Why derivatives don't reduce FX risk

Thomas E. Copeland

Yash Joshi

For a hedging program to work, it must increase the "time to ruin"

The goal is to reduce the variability of cashflows

A new study shows that few companies succeed

THE SAD TRUTH about foreign exchange risk management programs is that most would not pass the doctor's basic test, "First, do no harm." A study of nearly 200 large companies yielded enough evidence to cast serious doubt on the economic benefits of FX hedging programs. Even the most superbly designed and executed programs seem not to reduce cashflow volatility significantly for most firms. Given the scarce management time and substantial capital sums currently devoted to hedging, it is clear that many programs destroy value instead of protecting it.

How can it be that hedging programs that appear so elegant in theory don't work in practice? The reason, we believe, is that the theory assumes a static world in which all factors apart from FX rates stay exactly the same. In real life, however, a host of other variables – demand for parts and products, supply of raw materials, regulatory frameworks, cost and productivity of labor and capital – all change just as FX rates change.

Tom Copeland is a partner in McKinsey's New York office and *Yash Joshi* was formerly a consultant in the New Jersey office.

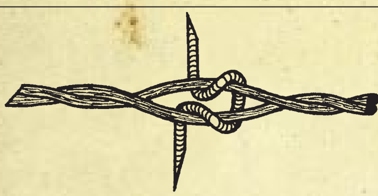

113. Baker and Bestor's Staple

114. Brotherton's Arched Barb

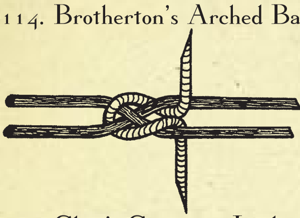

115. Claw's Crossover Lock

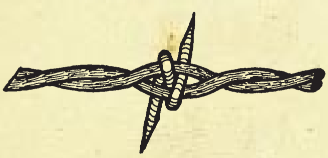

116. Edenborn's Locked-in Barb

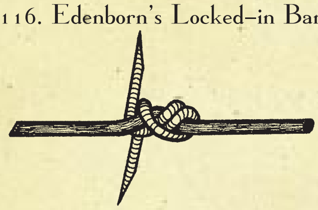

117. Ford's Kink and Wrap Barb

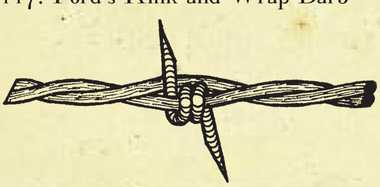

118. Glidden's Heavy Duty

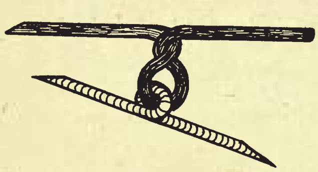

119. Kelly's Swing Barb

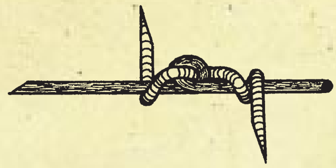

120. Locke's Loop Lock

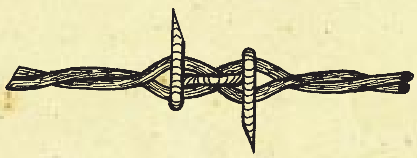

121. Nadelhoffer's Crossover

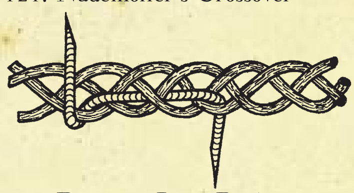

122. Preston's Long Barb

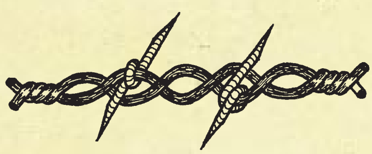

123. Riter's Visible

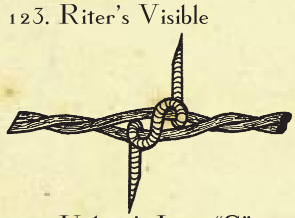

124. Upham's Lazy "S"

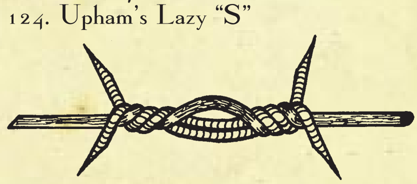

125. Waco Twist

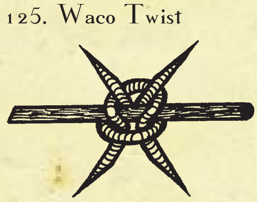

126. Wilkes' Two Staple

ROBERT T. CLIFTON, *BARBS, PRONGS, POINTS, PRICKERS AND STICKERS,* UNIVERSITY OF OKLAHOMA PRESS

In the nineteenth century there were hundreds of different types of barbed wire available, each designed to deal with a particular kind of risk.

Moreover, the relationships between all these factors are constantly shifting. Hard enough to understand in hindsight, they are virtually impossible to predict in advance. All told, FX winds up as only one of many drivers of total cashflow – and a small one at that.

Of all the variables that influence a firm's financial performance, FX is but a minor contributor to total risk, except in the event of catastrophic currency failures like that of the Mexican peso in 1995, which standard hedging programs are not in any case designed to cope with. This means that even eliminating FX risk completely – a practical impossibility – would hardly budge the needle on the total risk faced by a firm. A notable exception may be commodity and commodity-related businesses, where total cashflow is far more closely correlated to exchange rates or easily hedgeable commodities.*

Below, we evaluate FX risk management programs that are based on the use of financial derivatives. The details apply only to corporations looking to manage FX risk. However, the principles and methodology are also relevant to other kinds of risk management, such as interest rate risk management or commodity hedging.

Why manage risk, anyway?

The logic behind FX hedging is intuitively appealing. Consider the Japanese unit of a US firm that buys its goods from the United States for dollars and sells them in Japan for yen. When the yen weakens, it needs to pay more in yen for the US goods. This reduces its profits unless prices are raised. In addition, the weaker yen means that these smaller yen profits get translated into even fewer dollars. This one-two punch can have a dramatic effect on dollar profits (Exhibit 1).

Exhibit 1

Impact of weaker yen on Japanese unit's operating profit

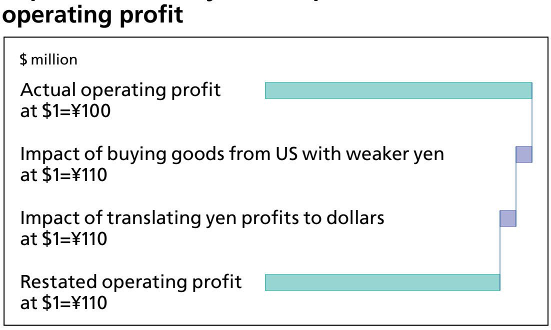

Thus, when the yen FX rate changes, the firm's profits and cashflows fluctuate. The higher the costs in dollars, the bigger the cashflow fluctuation will be.

The Japanese affiliate could limit its own FX risk by buying goods from the US firm in yen rather than dollars. This would eliminate the second bar in Exhibit 1. However, it would also increase risk for the US company, which

* *See* Timothy M. Koller, "Are you taking the wrong FX risk?" pp. 80–89, for a framework for understanding the foreign exchange risk in these businesses.

would receive revenues in yen but still incur costs in dollars. So a change in billing policy simply shifts the risk from the Japanese affiliate to the US parent, without altering the total FX exposure.

Value creation

The primary purpose of FX hedging is to reduce the cashflow volatility caused by precisely these kinds of FX movements. In turn, the reduction in FX-induced volatility is expected to dampen the volatility of a firm's total cashflow. This smoothing effect may create value in a number of ways:

◆ **By reducing the probability of business disruption costs.** Large swings in cashflows can lead to liquidity crises when cashflows turn negative unexpectedly. The cost of doing business will then rise because suppliers are slow to deliver when dealing with a customer in distress, customers shy away from a firm that may not be in business when its products need servicing, and workers depart or demand extra pay from an employer that may be gone tomorrow. Hedging to smooth cashflow can reduce these potential costs.

◆ **By creating new business opportunities.** Firms with smoother cashflows can gain competitive advantage over other companies in their industry. Merck, for example, suffered cuts in its value-creating R&D budget when its operating cashflow fell. By hedging to reduce cashflow troughs, it hoped to maintain R&D levels and thus keep up the supply of new products.

◆ **By reducing taxes.** Since tax rates are progressive – that is, a higher income attracts a higher percentage tax – smoothing cashflows across tax years reduces the tax percentage, and thus the tax liability.

◆ **By increasing debt capacity.** Lenders are more willing to deal with firms that have stable cashflows. If a company reduces the probability of a cash crunch, it improves its ability to borrow.

But even though the volatility of cashflows has a pronounced effect on these four ways of creating value, it is not the only factor at work.

"Time to ruin"

Other important variables include the volatility of unhedged operating cashflows and their trend over time, the volatility of FX contracts and their trend, the correlation between cashflows from operations and those from the FX position, and the relative distance between the level of operating cashflows and the claims against them (a coverage ratio). But what matters about cashflows in particular is the probability that they will decline to a point where business disruption costs are incurred. Mathematicians call this the "expected time to ruin." For a hedging program to be effective, it must increase the time to ruin.

Sometimes a hedge is totally unnecessary. If cashflows are well above and trending upward faster than fixed cash charges, and if the variability of operating cashflows is low, then the expected time to ruin may already be virtually infinite. This is why most large companies self-insure against minor unexpected losses. Exxon, for example, can self-insure against refinery explosions because the company as a whole experiences little or no business disruption. Smaller or less profitable companies might, however, need to insure against the same risk.

Even the most superbly designed and executed programs seem not to reduce cashflow volatility significantly for most firms

Of course, when the expected time to ruin is short, FX hedging should be investigated. Risk can be greatly reduced if cashflows from the FX position offset those from operations well enough to dampen their volatility appreciably. As we shall see, this depends on the correlation between operating cashflows and FX cashflows. It must be high for a hedge to be effective.

Most hedges are accompanied by a cost: namely, that the trend in FX positions, which is usually lower than the trend in operating cashflows, reduces an anticipated upward trend in the latter. Because the effect of hedging on the variability of operating cashflows is the primary variable that can be controlled, it will form the focus of the discussion below.

Exhibit 2

Impact of FX hedging on pre-tax income volatility

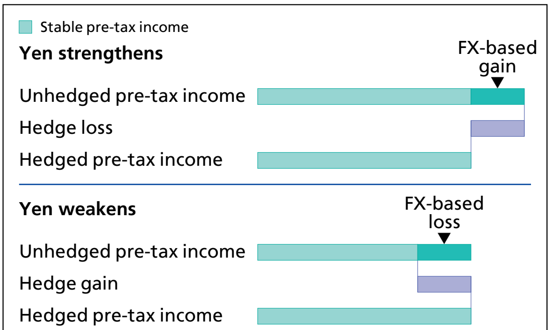

Question: Will this hedge reduce the firm's entire cashflow volatility?

Note: Assumes that cashflow from the hedge equals the negative of the increase in cumulative translation adjustment

Smoothing flows – or not?

To understand how hedging might smooth cashflows, let us return to the firm that sells in Japan but has sizable dollar-denominated costs. It loses when the yen weakens, and gains when the yen strengthens. FX hedging would simply neutralize the operational losses caused by FX movements – and, often, any gains too (Exhibit 2).

Reducing FX risk is usually thought to stabilize a firm's total cashflow and stock price (because this price is the capitalized value of future cashflows).

In a typical FX hedging program, a firm would estimate the foreign-currency-denominated cashflows it expects to receive at a future date. Usually these cashflows will be tied to expected transactions, such as foreign currency payments due on invoiced goods or currency receipts from receivables; thus, the hedge is often referred to as a "transaction hedge."

The company would then sell forward contracts denominated in the foreign currency to hedge the FX exposure. A firm that expects to receive a ¥1 billion cashflow in September 1996, for example, would hedge by selling 1 billion September yen forward.

Exhibit 3

An example of transaction hedging

	Base case ($1 = ¥100)			Strong yen ($1 = ¥50)			Weak yen ($1 = ¥150)		
	US parent	Japanese affiliate		US parent	Japanese affiliate		US parent	Japanese affiliate	
Revenue (bill in yen)	$1.00	¥400		$2.00	¥400		$0.67	¥400	
Cost: United States	–$0.90	–¥100		–$0.90	–¥100		–$0.90	–¥100	
Cost: Japan	$0.00	–¥200		$0.00	–¥200		$0.00	–¥200	
Profit	$0.10	¥100	$1.00	$1.10	¥100	$2.00	–$0.23	¥100	$0.67
Gain/loss on forwards	$0.00			–$1.00			$0.33		
Profit net of forwards	$0.10		$1.00	$0.10		$2.00	$0.10		$0.67

Using our imaginary multinational firm as an example, Exhibit 3 shows the numerical impact of FX fluctuations on cashflow and the benefits of hedging to reduce volatility. The firm sells goods produced in the United States to its Japanese affiliate at a transfer price of ¥100 per piece. It simultaneously sells forward the equivalent of $1.00 in yen; that is, it sells ¥100 to receive $1 on the settlement date.

In the base case, at $1 equals ¥100, the profit totals $1.10: $0.10 at the US parent company and $1.00 at the Japanese affiliate. When there is a strong yen, at $1 equals ¥50, the profit rises to $3.10 without hedging; with a hedge, the profit is lower at $2.10. More important, when the yen weakens to $1 equals ¥150, the profit drops to $0.44 without a hedge; but it is higher, at $0.77, with the hedge gain added in. The hedge thus limits the gains and losses caused by FX movements, dampening the volatility of profits.

> For a hedging program to be effective, it must increase what mathematicians call the "expected time to ruin"

Clearly, hedging will help only those firms whose cashflows are affected by (or correlated with) FX fluctuations. For other firms, it will have either no impact at all or an adverse impact (Exhibit 4).

Paradoxically, hedging can actually increase volatility. Suppose our imaginary firm is selling yen forward to hedge anticipated yen cashflows – a typical transaction hedge. What happens if the yen strengthens and is

Exhibit 4

Cashflow change versus FX change

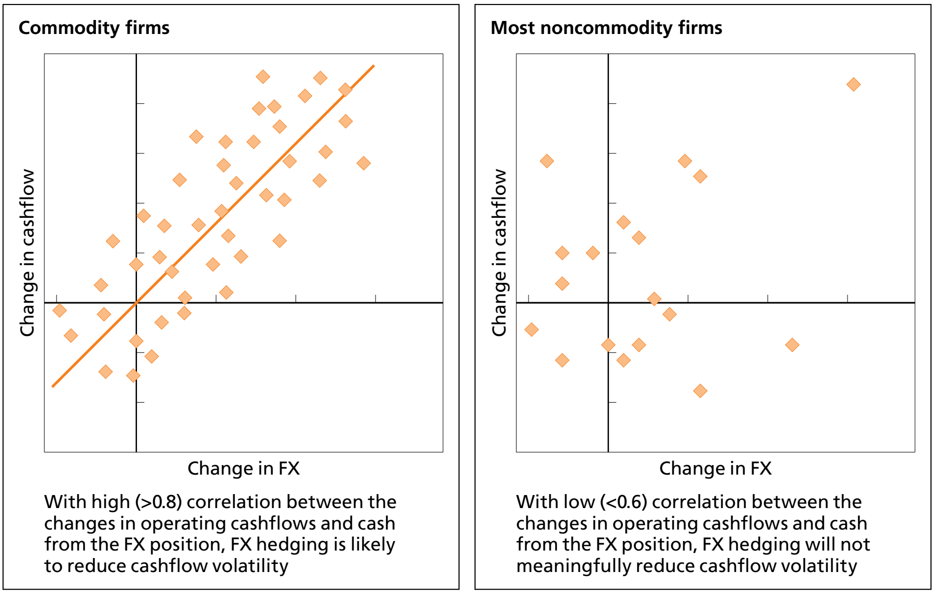

Exhibit 5

Cashflow impact of strong yen and recession in Japan

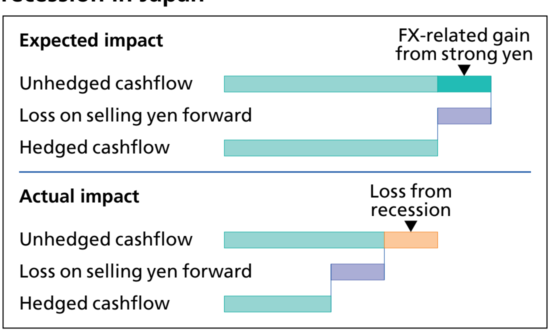

Exhibit 6

Cashflow impact of weak yen and economic growth in Japan

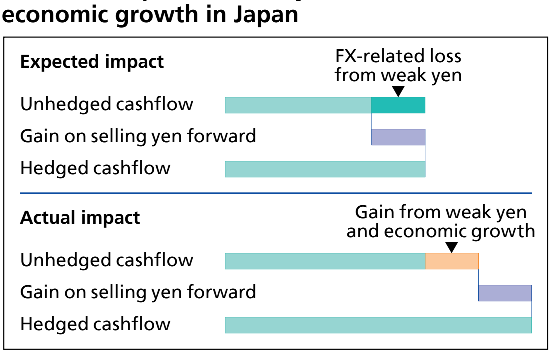

accompanied by a recession in Japan? Unhedged cashflows will fall – and losses on the yen sold forward will make things worse (Exhibit 5). Conversely, if the yen weakens and the Japanese economy grows, unhedged cashflows will rise, hedging will produce further gains, and volatility will again increase (Exhibit 6).

A case in point

An actual example illustrates this possiblity further. In the mid-1980s, a European airline contracted for several billion dollars' worth of Boeing 747s and 767s. The US dollar was strong at that time, and consensus was that it could go higher. Being short the dollar on its Boeing transaction, the airline bought roughly the same amount of dollars using forward contracts. If the dollar strengthened, the airline would lose on its Boeing contracts but win on its forward position. Therefore the airline had created a transaction hedge.

Unfortunately, even though the airline is European, its cashflows are positively correlated with the dollar, especially on transoceanic routes. If the dollar weakens, ticket prices in the home currency need to fall in order to keep the quantity sold relatively constant; therefore sales revenues decline. Furthermore, equipment and fuel costs fall in the local currency, while personnel, overhead, and other costs remain constant. In sum, the company takes in less cash if the dollar weakens, and more cash if the dollar strengthens.

Consequently, the Boeing contract was a natural hedge. If the dollar strengthened, the cashflows of the airline excluding the contract would go up, while the contract cost would also go up. Conversely, if the dollar weakened, the cashflows of the airline excluding the contract would go down, but the contract cost would also go down. The FX position, a transaction hedge, had the effect of removing the natural hedge, and was directionally wrong. In effect, hedging the transaction increased the company's risk.

The actual course of events was that the dollar peaked in February 1985, and by year-end it had fallen about 40 percent from its peak. When the hedge was marked to market, the airline's operating cashflows were down, and it suffered a large loss on its FX position as well. Thus transaction hedging does not always smooth operating cashflows.

Is hedging worthwhile?

To answer this question, we estimated the potential benefits of FX hedging for a sample of large firms. The sample was based on the 500 firms with the largest 1994 sales in the Compustat database. Firms with no adjustment for FX translation were excluded, as were those that had missing data for pre-tax income or FX translation adjustments over the previous 10 years. The final sample comprised 198 firms.

Cashflows from a hypothetical FX hedging program were estimated and their volatility was compared with that of the firms' actual cashflows. If hedging produced large reductions in volatility, this would suggest high potential benefits; small reductions would indicate low potential benefits.

A simple approximation

Estimating hedged cashflow involves estimating the changes in monthly or quarterly cashflow caused by FX fluctuations – not information that firms generally report. So a simple approximation was used instead. Firms usually report FX translation adjustments on their balance sheet. The impact of FX fluctuations on the balance sheet tends to resemble that on the income statement: both are likely to be positive or both negative, and the order of magnitude of the two figures will probably be similar.

Exhibit 7

Impact of temporary FX fluctuations on balance sheet

	Base case ($1 = ¥100)		Weak yen ($1 = ¥120)	Strong yen ($1 = ¥80)
Income statement	¥	$	$	$
Revenue	100,000	1,000	900	1,110
Expenses (half in $)	(70,000)	(700)	(665)	(740)
Tax	(15,000)	(150)	(120)	(185)
Net income	15,000	150	115	185
Change in net income			–35	+35
Balance sheet				
Net assets	50,000	500	415	625
Debt	20,000	200	165	250
Equity	30,000	300	300	300
Cumulative translation adjustment	0	0	–50	+75
Change in cumulative translation adjustment			–50	+75

Consider the example illustrated in Exhibit 7. When the yen weakens, net income falls by $35. The change in translation adjustment is –$50. If this is adjusted for tax at 50 percent, the after-tax change of –$25 is of the same order of magnitude as the change in net income. A similar relationship is seen in a period when the yen strengthens.

Therefore, we used the change in FX translation adjustment as our proxy for change in cashflow caused by FX. This approximation provided what we believe is directionally correct data which can form the basis for conclusions that are valid in the aggregate, though not at the level of individual firms.

Testing the proxy

A close look at a real-life FX hedging program suggests that the proxy is a reasonable one. Sara Lee uses forward foreign currency contracts to hedge its FX risk.* We estimated the gain on the company's forward foreign currency positions by multiplying the year-end FX contracts outstanding by the gains or losses on six-month forward foreign currency contracts for the year.

* The company's 1994 annual report indicates that it enters into contracts to sell forward foreign currencies for US dollars. These contracts generally have maturities of less than a year. The report states that "gains and losses in foreign exchange contracts and currency swaps used to hedge long-term foreign investments are recorded as a separate component of common stockholders' equity," which suggests that gains from FX contracts are not part of the reported operating profits.

The resulting estimated FX-induced gains and losses were directionally similar to our proxy (the change in cumulative foreign currency translation adjustments). In addition, the correlation between the two was reasonably high at 0.9, which, notwithstanding the small number of data points, appears to support our methodology (Exhibit 8).

Exhibit 8

Estimated gain on FX forward position compared with crude hedge gain*

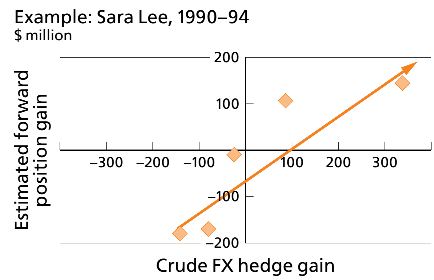

* Crude hedge gain equals the negative of the increase in cumulative translation adjustment

Source: Sara Lee annual reports; McKinsey analysis

Designing a hedging program

Next, a crude hedging program was designed. It simply neutralized FX-induced gain (or loss) by subtracting our proxy for gain from (or adding our proxy for loss to) reported pre-tax income. The result of this calculation represents hedged cashflow. Comparing the standard deviation of the hedged cashflow with that of the unhedged cashflow indicates the degree of benefit a firm would have received had it had in place an FX hedging program to neutralize the impact of FX translation.

All this assumes, of course, that firms do not hedge. This is probably true in general, since most firms do not manage FX risk. Only four companies out of a sample of 20 indicated in their annual reports that they used currency swaps to convert foreign-denominated debt to dollar debt.* And a recent survey by the Wharton School and Chase Manhattan found that three out of four manufacturers, four out of five service firms, and two out of three public utilities never use derivatives.†

> Clearly, hedging will help only those firms whose cashflows are affected by (or correlated with) FX fluctuations

Nevertheless, some firms, like Sara Lee, do hedge. Unlike Sara Lee, some of them may include FX gains in their operating income. In such cases, the reported cashflow is the hedged cashflow. In order to get the unhedged cashflow, our proxy gain would need to be added to (not subtracted from) the reported cashflow. As a second pass, we did exactly this, making an adjustment that implicitly assumes that all firms hedge.

Each firm in the sample can belong to only one group: either it hedges, or it doesn't. We gave each firm the benefit of the doubt by adopting the assumption that conferred the most advantage on it. In other words, we overestimated the benefit firms would derive from FX hedging.

* S. Waite Rawls, III and Charles W. Smithson, "Strategic risk management," *The New Corporate Finance: Where theory meets practice,* McGraw-Hill, New York, 1993.

† *Barron's,* May 1, 1995.

Applying the program to the sample

Using annual data for the previous 10 years, we estimated the volatility of each firm's pre-tax income and compared it to the volatility of hedged pre-tax income, defined as pre-tax income less the increase in cumulative FX translation adjustment. Thus, we estimated the actual change in volatility that would have occurred if the hedge had resulted in cashflows equal to the increase in cumulative FX translation adjustment.

Surprisingly, we found that just one firm out of the sample of 198 would have reduced its 10-year income volatility by more than 20 percent, and a mere 20 firms would have reduced their income volatility by more than 10 percent. Only a very small percentage of firms, then, appear to benefit from the hypothetical hedging program.

Benefits of an optimal program

However, our rather crude program may not represent the full benefits that might accrue from a better-designed hedging program. While it is impossible to design an optimal hedging program for individual firms without more information, we can calculate the maximum reduction in volatility that a hedging program could possibly achieve by using the following formula:

$$\text{minimum volatility} = \text{unhedged volatility} \times \sqrt{(1 - \text{correlation}^2)}$$

where correlation is the correlation coefficient between unhedged cashflows and change in FX rates.

As we have seen, the cumulative translation adjustment is a directionally correct approximation for the change in cashflow caused by FX fluctuations. Thus, the correlation between unhedged pre-tax income and translation adjustment changes stands as our proxy for the correlation required by the formula. The correlations we calculated for the firms in the sample are presented in Exhibit 9.

Exhibit 9

Correlation between reported income and change in cumulative FX translation adjustment

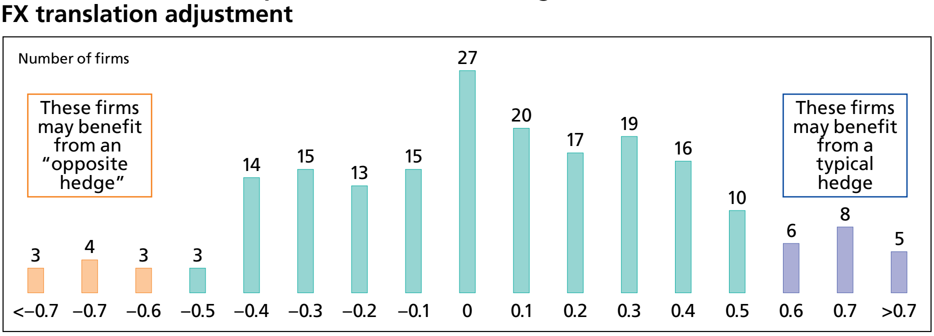

According to our formula, a correlation of 0.6 is required for a volatility reduction of about 20 percent. Only 9.6 percent of the firms in our sample had a correlation of 0.6 or higher. Consequently, few firms are likely to benefit even from an optimally designed FX hedging program.

A closer look

Using inside information to test our findings in a specific case, we also studied the FX hedging program of an actual multinational firm which we will call ABC Company. It produces goods in the US and sells them in Japan, and its dollar-denominated costs and yen-denominated revenues lead to the typical FX risk we have described. ABC decided to hedge its exposure by using forward contracts. Its US unit sold yen forward to tie in with budgeted cash inflows from Japan.

> Even eliminating FX risk completely would barely budge the needle on the total risk faced by a firm

So as to assess the effectiveness of ABC's FX hedging program, we first estimated the volatility of the hedged and unhedged quarterly (and monthly) cashflows. We found that the actual hedged cashflows were just as volatile as the unhedged cashflows (Exhibit 10).

Exhibit 10

Quarterly operating cashflows* with and without hedges

Example: ABC Company

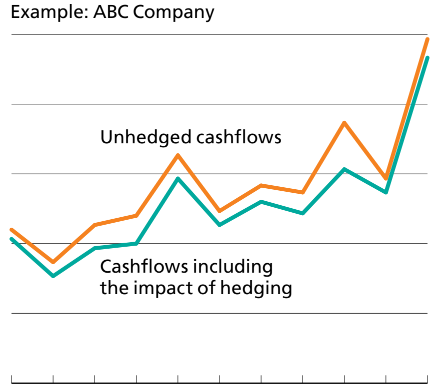

* Net income plus depreciation; including working capital changes makes cashflow more volatile, but does not affect the conclusion

Exhibit 11

Change in quarterly cashflow* compared with percent change in FX

Example: ABC Company, 1992–95

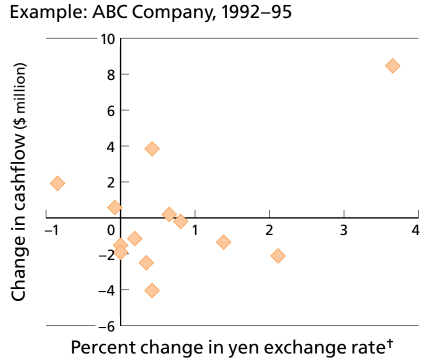

* Net income plus depreciation, less gain on closed hedge positions
† Equivalent to the percent gain on a short yen forward position with a constant size in dollars

Source: ABC Company financials; DRI/McGraw-Hill; Bloomberg; McKinsey analysis

To ascertain whether a different FX hedging program would have suited ABC better, we looked at the correlations of the firm's cashflows with FX rates. It emerged that these correlations were very low, at 0.4 (Exhibit 11).

Exhibit 12

Estimating optimal hedge position

Procedure using simulation

1. For the period chosen, estimate the unhedged quarterly cashflow in dollars, which equals reported cashflow less gain on closed forward position.
2. Estimate the standard deviation of the unhedged cashflow. This is the unhedged cashflow volatility.
3. Assume a hedge position (short six-month yen forward) equivalent to $1 million for the first simulation. Estimate the gain or loss on this hedge position in each quarter. Add this hedge gain to the unhedged cashflow to determine the hedged cashflow in each quarter. Estimate the standard deviation of the quarterly hedged cashflow. This is the hedged cashflow volatility.
4. Repeat step 3 several times, increasing the size of the hedge position by $1 million every time. As the hedge size is gradually increased, the volatility of cashflows will drop, hit a minimum, and then start rising.
5. The simulated hedge position that minimizes hedged cashflow volatility is the optimal hedge position.

Exhibit 13

Standard deviation of quarterly cashflows*

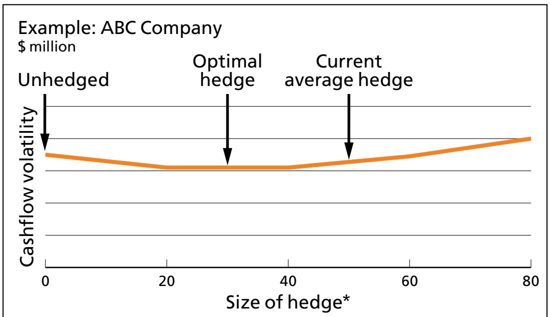

* Net income plus depreciation and amortization, less gain on closed hedges, plus gain on different hedge positions; using gains on open instead of closed positions does not alter the conclusion
† Short yen forwards
Source: ABC Company financials; McKinsey analysis

We then simulated cashflows for optimal hedges, following the steps shown in Exhibit 12. Such a procedure can be used by any company. It answers a simple question: were hedged cashflows less volatile than unhedged cashflows, and if so, by how much?

With all the benefits of hindsight, such as the elimination of forecasting errors in revenues and costs, we found that the optimal hedge would have reduced ABC's actual quarterly cashflow volatility by just 10 percent – not a meaningful amount (Exhibit 13). The impact of the hedge is so disappointingly low because FX fluctuations made only a small contribution to total cashflow volatility. As a result, even eliminating FX-induced volatility altogether would not significantly reduce total cashflow volatility.

Moreover, the firm's current hedge positions were two to three times the optimal hedge positions. Hence the second problem with hedging: it is easy to overhedge. And, unless correlations are stable over time, real hedges are unlikely ever to be optimal. In other words, even if ABC revised its hedge to match what we estimate is today's optimal position, there is no guarantee that the new hedge would remain optimal. Future correlations may differ so much from past correlations that the best possible hedge position this time next year may be nothing like the current one. The conclusion is clear: FX hedging did not and could not reduce cashflow volatility in this real-life case.

Although in general derivatives are, as we have shown, ineffective in managing foreign exchange risk, senior managers should not simply throw up their hands and resign themselves to being pummeled by the currency markets. Hedging individual transactions may not work, but overall foreign exchange exposure at the company cashflow level can and should be measured and managed. But again, don't look to derivatives for the answer. As successful Japanese auto factories in the United States attest, relocating plants and adjusting pricing often provide the best hedge against foreign exchange risk. Q