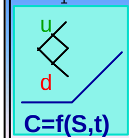

Derivatives Performance Attribution

Mark Rubinstein
Paul Stephens Professor of Applied Investment Analysis
University of California at Berkeley
rubinste@haas.berkeley.edu

Abstract

This paper shows how to decompose the dollar profit earned from an option into two basic components:

- ❶ mispricing of the option relative to the asset at the time of purchase, and
- ❷ profit from subsequent fortuitous changes or mispricing of the underlying asset.

This separation hinges on measuring the "true relative value" of the option from its realized payoff. The payoff from any one option has a huge standard error about this value which can be reduced by averaging the payoff from several independent option positions. It appears from simulations that 95% reductions in standard errors can be further achieved by using the payoff of a dynamic replicating portfolio as a Monte Carlo control variate. In addition, it is shown that these low standard errors are robust to discrete rather than continuous dynamic replication and to the likely degree of misspecification of the benchmark formula used to implement the replication.

The first basic component, the option mispricing profit, can be further decomposed into profit due to superior estimation of the volatility (**volatility profit**) and profit from using a superior option valuation formula (**formula profit**). In order to make this decomposition reliably, the benchmark formula used for the attribution needs to be similar to the formula implicitly used by the market to price options. If so, then simulation indicates that this further decomposition can be achieved with low standard errors.

The second basic component can be further decomposed into profit from a forward contract on the underlying asset (**asset profit**) and what I term pure option profit. The asset profit indicates whether or not the investor was skillful by buying or selling options on mispriced underlying assets. However, asset profit could also simply be just compensation for bearing risk -- a distinction beyond the scope of this paper. Although simulation indicates that the attribution procedure gives an unbiased allocation of the option profit to this source, its standard error is large -- a feature common with attempts by others to measure performance of assets.

May 3, 1998

Mark Rubinstein is the Paul Stephens Professor of Applied Investment Analysis at the University of California at Berkeley. The author would like to thank Barra, Inc. for encouraging this research.

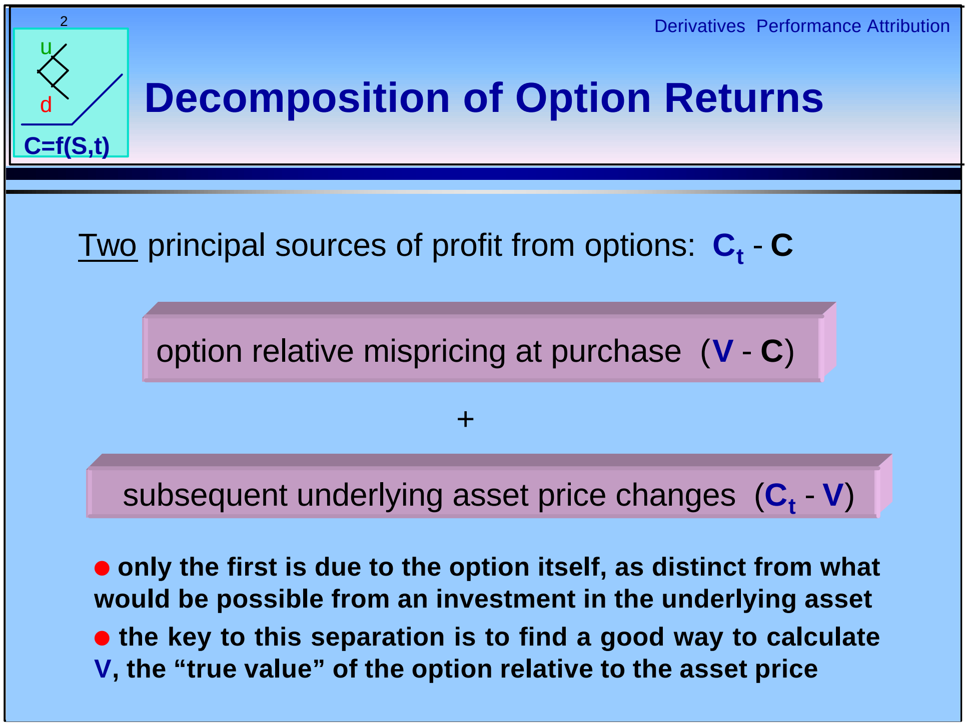

I. Introduction

The challenge of performance attribution is to find a way to decompose the realized return of a portfolio, which may contain derivatives, attributing its components to particular causes. As a simple and concrete example that we shall use throughout this paper, suppose that the portfolio contains a single call originally purchased for **C** on an underlying asset. We will measure performance after elapsed time **t** when the call has a market price of $\mathbf{C_t}$. The realized profit is then $C_t - C$.

The basic decomposition is to divide this profit into two parts:*

- ❶ mispricing of the option relative to the asset at the time of purchase, and
- ❷ profit from subsequent fortuitous changes or mispricing of the underlying asset.

If we let **V** be our estimate of the initial "true relative value" of the option (to be exactly defined shortly), then V - C measures the first component, and C_t - V measures the second component. Together these components obviously satisfy the adding-up condition:

$$C_t - C = (V - C) + (C_t - V)$$

This decomposition seeks to isolate the marginal contribution of derivatives to portfolio performance. The first component isolates the contribution that we would want to attribute to the skills of the investor in selecting underpriced options. If V > C, then his realized performance is positive. The second component, C_t - V, as we shall see, is primarily dependent on selecting the right asset on which to buy the option. Performance in this dimension we attribute to selecting calls on underpriced assets.

A problem in measuring performance, encountered here and elsewhere, is that profits and losses, whether they be from components ❶ or ❷, may be due to luck or skill. We need to find some way to isolate profits due to each source. To do this, we will need to examine a series of different investments, over different periods of time, so that the law of large numbers can begin to work to sort this out. One of our challenges will be to design a way of measuring these components, particularly for this paper, component ❶, that allows us to distinguish between luck and skill with as few observations and as short a time period as possible. It is widely known that attempting to make this separation between luck and skill for equity portfolios can take many years of observations. As we will see, for the purpose of measuring the marginal contribution of options, far fewer observations will be required.

* An earlier approach to this problem can be found in Galai (1983).

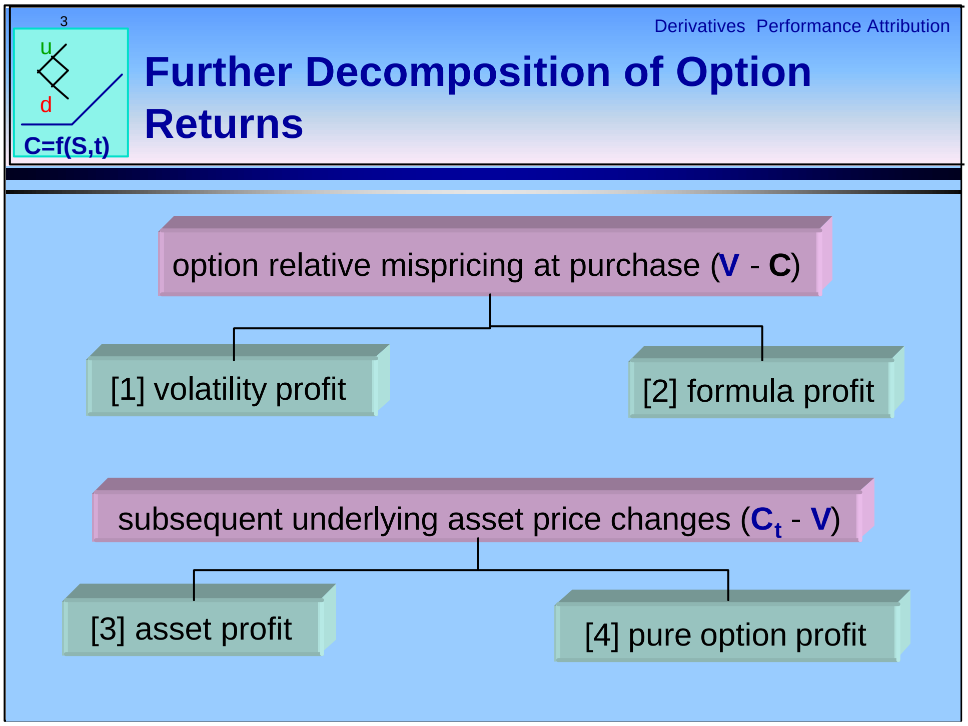

To provide greater insight, we will want to decompose the profits from the option even further.

It is commonly believed that the two major tasks in identifying mispriced options relative to their underlying asset price are estimation of parameters, particularly volatility, and use of the right option pricing formula. Thus, one of our goals will be to further decompose

❶ option relative mispricing at the time of purchase into:

- volatility profit (arising from superior volatility estimates)
- formula profit (arising from a superior option valuation model)

These two components will satisfy the additivity requirement since their sum is equal to ❶.

It is difficult to separate out without ambiguity the second component of profit, the profit due to fortuitous subsequent movement or mispricing of the underlying asset price. Hence we will break it into two parts, one that can be solely attributed without question to the underlying asset, one that is a joint result of selecting the underlying asset and the option. That is, we further decompose

❷ profit due to subsequent fortuitous movement or mispricing of the underlying asset into:

- asset profit (from a non-option strategy such as forward contract on the asset)
- pure option profit (additional profit arising by chance in an efficient option market from buying an option rather than a forward)

Again, these two components will add up to ❷.

4 Derivatives Performance Attribution

C=f(S,t)

Three Formulas

- **"true formula": the option valuation formula based on the actual risk-neutral stochastic process followed by the underlying asset**
- **"market's formula": the option valuation formula used by market participants to set market prices**
- **benchmark formula: the option valuation formula used in the process of performance attribution to**
  - **(1) help determine the "true relative value" of the option**
  - **(2) decompose option mispricing profit into components**

We distinguish between these formulas and their riskless return and volatility inputs, for which there are also three estimates -- true, market, and benchmark. Essentially, by the "formula" we mean the levels of all the other higher moments of the risk-neutral distribution, where each moment may possibly depend on the input riskless return and volatility.

We will need to distinguish between three option valuation formulas. The **"true formula"** captures the actual risk-neutral stochastic process of the underlying asset price. Although we do not know this formula, we can hope to learn something about it by observing the realized option payoffs. Using simulation in Part IV of this paper, the performance attribution approach does not generally know what the true formula is. But, since the simulation itself will know the true formula, we can ask how quickly our performance attribution method learns about this formula.

To separate components of performance, we also need to know the formula used by the market in setting option prices, which we call the **"market's formula"**. In an inefficient market, this may not coincide with the "true formula".

A third formula, used directly in the methodology, is called the **benchmark formula**, such as the standard binomial option pricing model. We will use the benchmark formula for two purposes:

(1) to help determine the "true relative value" V of the option

(2) to decompose the profit due to option mispricing into components

For the first, we will want to use a benchmark formula that comes as close as possible to the "true formula." For the second, we want a benchmark formula that comes as close as possible to the "market's formula." In particular, if the market uses the wrong formula, the same formula will not satisfy both these purposes, so more generally we would want to consider using different benchmark formulas for each purpose. It is one thing to hope we understand how the market values options, but it is quite another to know the "true formula" if this is different. So, in this paper, we will use the same formula -- our best guess about the "market's formula" -- for both purposes.

Fortunately, as we shall see, even if our benchmark formula is a poor approximation of the "market's formula," our calculation of V may not be seriously affected since it can be cured by the law of large numbers. However, failure to use the "market's formula" will affect the way we decompose option mispricing profit and will not be cured by a large sample.

In this paper, we will use the standard binomial model as the benchmark formula. But in practical applications we may prefer using the constant elasticity of variance, jump-diffusion, or an implied binomial tree model depending on the nature of the stochastic process of the underlying asset and our beliefs about how the market prices options.

C=f(S,t)

[1] Volatility Profit

The profit from option relative mispricing at purchase V - C, can be decomposed into 2 parts:

- **Payoff attributed to difference between the realized (s) and implied volatility at purchase based on benchmark formula:**

C(s) - C

- **Positive results indicate that the investor was clever enough to buy options in situations where the market underestimated the forthcoming volatility.**

(Caveat: Although we can know that an option is mispriced, we will not be able to tell why it is mispriced if the benchmark formula used to calculate C(s) is not a good approximation of the formula used by the market to set the option price C.)

II. Four Components of Performance

No doubt our attempt to describe this decomposition of option profit in succinct English is not as clear as one would like. So now, with the aid of mathematics, we take each of the components up one by one.

Profit results if an investor can successfully identify mispriced options. We will decompose this basic component of profit into two parts: profit made because the investor was good at predicting volatility and profit made because the investor was good at finding a superior option valuation formula to the one apparently used by the market, reflected in the initial market price of the option. Academic research that investigates how well a model's implied volatility forecasts realized volatility concerns itself with the first of these. Whereas academic research that focuses on how well alternative option pricing models forecast future option prices, conditional on the future underlying asset price, is primarily concerned with the second of these.

To isolate the profits made from superior forecasts of volatility, we calculate the initial value the option should have had, using the benchmark formula, had the realized volatility of the asset over the assessment period (elapsed time t) been known in advance. We use **s** to represent annualized realized volatility over the assessment period. It is not immediately clear how this should be measured, but a first cut is to measure the sample volatility of realized daily asset returns. The value of the option measured by the benchmark option pricing formula with volatility parameter s is denoted by C(s). In contrast, the current option price C can be interpreted as the value of the option measured by the benchmark option pricing formula with volatility parameter s, which is the implied volatility. So we can interpret the difference, **volatility profit**,

$$C(s) - C$$

as the mispricing of the option due to the fact that under the benchmark formula, the market mistakenly thought that the volatility was σ rather than the volatility s that actually was realized.

If this is positive, it suggests that the investor may have skill in forecasting volatility.

C=f(S,t)

[2] Formula Profit

- **Profit attributed to using a formula superior to the benchmark formula, assuming realized volatility were known in advance:**

V - C(s)

- **Positive results indicate that the investor was clever enough to buy options for which the benchmark formula, even with foreknowledge of the realized volatility, undervalued the options.**

(Caveat: Although we can know that an option is mispriced, we will not be able to tell why it is mispriced if the benchmark formula used to calculate C(s) is not a good approximation of the formula used by the market to set the option price C.)

It is important to realize that our measure of volatility profit can be quite sensitive to the benchmark formula. For this purpose, we want to choose as a benchmark our best guess for the formula used by the market to price options. If the benchmark is incorrect, then we will mistakenly confuse volatility profit with formula profit. If we do not know the market's formula, then although we can know that an option tends to be mispriced, we will not be able to tell why it is mispriced.

The second mispricing component isolates the profit from using an option pricing formula superior to the benchmark. If V is the relative value of the option based not only on the true formula but also on the realized volatility, then the difference we call the **formula profit**,

$$V - C(s)$$

is the profit due to using a superior option valuation formula to the benchmark. Since both V and C(s) are measured using the realized volatility, this difference isolates the role of the option pricing formula from skill in forecasting volatility.

Again, for this decomposition to work, we need to be using a benchmark that is a good approximation for the formula used by the market to set the option price C.

Note: In practice, to reduce the standard error of the decomposition of option mispricing into volatility and formula profits, we will average the calculations over several option investments. Averaging C(s) – the benchmark formula values based on realized volatility – produces a biased estimate because the formula value is generally a non-linear function of the volatility. However, using benchmark formulas similar to Black-Scholes induces a trivial bias since this formula is almost linear in volatility over the relevant range. For example, using the parameter inputs we will later use in our simulation, at-the-money European calls using the Black-Scholes formula have the following values associated with their annualized volatility:

volatility	option value
15%	$2.7509
20%	$3.5549
25%	$4.3599

Given the option values at 15% and 25% volatilities, if the formula were exactly linear in volatility, then the option value at 20% volatility would be $3.5553.

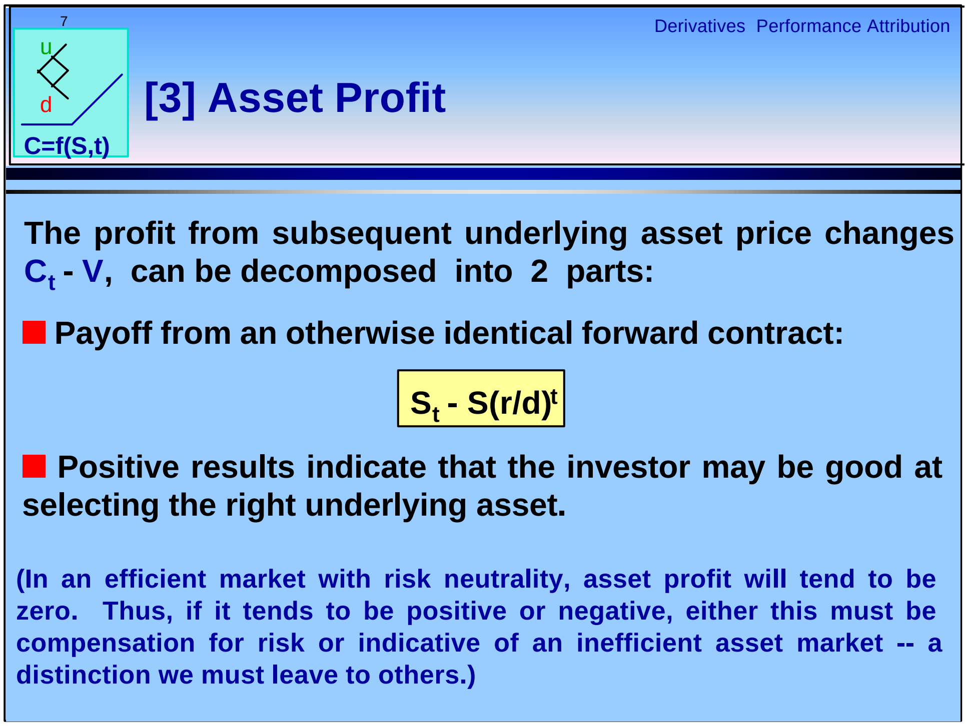

The next component of profit is simply the performance of a benchmark strategy, assumed to be a forward contract on the underlying asset, which we call the **asset profit**:

$$S_t - S(r/d)^t$$

where:

- **S** is the price of the asset at the time the option was purchased
- $\mathbf{S_t}$ is the price of the asset after elapsed time t (when performance is being assessed)
- **r** is the annualized riskless return over the period
- **d** is the annualized payout return over the period for the asset

r and d are two parameters, along with volatility, that determine the value of options.*

$S(r/d)^t$ is the formula for the fair value of a forward contract on the asset with a time-to-delivery of t. Therefore, the difference $S_t - S(r/d)^t$ is the realized profit from this forward contract. In an efficient asset market, the present value of this is zero.

Note that the choice of the benchmark strategy is to some extent arbitrary. Another benchmark could easily be used instead. The benchmark hides the factors that will be unexplained by our analysis. While our calculation $S_t - S(r/d)^t$ measures the profit from selecting the asset underlying the option, in this paper we will not in turn decompose that into its sources -- a problem which has many commercially available solutions.

In an efficient asset market with risk neutrality, the asset profit will on average be zero. If it tends to be unequal to zero, then either the market is risk averse (or risk preferring) or the underlying asset is mispriced. This is not a distinction which this paper can help sort out.

* While we assume that the true volatility is not known by the investor in advance, to simplify the paper and address the most important sources of option profit, we assume that r and d are known in advance.

[4] Pure Option Profit

- **Difference in profit from an efficiently priced call and the profit from a forward on the same underlying asset maturing at t:**

$$C_t - V - [S_t - S(r/d)^t]$$

- **For example, if the option expires after time t with striking price $S(r/d)^t$, then since $C_t = \max[0, S_t - S(r/d)^t]$, the pure option profit is:**

$$\text{if } S_t \geq S(r/d)^t: \quad -V < 0$$

$$\text{if } S_t < S(r/d)^t: \quad -V - [S_t - S(r/d)^t]$$

which is positive whenever $S_t < S(r/d)^t - V$, illustrating the idea that compared to a forward, the advantage of a call is that it places a floor on losses.

The final profit component attempts to capture the portion of the profit that is due to the use of efficiently priced options rather than forwards, which we call the **pure option profit**:

$$C_t - V - [S_t - S(r/d)^t]$$

With striking price K, if the option expires after time t, then since $C_t = \max[0, S_t - K]$, the pure option profit becomes:

$$\max[0, S_t - K] - V - [S_t - S(r/d)^t]$$

To interpret this, suppose the option is initially "at-the-money" in the sense that $K = S(r/d)^t$. If S_t ends up less than K, the pure option profit is $-V - [S_t - S(r/d)^t]$; and if S_t ends up greater than K, the pure option profit is $-V$.

In the latter "in-the-money" case, the pure option profit is negative because the investor would have been better off simply with a forward contract. As compensation, in the former "out-of-the-money" case, if S_t ends up less that K by more than V, the pure option profit will be positive since the investor does better than had he instead purchased a forward contract.

Using μ to represent the annualized expected ex-dividend return of the underlying asset, we can write $E(S_t) = S\mu^t$.

Consider two extreme cases. First, suppose the option is certain to finish in-the-money. In that case, $V = Sd^{-t} - Kr^{-t}$. Then the expected pure option profit is:

$$[S\mu^t - S(r/d)^t] - [Sd^{-t} - S(r/d)^t r^{-t}] - [S\mu^t - S(r/d)^t] = 0$$

In this case, whether or not the underlying asset price S is determined in an efficient market, since it cancels out of the expression, the expected pure option profit is zero.

In the other extreme, suppose the option starts out so much out-of-the-money that it is certain to finish out-of-the-money. In that case, since both $\max[0, S_t - K]$ and V equal zero, the pure option profit is $-[S_t - S(r/d)^t]$. This exactly offsets the asset profit leaving the second basic component of profit (profit from subsequent fortuitous changes or mispricing of the underlying asset) equal to zero. In this extreme case, even though the investor may have bought a call on an underpriced asset, it did not help him since the call he bought was too far out-of-the-money.

Adding-Up Constraint

$C_t - C =$

[1] $C(s) - C$	(volatility profit)
+ [2] $V - C(s)$	(formula profit)
+ [3] $S_t - S(r/d)^t$	(asset profit)
+ [4] $C_t - V - [S_t - S(r/d)^t]$	(pure option profit)

For efficiently priced options: $C = E[V]$

If options are also correctly benchmarked:

$E[C(s) - C] = E[V - C(s)] = 0$

For efficiently priced assets: $PV[S_t - S(r/d)^t] = 0$

The attached picture summarizes our decomposition of option profit. The profit of the option equals:

[1]		[2]		[3]		[4]
volatility profit	+	**formula profit**	+	**asset profit**	+	**pure option profit**

For <u>efficiently priced options</u>, on average **[1]** + **[2]** = 0, whether or not the benchmark formula closely approximates the formula the market uses to value options.

If, in addition to the option being efficiently priced, the benchmark formula captures the market's approach to option valuation, then on average **[1]** is zero and on average **[2]** is zero.

On any one option investment, due to sampling error, the realized volatility will be different than the true population volatility, so $C(s) \neq C$, but on average (ignoring slight nonlinearity effects), $C(s) = C$.

If the benchmark formula approximates the market's formula, for <u>inefficiently priced options</u>, the magnitude of **[1]** should on average indicate the value of an investor's ability to make superior volatility forecasts, and the magnitude of **[2]** should isolate on average the extent an investor is using a superior formula to the market's formula.

For <u>efficiently priced assets</u>, the present value of **[3]** should be zero while its magnitude indicates appropriate compensation for bearing risk. For <u>inefficiently priced assets</u>, its present value will not be zero and measures the extent an investor is skillful in selecting options with inefficiently priced underlying assets.

For <u>efficiently priced assets</u>, the present value of **[4]** equals zero.

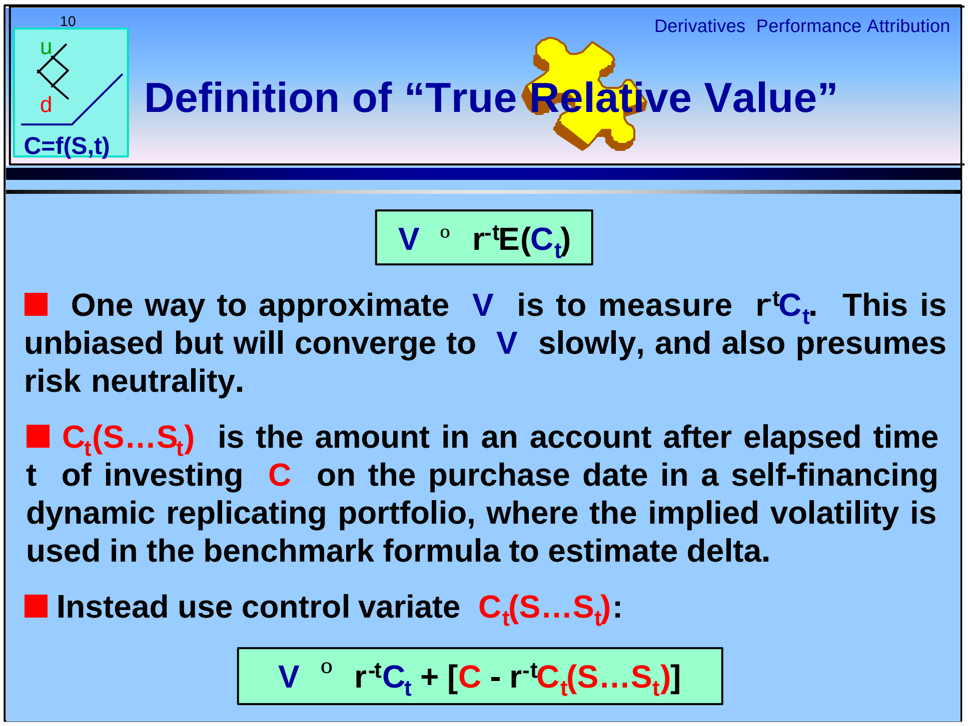

III. Estimating the "True Relative Value" of the Option

The trick to separating option mispricing profit from profit due to fortuitous underlying asset price movements or asset mispricing is to find some way to estimate the "true relative value" V of the option.

In an efficient market for the asset, ignoring risk aversion, $V \equiv r^{-t}E(C_t)$. So one way to estimate V would simply be to observe a single $r^{-t}C_t$. Any one $r^{-t}C_t$ would overstate or understate V. But over many realizations of sample paths for the underlying asset, V would be approximated by the average outcome for $r^{-t}C_t$ over these paths.

One problem with this simple Monte Carlo approach is that it may take a very large number of realized sample paths for the average outcome of $r^{-t}C_t$ to come close to V. We want to be able to measure the performance of the investor more quickly. A way to speed up the process is to use a control variate.

Let $C_t(S\ldots S_t)$ be the amount in an account after elapsed time t from investing C on the purchase date in a portfolio comprised of the underlying asset and cash, and then subsequently attempting to replicate dynamically (with self-financing) the payoff of a call option with the same time-to-expiration and striking price as the purchased option. The option is replicated by using the benchmark formula together with the implied volatility to estimate the delta. In practice, it may be sufficient in most markets to assume daily rebalancing at the close to correct the delta. The realized difference $C_t - C_t(S\ldots S_t)$ measures the extent by which the benchmark formula, using the implied volatility, fails to replicate the option.

The natural control variate is the value of the replicating portfolio $C_t(S\ldots S_t)$, so that

$$V \equiv C + r^{-t}[C_t - C_t(S\ldots S_t)]$$

Although, I refer to $C_t(S\ldots S_t)$ as a "control variate," in real life application, it is simply the result of running a parallel paper replicating portfolio for each option position under analysis.

Note: One might have thought that a better replicating strategy for our purpose would have been to base the strategy on the realized volatility along the sample path, rather than the beginning implied volatility. Unfortunately, this can lead to biased measures of value. For example, suppose the true stochastic process implies that realized volatility is inversely correlated with asset price. Then knowing at the beginning that the realized volatility will be high leads one to expect a decline in the asset price. This information could then be used to almost assure that $C_t(S\ldots S_t) > C_t$, along every path with the given realized volatility.

C=f(S,t)

The Monte Carlo Logic

$$V \equiv r^{-t}C_t + [C - r^{-t}C_t(S \ldots S_t)]$$

simplifying : $V^* \equiv r^{-t}C_t$ **and** $C^* \equiv r^{-t}C_t(S \ldots S_t)$

$$Var(V) = Var(V^*) + Var(C^*) - 2\,Cov(V^*, C^*)$$

Suppose that $Var(V^*) = Var(C^*)$**, then**

$$Var(V) = 2\,[Var\ V^*]\,[1 - r(V^*, C^*)]$$

Suppose that $r(V^*, C^*) = .9$ **(a fair benchmark), then**

$$Var(V) = .2[Var(V^*)]$$

To see why the proposed measure $V \equiv r^{-t}C_t + [C - r^{-t}C_t(S \ldots S_t)]$ compares favorably to $V \equiv r^{-t}C_t$, simplify the notation and set

$$V^* \equiv r^{-t}C_t$$

$$C^* \equiv r^{-t}C_t(S \ldots S_t)]$$

So we have:

$$V = V^* + [C - C^*]$$

Taking variances of both sides:

$$Var(V) = Var(V^*) + Var(C^*) - 2\,Cov(V^*,C^*)$$

To get a rough idea of the magnitudes involved, to a first approximation we could well expect

$$Var(V^*) = Var(C^*)$$

After all, in a Black-Scholes world of continuous trading, by using Black-Scholes also as the benchmark formula, V^* is exactly equal to C^*. In that case, this approximation is exact.

Then, we can simplify to get:

$$Var(V) = 2\,Var(V^*)\,[1 - \rho(V^*,C^*)]$$

where $\rho(V^*, C^*)$ is the correlation coefficient measuring association between V^* and C^*.

Experience with replicating strategies suggests that if the benchmark formula is Black-Scholes and the volatility is implied, the benchmark-replicating strategy C^* might be expected to produce outcomes highly correlated with the realized option payoff. In light of this, assume that $\rho(V^*,C^*) = .9$. In that case,

$$Var(V) = .2\,Var(V^*)$$

This shows why, for this purpose, we want to use a benchmark formula that has the highest $\rho(V^*,C^*)$.

Since $Var(V)$ is much smaller than $Var(V^*)$, we can afford to use a much smaller sample to estimate V than V^*. This translates into being able to identify ability to select mispriced options much more quickly.

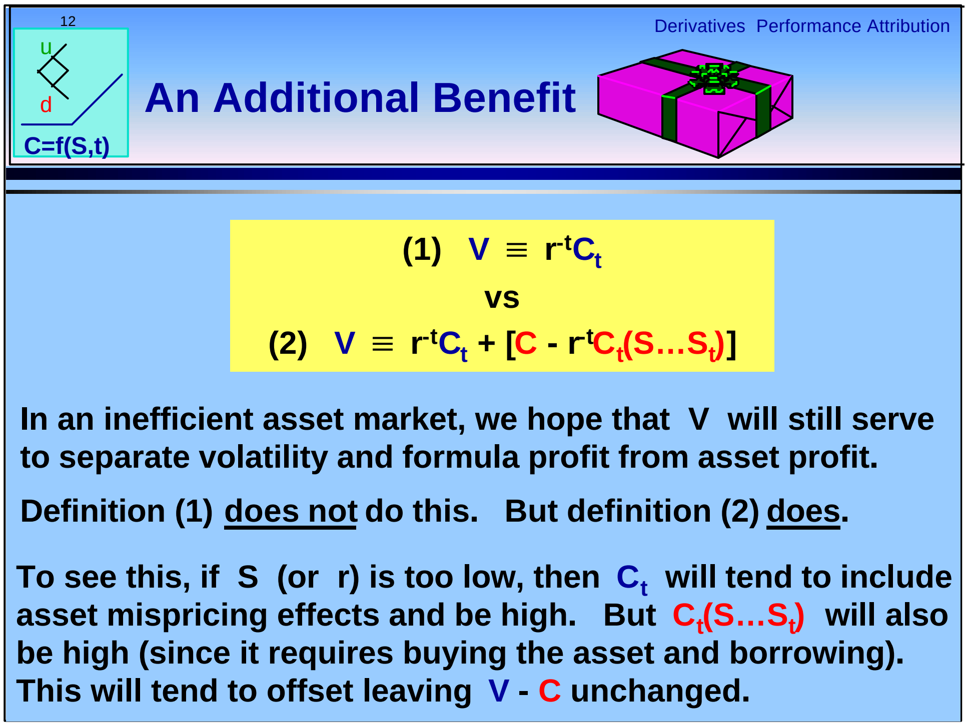

Recall that we have defined the "true relative value" as:

$$V \equiv r^{-t}C_t + [C - r^{-t}C_t(S\ldots S_t)]$$

What happens to V if the market misprices the underlying asset itself at the purchase of the option, but we assume the asset is correctly priced at the end of the assessment period? In that case, had we simply defined

$$V \equiv r^{-t}C_t$$

we would have a problem. Then, in a risk-neutral market, since on average $V = r^{-t}E(C_t)$ and C_t would reflect the correct asset price (which is the realized asset price at expiration), V would then be the "true <u>absolute</u> value" of the option inclusive of asset mispricing. This would be unfortunate for our purposes since we are hoping to use V to separate the effects of option mispricing (volatility and formula profit) from asset mispricing (asset profit).

Fortunately, the Monte Carlo control-variate approach we have proposed, in addition to reducing standard errors, also holds out the promise of correcting this problem. To see this, suppose that on the option purchase date, the underlying asset is underpriced. What this really means is that the expected risk-neutral return of the asset is greater than the riskless return r available in the market. This means that the dynamic replicating portfolio strategy will do better than we would have expected in an efficient asset market. To replicate a call, we always need to be long the underlying asset partially financed by borrowing. Since the asset will appreciate faster than it should in an efficient market and we are long the asset (or alternatively, since the riskless interest rate is lower than it should be in an efficient market and we are borrowing), $r^{-t}C_t(S\ldots S_t)$ will tend to exceed C. In fact, it will tend to exceed C by the amount that $r^{t}C_t$ is higher than it should be because the asset appreciates faster than it would in an efficient market.

C, of course, will be low because of the asset underpricing. Because of this offset, V will tend to be low by the same amount, leaving the difference $V - C$ unaffected, just as we would wish. So as long as we measure V using the dynamic replicating strategy as the control variate, V should live up to its billing as the "true <u>relative</u> value" of the option.

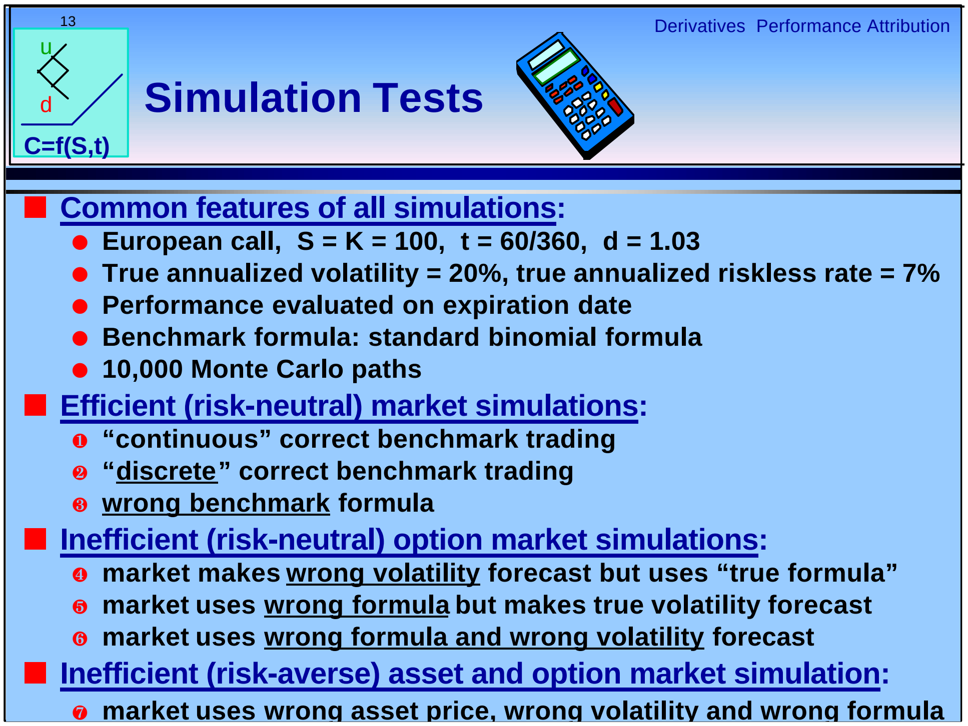

IV. Simulation Results

Ascertaining from observed performance whether or not an investor has earned a sufficiently high rate of return to be considered skillful to a high probability is difficult, even within an investor's lifetime. So a key challenge of performance measurement and attribution is to get the job done quickly.

To check this out for our proposed attribution, we will run several simulations. In each case, we will assume an investor has purchased a European call with underlying asset price of 100, a striking price of 100, and 60 days-to-expiration in a 360-day year. The "true" annualized volatility is assumed to be 20%, the "true" annualized riskless rate is assumed to be 7%, and the annualized payout rate 3%. In each case, we will examine performance attribution on the expiration date. Our benchmark formula will in every case be the standard binomial option pricing model [Cox, Ross and Rubinstein 1979]. Each reported simulation will use 10,000 Monte Carlo paths. The simulations differ as follows:

❶ **efficient (risk-neutral) market with "continuous" and correct benchmark trading:** market knows the true population volatility and the correct formula; benchmark formula also correct with benchmark trading to target delta taking place at every binomial move

❷ **efficient (risk-neutral) market with "discrete" but correct benchmark trading:** like ❶ except benchmark trading takes place after more than one binomial move

❸ **efficient (risk-neutral) market with incorrect benchmark formula:** like ❶ except benchmark formula wrong since it underestimates leptokurtosis and underestimates left-skewness

❹ **inefficient (risk-neutral) option market (wrong volatility but correct formula):** market knows the correct formula but misprices the call because it underestimates or overestimates true population volatility

❺ **inefficient (risk-neutral) option market (correct volatility but wrong formula):** market knows the true population volatility but is using a formula that overprices the call

❻ **inefficient (risk-neutral) option market (wrong volatility and formula):** market misprices the call because it uses both the wrong volatility and the wrong formula

❼ **inefficient (risk-averse) asset and option markets:** market misprices the call because it uses the wrong underlying asset price, wrong volatility forecast and wrong option valuation formula.

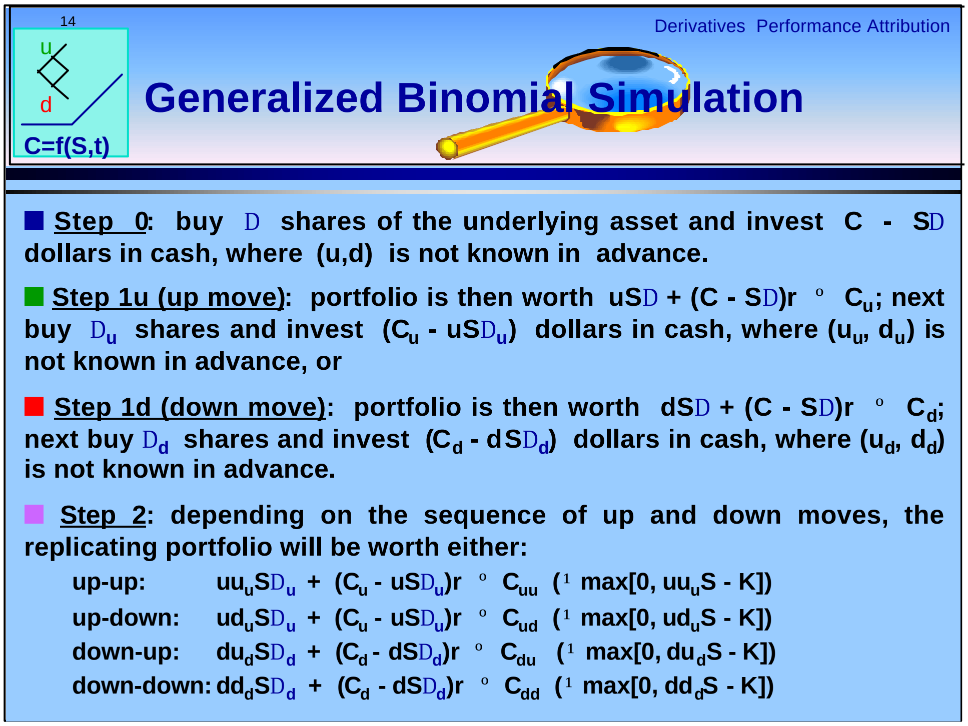

For simulation purposes, the underlying asset price is assumed to have a stochastic process conforming to a generalized but recombining binomial tree. Consider an example in which the call expires at the end of the second move. The benchmark strategy is implemented as follows:

Step 0: buy D shares of the underlying asset and invest C - SD dollars in cash.

Step 1u (up move): portfolio is then worth $uSD + (C - SD)r \equiv C_u$; next buy D_u shares and invest $(C_u - uSD_u)$ dollars in cash, or

Step 1d (down move): portfolio is then worth $dSD + (C - SD)r \equiv C_d$; next buy D_d shares and invest $(C_d - dSD_d)$ dollars in cash.

Step 2: depending on the sequence of up and down moves, the replicating portfolio will be worth either:

up-up: $uu_uSD_u + (C_u - uSD_u)r \equiv C_{uu}$ up-down: $ud_uSD_u + (C_u - uSD_u)r \equiv C_{ud}$

down-up: $du_dSD_d + (C_d - dSD_d)r \equiv C_{du}$ down-down: $dd_dSD_d + (C_d - dSD_d)r \equiv C_{dd}$

Here u_d, for example, is the up move in the second step conditional on having had a down move in the first step. Although we will assume a recombining tree so that $ud_u = du_d$, we do not usually assume that $d_u = d_d$ or that $u_d = u_u$.

Note that the benchmark strategy, like the call, requires an initial investment of the market price and is self-financing.

In the usual development of the binomial option pricing formula, it is assumed that D, D_u and D_d are selected knowing in advance what u, d, u_u, u_d, d_u and d_d are. Moreover, it is assumed that the boundary conditions C_{uu} = max[0, uuS - K], $C_{ud} = C_{du}$ = max[0, udS - K] and C_{dd} = max[0, ddS = K] are satisfied.

However, the benchmark formula is assumed to be implemented by only guessing and without knowing the underlying stochastic process. That is, D , D_u and D_d must be chosen without knowing in advance what the sizes of the subsequent up and down moves will be. As a result, it can be shown that the resulting payoff of the benchmark strategy will be incorrect as well as path-dependent (assuming it it not static: $D = D_u = D_d$). That is, in our 2-step example, not only will it generally be the case that $C_{uu} \neq$ max[0, uuS - K], $C_{ud} \neq$ max[0, udS - K], $C_{du} \neq$ max[0, duS - K], and $C_{dd} \neq$ max[0, ddS - K], but also that $C_{ud} \neq C_{du}$ (path-dependence).

Of course, as we have argued, we will try our best to choose a benchmark strategy (D, D_u, D_d) where the payoff at expiration is as close as possible to a call, but whatever we do, we cannot expect to replicate the call perfectly.

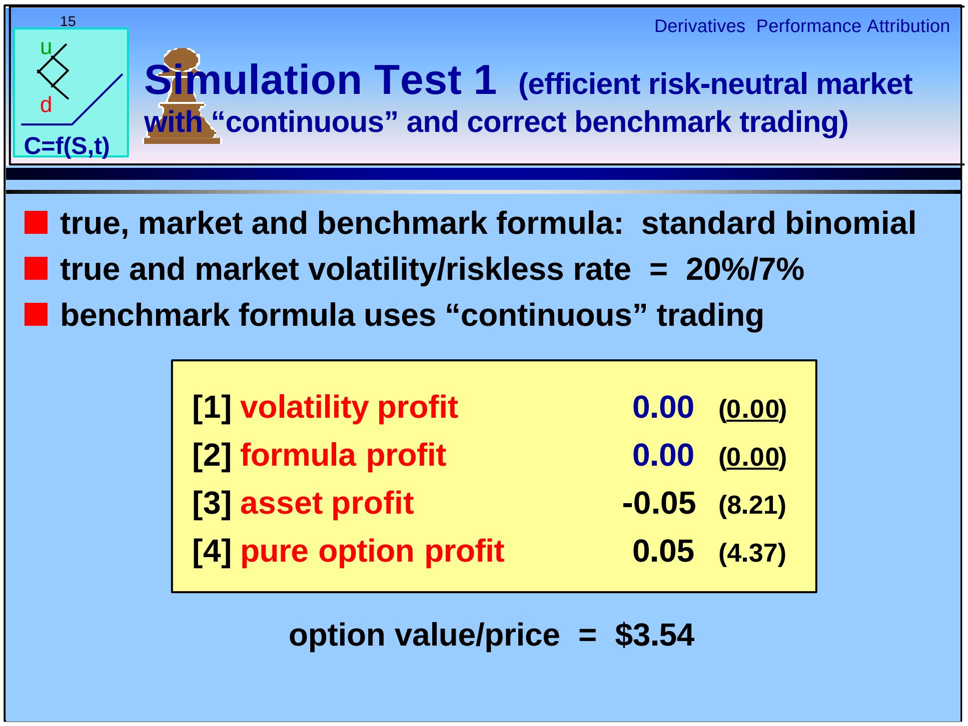

The first simulation is designed to test whether, if both the asset and options markets are efficient, the call will show up as properly priced. In particular, volatility, formula and asset profits should each be zero. In this simulation, the benchmark, market's and true formulas are the standard binomial option pricing model estimated using the true population volatility and true riskless rate.

To implement the simulation, we construct a 60-move (one move per day) standard binomial tree based on an annualized volatility of 20%. So that the up move $u = e^{2\sqrt{(.1667/60)}}$, the down move $d = 1/u$ and the up-move risk-neutral probability is $p = .502385$. With efficient risk-neutral markets, p is the same as the consensus market subjective probability. In this situation, the call is priced by the market correctly at $3.54. We then select 10,000 sample paths through the tree. To construct a single path, for each move in the path, we select a number from 1 to 100,000,000 under a uniform distribution. If p times 100,000,000 is greater than this number, an up move occurs; otherwise a down move occurs. As we move along each path, at each node we use the benchmark formula to calculate the benchmark delta (number of shares in the replicating portfolio) at that node. With these deltas in hand along a path, we can calculate the expiration-date value of the replicating portfolio, where it is assumed to start with an initial investment equal to the beginning call market price ($3.54) and be self-financing thereafter.

The attached picture lists the average component profits (with standard errors in parentheses).

In this efficient risk-neutral market environment, we expect the asset profit to be zero, with a small error due to the finite Monte Carlo sample. The pure option profit should also be zero. We also expect the volatility profit and the formula profit to be zero and have zero standard error since the benchmark formula works perfectly. As we see, we need not worry that in this efficient market environment, our procedure will will end up attributing volatility forecasting or formula selection skill to any investor.

However, it will not be as easy to decide if the investor has skill in selecting the right underlying asset. Although the average asset profit is almost zero, its standard error is $8.21, more than twice the cost of the option. Even if we can observe 100 such independent investments, the standard error of the average of the 100 asset profits will be about $8.21/\sqrt{100} = \$0.821$, still significant. Here, as elsewhere, it is difficult to distinguish luck from skill in selecting assets.

16 Derivatives Performance Attribution

Simulation Test 2 (efficient risk-neutral market with "discrete" but correct benchmark trading)

- true, market and benchmark formula: standard binomial
- true and market volatility/riskless rate = 20%/7%
- benchmark formula uses "discrete" trading (once a day)

	move every 1/2 day	move every 1/8 day
[1] volatility profit	-0.005 (0.21)	-0.011 (0.28)
[2] formula profit	0.008 (0.15)	0.013 (0.17)
[3] asset profit	-0.10 (8.15)	0.08 (8.28)
[4] pure option profit	0.06 (4.14)	0.02 (4.05)

option value/price = $3.54

The next simulation checks the robustness of the benchmark replicating strategy to non-continuous observations. In the previous simulation, it was assumed that every time the underlying asset made a binomial move, the benchmark strategy revised its position based on a freshly calculated delta. This led to zero volatility sampling error since all paths through a standard binomial tree have the same sample path volatility.

More realistically, we assume now that we only observe the asset price at "discrete" intervals, that is, after more than one binomial move. Moreover, we only calculate the realized volatility based on this sample of observations. Since we are not really trading, our reason for this discrete revision is not trading costs, so we can indeed afford to sample the asset price quite often, perhaps several times intra-day. Nonetheless, it is impractical to sample truly continuously. We also do not want our sample to be significantly influenced by bid-ask bounce; and we also might want to force discrete sampling on our simulation to crudely capture mild jump risk.

So we revised the simulation by sampling every half day for a total of 120 moves, every quarter day for a total of 240 moves, and every 1/8 day for a total of 480 moves, in every case covering the 60 days to expiration. But in each case, we assumed that the benchmark replicating portfolio was revised only at the end of each day and the realized volatility was calculated based only on end-of-day prices. The component average profits (with standard errors in parentheses) can be found in the attached picture. The simulation with a move every 1/4th day shows almost the same standard errors for volatility and formula profit as the results of moves every 1/8th day.

In contrast to our earlier simulation, we now have positive standard errors for volatility profit and for formula profit. If we allow ourselves at least a sample of 100 independent option investments, these standard errors are quite small (about 1/10th of the errors indicated), so that our techniques are quite robust to discrete-time benchmark calculations. For example, formula profit is estimated with a standard error of about two cents.

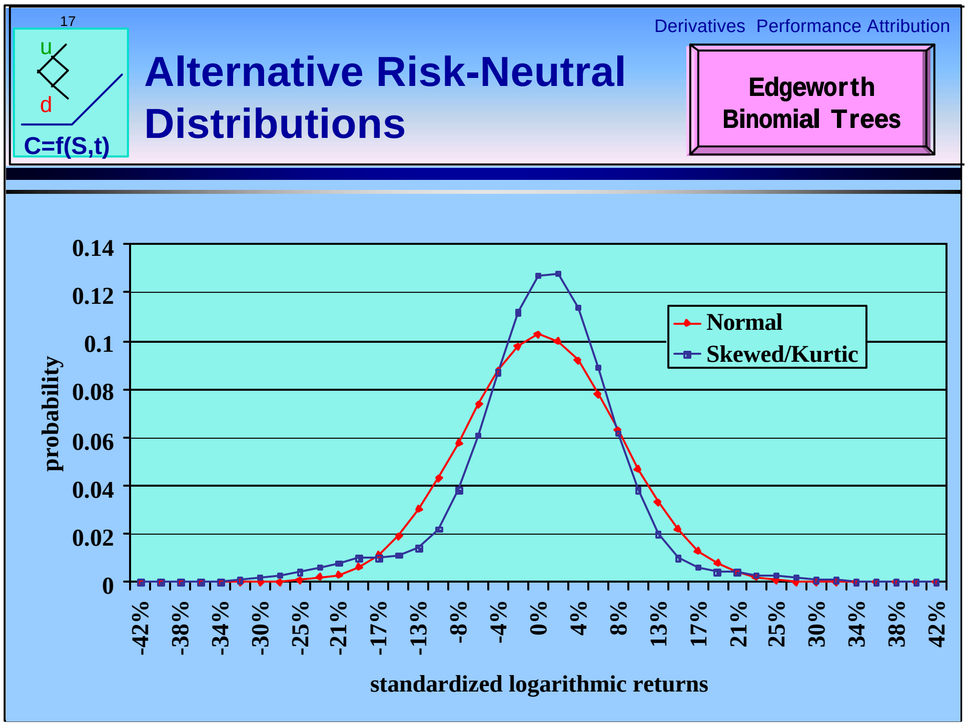

In several of the remaining simulations, the "true formula" is the risk-neutral discounted value of the attached left-skewed and leptokurtic standardized distribution of annualized logarithmic returns. The "normal" distribution describes logarithmic returns which are 60-move standard binomial (with a skewness of 0 and a kurtosis of 3). The "skewed/kurtic" distribution describes logarithmic returns generated by an Edgeworth expansion of the 60-move standard binomial with a skewness of -.398 and a kurtosis of 4.86.*

This was chosen to match the option-implied distribution from S&P 500 Index options in the years after the 1987 stock market crash. In these simulations, the true stochastic process is the unique implied binomial tree consistent with this expiration-date distribution following techniques developed in Rubinstein (1994).

By contrast, in all but the next simulation, the "market's formula" is the risk-neutral discounted value of the above distribution with a skewness of 0 and kurtosis of 3. This will allow us to examine how the procedures proposed for performance attribution work when the option market is inefficient in the sense that it uses the wrong formula (essentially ignoring the left-skewness and leptokurtosis of the true expiration-date distribution when it prices options).

---

*See Rubinstein (1998) for development of techniques to generate unimodal standardized distributions with prespecified third and fourth central moments by transforming a standard binomial distribution via an Edgeworth expansion.

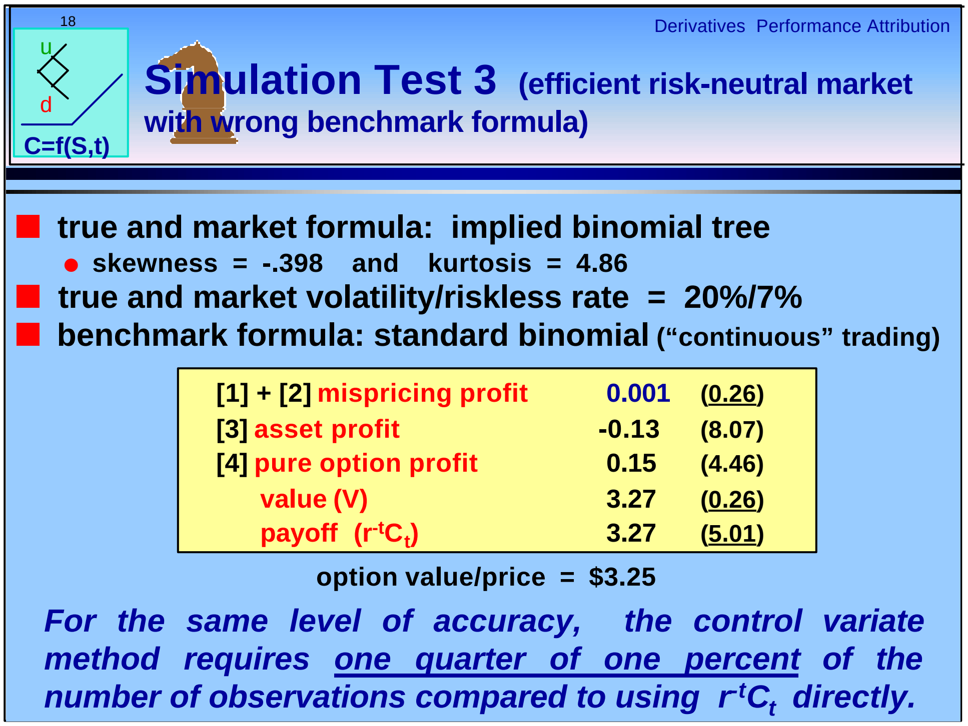

Realistically, we do not have the luxury of being able to set our benchmark formula to the one which actually determines option values and prices. To test the robustness of the benchmark formula to this misspecification, we continue to assume a risk-neutral efficient market and that the benchmark is the standard binomial model; but now suppose that the true stochastic process is derived from an implied binomial tree where the expiration-date distribution has more leptokurtosis and more left-skewness than allowed by a standard (constant move size) binomial tree.

In this efficient market, the option mispricing profit is zero. With this now inferior dynamic replicating strategy, as expected, the average mispricing profit is virtually zero. Because we are using the wrong benchmark formula, however, the standard error (compared to our first simulation) is now positive, but nonetheless not large. Again aggregating across 100 independent option investments produces a standard error of about 3 cents.

We also can directly measure the success of our control variate approach in reducing the standard error in estimating the "true relative value" of the option. The standard error of the payoff r^tC_t is $5.01 for a single option, while the standard deviation of the value V for a single option estimated with the control variate is only $0.26, representing about a 95% improvement. *To reduce the standard error to about 2.5 cents, would require 100 independent observations for V (.26/Ö100) and 40,000 independent observations for r^tC_t (5.01/Ö40000). Thus the control variate method requires one quarter of one percent of the number of observations compared to using $r^{-t}C_t$ directly.*

For this simulation, we have not reported separate results for volatility and formula profits since this decomposition only works properly if the benchmark formula we are using is close to the market's formula. Since, by assumption for this simulation, these are not the same, this attribution will be incorrect.

19 Derivatives Performance Attribution

C=f(S,t)

Simulation Test 4 (inefficient risk-neutral option market because market uses wrong volatility)

- true, market and benchmark formula: standard binomial
- true and market riskless rate = 7%
- true volatility = 20% market volatility = 15%/25%

	market vol = 15%		market vol = 25%	
[1] volatility profit	0.801	(0.00)	-0.802	(0.00)
[2] formula profit	0.007	(0.32)	-0.004	(0.30)
[3] asset profit	-0.04	(8.12)	-0.09	(8.30)
[4] pure option profit	0.00	(4.04)	0.11	(4.07)
value (V)	3.55	(0.33)	3.55	(0.29)
payoff (r^tC_t)	3.48	(5.08)	3.53	(5.19)

option value = $3.54 and option price = $2.74/$4.34

In the attached simulation results, the market values the option based on an incorrect estimate of volatility. Otherwise, the market makes no mistakes. In particular, it correctly values the asset and uses the true option valuation formula. With the true volatility, the option is worth $3.54; but the market errs in one case by underpricing the option by $0.80, and in another case overpricing the option by $0.80. In this case, we would hope that our attribution procedures would calculate an average option mispricing profit of plus or minus 80 cents.

Since we also assume that the benchmark formula is the same as the "market's formula", we would also hope that the procedure would allocate the entire mispricing profit to volatility profit and none of it to formula profit.

As the attached picture shows, we are right on target, with the added plus of low standard errors for volatility and formula profit.

Simulation Test 5 (inefficient risk-neutral option market because market uses wrong formula)

- true formula: implied binomial tree
  - skewness = -.398 and kurtosis = 4.86
- true and market volatility/riskless rate = 20%/7%
- benchmark and market formula: standard binomial

[1] volatility profit	-0.000	(0.53)
[2] formula profit	-0.284	(0.27)
[3] asset profit	-0.24	(7.99)
[4] pure option profit	0.00	(4.31)
value (V)	3.26	(0.33)
payoff ($r^{-t}C_t$)	3.25	(4.96)

option value = $3.25 and option price = $3.54

In the prior simulation, we assumed that the market estimated the volatility incorrectly but got the formula right. In this case, we assume the reverse. Using the standard binomial formula, because the market fails to account for the left-skewness and leptokurtosis of the true underlying stochastic process, it overprices the option by about 29 cents.

We would hope that the our procedure would not only capture this overpricing but also attribute it correctly all to formula profit and none to volatility profit. Again the attribution procedure comes through, with low standard errors to boot.

C=f(S,t)

Simulation Test 6 (inefficient risk-neutral option market because market uses wrong volatility/formula)

- **true formula: implied binomial tree**
  - **skewness = -.398 and kurtosis = 4.86**
- **true and market riskless rate = 7%**
- **true volatility = 20% market volatility =15%/25%**
- **benchmark and market formula: standard binomial**

	market vol = 15%		market vol = 25%	
[1] volatility profit	0.793	(0.54)	-0.802	(0.53)
[2] formula profit	-0.282	(0.53)	-0.287	(0.36)
[3] asset profit	-0.13	(8.03)	-0.06	(8.19)
[4] pure option profit	0.02	(4.50)	0.06	(4.21)
value (V)	3.25	(0.25)	3.25	(0.54)
payoff ($r^{-t}C_t$)	3.25	(4.90)	3.34	(5.10)

option value = $3.25 and option price = $2.74/$4.34

Now we consider simultaneously errors in the market's estimate of volatility and use of the wrong formula. The market continues to use the standard binomial model, failing to account for non-normality and in addition under- or over-estimates the volatility. In this case, the option value is $3.25 (as in the prior simulation). We know that the market has made an underpricing error of 29 cents due to use of the wrong formula. In addition, by underestimating (overestimating) the volatility, the market underprices (overprices) the option by an additional 80 cents.

As the attached picture shows, again the attribution procedure works almost perfectly on average with low standard errors. Aggregating over 100 independent option investments, the standard errors for volatility and formula profit are from three to five cents.

22 Derivatives Performance Attribution

C=f(S,t)

Simulation Test 7 (inefficient risk-averse market because market uses wrong asset price, volatility, formula)

- true formula: implied binomial tree
  - skewness = -.398 and kurtosis = 4.86
- true volatility = 20% market volatility = 25%
- true riskless rate = 7% market riskless rate = 5%/9%
- benchmark and market formula: standard binomial

	riskless rate = 5%		riskless rate = 9%	
[1] volatility profit	-0.810	(0.54)	-0.794	(0.54)
[2] formula profit	-0.286	(0.24)	-0.281	(0.28)
[3] asset profit	0.38	(8.03)	-0.36	(8.04)
[4] pure option profit	-0.14	(4.14)	0.20	(4.21)
value (V)	3.09	(0.50)	3.42	(0.58)
payoff ($r^{-t}C_t$)	3.29	(4.99)	3.22	(4.89)

option value = \$3.25 (\$3.07/\$3.44) and option price = \$4.19/\$4.50

This final simulation is like the ending fireworks display on the Fourth of July: we pull out all the stops. To the errors made by the market in the prior simulation, we add potential mispricing of the underlying asset itself. Alternatively, the simulation can be interpreted as allowing consensus market risk aversion (where before we assumed risk-neutrality).

This modification can be introduced into the simulation quite easily by allowing for the riskless return available in the market to be more or less than the riskless return built into the true stochastic process of the underlying asset. We now have two riskless returns. One, the "true" riskless return, is used only to determine the implied binomial tree describing the actual behavior of the underlying asset. As before, we continue to assume this is 7%. So at each node in the implied binomial tree, the up and down move sizes for the next move are chosen so that the annualized risk-neutral expected return of the underlying asset equals 7%. The market, however, does not understand this. It believes that the riskless return is 5% or 9%, and sets the interest rate available to investors accordingly. In particular, this is the interest rate paid on borrowing in the benchmark dynamic replicating strategy. It is also the interest rate used to measure the various components of performance (since the agency measuring the attribution also doesn't know any better either).

For example, consider the effects of this on asset profit: $S_t - S(r/d)^t$. In one simulation, the r in this formula is incorrectly set by the market at 5%. But the asset price actually appreciates at a risk-neutral expectation of 7%. In a risk-neutral market, we would interpret this as a 2% positive mispricing "alpha" (in which case, the option value is \$3.25). In a risk-averse market, we could interpret this as compensation for bearing risk (in which case, the option value is \$3.07). It lies beyond the scope of this paper to make this distinction. This is the critical issue of asset performance measurement. But under either interpretation, this should show up as positive asset profit approximately equal to $100(1.07/1.03)^{1/6} - 100(1.05/1.03)^{1/6} = 0.32$.

So bottom line we hope that our simulation would as in the prior simulation attribute about -80 cents to volatility profit, and -29 cents to formula profit. Hence, we hope it will show 32 cents asset profit. As the attached picture shows, even with this more complex economy, it continues to sort through the results and attribute the correct volatility and formula profits with low standard errors. In addition, it comes reasonably close on average to capturing the asset profit, although the large standard error makes this difficult to rely upon.

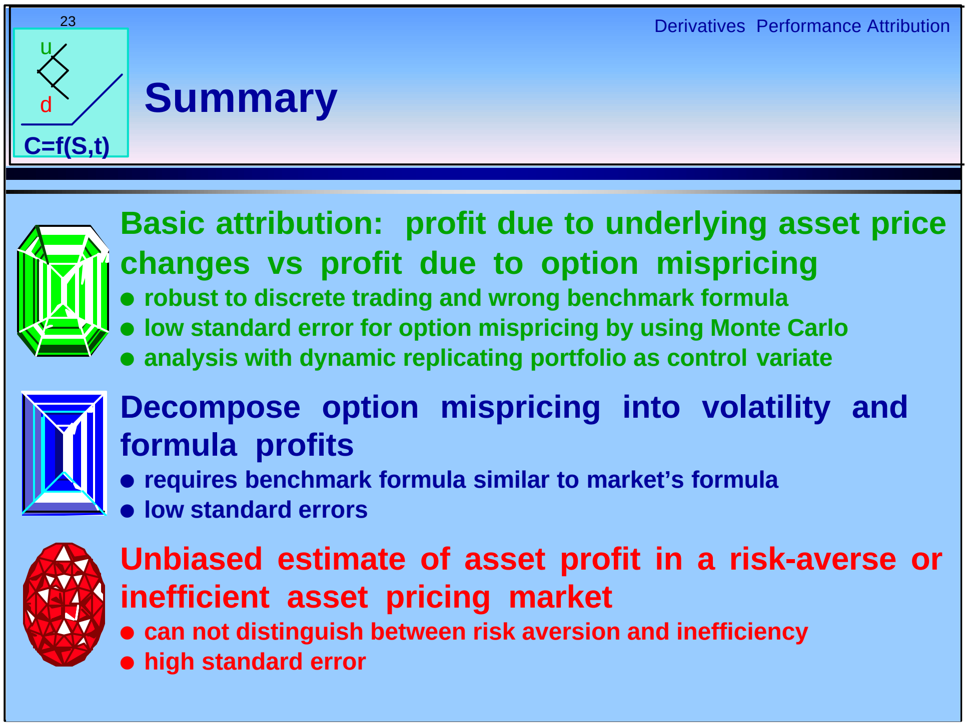

V. Summary

We have decomposed the realized profit from an option into two principal components: (1) the mispricing of the option at the time of purchase and (2) the profit from subsequent fortuitous changes or mispricing of its underlying asset. We found that the first is relatively easy to isolate, requiring in the simulations about one-quarter of one percent of the number of observations needed for the second to achieve the same level of accuracy. The trick to this variance reduction is to estimate the "true relative value" of the option by using the results of a dynamic replicating strategy as the control variate. In addition, this has the further benefit of separating out the profit from mispricing of the underlying asset. The results apply generally irrespective of market risk-aversion.

To separate the mispricing of the option at the time of purchase into volatility and formula profit requires knowledge of the formula used by the market to price options. But if this formula is known, then this separation can be accomplished with few observations.

Although this paper illustrates the attribution approach for a single European call, it can be also be used for American options. To apply it to portfolios of derivatives, it will be necessary to consider the correlation of their prices to estimate standard errors.

Bibliography

Cox, J.C., S.A. Ross and M. Rubinstein, "Option Pricing: A Simplified Approach," *Journal of Financial Economics* 7, No. 3 (September 1979), pp. 229-263.

Galai, D., "The Components of the Return from Hedging Options against Stocks," *Journal of Business* 56, No. 1 (January 1983), pp. 45-54.

Rubinstein, M. "Implied Binomial Trees," *Journal of Finance* 49, No. 3 (July 1994), pp. 771-818.

Rubinstein, M. "Edgeworth Binomial Trees," *Journal of Derivatives* 5, No. 3 (Spring 1998), pp. 20-27.