

To Hedge or not to Hedge?

Theory, ideas and mechanics for hedging in multilateral netting systems

by EuroNetting Inc.

The underlying purpose of hedging intercompany payables is quite simple: consider a CAD 2 million payable to a U.S. receiver, identified in October 1997 and payable in January 1998. With the Canadian Dollar at 1.3690 in October, the U.S. receiver expected USD1,460,920, but when it is settled at spot through the netting in January at 1.4680, only USD 1,362,398 is received. The result is a realized FX loss of about USD100,000 for the receiver and therefore to the company as a whole.

This type of situation often becomes a contentious issue when a company is implementing a new netting system. Prior to netting, participants have been responsible for managing their own payables and may have been hedging significant non-base currency payables and receivables locally. When the parent comes in with a netting system, participants may be reluctant to yield control over their effective conversion rates, particularly if the parent has not given thought to hedging or is using a netting system that cannot handle hedging. If the system cannot handle hedges, participants may not only lose the ability to hedge in future, but also may have difficulty transitioning the now-orphaned hedges they already have on their books.

Centralized hedging of netting cashflows is always more efficient than participant-level hedging for two reasons. First, just as a netting system offsets cashflows so that only the final net position is traded, some companies are able to aggregate participant netting exposures and then hedge only the net currency positions externally, resulting in both a reduction of aggregate FX trade volume and also elimination of situations in which the company is simultaneously buying and selling the same currency in different participant locations. Secondly, the corporate treasury department will probably be able to trade FX more efficiently (that is to say, at better spreads) than individual participants because of better trading tools and FX relationships.

Here are some different hedging methodologies that are currently being used by netting centers in North America and Europe:

Transaction-level hedging

Perhaps the simplest mode is that in which participants advise the netting center of specific significant future cashflows that will settle through the netting system. This may be just for the current month, or may stretch out across future netting cycles according to billing practices. Depending on corporate FX policy, the participant may request a hedge or the netting center

may automatically identify the need to hedge individual items. The netting center then locks in an external hedge contract with one of its banks to settle on the relevant netting cycle's settlement date, and guarantees this hedge rate to the participant for identified items. For an individual item the hedge contract will be between the participant's base currency and the payable currency. Then, in the relevant netting cycle, the external hedge contract settles with the bank and participant settles at the hedge rate instead of the default netting rate. The net result is that the participant has settled the cashflow at the hedge rate, and the netting center's FX P&L position remains flat.

How the cashflow settlement with the participant is achieved at the hedge rate instead of the netting rate depends on the capabilities of the netting software being used. From the participant's point of view, it would be preferable to simply see the cashflow converted in the netting report at the hedge rate, perhaps highlighting the fact that a hedge rate has been applied instead of the default netting rates. Unfortunately, not many netting systems can handle this process, and they must instead channel the two sides of the actual FX trade settlement through the participant's netting report to yield the correct FX profit or loss (denominated in the participant's base currency) to offset the apparent loss or profit from settlement of the hedged cashflow which has been reported at the netting rate. The net P&L result of this more cumbersome operation is identical to the desired result, but is more difficult to identify.

Note that only one of the two participants to the cashflow has been affected by the hedge, either the payer or receiver according to who requested the hedge. The other participant will settle at the default hedge rate. In practice this is not particularly difficult to envision because the payable currency will likely be the base currency of either the payer or receiver, and so only one participant will have any interest in hedging. But consider a '3-currency' payable, one in which the payable currency is neither the payer's or the receiver's base currency, for example a EUR payable from a GBP participant to a USD participant. This is not very common but does happen occasionally. The GBP-based payer may wish to hedge the GBP cost of the EUR, while the USD-based receiver wants to hedge the amount of EUR that will be received.

Hedging in this situation becomes a little more complex. The easy way is to lock in two separate hedges, each between a participant's base currency and the payable currency. In this way, one side of each hedge contract can be set to equal the known amount of the USD payable. But it would be more efficient to eliminate one trade and just lock in a GBP/USD hedge, then split the cross rate apart into GBP/EUR and EUR/USD rates for the two participants. Careful calculation of the two crosses can still result in a neutral P&L for the netting center.

Position-level hedging

Transaction-level hedging is fine for lower volume netting systems or FX policies that decree a selective mode of hedging. But some companies want to implement a more systematic methodology that is less subjective and which can handle higher volumes of cashflows over multiple cycles. In this case, it may be better to hedge positions rather than transactions.

However the process of feeding netting data to the netting center will have to be modified so that the data is received for integration into the position model as soon as it becomes available. This should ideally take place at the invoice level by taking a feed from a centralized data warehouse responsible for tracking all booked invoices. By applying simple position analysis techniques to the invoice database grouped by exposure horizons matching the netting cycles, plus any hedges already in place, it is then relatively simple to calculate required hedges that will deliver an exposure profile meeting corporate FX policy.

The hedge data should be viewable at both the participant level and at the aggregated netting center level, so that required hedges can be set up as efficiently as possible using a single currency trade per exposure horizon. The netting center then splits apart the aggregated hedges into component internal hedges which are settled through the netting center.

Not everybody is so lucky as to have a centralized data warehouse. In this event it may be necessary to solicit cashflow forecasts from participants, again grouped by currency and exposure horizon. Implementing this type of process is not dissimilar from remote entry in a netting system: participants submit monthly exposure updates which are imported into the exposure tracking system. One drawback of this process which uses forecasts rather than actual invoice data, however, is that it is somewhat subjective, and may therefore be subject to whipsawing from cycle to cycle as participants adjust their forecasts. The benefits of hedging forecasted positions may then be tempered by a certain amount of over-trading which may have occurred in an attempt to adjust hedge levels from cycle to cycle.

Accounting for the P&L resulting from position-level hedging is going to be more challenging than for simple transaction-level hedging and should be thoroughly investigated before committing to this approach.

One final consideration for position-level hedging is that analytical tools must be on hand to model positions and calculate hedge requirements, and to also track the number of hedge contracts that must be administered. The source data should be stored locally so that it is readily available for analysis.

Projecting positions from prior data

There is a third way to hedge netting systems that is quite rare but which may still yield beneficial results. Intercompany cashflows in many businesses are quite consistent from cycle to cycle, resulting in relatively stable net currency positions in the netting system from one cycle to the next. In this case, one can look back at the net FX positions in previous cycles and project positions for future cycles, potentially locking in hedges where exchange rate forecasts suggest that future rates may go against a position.

The challenge then becomes how to apply a hedge contract when the relevant cycle comes up for settlement, as it will almost certainly not have been set up for the exact net position. One possible approach is to trade the residual (under- or over-hedged) net position at spot rates, calculate the melded rate resulting from the two (or more) trades, and then apply it across all

relevant cashflows. This will guarantee a P&L neutral result at the netting center. On the other hand, if you are running the netting system as a profit center, the net could be settled at current spot rates and the residual P&L from the hedge retained at the netting center level.

Netting at book rates

One other trend alluded to in my previous article is to abandon netting with the spot rates available on the netting date in favor of predefined conversion rates, which will be applied to all cashflows in the net. At one extreme these may be month-end reporting rates established just a few days before the net is run, while at the other extreme one might use budget rates which may have been set up to 16 months before. In each case, the objective is to shift the burden or responsibility for hedging from participants to the netting center which is presumed to be better positioned to manage the exposure inherent in the netting data.

This process has some interesting side effects on the netting data. By using predefined netting rates each participant's net position is immediately known as soon as the data has been fully logged - in effect, each payer to receiver cashflow becomes two discrete cashflows between the netting center and the payer and receiver denominated in their respective base currencies, and running the net on the netting data becomes a simple matter of summing the participants' already known net positions.

Of course, this means that the netting center is then handed net positions in a set of currencies which will almost certainly not net out to zero when traded, as they would have done if the net had been settled using the traditional live spot rate method. However this does not mean that any more or less FX P&L has impacted the company as whole, it has merely been centralized instead of being scattered across the participants.

One final permutation on the book rate theme is the opposite of the twin base currencies scenario: in this permutation, payables are split apart into pairs of cashflows to and from the netting center, all converted at book rates to the parent company's base currency. I'll leave you to figure out the implications of this practice on your own!

Choosing the right approach

While hedging of one sort or another will almost always be beneficial, there is no single 'right' approach to implementing a hedging component within a netting system. The factors touched on above include the volume of items to be netted, corporate FX policy, availability of systems and the nature of the cashflows which are being netted. The manpower required to implement an effective program may not be as great as it seems, however. This is because netting is a cyclical operation that you will not be required to administer every day. Even in the most complex netting operations, judicious selection of the right systems can still make data import, analysis, hedging execution and netting settlement a one or two man-day operation each month. Try it!

