

Efficient XVA management: pricing, hedging and allocation

Banks must manage their trading books, not just value them. Valuing includes valuation adjustments collectively known as XVA (credit, funding, capital and tax, at least). Here, Chris Kenyon and Andrew Green show how three technical elements can be combined to radically simplify XVA management, both for calculation and implementation

Banks must calculate and manage valuation adjustments (XVA) across their entire trading portfolio. XVA includes the effects of credit (CVA, DVA), funding (FVA, MVA) (Burgard & Kjaer 2013; Green & Kenyon 2015a), capital (KVA) (Green, Kenyon & Dennis 2014) and tax (TVA) (Kenyon & Green 2015). XVA management includes allocation, hedging and pricing. Allocation refers to the allocation of XVA, and XVA hedging costs, to desks. Hedging costs require the computation of both first-order sensitivities and second-order sensitivities, such as interest rate-credit cross-gamma. Incremental allocation is required for daily trading.

Here, we provide an analytically rigorous method for managing XVA efficiently, which combines three elements: trade-level regression, analytic computation of sensitivities and global conditioning. Regression has been demonstrated in MVA to speed things up by one-to-two orders of magnitude, even on vanilla instruments for medium-sized portfolios (1,000–10,000 swaps) (Green & Kenyon 2015a). Analytic computation of sensitivities gives one-to-two orders of magnitude (Capriotti, Lee & Peacock 2011; Giles & Glasserman 2006) (the more sensitivities required, the more improvement). Global conditioning can give a 40:1 improvement for MVA. These are point examples; our contribution is to make these improvements systematically available and extend them generally across trading book XVA management. Technically, this paper generalises Green & Kenyon (2015a) from MVA to XVA, and adds sensitivities and allocation; it also makes explicit elements that are implicit in Green & Kenyon (2014).

Regression-based CVA was developed in Cesari *et al* (2010). We apply regression to all trades; ie, including non-callable trades and European callable trades. Wang & Caflish (2009) used regressions for sensitivities, but were limited to the regression variables themselves. We cover all sensitivities by including sensitivities of the underlyings to hedging instruments via the chain rule. Sensitivities of underlyings to hedging instruments can be calculated using analytic derivatives (AD) (Broadie & Glasserman 1996) or adjoint algorithmic differentiation (AAD) (Capriotti, Lee & Peacock 2011) – we use A/AD as a label for both. Examples of allocation methods for capital and CVA include Tasche (2008). We extend to other XVAs and their sensitivities, which requires a general global conditioning approach to cover MVA when based on expected shortfall (ES) as well as CVA, DVA and FVA.

The contribution of this paper is to show how to systematically combine trade-level regression, analytic computation of sensitivities and global conditioning to make XVA management computations radically more efficient. Each element alone is useful, but only in combination does the step-change in efficiency of computation and implementation appear. This methodology naturally handles wrong-way risk (see, for

A. Trade-level regression example

	Basis functions			Regression equation $f(x)$
	x^0	x^1	x^2	
	Coefficients of basis functions			
Trade #1	−0.2	1	0.1	$−0.2 + 1x + 0.1x^2$
Trade #2	1	2	0.2	$1 + 2x + 0.2x^2$
Trade #3	2	3	−0.2	$2 + 3x − 0.2x^2$
Portfolio	2.8	6	0.1	$2.8 + 6x + 0.1x^2$

	Scenario		
	A	B	C
x	−0.6	0.1	1.1
Trade #1	−0.764	−0.099	1.021
Trade #2	−0.128	1.202	3.442
Trade #3	0.128	2.298	5.058
Total	−0.764	3.401	9.521
Portfolio	−0.764	3.401	9.521

When regressions are linear in their coefficients, calculating from the sum of the basis function coefficients (the portfolio regression) is identical to calculating the sum of each trade regression

example, (2)) and is suitable for parallel implementation on graphics processing units (GPUs).

Examples

We start with a set of examples of XVA management cases on toy problems.

■ **Values with trade-level regression.** We show that computing portfolio prices from a portfolio regression is identical to computing the sum of the individual trade regressions. The basis function coefficients of the portfolio regression are the sums of the trade basis function coefficients.

We start with three trades, with a regression equation for each one; these are shown in table A. Each equation has three basis functions, $\{x^0, x^1, x^2\}$, and so is quadratic in the underlying x but linear in the coefficients of the basis functions. Table A considers three scenarios $\{A, B, C\}$. These might be Monte Carlo realisations and there would generally be many thousands of them. Pricing with individual trade-level regressions and pricing with their sum are identical. Each scenario $\{A, B, C\}$ is distinguished by the value of the underlying x , which might, for example, be a stock price or the price of an interest rate swap.

■ **First-order sensitivities with trade-level regression.** This example has three objectives: (1) to show how to compute the first-order sensitivity of each trade with respect to any calibration instrument s ; (2) to demonstrate the resulting trade-level sensitivity regressions sum, as before, to the portfolio sensitivity regression; and

B. First-order sensitivity example, s , is any calibration instrument				
	Basis functions			$df(x)/dx$ equation
	x^0	x^1	x^2	
	Coefficients of basis functions			
Trade #1	1	0.2	0	$1 + 0.2x + 0x^2$
Trade #2	2	0.4	0	$2 + 0.4x + 0x^2$
Trade #3	3	-0.4	0	$3 - 0.4x + 0x^2$
Portfolio	6	0.2	0	$6 + 0.2x + 0x^2$

	Scenario		
	A	B	C
x	-0.6	0.1	1.1
dx/ds	0.2	-0.3	0.1
Trade #1	0.176	-0.306	0.122
Trade #2	0.352	-0.612	0.244
Trade #3	0.648	-0.888	0.256
Total	1.176	-1.806	0.622
Portfolio	1.176	-1.806	0.622

Calculation from trade sensitivity equations or portfolio sensitivity equation is identical. Portfolio sensitivity equation is identical to the sum of the individual trade equations

(3) to demonstrate the implementation effort is radically reduced with respect to non-regression-plus-A/AD approaches. We reuse the setup from the previous example.

To get the first-order sensitivity of any trade regression $f(x)$ to the calibration instrument s we use the chain rule:

$$\frac{\partial f(x)}{\partial s} = \frac{\partial f(x)}{\partial x} \frac{\partial x}{\partial s} \quad (1)$$

This is also valid for the portfolio regression. In addition, the coefficients of the basis functions of the sensitivity regressions sum, as before, to give the coefficients of the basis functions of the portfolio regression. Equation (1) is valid for *all* first-order sensitivities. Each separate first-order sensitivity is distinguished by a different s and therefore different $\partial x/\partial s$.

The derivative of the regression with respect to the underlying is generally trivial to compute analytically. The derivative of the underlying with respect to the calibration instrument is generally more involved and can be tackled using A/AD.

By separating the derivative into two parts, we radically reduce the implementation effort of A/AD, because the second part $\partial x/\partial s$ is the same for all trades, while the first part $\partial f(x)/\partial x$ is generally trivial analytically. Both a regression approach and A/AD are required.

Second-order sensitivities work just like first-order sensitivities.

■ **CVA calculation and exact allocation.** This example introduces global conditioning, and demonstrates its use for CVA and exact allocation. Global conditioning means we first identify the scenarios that contribute to a computation using the portfolio regression. We subsequently calculate using only those scenarios. Within the selected scenarios, computations are additive because conditional expectation is a linear operator (Shreve 2004).

We use the same portfolio and scenarios as before. Looking at the portfolio values in table A, we see that only scenarios B and C contribute to CVA because that is where the portfolio has positive value.

C. CVA example with exact allocation via global conditioning					
	Scenario				
	A	B	C		
Portfolio	-0.764	3.401	9.521		
PD	0.2	0.2	0.2		
LGD	0.6	0.6	0.6		
CVA by scenario	0	0.40812	1.14252	0.51688	

	Scenario				CVA by trade
	A	B	C		
Trade #1 CVA	0	-0.01188	0.12252		0.036880
Trade #2 CVA	0	0.14424	0.41304		0.185760
Trade #3 CVA	0	0.27576	0.60696		0.294240
CVA by scenario	0	0.40812	1.14252	0.51688	0.516880

PD is probability of default; LGD is loss given default. Global conditioning means we only use scenarios identified as contributing to CVA in the trade-level calculations (ie, B and C) and set the contributions from scenario A to zero. So the calculation of CVA by trade using global conditioning provides exact, and additive, allocation

Table C shows the CVA calculation. Trade-level contributions to CVA sum exactly. Global conditioning therefore provides exact trade-level CVA allocation.

Reallocation of trade-level CVA to different desks, or for different reports, is trivial because trade-level contributions add up exactly.

■ **Incremental CVA and exact allocation.** Suppose we have an incremental trade (trade #4). Now, for CVA, we know that provided its value in scenario A is less than 0.764, then scenario A will not become relevant for CVA. We also know that if the value of the new trade is greater than -3.401 in scenario B and greater than -9.521 in scenario C, then these two scenarios will remain relevant for CVA. Thus, within these (one-sided) bounds the trade-level allocation of CVA will be *unchanged* by the new trade. Its own contribution can be computed independently (conditioned on scenarios B and C).

If the incremental trade value is outside the bounds of the portfolio value in each scenario, then we need to calculate trade values in the newly relevant scenario. The previous trade-level values in each previously relevant scenario will either remain valid or be set to zero. Thus recalculation is essentially trivial, and exact, both for CVA itself and trade-level allocation. The same result holds for trade-level contributions to first- and second-order sensitivities. We say *contributions* because sensitivities combine trade parts and parts from default probabilities and recovery rates, as we now show.

■ **CVA first-order sensitivities and exact allocation.** This example shows how global conditioning, regression and A/AD combine in the computation of first-order sensitivities for CVA. It also demonstrates exact trade-level allocation of first-order CVA sensitivities. We will see that the computational and implementation advantages previously observed also apply here.

In this example we take LGD as a function of two underlyings, x and y , and PD to be a function of a single underlying, y . Both $LGD(x, y)$ and $PD(x)$ are given by regression equations that are linear in the coefficients of their basis functions.

Since sensitivities are infinitesimal calculations, the scenarios that contribute to CVA sensitivities are exactly those that contribute to CVA value. Thus the scenarios identified by conditioning on positive portfolio value are still the ones we use for the computation of CVA sensitivity. Within each selected scenario we calculate first-order CVA

sensitivity with respect to a calibration instrument s as:

$$\begin{aligned} & \frac{\partial \{f(x) \text{LGD}(x, y) \text{PD}(y)\}}{\partial s} \\ &= \frac{\partial f(x)}{\partial x} \frac{\partial x}{\partial s} \text{LGD}(x, y) \text{PD}(y) \\ &+ f(x) \left(\frac{\partial \text{LGD}(x, y)}{\partial x} \frac{\partial x}{\partial s} + \frac{\partial \text{LGD}(x, y)}{\partial y} \frac{\partial y}{\partial s} \right) \text{PD}(y) \\ &+ f(x) \text{LGD}(x, y) \frac{\partial \text{PD}(y)}{\partial y} \frac{\partial y}{\partial s} \end{aligned} \quad (2)$$

Although $f(x) \text{LGD}(x, y) \text{PD}(y)$ is non-linear in x and y , it is linear in the coefficients of x and y . Essentially, it is a new function:

$$g_*(x, y) = f_*(x) \text{LGD}(x, y) \text{PD}(y)$$

where $*$ can be ‘trade’ or ‘portfolio’. This new function is linear in the coefficients of its basis functions provided that its components $f(x)$, $\text{LGD}(x, y)$ and $\text{PD}(y)$ were. This is valid for each trade individually and for the portfolio. Thus, within the scenarios chosen, this example is identical to the example in the section on first-order sensitivities with trade-level regression.

Putting this together, the first-order sensitivity of CVA to the calibration instrument s is, in our example:

$$\begin{aligned} \frac{\partial \text{CVA}}{\partial s} &= \frac{1}{3} \sum_{f_{\text{portfolio}}(x) \geq 0} \frac{\partial g_{\text{portfolio}}(x, y)}{\partial s} \\ &= \frac{1}{3} \sum_{\text{scenario A or B}} \left(\frac{\partial g_{\text{portfolio}}(x, y)}{\partial x} \frac{\partial x}{\partial s} + \frac{\partial g_{\text{portfolio}}(x, y)}{\partial y} \frac{\partial y}{\partial s} \right) \end{aligned} \quad (3)$$

where we compute over only the two scenarios where the portfolio, $f_{\text{portfolio}}(x)$, is positive. The factor of one-third comes from averaging over all scenarios, although only two scenarios provide any contribution. Note that we retain the separation of trivial differentiation for $g(x, y)$ and the complex derivatives $\partial y/\partial s$ and $\partial x/\partial s$. As before, the more complex derivatives will require A/AD, but they are only needed for the underlyings (x, y) . The derivatives of the underlyings are common for all trades and we therefore retain the implementation simplification.

We can expand (3) to see the trade contributions:

$$\begin{aligned} \frac{\partial \text{CVA}}{\partial s} &= \frac{1}{3} \sum_{f_{\text{portfolio}}(x) \geq 0} \sum_{\text{trades}} \frac{\partial g_{\text{trade}}(x, y)}{\partial s} \\ &= \sum_{\text{trades}} \frac{1}{3} \sum_{f_{\text{portfolio}}(x) \geq 0} \frac{\partial g_{\text{trade}}(x, y)}{\partial s} \\ &= \sum_{\text{trades}} \frac{1}{3} \sum_{\text{scenario A or B}} \left(\frac{\partial g_{\text{trade}}(x, y)}{\partial x} \frac{\partial x}{\partial s} + \frac{\partial g_{\text{trade}}(x, y)}{\partial y} \frac{\partial y}{\partial s} \right) \end{aligned} \quad (4)$$

Note that global conditioning is used on the portfolio value $f_{\text{portfolio}}(x)$ and that we can invert the order of summations. The trade sensitivity contributions are exactly their CVA sensitivity allocations.

Second-order sensitivities for CVA, and their exact allocation, work just like first-order sensitivities for CVA.

Methodology

Having given a set of motivating examples, we now develop the general theory. We start by formally introducing the three key technical

elements: trade-level regression, analytic sensitivity computation and global conditioning. We then combine all three for XVA management.

■ **Trade-level regression.** We express every trade across the whole portfolio in terms of a set of basis functions. That is, we regress the value of each trade across an expanded state space against the basis trade set, at every time point of interest (including $t = 0$), which we term stopping dates. For trades that are Bermudan callable, this is done using early-start Longstaff-Schwartz (Wang & Calfish 2009); for those that are not Bermudan callable, we can apply the simpler, augmented state-space approach (Green & Kenyon 2015a). We are interested in many time points for XVA calculations because these involve integrals over time (Burgard & Kjaer 2013; Green, Kenyon & Dennis 2014).

For each trade U_i , $i = 1, \dots, |\Pi|$, in the overall portfolio Π , we have:

$$U_i(t_k; \xi) = \sum_{l=1}^{|f_{*,k}|} a_{i,j,k} f_{l,k}(B_{\bar{j},k}(t_k; \xi)) + \epsilon_{i,k}(t_k; \xi) \quad \forall k = 1, \dots, |K|, \xi \in \mathcal{E}(t_k) \quad (5)$$

where $f_{l,k}(B_{\bar{j},k})$ are $|f_{*,k}|$ functions of the $|B_{*,k}|$ basis instruments at stopping date k , and there are $|K|$ stopping dates from $t = 0$ to the last date of interest: say, the last cashflow date of the portfolio, t_k . ϵ_* expresses the regression error at a point ξ within the augmented state space $\mathcal{E}(t_k)$ at time t_k . \bar{j} indicates that each $f(\cdot)$ may depend on an arbitrary subset of the basis instruments. Apart from standard regularity conditions, the $f(\cdot)$ have no restrictions. The augmented state space is created either by early-start for a simulation (Wang & Calfish 2009) or by direct augmentation (Green & Kenyon 2015a). The a_* are constants.

In (5) it is critical that only the $a_{i,j,k}$ depend on the trade. The $f_{l,k}$ functions and the $B_{\bar{j},k}$ basis instruments do not depend on the trade. This means that while the basis instruments and function may be arbitrarily complex, they are common for all trades. This is not a significant restriction because of the finite precision of computation.

With a sufficient number of basis functions, $\epsilon_{j,k}(t_k; \xi)$ can be made arbitrarily small and we do not include it further. A portfolio of swaps of 30 years maturity (Green & Kenyon 2015a) showed convergence for a few tens of basis functions for lifetime MVA calculation and general theoretical results are available.

■ **Computation of analytic sensitivities.** We use analytic derivatives from regressions, together with analytic or algorithmic derivatives of underlyings, to obtain sensitivities. Thus, as in Wang & Calfish (2009), we depend on the regression being a good representation of the value function. Convergence of the regression to the value function itself has been extensively studied for diffusions (Glasserman & Yu 2004). Convergence of the derivatives is covered in theorem 1 of Wang & Calfish (2009). Both AD and AAD may be used for derivatives of underlyings with respect to calibration instruments.

■ **Global conditioning.** We compute using global conditioning, meaning that we use global criteria to select scenarios and then compute only on those scenarios. One example of a global criterion is the sign of the value of the portfolio (eg, for CVA). A second example would be whether the value of the portfolio is among the n with the most losses (a typical MVA criterion). We might then use the scenarios

this selects to compute trade-level sensitivities. In computing the sensitivities of MVA, this is clearly a huge advantage as we only compute on a tiny fraction of the total scenarios.

Technically, we are using the linearity property of conditional expectation over filtered probability spaces. A filtered probability space is a probability space with a filtration; ie, an increasing information structure (see Shreve (2004) for details). Valuation adjustment are typically additive both *within* conditional expectations and *across* times. This depends only on the properties of conditional expectation: it is independent of the particular scenario or scenarios that may be selected.

Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is the universe of events, \mathcal{F} is a filtration on Ω and \mathbb{P} is a probability measure on \mathcal{F} . Let $X(t, \omega(t)), Y(t, \omega(t))$ be random variables defined on $(\Omega, \mathcal{F}, \mathbb{P})$, let \mathcal{G} be a sub-filtration of \mathcal{F} where $\omega(t) \in \mathcal{G}(t)$, and let $\alpha(t)$ be a deterministic scalar. An elementary result from definition 2.3.1 and theorem 2.3.2 in Shreve (2004) is then:

$$\begin{aligned} \mathbb{E}[X(t, \omega(t)) + \alpha(t)Y(t, \omega(t)) | \mathcal{G}(t)] \\ = \mathbb{E}[X(t, \omega(t)) | \mathcal{G}(t)] + \alpha(t)\mathbb{E}[Y(t, \omega(t)) | \mathcal{G}(t)] \quad \forall t \end{aligned} \quad (6)$$

So, given a discount bond price $D(t, \omega(t))$, also defined on $(\Omega, \mathcal{F}, \mathbb{P})$, it is obvious (with appropriate regularity conditions) that:

$$\begin{aligned} \int_{t=0}^T \mathbb{E}[(X(t, \omega(t)) + \alpha(t)Y(t, \omega(t)))D(t, \omega(t)) | \mathcal{G}(t)] dt \\ = \int_{t=0}^T \mathbb{E}[X(t, \omega(t))D(t, \omega(t)) | \mathcal{G}(t)] dt \\ + \int_{t=0}^T \alpha(t)\mathbb{E}[Y(t, \omega(t))D(t, \omega(t)) | \mathcal{G}(t)] dt \end{aligned}$$

\mathcal{G} may in turn have sub-filtrations within it; that is, $(\Omega, \mathcal{F}, \mathbb{P})$ may contain nested probability spaces. We have therefore actually demonstrated linearity of computation of the lifetime costs of conditional quantities; eg, exposures (including survival-weighted), value-at-risk (VAR), expected shortfall, etc. Again, Shreve (2004) is a convenient reference for the basic probability framework in the finance context. This is key for allocation and reallocation without resimulation using trade-level regression. It also means that allocation methods can be exact, given \mathcal{G} .

We now apply the technical points to XVA management and show how they simplify, and accelerate, computation.

■ **CVA, DVA and FVA: pricing, sensitivities and allocation.** We start from a generic valuation adjustment of an uncollateralised netting set between a bank B and its counterparty C . For simplicity of exposition we have not made the hazard rate or the discount factor depend on common factors with the exposure (unlike our earlier example with (2)). When this is the case this methodology naturally handles wrong-way risk.

The valuation adjustment of interest can be for CVA, DVA or FVA, so we label it xVA. Its equation from ‘Strategy I: semi-replication with no shortfall at own default’ in Burgard & Kjaer (2013) is:

$$\text{xVA}_*^\nabla = -\text{LGD}_* \int_t^T \lambda_*(u) D_q(t, u) \mathbb{E}_t[V(u)^\nabla] du \quad (7)$$

where ∇ determines which xVA this computes (see table D). $*$ is either B for the bank or C for the counterparty. ∇ can select the

D. Alternatives for $*$ and ∇ that select different xVA possibilities in (7) in the text (equation adapted from Burgard & Kjaer (2013))

xVA_*^∇	$*$	∇
CVA	C	$+$
DVA	B	$-$
FCA	B	$+$
FVA = DVA + FCA	B	(blank)

positive exposure, the negative exposure or do nothing, depending on the particular valuation adjustment. $D_q(t, u)$ is the discount factor between u and t for the rate q , $q = r + \lambda_B + \lambda_C$: r is the riskless rate and λ_* is the hazard rate of the bank or counterparty. V is the unadjusted value of the netting set.

Let the netting set V be made up of a set of trades π_i , so (7) becomes:

$$\text{xVA}_*^\nabla = -\text{LGD}_* \int_t^T \lambda_*(u) D_q(t, u) \mathbb{E}_t \left[\left(\sum_i \pi_i(u) \right)^\nabla \right] du \quad (8)$$

Now suppose the expectation and time integral are both computed using sets of observations (eg, generated by simulation). We then have (using the simplest possible time-integration scheme):

$$\begin{aligned} \text{xVA}_*^\nabla = -\text{LGD}_* \sum_{k=1}^{n_k} (t_k - t_{k-1}) \lambda_*(u_k) D_q(t, u_k) \\ \times \frac{1}{n_j} \sum_{j=1}^{n_j} \left(\sum_i \pi_i(u_k; \omega_{j,k}) \right)^\nabla \end{aligned}$$

where n_k is the number of time steps and n_j is the number of scenarios at each time step. $\omega_{j,k}$ represents the realisation of the random factors at time t_k in scenario j . $\omega_{j,k} \in \mathcal{G}(t_k)$, where \mathcal{G} is the filtration.

Applying (5), ie, using regressions for the trade values, we obtain:

$$\begin{aligned} \text{xVA}_*^\nabla = -\text{LGD}_* \sum_{k=1}^{n_k} (t_k - t_{k-1}) \lambda_*(u_k) D_q(t, u_k) \\ \times \frac{1}{n_j} \sum_{j=1}^{n_j} \left(\sum_i \sum_{l=1}^{|f_{*,k}|} a_{i,l,k} f_{l,k}(B_{\bar{j},k}(t_k; \omega_{j,k})) \right)^\nabla \\ = -\text{LGD}_* \sum_{k=1}^{n_k} (t_k - t_{k-1}) \lambda_*(u_k) D_q(t, u_k) \\ \times \frac{1}{n_j} \sum_{j=1}^{n_j} \left(\sum_{l=1}^{|f_{*,k}|} a_{l,k} f_{l,k}(B_{\bar{j},k}(t_k; \omega_{j,k})) \right)^\nabla \\ a_{l,k} = \sum_i a_{i,l,k} \end{aligned}$$

Since the basis instruments are common for all trades we now have an equation involving only these instruments. We can expand the equation for xVA_*^∇ above to remove the non-linearity of the $(\cdot)^\nabla$ bracket as:

$$\begin{aligned} \text{xVA}_*^\nabla = -\text{LGD}_* \sum_{k=1}^{n_k} (t_k - t_{k-1}) \lambda_*(u_k) D_q(t, u_k) \\ \times \frac{1}{n_j} \sum_{j=1}^{n_j} I_{V_{j,k}^\nabla} \sum_{l=1}^{|f_{*,k}|} a_{l,k} f_{l,k}(B_{\bar{j},k}(t_k; \omega_{j,k})) \end{aligned} \quad (9)$$

$$\begin{aligned}
&= -\text{LGD}_* \sum_{k=1}^{n_k} (t_k - t_{k-1}) \lambda_*(u_k) D_q(t, u_k) \\
&\quad \times \frac{1}{n_j} \sum_{j | V_{j,k}^\nabla} \sum_{l=1}^{|f_{*,k}|} a_{l,k} f_{l,k}(B_{\bar{j},k}(t_k; \omega_{j,k})) \quad (10)
\end{aligned}$$

The boxes indicate the key point: we compute with respect to global scenarios selected by our conditioning criteria. This is an example of global conditioning. We use $I_{V_{j,k}^\nabla}$ as the indicator function on the sign of the unadjusted future netting set value (see table D for choices). In (10) we have selected, at each time point, those scenarios such that the netting set value satisfies the criteria (ie, the same selection as the indicator function). These scenarios will be different for each time point.

It may appear that we have done more work, not less, but xVA pricing is a first step: our target is XVA *management*. We now demonstrate what we have achieved for sensitivities, allocation of XVA and allocation of XVA sensitivities. In general, all of these calculations are more costly than the initial xVA computation. Below we will use exactly the same regressions for MVA computation, sensitivities, allocation and allocation of MVA sensitivities. Thus, even for xVA pricing, we will demonstrate significant advantages.

■ **Sensitivities.** By sensitivity we mean sensitivity with respect to hedging; ie, calibration, instruments. We assume sensitivities are being computed analytically.

Obviously, we have immediately reduced the implementation cost of analytic derivatives from all the trade types found in the entire portfolio to the set of basis instruments of the regression. This will usually represent a major saving in implementation time.

The next key observation is that since analytic sensitivities are based on infinitesimal changes and we compute at finite precision, the set of scenarios we calculate over, $j | V_{j,k}^\nabla$, is unchanged. Hence, for a calibration instrument s that we want a sensitivity for:

$$\frac{\partial \text{xVA}_*^\nabla}{\partial s} = \dots \frac{\partial f_{l,k}(B_{\bar{j},k}(t_k; \omega_{j,k}))}{\partial B_{\bar{j},k}(t_k; \omega_{j,k})} \frac{\partial B_{\bar{j},k}(t_k; \omega_{j,k})}{\partial s}$$

Alternatively:

$$J_{\text{xVA}_*^\nabla, s} = \dots J_{f_{l,k}(\cdot), B_{\bar{j},k}(\cdot)} J_{B_{\bar{j},k}(\cdot), s}$$

where $J_{*,*}$ are the Jacobians.

We have pre-selected the scenarios $j | V_{j,k}^\nabla$ to calculate over so everything is linear, and since differentiation is a linear operator, this remains the case. Generally, $J_{f_{l,k}(\cdot), B_{\bar{j},k}(\cdot)}$ will be analytic because the $f_{l,k}(\cdot)$ will have been selected for that property.

■ **Trade-level allocation of xVA.** Allocation of valuation adjustment prices to desks is a core activity of XVA desks. Trade-level allocation of xVA is given directly from (10) by considering the contribution of each trade in terms of its regression coefficients. For example, for trade i :

$$\begin{aligned}
\text{xVA}_*^\nabla &= -\text{LGD}_* \sum_{k=1}^{n_k} (t_k - t_{k-1}) \lambda_*(u_k) D_q(t, u_k) \\
&\quad \times \frac{1}{n_j} \sum_{j | V_{j,k}^\nabla} \sum_{l=1}^{|f_{*,k}|} a_{i,l,k} f_{l,k}(B_{\bar{j},k}(t_k; \omega_{j,k}))
\end{aligned}$$

where the difference from (10) is that we now use trade i 's regression coefficients $a_{i,l,k}$ rather than the netting-set regression coefficients $a_{l,k}$. Allocation is both exact and additive using trade-level regression with global conditioning; ie, computing within the selected scenarios $j | V_{j,k}^\nabla$. Thus, reallocation – the reallocation of different trades' xVA to different groupings (eg, for reporting) – is trivial.

■ **Trade-level allocation of sensitivities.** Hedging costs are often derived from sensitivities and thus trade-level allocation of these sensitivities is a core activity of XVA desks. Since differentiation is a linear operator, we can combine the arguments of the previous two sections to observe that using our regression and conditioning approach with respect to a calibration instrument s :

$$\frac{\partial \text{CVA}(\Pi^C, t)}{\partial s} = \sum_i \frac{\partial \text{CVA}(\pi_i^C, t)}{\partial s}$$

where we only compute 'trade' sensitivities *within selected scenarios*. We put trade in quotes because we can arbitrarily create new trade groupings using the additivity of their regression coefficients. We can thus allocate, and reallocate, hedging costs freely. That is, the costs are linear and we only reallocate sums of scalar numbers to different pots (desks, groups, etc.). In addition, we only compute sensitivities at the coarsest level required using the appropriate regression coefficients.

■ **Incremental xVA.** During a trading day there will be continual changes to portfolios, and the XVA desk must provide prices for these changes to other desks. Portfolio changes can be expected to change the conditioning set $j | V_{j,k}^\nabla$. For xVA, the conditioning set is specific to each counterparty. First note that we have already calculated the unconditioned portfolio values in each scenario:

$$V_{j,k}^{\text{Unconditioned}}(\Pi) = V_{j,k}$$

We follow the same procedure for the new trades as for the existing portfolio by calculating their trade-level regressions. Now we calculate their values for the same set of overall scenarios as the existing portfolio, as well as calculating the updated conditioning scenarios:

$$\begin{aligned}
j | V_{j,k}^\nabla(\Pi(\text{original}) + \Pi(\text{changes})) \\
&= j | \{j | V_{j,k}(\Pi(\text{original})) + j | V_{j,k}(\Pi(\text{changes}))\}^\nabla \\
&= j | V_{j,k}^\nabla(\Pi(\text{updated}))
\end{aligned}$$

To compute the first line above we need only the scenario values of the original portfolio and the changes to the portfolio. No recomputation of the original portfolio is required. Thus we can recompute the CVA without recomputing the original portfolio; we just include the previous values for the additional scenarios.

Again, it appears we have computed the regression of the changes to the portfolio as extra work. However, this extra work makes the other XVA elements and their management (sensitivities and allocation) orders of magnitude faster.

For incremental sensitivities the arguments of the previous sections apply directly. This is also true for both incremental trade-level allocation and incremental trade-level allocation of sensitivities.

■ **MVA pricing, sensitivities and allocation.** Central counterparties often require posting of initial margin (IM). The lifetime cost of funding this IM is termed margin valuation adjustment (MVA).

As indicated in the examples section, our technique also applies to the lifetime cost of funding IM (ie, MVA). IM for trades with central counterparties is often based on VAR and/or ES, so we consider these next. As shown in the examples, the use of global scenario selection makes VAR and ES computation additive.

ES is a conditional expectation by definition. Thus the derivation from CVA above applies exactly. Furthermore, since we have already calculated the trade-level regression functions and analytic derivatives for CVA, there is no need to recompute them for ES. The only change between CVA and ES is a different condition:

$$j \mid \mathbb{P}_{\text{ES}}(V_{j,k}) \leq \alpha$$

instead of:

$$j \mid V_{j,k}^{\nabla}$$

Once the scenarios are identified, the same computations apply. We have used \mathbb{P}_{ES} for the distribution of portfolio values as used for ES. This distribution will usually be on a sub-filtration of Ω equivalent to that used for the risk-neutral measure (but obviously with a different measure). We use α for the percentile of interest: typically around 97.5%.

We approach VAR as a limit of ES definitions, ie:

$$j \mid \mathbb{P}_{\text{ES}}(V_{j,k}) = \alpha = \lim_{\beta \rightarrow \alpha} \{j \mid \beta \leq \mathbb{P}_{\text{ES}}(V_{j,k}) \leq \alpha\}$$

Thus the development above also applies to VAR. Only conditioning has changed, so the previous arguments apply exactly to all management cases.

Conclusions

We have shown how XVA (CVA, FVA, DVA and MVA) management is radically more efficient using a combination of three technical elements: trade-level regression, analytic derivatives and global conditioning. Regression for KVA is covered in Green & Kenyon (2014). This methodology naturally handles wrong-way risk. Furthermore, as trade prices are computed separately from the underlyings, the underlying dynamics can be almost arbitrarily complex to match time-zero pricing. Thus, only a single system may be required for both real-time trade pricing and XVA.

The main limitation of the technique in this paper is that it is essentially a first-order approach, excluding option exercise boundary changes from XVA interactions. This is a topic of further research (Green & Kenyon 2015b). Second-order sensitivities present more numerical challenges than first-order ones, and many implementation- and portfolio-specific details will matter. **R**

Andrew Green is head of the CVA/FVA quantitative research team at Lloyds Banking Group in London and Chris Kenyon is a director in the team. The views expressed are those of the authors only; no other representation should be attributed. Email: andrew.green2@lloydsbanking.com, chris.kenyon@lloydsbanking.com.

REFERENCES

- | | | | |
|---|--|--|--|
| Broadie M and P Glasserman, 1996
<i>Estimating security price derivatives using simulation</i>
<i>Management Science</i> 42(2), pages 269–285 | Giles M and P Glasserman, 2006
<i>Smoking adjoints: fast Monte Carlo Greeks</i>
<i>Risk</i> January 1, pages 92–96 | Green A and C Kenyon, 2015a
<i>MVA by replication and regression</i>
<i>Risk</i> May, pages 82–87 | Shreve, S, 2004
<i>Stochastic Calculus for Finance, Volume II: Continuous-Time Models</i>
Springer |
| Burgard C and M Kjaer, 2013
<i>Funding strategies, funding costs</i>
<i>Risk</i> December, pages 82–87 | Glasserman P and B Yu, 2004
<i>Number of paths versus number of basis functions in American option pricing</i>
<i>Annals of Applied Probability</i> 14, pages 2,090–2,119 | Green A and C Kenyon, 2015b
XVA at the exercise boundary
Forthcoming | Tasche, D, 2008
Capital allocation to business units and sub-portfolios: the Euler principle
In <i>Pillar II in the New Basel Accord: The Challenge of Economic Capital</i> , pages 423–453 |
| Capriotti L, J Lee and M Peacock, 2011
<i>Real-time counterparty credit risk management in Monte Carlo</i>
<i>Risk</i> June, pages 86–90 | Green A, 2015
XVA: Credit, Funding and Capital Valuation Adjustments
Wiley (forthcoming) | Green A, C Kenyon and CR Dennis, 2014
<i>KVA: capital valuation adjustment by replication</i>
<i>Risk</i> December, pages 82–87 | Wang Y and R Calfish, 2009
<i>Pricing and hedging American-style options: a simple simulation-based approach</i>
<i>Journal of Computational Finance</i> 13(4), pages 95–125 |
| Cesari G, J Aquilina, N Charpillon, Z Filipovic, G Lee and I Manda, 2010
<i>Modelling, Pricing, and Hedging Counterparty Credit Exposure: A Technical Guide</i>
Springer Finance | Green A and C Kenyon, 2014
<i>Portfolio KVA: I Theory</i>
Preprint (SSRN, ssrn.com/abstract=2519475) | Kenyon C and A Green, 2015
<i>Warehousing credit risk: pricing, capital and tax</i>
<i>Risk</i> February, pages 70–75 | |