REPRINTED FROM

Risk

RISK MANAGEMENT • DERIVATIVES • REGULATION

Risk.net January 2015

Cutting edge

Valuation adjustments

Backward induction for future values

Alexandre Antonov, Serguei Issakov and Serguei Mechkov generalise the American Monte Carlo method to efficiently calculate future values (or exposures) of derivatives using an arbitrage-free model. The resulting algorithm is especially attractive for exotic portfolios.

Backward induction for future values

Here, Alexandre Antonov, Serguei Issakov and Serguei Mechkov generalise the American Monte Carlo method to efficiently calculate future values (or exposures) of derivatives using an arbitrage-free model. The resulting algorithm is especially attractive for exotic portfolios

The American (or least squares) Monte Carlo method[1] in its original formulation (see, for example, Carriere 1996; Longstaff & Schwartz 2001; and see also Glasserman (2003) for a more complete list of references) uses a backward induction to compute the continuation value of a derivative. In this article we generalise the backward induction to compute a future value of a derivative that corresponds to the full instrument value on future dates with effects of exercises and triggers included.

This method allows for an efficient simulation of exposures (distributions of future values) in the contexts of:

- various valuation adjustments (XVA) due to counterparty risk, funding, capital, etc,
- calculation of risk measures that use averages of future values, such as value-at-risk and expected shortfall for market risk, and potential future exposure (PFE), expected positive exposure/expected negative exposure (EPE/ENE), etc, for counterparty risk,
- scenario generation, also in the real-world measure.

We refer to the pure backward induction for computing future values as the algorithmic exposure method. The word 'algorithmic' indicates that the computation is done essentially within the standard pricing procedure, provided that backward induction is enhanced with some extra terms described below.

Our approach can easily be applied to complicated instruments (eg, flexi caps, target accrual redemption notes (Tarns), barrier triggers, path-dependent instruments, etc) where a 'manual' future values calculation can be very tricky.

The presented method of computation of future values can be formulated in arbitrary probability measure. We propose calibration of the Radon-Nikodym process corresponding to the change of measure to projections of market rates and indexes.

The article is organised as follows. First we define the future value of financial instrument, then explain our methodology of the algorithmic exposure calculation for exotic portfolios. Finally, we give insights on real-world measure calibration, with numerical results.

Future value

In this section we define the future value of a generic instrument that can be priced by backward Monte Carlo methods (see, for example, Andersen & Piterbarg 2010; Fries 2007). That includes callable instruments, instruments with barrier triggers, and callable path dependent deals.

Consider a payment of a certain amount $\mathcal{A}$ on a date τ. The present value (PV) of the payment is the discounted expectation $A = \mathbb{E}[\mathcal{A}/N(\tau)]$, where $N(t)$ is the model numeraire and $\mathbb{E}[\cdots]$ is the pricing expectation in the model measure. We refer to the numeraire currency as the domestic one and require all payments to be expressed in this currency. The payment seen from some observation date $t < \tau$ is the conditional expectation (CE), $A(t) = N(t)\,\mathbb{E}[\mathcal{A}/(N(\tau)) \mid \mathcal{F}_t]$. For any regular path-independent payoff $\mathcal{P}(T)$ at time T, the discounted CE at time t is a continuous function of model states $x_i(t)$:

$$N(t)\mathbb{E}\bigg[\frac{\mathcal{P}(T)}{N(T)}\,\bigg|\,\mathcal{F}_t\bigg] = f(t; x_i(t))$$

For example, in a hybrid cross-currency setup, the model states can be short rates and foreign exchange rates. We suppose, of course, that the model is Markovian in finite dimensions. Below we assume the possibility of a CE calculation embedded in the Monte Carlo (MC) simulation. A typical numerical method for CE calculation is regression to state variables[2] (Carriere 1996; see also Andersen & Piterbarg 2010; Fries 2007).

A general instrument is defined as a payment stream of amounts $\mathcal{A}_j$ paid at τ_j for $j = 1, \ldots, M$. The amounts are expressed in domestic currency units. They are linked with financial market rates (eg, Libor, constant maturity swaps (CMSs)) and, eventually, obey exercise or barrier conditions. The instrument's price today is the PV (discounted expectation) of the payments:

$$V = \mathbb{E}\bigg[\sum_{j=1}^{M} \frac{\mathcal{A}_j}{N(\tau_j)}\bigg] \tag{1}$$

The instrument's future price or future value (FV) at a certain observation date t_{obs} is a conditional expectation of future payments:

$$v = N(t_{\text{obs}})\,\mathbb{E}\bigg[\sum_{j=1,\tau_j>t_{\text{obs}}}^{M} \frac{\mathcal{A}_j}{N(\tau_j)}\,\bigg|\,\mathcal{F}_{t_{\text{obs}}}\bigg] \tag{2}$$

Let us stress the difference between the product PV (a number) and the FV (a stochastic variable). Here and below, we use the same letter for the PV and the FV but capitalise the PV notation.

It is important to note that the cashflows $\mathcal{A}_j$ in the formula above may depend on exercises before the observation date. For example, an FV of a Bermudan swaption giving the right to enter into a swap will include the swap payments after the observation date but only for the paths on which we exercised before t_{obs}. This means that for the FV calculation we need to know all the exercise events that happened before the observation date. Thus, a direct simulation of the Bermudan swaption FV will include a backwards pricing procedure which will determine exercise indicators (or exercise boundary) with its subsequent forward pass aggregating cashflows with exercises. This two-direction (or 'manual') procedure for Bermudan options was first

[1] We suggest referring to the American Monte Carlo method as the Las Vegas method for its brevity and geographical consistency.

[2] Note that, for path-dependent payoffs, the model space is augmented and the regression functions should cover the additional dependencies.

1 Instrument tree representation

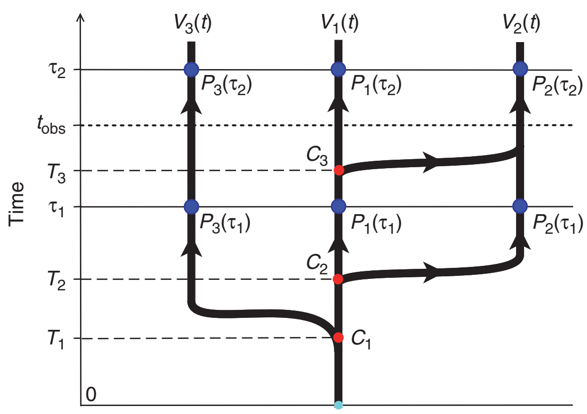

proposed in Cesari *et al* (2010). In this paper we present a generalisation of the backward induction that allows us to calculate the future value in a purely algorithmic way. This is done by extending the standard backward pricing procedure by a number of adjoint calculations on every step.

Algorithmic computation of FVs in backward induction

Algorithmic rules In this section we present the backward iterative procedure for future values calculation of an arbitrary backward instrument. To introduce notation, consider the multi-callable instrument depicted in figure 1.

The instrument described by the graph (a rooted and ordered tree[3]) corresponds to a Bermudan option with an extra possibility to enter a different payment stream. The instrument starts with the node at the origin. The other nodes designate the instrument events: blue nodes (P_i) mean payments, and red ones (C_i) mark exercises. There are three possible 'branches' in the instrument: the middle branch V_1 is associated with the 'default' payments stream (ie, without exercise); the two other branches correspond to payment streams after eventual exercises. The right branch (payment stream V_2) can be entered twice, on exercise dates T_2 and T_3. The left branch (payment stream V_3) can be entered once on T_1. The payments happen on payment dates τ_1 and τ_2: for the ith branch we pay/receive amounts $\mathcal{P}_i(\tau_k)$.

Turn now to the general case of an arbitrary number of branches. The main component of the backward pricing procedure is a branch continuation value (CV) at time t expressed in domestic currency units, $V_j(t)$, which is a certain function of the model state variables at t. Financially, $V_j(t)$ means a hold value: the price of staying with the branch j at time t with all possible payments and exercises after t. For coherence, we express all the CVs in domestic currency units.

A simple CV example is a leg with a single unit payment at time T in some currency c. Its value in domestic currency with forex rate $X_c(T)$ is simply $V_{X_c}(T) = X_c(T)$. The CV can be seen from an earlier time t as a currency-c zero-coupon bond in domestic currency:

$$V_{X_c}(t) = P_c(t,T) = N(t)\,\mathbb{E}\left[\frac{X_c(T)}{N(T)}\,\middle|\,\mathcal{F}_t\right] \tag{3}$$

[3] The recombining character of the tree reflects multiple exercises into the same instrument.

The backward pricing procedure consists of transformation of the CVs on instrument dates (including payment, exercise and trigger dates) and their backwards propagation (a discounted conditional expectation) between the dates.

The transformations may include payments and exercises at an instrument date T. They reduce to a linear combination:

$$V_n(T^-) = \sum_m \gamma_{m,n} V_m(T) \tag{4}$$

with dimensionless coefficients $\gamma_{m,n}$.

Here and below we will denote by T^- a moment in time preceding T by a 'quant' of business time: $T^- = T - \mathrm{d}T$. Thus, the CV prices are continuous from above. By avoiding special notations for the CVs in the left-hand side, this convention serves to simplify a writing of exercise and payment conditions.

The discounted conditional expectation between two instrument dates ($t < T$) is:

$$V_n(t) = N(t)\,\mathbb{E}\left[\frac{V_n(T^-)}{N(T)}\,\middle|\,\mathcal{F}_t\right] \tag{5}$$

As mentioned above, the most common numerical method to compute the conditional expectation is regression.

The coefficients $\gamma_{m,n}$ are the other major component of the backward pricing. We call them generalised rates because they are closely related to the market rates (eg, Libor, CMS, etc) multiplied by an appropriate daycount fraction. In addition, if a barrier or exercise condition is present, the generalised rates include the corresponding indicators. Generally speaking, these rates are dimensionless coefficients absorbing all the nonlinearities that appear in the instrument pricing. The coefficients are determined separately for every transformation date with no direct relation to other such dates. This is the exact opposite of the CV objects, which only allow linear combinations (with the coefficients $\gamma_{m,n}$) and connected in time via the conditional expectation (5).

For the moment we assume a single observation date t_{obs}. In the next section, we generalise the algorithm to multiple observation dates and show that it is suitable for the XVA calculation related with FVs on multiple horizons.

To calculate future values u_m at t_{obs} for all branches by backward induction, we use the corresponding observation values (OVs) $v_m(t)$. The OV is computed gradually back in time, together with the CV. Once at the origin, the OV $v_n(0)$ gives the FV v_n. Note the difference between the product PV (a number) and the FV (a stochastic variable). As both FV and OV depend on the observation date, strictly speaking we should write $v_m(t; t_{\mathrm{obs}})$, but we omit the last argument.

The OVs are initialised on the observation date by the corresponding CVs:

$$v_m(t_{\mathrm{obs}}) = V_m(t_{\mathrm{obs}}) \tag{6}$$

Now, the backward-propagation rules (4) and (5) algorithmically induce the following concise rules for the OVs for $t < t_{\mathrm{obs}}$:

- Linear update rules for OVs are the same as for CVs:

$$V_k(T^-) = \sum_m \gamma_{m,n} V_m(T) \quad \Rightarrow \quad v_k(T^-) = \sum_m \gamma_{m,n} v_m(T) \tag{7}$$

where $\gamma_{m,n}$ are generalised market rates (for example, numbers or barrier exercise indicators).

■ No changes for the OVs during the backwards propagation of the CVs:[4]

$$V_n(t) = N(t)\mathbb{E}\left[\frac{V_n(T^-)}{N(T)}\,\middle|\,\mathcal{F}_t\right] \quad\Rightarrow\quad v_n(t) = v_n(T^-) \tag{8}$$

This way, the exposure is finalised once the backward procedure reaches the origin, $v_m = v_m(0)$.

The procedure of the future value calculation is purely backward and automatic: it avoids complications of a manual method caused by cumbersome exercises aggregation (below we will comment on the manual procedure on examples of a Bermudan swaption and a flexi cap). Based on this backward induction, we will present a new algorithm, a robust software design suitable for simultaneous valuation of instrument prices and their future values.

As a simple illustration, consider a CV of a unit payment in currency c at $T > t_{\text{obs}}$. When the backward pricing time t reaches the observation date, the OV is initialised by the running CV (3), $v_{X_c}(t_{\text{obs}}) = V_{X_c}(t_{\text{obs}}) = P_c(t_{\text{obs}}, T)$, and remains unchanged for $t \leqslant t_{\text{obs}}$. As a consequence, this payment added to a leg V at time $t < t_{\text{obs}}$ will induce the following OV transformation:

$$V(t^-) = V(t) + \alpha\, P_c(t,T) \quad\Rightarrow\quad v(t^-) = v(t) + \alpha\, P_c(t_{\text{obs}}, T)$$

Instrument examples A simple but didactic instrument example is a Bermudan swaption giving a right to enter into a swap on exercise dates T_i. Other examples, including a barrier option, can be found in Antonov, Issakov & Mechkov (2011). Suppose the underlying swap pays a generalised rate α_j in some currency c_j at a date τ_j for $j = 1,\dots,M$. Then the swaption PV and FV are built with payments subjected to the exercise conditions

$$\left.\begin{aligned} V &= \mathbb{E}\left[\sum_{j=1}^{M} \frac{\mathcal{I}(\tau_j)\alpha_j X_{c_j}(\tau_j)}{N(\tau_j)}\right] \\ v &= N(t_{\text{obs}})\mathbb{E}\left[\sum_{j=1,\,\tau_j>t_{\text{obs}}}^{M} \frac{\mathcal{I}(\tau_j)\alpha_j X_{c_j}(\tau_j)}{N(\tau_j)}\,\middle|\,\mathcal{F}_{t_{\text{obs}}}\right] \end{aligned}\right\} \tag{9}$$

where indicator $\mathcal{I}(\tau_j)$ equals 1 if we have entered into the swap before the payment date τ_j and zero otherwise.

Let $V(t)$ and $S(t)$ denote the CVs of the swaption and the swap, respectively. Their OVs for observation date t_{obs}, denoted $v(t)$ and $s(t)$, respectively, are initialised as prescribed by (6). We will refer to instrument dates as a union of exercise and payment dates. The swaption's backward pricing and the induced rules for the OVs are performed by two repeated steps.[5]

[4] In Antonov, Issakov & Mechkov (2011), we considered a scaled version of the current OV by dividing by the corresponding numeraire. This slightly modifies the propagation rule.

[5] Recall that the CVs are modified on all instrument dates; the OVs are active only before the observation date, $t < t_{\text{obs}}$.

The first step, (7), includes transformations on the instrument dates:

$$\begin{aligned} V(T_j^-) &= \max(V(T_j), S(T_j)) \\ &\Rightarrow\quad v(T_j^-) = v(T_j)\,(1-\mathcal{C}_j) + s(T_j)\,\mathcal{C}_j \quad \text{for } T_j < t_{\text{obs}} \end{aligned} \tag{10}$$

$$\begin{aligned} S(\tau_k^-) &= S(\tau_k) + \alpha_k\, X_{c_k}(\tau_k) \\ &\Rightarrow\quad s(\tau_k^-) = s(\tau_k) \quad \text{for } \tau_k < t_{\text{obs}} \end{aligned} \tag{11}$$

where a conditional exercise indicator at time T_j:

$$\mathcal{C}_j = \mathbb{1}_{S(T_j) > V(T_j)}$$

is equal to 1 if we enter the swap at time T_j provided that the swaption was not exercised before. Note that there is no contribution from the kth payment to $s(\tau_k)$ because $\tau_k < t_{\text{obs}}$.

The second step, (8), is a propagation between the instrument dates; the discounted conditional expectation of the CVs induces the invariability of the OVs.

The backward procedure delivers the swap's today price $S(0)$ and the swaption's today price $V(0)$ along with the exercise indicators. As a by-product, we also obtain FVs of the swap and the swaption. Indeed, we can formally prove (see Antonov, Issakov & Mechkov 2011) that the result of the OV recursion, $v(0)$, does coincide with the definition (9) for v provided that the (unconditional) exercise indicators are $\mathcal{I}(t) = \mathcal{I}_j$ for $T_j \leqslant t < T_{j+1}$ and $\mathcal{I}_j = 1 - \prod_{i=1}^{j}(1-\mathcal{C}_i)$. We note that the indicator is path-dependent.

The FV recursion becomes transparent: the differences between PV and FV in (9) are the presence of a conditional expectation in the FV and the absence of the past payments with respect to the observation date. That is why we have stopped the regression in the OVs calculation before the observation date along with the payment updates.

The backward price $V(0)$ is not the only numerical way to calculate PVs of exotics. Alternatively, we can represent the instrument PV as in (9), simulate the payments $\alpha_j X_{c_j}(\tau_j)$, reuse the exercise indicators $\mathcal{I}_j$ from the backward pricing and take the simple average of the resulting sum. This way of Bermudan option evaluation has less bias than the traditional backward-induction result $V(0)$ because the regression error appears only in exercise indicator calculations and not in the swap payments. Note that slightly modifying the OV logic allows us to express the swaption PV in this more advantageous form. See also Fries (2007, Section 15.4), Egloff, Kohler & Todorovic (2007) and Andersen & Piterbarg (2010, Section 18.3.4).

Alternatively (see Antonov, Issakov & Mechkov 2011), the swaption FV at t_{obs} can be written as:

$$v = V(t_{\text{obs}})\,(1-\mathcal{I}(t_{\text{obs}})) + S(t_{\text{obs}})\,\mathcal{I}(t_{\text{obs}}) \tag{12}$$

and evaluated in a manual two-step way. The first step is a backward pricing coupled with storage of the CVs $V(t_{\text{obs}})$ and $S(t_{\text{obs}})$ as well as the indicators $\mathcal{C}_k$ for all k. The second step is a forward pass while aggregating $\mathcal{C}_k$ into the global exercise indicators $\mathcal{I}_k$ and assembling in the final formula (12).

Consider now a flexi cap, an instrument with multiple (a limited number of) exercises. A flexi cap gives the right to enter N times into floating-fixed exchanges in domestic currency on a set of dates T_i.

Each payment index at T_{i+1} is defined as $L_i - K$, where L_i is a Libor fixed at time T_i and K is a fixed rate. Thus, the payment as seen at time T_i is given by $H(T_i) = (L_i - K)\,P_d(T_i, T_{i+1})$ where $P_d(T_i, T_{i+1})$ is a zero bond in domestic currency. For transparency we limit ourselves to two exercises $N = 2$ (see Antonov (2004) for more exercises).

The flexi cap pricing is done via a backward induction. Introduce two CVs: $V^{(1)}$ for the product with one exercise and $V^{(2)}$ for the product with two exercises (the instrument to evaluate). The pricing logic:

$$V^{(1)}(T_j^-) = \max(V^{(1)}(T_j), H(T_j)) \tag{13}$$

$$V^{(2)}(T_j^-) = \max(V^{(2)}(T_j), H(T_j) + V^{(1)}(T_j)) \tag{14}$$

is transparent: having a two-exercise product we choose to stay with it or exercise into the floating-fixed payment and one-exercise instrument (14). If we have already exercised once, we have $V^{(1)}$ and can exercise into the floating-fixed payment (13). The iterative procedure which is initialised on T_M by $V^{(n)}(T_M) = 0$ for $n = 1, 2$ gives the flexi cap PV at origin $V = V^{(2)}(0)$.

In order to calculate the future values we denote by $v^{(i)}(t)$ the corresponding OVs for observation date t_{obs} and initialise them as prescribed by (6). The flexi cap backward pricing and the induced rules for the OVs are performed by two repeated steps.

Given the CV transformations (13)–(14) the first step (7) leads to the following OV ones:

$$v^{(1)}(T_j^-) = v^{(1)}(T_j)\bar{\mathcal{C}}_j^{(1)} + \delta_{jk}\,H(T_k)\,\mathcal{C}_j^{(1)} \tag{15}$$

$$v^{(2)}(T_j^-) = v^{(2)}(T_j)\bar{\mathcal{C}}_j^{(2)} + (\delta_{jk}\,H(T_k) + v^{(1)}(T_j))\mathcal{C}_j^{(2)} \tag{16}$$

where the Kronecker delta-symbol $\delta_{kj} = 1_{k=j}$ in front of the floating-fixed payment reflects zero OV of the CV $H(T_j)$ for $j < k$. Here, the indicator:

$$\mathcal{C}_j^{(2)} = 1_{V^{(2)}(T_j) < H(T_j) + V^{(1)}(T_j)}$$

is equal to 1 if we exercise at T_i provided that we have not exercised before T_i (and zero otherwise). Analogously:

$$\mathcal{C}_j^{(1)} = 1_{V^{(1)}(T_j) < H(T_j)}$$

is equal to 1 if we exercise at T_i provided that we have already exercised once before T_i (and zero otherwise).

The second step (8) is a propagation between the instrument dates; the discounted conditional expectation of the CVs induces the invariability of the OVs.

The iterative backward procedure delivers the flexi cap today price $V^{(2)}(0)$ along with the exercise indicators. As a by-product, we also obtain the corresponding FVs. We can formally prove (see Antonov, Issakov & Mechkov 2011) that the algorithmic result of the OV recursion, $v(0)$, does coincide with a manual construction for v described below in (17).

Indeed, denote by $\mathscr{I}_k^{(2)}$ an indicator that there was no exercise up to T_k, inclusively (ie, it rests two exercises after T_k). Similarly, $\mathscr{I}_k^{(1)}$ is an indicator of only one exercise up to T_k, inclusively (ie, it rests one exercise after T_k). We call $\mathcal{E}_k$ an indicator of an exercise at T_k (the first or the second). These global (unconditional) exercise indicators can be obtained using the conditional ones $\mathcal{C}_j^{(i)}$ (see Antonov, Issakov & Mechkov 2011 for details).

It is clear that the FV at observation date $t_{\text{obs}} = T_k$ is composed of the CV with two exercises $V^{(2)}(T_k)$ provided that we have not exercised before ($\mathscr{I}_k^{(2)} = 1$), plus the CV with one exercise $V^{(1)}(T_k)$ provided that we have exercised once before ($\mathscr{I}_k^{(1)} = 1$), plus the floating-fixed payment at T_{k+1} (seen at T_k as $H(T_k)$) if we have exercised exactly on T_k, ie:

$$v = \mathscr{I}_k^{(2)}\,V^{(2)}(T_k) + \mathscr{I}_k^{(1)}\,V^{(1)}(T_k) + \mathcal{E}_k\,H(T_k) \tag{17}$$

Software design and numerical performance Based on backwards induction for FVs, we present here our software design permitting algorithmic calculation of the FVs as by-product of backward pricing. Below we describe the methodology of dealing with multiple observation dates $\{t_{\text{obs}}^{(n)}\}$, $n = 1, \dots, N$. The resulting FVs can be used for the XVA calculations where the exposure should be aggregated across multiple time horizons.

Suppose that during the backward pricing procedure, we manipulate the CVs ($V_m(t)$, $m = 1, \dots, M$) of multiple products on the instrument dates.[6] To calculate the FVs of the products, we add extra entities (OVs corresponding to the observation dates) to each CV:

$$V_m(t) \to \{V_m(t); v_m(t; t_{\text{obs}}^{(1)}), \dots, v_m(t; t_{\text{obs}}^{(N)})\}.$$

Then, each time we modify the CV, its observation counterparts are also modified according to the simple algorithmic rules (7) and (8). It is important to stress that the initialisation rule (6) requires calculation of the CVs on the observation dates. This results in numerical overheads compared with the pure pricing procedure, where we need the CVs only on instrument dates. (See also the end of this section.)

To calculate the FVs of the product on all the observation dates, we repeat the general rules (7) and (8) independently for each OV. For example, the OV $v_m(t; t_{\text{obs}}^{(n)})$ becomes initialised when the CV backward pricing comes to $t_{\text{obs}}^{(n)}$:

$$v_m(t_{\text{obs}}^{(n)}; t_{\text{obs}}^{(n)}) = V_m(t_{\text{obs}}^{(n)})$$

When the CV continues its backwards propagation to the origin, we apply the induced rules (7)–(8) to the OV $v_m(t; t_{\text{obs}}^{(n)})$ for $t < t_{\text{obs}}^{(n)}$. Finally, the OVs at the origin give the future values for all observation dates and products.

As we have seen, the pricing procedure combines the CVs (expressed in domestic currency units) linearly, with the dimensionless generalised rates γ. Note that these rates are strictly specific to the operations performed on them and thus do not have observation counterparts. To be adaptable to the algorithmic FV calculations, the backward pricing procedure must clearly distinguish the CV bearing currency unit from the dimensionless generalised rate serving as coefficient in the linear combinations of CVs. Technically, it is wise to program CVs and generalised rates as different types of objects, in order to exclude any currency non-preserving operations.

[6] For example, adding payments, applying exercises, extracting values, etc.

Our algorithmic approach is a universal alternative to the manual one first proposed in Cesari *et al* (2010) for single-exercise instruments. The authors have treated a callable deal as a pair of payment streams before and after the exercise. However, the authors do not consider multi-exercise instruments, such as a call with a barrier trigger or multi-call features. For such instruments, the manual approach requires substantial modifications, including both backward and forward steps with cumbersome and complicated logic of exercise indicators calculation and aggregation.

As mentioned above, in order to apply the algorithmic exposure method to forward instruments, we should convert them into backward ones. For example, a non-callable, path-dependent product can be transformed into the backward one using a preliminary forward simulation of the additional state variables (eg, cumulative coupons for Tarns) and a final backward propagation such that the regression is done on the extended state space.

The computational effort is split between Monte Carlo simulations of model rates and conditional expectation (CE) calculation using the least-squares Monte Carlo. Quite regularly, the latter is much more time consuming than the former. Indeed, it implies intensive calculus with a large number of basis functions, heavy linear algebra and possibly nonlinear search procedures. Thus, most time is spent on the least-squares Monte Carlo related calculations. As we have seen in the previous section, a calculation of FVs on a set of observation dates $\{t^{(n)}_{\mathrm{obs}}\}$, $n = 1, \ldots, N$, requires an extra N conditional expectations for each product.[7] For many instruments, the number of payment/exercise dates is much smaller than the number of FV observation points. Thus, for these cases the XVA calculation is much slower than the pure pricing.

Note that both algorithmic FV calculation and the manual one have the same numerical performance, but that the former has an advantage in its robust design.

2 6M Libor averages for different measures

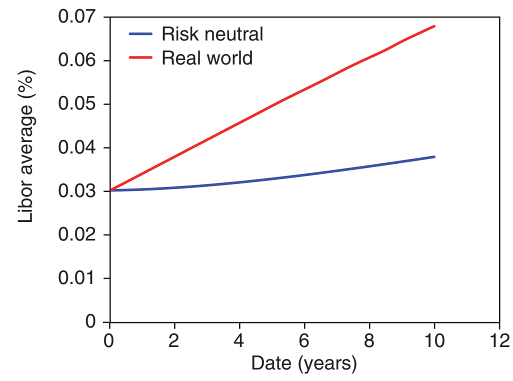

3 CDF of 6Y future values for different measures

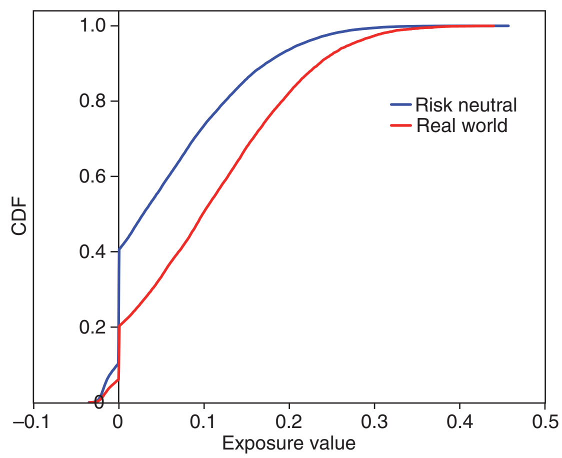

Numerical modelling of FVs in different probability measures

In this section, we discuss future values in different probability measures and present related numerical experiments. We calibrate the Radon-Nikodym process responsible for the measure change to expected values of market rates and indexes on future dates.

While valuation adjustments (XVA), being adjustments to price, do not depend on the probability measures, the distribution of future values does. Hence the risk measures computed as statistical characteristics of the distributions of future values through non-discounted averages are measure dependent, in contrast to the measure-independent price and valuation adjustments expressed though discounted expectations. Examples of such risk measures include VAR and expected shortfall for market risk, and potential future exposure and expected positive/negative exposure for counterparty risk (see, for example, Brigo, Morini & Pallavicini (2013), Cesari *et al* (2010), Gregory (2012), Canabarro & Pykhtin (2014) and references therein).

We refer to any probability measure different from the risk-neutral measure as the real-world measure. We propose that a Radon-Nikodym process corresponding to the transformation from the risk-neutral measure to a real-world measure be calibrated to the expected future values (averages and/or ranges) of market rates and indexes. The general calibration setup is as follows.

Let our future prices be simulated in the model risk-neutral measure. To calculate their distributions in the real-world measure we suppose that we have a set of rates R_j fixed at some times t_j and their projections in the future, $p_j = \mathbb{E}_{\mathrm{RW}}[R_j]$, as expectation in the real-world measure. The projections are either inferred by traders and risk managers from their views of expected evolution of markets or calculated using the historical time-series. In both cases, we consider them an input permitting us to define the real-world measure.

Formally, expectations in the real-world and the risk-neutral measures are related via a positive risk-neutral martingale process $\xi(t)$ called the Radon-Nikodym derivative, $\mathbb{E}_{\mathrm{RW}}[X] = \mathbb{E}_{\mathrm{RN}}[X\xi(t)]$ for a t-measurable random variable X. Following Cesari *et al* (2010), we compute a real-world measure distribution of the portfolio future value v as $\mathbb{E}_{\mathrm{RW}}[1_{x>v}] = \mathbb{E}_{\mathrm{RN}}[1_{x>v}\xi(t_{\mathrm{obs}})]$, for observation date t_{obs}.

We propose to define and calibrate the Radon-Nikodym derivative process to reflect measure change in simulation of exposures. As the simplest example we consider a log-normal process $\mathrm{d}\xi(t) = \xi(t)\sigma_\xi(t)\,\mathrm{d}W_\xi(t)$ with certain volatility $\sigma_\xi(t)$ and Brownian motion $W_\xi(t)$ correlated with the initial model. Treating the volatility and correlations as free parameters, we can use a numerical solver to fit

[7] Unless the observation dates intersect with the instrument ones.

4 PFE 97.5% for different measures

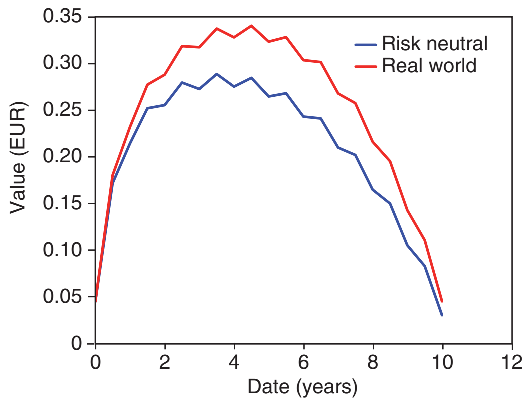

the projections $\mathbb{E}_{\mathrm{RN}}[R_j\,\xi(t_j)] = p_j$. This calibration procedure does not change the initial arbitrage-free model but explicitly determines the Radon-Nikodym derivative and the real-world measure. More complicated Radon-Nikodym evolution (eg, the driftless Heston model) gives a richer family of the measure change and provides more degrees of freedom for the calibration to real-world projections.

For numerical experiments we use an example of a 10Y cancellable swap: receive semi-annually a 6M Libor, pay annually a fixed rate (2.57%, a swap rate at origin) on 1 EUR notional, with the right to cancel the swap annually from 4Y. The pricing model is the Hull-White interest rate model with 3% rate, 5% mean-reversion and 1.5% volatility (we assume that the model parameters are obtained by calibrating to the implied market). The target projections of a 6M Libor are as in figure 2.

Calibration of the Radon-Nikodym derivative volatility leads to $\sigma_\xi = 25\%$ with its Brownian motion W_ξ perfectly correlated with the HW one. Of course, in real-world models the volatility σ_ξ is time-dependent. In figures 3 and 4 we can observe typical shapes of the risk profiles with significant differences in the distributions of future values in different measures.

Conclusion

In this article, we presented efficient calculations of the portfolio future values (exposure) in a self-consistent way using an arbitrage-free model that is calibrated to both implied market and real-world projections. We proposed a new algorithmic method of simulation of exposures (distributions of future values) based on an iterative backward induction, a generalisation of backward induction, especially attractive for exotic portfolios.

We applied this generalisation to simulation of exposures (distributions of future values) in the contexts of

- various valuation adjustments (XVA) due to counterparty risk, funding, capital, etc,
- calculation of risk measures that use averages of future values, such as VAR and expected shortfall for market risk, and PFE, EPE/ENE, etc, for counterparty risk,
- scenario generation, also in real-world measure. **R**

Alexandre Antonov and Serguei Mechkov are senior vice-presidents of the quantitative research group at Numerix. Serguei Issakov is senior vice-president and global head of the group. **In preparing this work, the authors greatly benefited from the insights of Vladimir Piterbarg. They are grateful to their colleagues at Numerix including Ion Mihai, Nic Trainor and especially Gregory Whitten for providing a stimulating research environment and support. They also would like to thank the anonymous referees for their insightful comments and suggestions.** **Email: antonov@numerix.com, isakov@numerix.com, meshkov@numerix.com**

REFERENCES

Andersen L and V Piterbarg, 2010
Interest Rate Modeling
Atlantic Financial Press

Antonov A, 2004
Generic instrument pricing
Working Paper, Numerix

Antonov A, S Issakov and S Mechkov, 2011
Algorithmic exposure and CVA for exotic derivatives
Available at SSRN

Brigo, D, M Morini and A Pallavicini, 2013
Counterparty Credit Risk, Collateral and Funding: With Pricing Cases for All Asset Classes
Wiley

Carriere J, 1996
Valuation of the early-exercise price for options using simulations and nonparametric regression
Insurance: Mathematics and Economics 19(30)

Cesari G, J Aquilina, N Charpillon, Z Filipovic, G Lee and I Manda, 2010
Modelling, Pricing, and Hedging Counterparty Credit Exposure: A Technical Guide
Springer Finance

Egloff D, M Kohler and N Todorovic, 2007
A dynamic look-ahead Monte Carlo algorithm for pricing Bermudan options
Annals of Applies Probability 17(4), pages 1138–1171

Fries C, 2007
Mathematical Finance: Theory, Modeling, Implementation
Wiley

Glasserman P, 2003
Monte Carlo Methods in Financial Engineering
Springer

Gregory J, 2012
Counterparty Credit Risk and Credit Value Adjustment: A Continuing Challenge for Global Financial Markets
Wiley Finance

Longstaff FA and ES Schwartz, 2001
Valuing American options by simulation: a simple least-squares approach
Review of Financial Studies 14(1), pages 113–147

Canabarro E and M Pykhtin, 2014
Counterparty Risk Management: Measurement, Pricing and Regulation
Risk Books