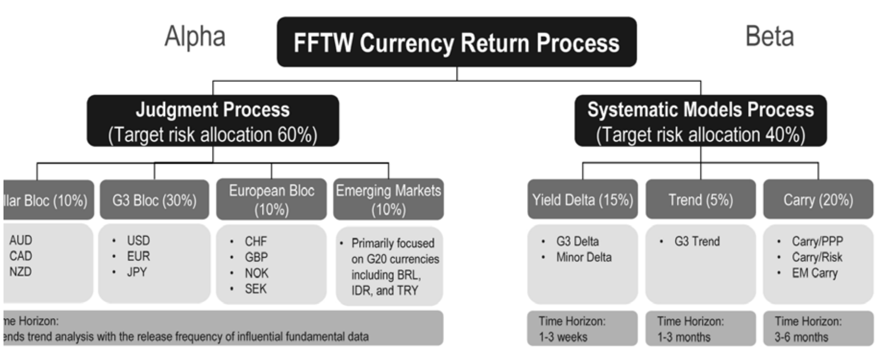

The models were developed over a decade of proprietary research at FFTW and have been allocated a risk budget of 40 percent of the total currency process since October 2006. The remaining 60 percent of the risk budget is allocated to a judgment-based alpha process that seeks to generate return by anticipating capital flows resulting from global macro and country specific factors.

There are many factors that can influence currency markets and capital flows. At FFTW, we focus on five areas to make judgments about alpha opportunities, as illustrated in Figure 9.

Two of these areas are directly linked to interest rates: First, yield spreads, which incorporate the foundation of expected return differentials that create incentives for capital flows from areas with surplus savings to deficit countries. Second, monetary policy expectations, which have a large impact on the expected evolution of interest rate spreads, and can change the risk/ reward trade-off of investments in one currency versus another.

A third factor is related to the terms of trade dynamics. Improving terms of trade (ratio of export prices over import prices) generally have two positive impacts on capital flows: (1) an improving trade balance, and (2) foreign direct investment inflows into the country's export sector.

A fourth area of importance is fiscal and regulatory policy: tax policy changes can significantly alter the expected return of holding a particular currency. For example, a shift in corporate taxes can cause large inflows into a country, as happened in the United States and Japan at various times in the last five years (e.g., the Homeland Investment Act of 2004 allowed a reduced tax rate for repatriated U.S. corporate profits for one year, which contributed to a strengthening of the U.S. dollar in 2005).

Finally, a fifth factor which is very difficult to model but extremely important in affecting capital flows is risk-aversion. Unexpected developments, such as the Lehman bankruptcy of 2008, the 9/11 attacks of 2001, the LTCM bankruptcy of 1998, the Japanese Tsunami of 2011, and many lesser events or investor sentiment swings can cause an increase in risk-aversion. High levels of risk aversion tend to drive a reversal of "usual" capital flows, as risky holdings are pared back in favor of cash. Gauging risk-aversion is best done using a discretionary judgment-based approach, as every episode tends to be different and appears with little warning.

Figure 9. Major Factors influencing Currency Markets

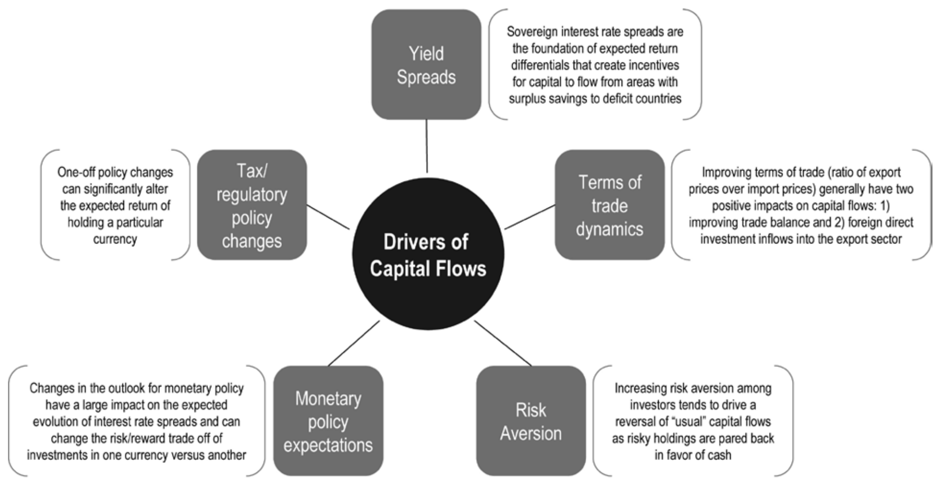

The FFTW currency alpha process is based on a flexible way of analyzing the factors described above and utilizing judgment to assign a proper weight to the opportunities and risks presented, as opposed to following the rigid styles incorporated in the beta models.

As described earlier, the models operate systematically and are designed to capture various beta factors. However, the judgment process led by the team head can always dominate model positions given the 60/40 risk budget split in favor of judgment. In practice, models are rarely interfered with, as they were developed within the currency team, which strongly believes in these beta factors (models were briefly stopped during the Lehman crisis -- with little overall impact on performance). At the end of the day, however, the currency team leader retains discretion to override models at any time using judgment.

The currency team continuously studies the characteristics of model performance to retain confidence in the structure of the models over time. The model styles selected at FFTW are well understood currency factors that have worked well for decades (especially trend and carry). The model suite is meant to complement the judgment process with a portion of the risk budget and create a diversified portfolio of investment styles. The critical element is the combination of models, and the additional diversification benefit of using judgment, which remains the most flexible approach in times of stress, or changing conditions.

THE CASE FOR ACTIVE CURRENCY: CORRELATIONS WITH MAJOR INDICES

We have discussed the evidence that active currency management can be expected to provide attractive excess returns over time. What makes the case for including active currency as a separate source of return in longer term global portfolios particularly compelling is the low level of correlation active currency shows with major stock, bond, and commodity indices. The data in Table 3 (Pg. 47) was computed utilizing FFTW's active currency returns over a period of over four years from October 2006 through December 2010, which covers the period for which FFTW's currency composite data is available.

CONTINUED ON **PAGE 46**

As can be seen in the Table, FFTW's active currency returns show 27 percent or lower levels of correlation with all the selected passive asset indices. The third to last column also shows a low correlation of 28 percent with the HFRI fund weighted composite of hedge funds, an index of active returns.

Furthermore, over this more than four-year period, the annualized excess return of 4.3 percent over one month LIBOR achieved by the FFTW active currency composite exceeds the excess return generated by all the other selected indices, with an attractive information ratio of 0.62.

The correlation and return data here suggest that long-term investors should seriously consider allocating some of their portfolio risk to active currency programs as they now do to equities, bonds, and other alternatives such as hedge funds, commodities, and real estate.

Table 3. Excess Return and Correlations for Select Indices

Correlations

	S&P 500	MSCI World	BarCap Global Agg	BarCap US Agg	BarCap 3-6 Months US T-Bills	BarCap World IL	HFRI Fund Weighted Composite	GS Commodities Index	FFTW Currency Alpha
S&P 500	100%								
MSCI World	98%	100%							
BarCap Global Agg	2%	0%	100%						
BarCap US Agg	16%	12%	93%	100%					
BarCap 3-6 Months US T-Bills	-14%	-16%	6%	3%	100%				
BarCap World IL	34%	30%	67%	77%	-3%	100%			
HFRI Fund Weighted Composite	82%	87%	-7%	8%	-20%	32%	100%		
GS Commodities Index	56%	56%	-17%	1%	-7%	32%	74%	100%	
FFTW Currency Alpha	27%	25%	10%	18%	12%	17%	28%	13%	100%

Returns

	S&P 500	MSCI World	BarCap Global Agg	BarCap US Agg	BarCap 3-6 Months US T-Bills	BarCap World IL	HFRI Fund Weighted Composite	GS Commodities Index	FFTW Currency Alpha
Annualized Excess Return over 1-Mo Libor	**-1.7%**	**-5.1%**	**2.6%**	**3.7%**	**-0.3%**	**2.7%**	**2.8%**	**-6.5%**	**4.3%**
Information Ratio	-0.09	-0.28	0.85	0.92	-1.37	0.41	0.34	-0.22	0.62

** A 2% cash allocation is levered up by 10 to achieve a 16.95% (or 20/118) allocation to currency and 83.05% allocation to equity (1–0.1695)*

Source: Pojarliev, M., and Levich, R.M. 2010. "Are All Currency Managers Equal?" NYU finance working paper. Portfolios 1–4 determined based on the in-sample period of April 26, 2005 to March 26, 2008

CONCLUSIONS

Global investing implies foreign currency exposure. Yet currency management has often been an after-thought in the decision process of global investors. This is both a serious peril as well as a missed opportunity. Many of those focused on the peril of currency risk have taken a passive hedging approach. Such an approach is usually better than doing nothing, especially as the amount of international investment grows as a share of the overall portfolio. However, taking a step further and embracing a program of active currency management can be a rewarding strategy that transforms currency risk into a significant opportunity for excess return. Academic research suggests that currency returns are hard to explain using economic data. In spite of this, currency managers have been successful in adding value in practice. We discussed some recent analysis of long-run currency manager performance that shows evidence of alpha and beta factors, and suggests approaches for selecting active currency managers. We also examined what makes currency alpha possible and gave an account of how FFTW, as an example of an active currency manager, has structured a process which includes both a judgment-based as well as a systematically-driven approach. Finally, we reviewed some evidence showing that active currency is an excellent source of uncorrelated alpha in global portfolios.

DISCLAIMER

either express or implied, that any investment strategies work under all market conditions. Past performance is not a guarantee of future results. Readers should independently evaluate the information presented and reliance upon such information is at their sole discretion. Investment advice or recommendations, if any are set forth in this document or elsewhere, may not be suitable for readers or any portfolio.

The information contained herein includes estimates and assumptions and involves significant elements of subjective judgment and analysis. No representations are made as to the accuracy of such estimates and assumptions, and there can be no assurance that actual events will not differ materially from those estimated or assumed. In the event that any of the estimates or assumptions used in this presentation prove to be untrue, results are likely to vary from those discussed herein.
Fischer Francis Trees & Watts, Inc. is registered with the US Securities and Exchange Commission as an investment adviser under the Investment Advisers Act of 1940.

Fischer Francis Trees & Watts UK Limited is authorized and regulated by the Financial Services Authority. Registered in England No: 979759, registered office: 5 Aldermanbury Square, London, England, EC2V 7BP, United Kingdom. FFTW UK is also registered with the US Securities and Exchange Commission as an investment adviser under the Investment Advisers Act of 1940, as amended.

Fischer Francis Trees & Watts Singapore Limited is registered with the Monetary Authority of Singapore for the conduct of Fund Management business and is a holder of a Capital Markets Services License. FFTW Singapore is also registered with the US Securities and Exchange Commission as an investment adviser under the Investment Advisers Act of 1940.

END NOTES

1. Meese, Richard and Kenneth Rogoff. "Empirical Exchange Rate Models of the Seventies." Journal of International Economics. 14 (1983): 3-24. North–Holland Publishing Company.
2. Rogoff, Kenneth. "The Failure of Empirical Exchange Rate Models: No Longer New, But Still True." Blackwell. 2002.

Adnan Akant, PhD, is head of Currencies at Fischer Francis Trees & Watts (FFTW), a BNP Paribas Investment Partner. He can be contacted at adnan.akant@bnpparibas.com

SOCIETY OF ACTUARIES Investment Section

475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
p: 847.706.3500 f: 847.706.3599
w: www.soa.org

Risks & Rewards