

Market RISK

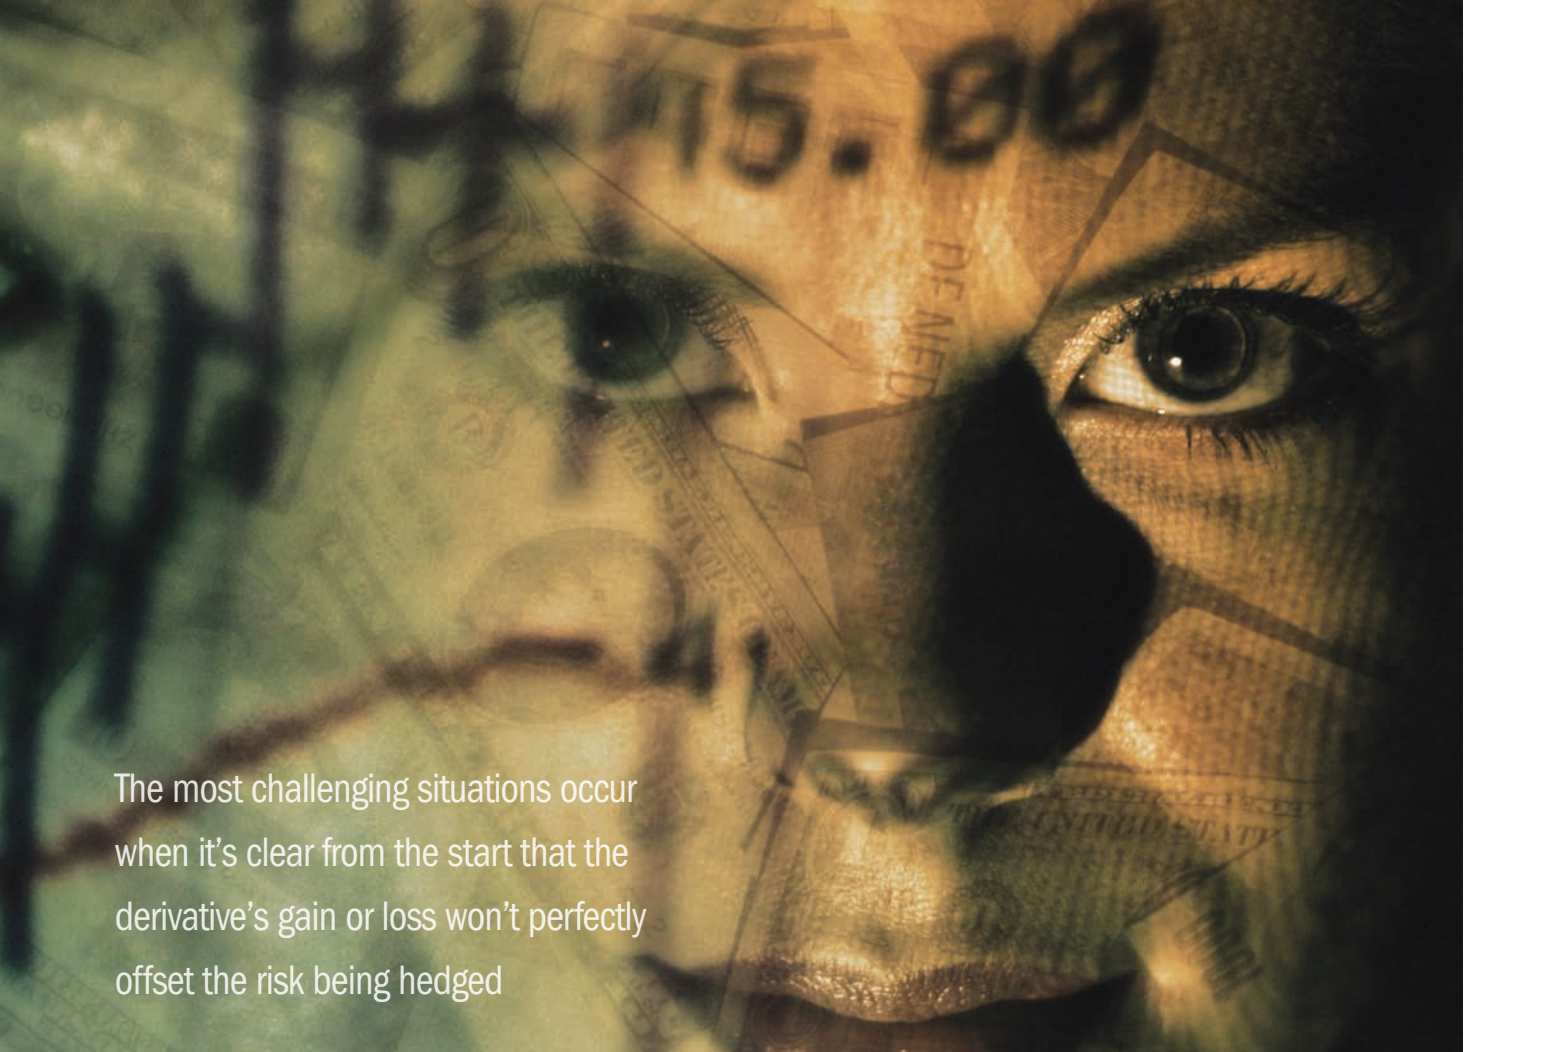
*What hedgers (and auditors)
should know about regression*

Ira G. Kawaller



It's hardly controversial to expect financial statements to provide accurate and meaningful information about company performance, but even the standard-setters themselves seem to recognize that something's amiss in connection with the way derivatives and hedging transactions are reported. Since its issuance in June of 1998, the Financial Accounting Standard No. 133 (FAS 133) has been widely cited as being one of the most difficult standards to apply; but now, even after two amendments, the FASB appears ready to make some significant and some not-so-significant changes to this accounting treatment.

Under both the current and prospective rules, the FASB provides for two distinct types of accounting—regular accounting and “special hedge accounting.” Regular accounting is easy: Derivatives are recorded as assets or liabilities, carried at fair market value, with gains or losses recorded in current income. The problem with this treatment, however, is that it might not result in an accounting presentation that reflects the objective for using the derivative. In particular, when derivatives are used as hedging instruments, most hedgers prefer that their financial statements reflect their hedging objective; that is, they would hope to recognize earnings impacts of hedging derivatives concurrently with the earnings impacts associated with the risks being hedged.



The most challenging situations occur when it's clear from the start that the derivative's gain or loss won't perfectly offset the risk being hedged

Special hedge accounting achieves this concurrent pairing of the two income effects, while regular accounting generally doesn't. Thus, reported earnings would likely be more volatile under regular accounting and less volatile under hedge accounting. And because analysts and investors tend to reward companies that have lower income volatility with higher valuations (all else being equal), applying hedge accounting becomes a priority at the highest level of corporate governance. Hedge accounting, however, isn't an automatic election. Rather, companies have to qualify for this treatment by satisfying some rather demanding protocols.

While some circumstances are easier to deal with than others, the most challenging situations occur when it's clear from the start that the derivative's gain or loss won't perfectly offset the risk being hedged. This assessment might be made because (a) the hedging derivative incorporates value dates or settlement dates that differ from those of the exposure being hedged, (b) the derivative happens to be an option contract, or (c) the underlying price (rate) variable pertaining to the derivative differs from the price (rate) that underlies the exposure being hedged. In these situations, qualifying for hedge accounting requires devising (and passing) prospective and retrospective hedge-effectiveness tests that specifically address the facts and circumstances that give rise to the less-than-perfect hedge results.

So, what's required to pass these prospective effectiveness tests? According to the standard itself, all that's required to satisfy this prospective test is the ability to demonstrate that the price (rate) underlying the derivative is highly correlated with the price underlying the hedged item. (See Paragraph 75.) This requirement has been universally understood to mean that the R-square statistic generated by the regression must be greater than equal to 0.8.

The R-square statistic measures the portion of variance observed in one variable that can be explained by the other variable. For example, an R-square statistic of 0.8 would mean that 80% of the variability of the price (rate) risk had been covered by the derivative during the sample period. Clearly, the higher the R-square the better, but the upper bound of 1.0 for the R-square statistic is unity, which would only be generated in connection with perfect hedges. Beyond that, it deserves mention that the R-square statistic is equal to the square of the correlation between the two variables involved in the regression. Hence, the requirement for R-square to be greater than or equal to 0.8 amounts to a requirement that the correlation between the two variables be greater than or equal to approximately 0.9.

Unfortunately, several of the major auditing firms have raised the bar beyond the requirements in Paragraph 75, stipulating

Risk Management

additional qualifying criteria (besides a high R-square) and/or specifying explicit testing methodologies. For instance, they might dictate whether the regression variables should reflect price (or rate) *levels* or price (or rate) *changes*, or they might require using some number of observations to use in the analysis, or they might impose some conditions relating to the statistical significance of the regression results.

Although FASB is currently evaluating whether to liberalize the requirements relating to prospective effectiveness testing—a direction that would likely obviate much, if not all, of the supplemental guidance being put forth by the auditing firms—until such time as FAS 133 is amended, the directives of the auditing firms will still carry considerable weight. Thus, at this point, the potential costs and benefits of requiring each of these “suggestions” deserves further scrutiny.

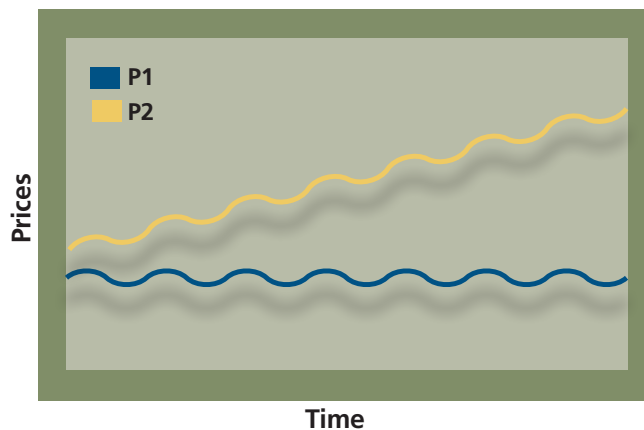
Levels versus changes

As mentioned above, to qualify for hedge accounting, Paragraph 75 stipulates that it is sufficient to demonstrate that price levels are highly correlated. Auditing firms feel some license to override this guidance, probably because the paragraph asserts something that isn't necessarily true. That is, the paragraph states that if price levels are highly correlated, it would be “reasonable to expect” the derivative's results to closely offset the changes in fair values or cash flows associated with the risk being hedged. Put another way, the standard suggests that if price *levels* are highly correlated, one should be able to infer that price *changes* should be highly correlated as well. Unfortunately, this guidance has no statistical basis. It's true only because FASB says it's true.

With this shortcoming in mind, some audit firms recommend or require ignoring Paragraph 75 and instead rely on regressions that are designed to explicitly address the question of offset. In other words, auditing firms may recommend or require that the test shouldn't ask whether the two prices (rates) are correlated, but that it should ask whether the derivative can be expected to deliver the desired offset to the risk being hedged. In constructing this test, the interval of the price changes is critical. Using monthly changes would be useful in evaluating the offset capabilities over a month; using quarterly changes would be useful in evaluating offset capabilities over a quarter. Equally important, these are independent tests. It would be inappropriate to assume that the conclusions that are appropriate for results based on one interval would carry over to intervals of different lengths.

Arguably, the regression using price levels gives the better indication as to how the hedge might be expected to work over the long run, but it says nothing about how it should work over shorter periods; similarly, a regression with a high R-square using changes in prices might be associated with a hedge that would perform abysmally over the long run. The chart below, showing two data series where price *changes* are perfectly correlated (R-square using price changes equals 1.0), demonstrates just this situation. It should be obvious to the most statistically naïve that a derivative that is reliant on one of

Price Levels Uncorrelated; Price Changes Perfectly Correlated



these series would serve as a horrible hedge to an exposure that is reliant on the other.

Number of observations

Statistical analysis involves constructing a data set from which inferences may be made. Typically, for effectiveness-testing purposes, this exercise requires collecting historical data. In assembling these data, two key questions arise: (1) What's the correct frequency of the data (i.e., daily, weekly, monthly, or quarterly)? and (2) How far back should the data extend?

As a rule of thumb, more (as opposed to fewer) observations are generally preferred, with an important caveat: Data should be collected from a sample that has a consistent nature; that is, there shouldn't be a structural event in the marketplace that makes the behavior of earlier data inconsistent with latter data. The kind of structural event that would have a bearing in this regard might be associated with changing regulation, technology, or innovation.

While there's no *universally correct* answer to either of the above questions, it's certainly wrong to dictate some finite number of observations that artificially restricts the data set and thereby ignores potentially relevant, accessible information. It may be reasonable to require a minimum sample size, but an arbitrary *upper* bound on the sample size—barring any legitimate structural concern—would be wholly inappropriate. Stated more positively, the preferred answers to questions about frequency and historical length of the data used in the regression will likely depend on data availability and the nature of the hedge relationship being assessed; in general, the more data, the better. A one-size-fits-all directive for a finite sample size with a specified frequency deserves to be challenged—and even ridiculed!

Statistical significance

The issue of statistical significance pertains to the capacity to make inferences regarding the linear equation that the regression generates. Along with the R-square statistic, the regression provides a linear equation that relates the Y (dependent) variable to the X (independent) variable, in the form of $Y = aX + b$. The regression gives explicit estimates

for the X-coefficient (a) and the Y intercept (b). It will also generate standard error and t-statistics that are used in making probabilistic statements about a and b , respectively. These statistics allow us to develop confidence intervals that surround the a and b estimates.

In the absence of perfect correlation (i.e., R-square identically equal to unity), we can't be sure that the regression-generated values for a and b are correct. The best we can do is to assert that the correct value is within some specified range around that estimate, with some level of confidence. For instance, if the coefficient were 1.02 and the standard error were 0.005, the 95% confidence interval around 1.02 would be found by adding and subtracting twice the standard error from the estimate. In this case, the interval would be from 1.01 to 1.03. In other words, we would be 95% sure that the correct value for a falls within the range between 1.01 and 1.03—and, of course, that means there's a 5% chance that the true value is outside of this range. Even this statement has its limitations in that it presumes that the regression's error terms (i.e., the deviations between actual and predicted values of our Y variable) have a normal distribution, which may or may not be the case. In any case, this assumption is the traditional starting point for statistical inferences we might like to make.

Why does this confidence interval matter? Frankly, it's not clear that it does. It is relevant to ask whether the hedge is sized correctly, and this X-coefficient may have some informational content in that regard. To extend the example, a value of 1.02 for a suggests that each dollar change in the X variable is generally, on average, associated with a \$1.02 change in the Y variable (or, analogously, a 100-basis-point rate change of the X variable is generally associated with a 102-basis-point rate change of the Y variable). Thus, it would seem that, for a given exposure, the size of the derivative should be a multiple of 1.02 times the size of the exposure. In this case, the confidence interval about the X-coefficient (a) provides the hedger with probabilistic information that indicates the degree of confidence or precision that the hedger should place in this estimated hedged ratio, *conditional on the sample used to estimate the hedge ratio*. That is, it's never clear whether this coefficient estimate is the "true" value or whether it's simply a reflection of the data that were examined.

In a good many cases, hedgers *know* the correct hedge size. They know that the notional size of the hedge should exactly equal the size of the exposure. For example, when swapping from floating to fixed interest rates (assuming LIBOR-based debt and a LIBOR-based interest rate swap), the correct size of the hedge is one-for-one—irrespective of any possible indications from the regression analysis to the contrary. If the regression-generated value for a happened to be some value *other than* unity (or if it were statistically different from unity), in all likelihood, the correct response would be to appreciate that the result were an artifact of the sample used in the regression, rather than any real justification for thinking that the hedge were improperly sized.

We should pay attention to the magnitude of this coefficient *only* if we're relying on it to tell us something that we don't

know. For instance, if the company were using a gasoline contract to hedge jet fuel, it would need to know how sensitive jet fuel prices have been (and presumably will be) to changes in gasoline prices. In such a case, the slope coefficient in the regression (a) tells us the best hedge ratio. We'd presumably want to be able to validate that our hedge will be sized in a manner consistent with this hedge ratio, but requiring statistical significance could wrongly disqualify hedge accounting if the actual hedge ratio is reasonably close to a , but outside of the 95% confidence interval surrounding a .

Returning again to the example where $a = 1.02$ and the standard error is 0.005, we have the situation where the statistical analysis indicates that the true value of a falls within 1.01 and 1.03, with a 95% level of confidence. In this case, if the effectiveness assessment criteria had required the value of the slope coefficient, a , to be "not statistically different from unity at the 95% confidence level," we would have failed this effectiveness test, and hedge accounting would have been denied. It defies reason to demand this level of statistical precision when it's questionable that the regression's estimate of this coefficient is even correct—but that's exactly what's happening in practice!

Looking ahead

FASB appears to be sensitive to at least some of the above concerns. The soon-to-be-released exposure draft is expected to propose constraining the prescriptive detail currently being required in prospective effectiveness assessments. The revised standard, if enacted, would likely adjust the current prospective testing to require, "at a minimum," a qualitative assessment to "demonstrate that (1) an economic relation exists between the hedging instrument and the hedged item or hedge forecasted transaction, and (2) changes in fair value of the hedging instrument would be reasonably effective in offsetting changes in the hedged item's fair value or the variability in the hedged cash flows attributable to the hedged risk."

Apparently, recognizing the potential shortcomings of a *qualitative demonstration*, the proposal allows that quantitative assessment may be needed, but whether qualitative or quantitative, it seems clear that the FASB's intent is to return to a more relaxed approach—one presumably in line with Paragraph 75 of the original standard, without embellishment. Additionally, whereas the current guidance requires prospective testing no less frequently than quarterly, under the proposal, this assessment would be required only at the inception of the hedge relationship, or if there were some reason to believe that the hedge would no longer be effective, going forward.

Until these changes come into effect, hedgers and auditors alike would be well served to understand some of the nuances of statistical testing to minimize the prospect that hedge accounting will be disallowed unreasonably or inappropriately.

Ira Kawaller is the President of Kawaller & Company and the Managing Director of the Kawaller Fund. He welcomes comments directed to kawaller@kawaller.com.