Getting CVA Up and Running

Frédéric Vrins* and Jon Gregory†

1 Introduction

An important consequence of the recent financial crisis is the growing interest (and more precisely, the necessity, both due to market volatility and regulatory constraints) for banks to include counterparty risk in their P&L. The adjustment made to the price of an OTC derivatives transaction as a result of the risky nature of the counterparty is known as the *credit value adjustment* (CVA). The CVA becomes an effective price factor of a deal, just like the interest rate or the exchange rate and can often have a significant impact on trading decisions. In terms of definition, CVA is essentially the expected loss faced on the trade occurring from the default of the counterparty one is trading with (see e.g. [Gregory (2010)], [Cesari (2009)], [Canabarro (2009)], [Pykhtin (2005)]). This would imply that the CVA is best defined as an *upfront amount*. In many circumstances, however, it may be more appropriate to use a *running CVA* or *CVA spread*. The reason is that clients may be more keen to adjust a running parameter (such as a swap rate for example) rather than making an up-front payment [1].

Given the complexity and computation requirements imposed by accurate CVA computation, a consistent CVA conversion framework is needed within a general CVA setup, where the CVA in question is either a stand-alone or an incremental CVA (i.e. the CVA of the trade taking into account netting and the impact of any other risk mitigants). In this paper, we detail that this conversion is not as trivial as it may seem, both for theoretical and practical reasons and requires iteration. Whereas an iterative procedure is possible, a simpler solution would be desirable in the typical framework where the CVA is calculated by Monte-Carlo simulation. For this reason, we investigate if simpler and computationally faster methods can be used, avoiding costly iterations. We describe upper and lower bounds on the CVA spread and discuss the relevance and accuracies of the simpler approaches which we illustrate with pricing examples. We will show that there is a simple method to determine the running CVA which requires no additional computation and gives results to within 2% accuracy in several test cases.

2 Running after CVA

We start by recalling how the (upfront) CVA on a risky derivative can be obtained. We then review several intuitive proposals to convert the upfront CVA into a running CVA. This involves changing a contractual rate in the derivative to incorporate the CVA component. The standard way of defining a CVA (in a general setting) is as follows:

$$\mathrm{CVA}(.) \quad = \quad \mathrm{V_D}(.) - \widetilde{\mathrm{V_D}}(.), \tag{1}$$

where $\widetilde{\mathrm{V_D}}(.)$ and $\mathrm{V_D}(.)$ respectively represent the value to the bank of the derivative with and without counterparty risk. A derivative featuring counterparty risk should always be worth less than the equivalent

*ING Financial Markets - Rates, Inflation and Credit Modelling Team, MII+4, Avenue Marnix 24, 1000 Brussels, Belgium, Frederic.Vrins@ing.be

†Solum Financial Partners, 12 Austin Friars, London. EC2N 2HE, jon@solum-financial.com

[1] We are referring here to derivatives, such as interest rate, basis and cross-currency swaps where a conversion is needed. For some products, such as options with up-front premiums, it is clearly not required.

risk-free derivative. Therefore, the upfront CVA is here always considered as a positive number; this is also the sign convention followed by Bloomberg, for instance.

2.1 Accurate method

For fixed-floating swaps, we take the convention to add the running CVA to the fixed rate as it helps customers to better compare the offers the get from the market, although other choices can be made (for basis swaps or floating-floating cross-currency swap, one needs to choose the leg to which the CVA spread should be added).

Suppose c is a rate in a derivative contract (such as LIBOR or a swap rate) and $c+\epsilon$ is that rate adjusted for CVA (ϵ may be positive or negative). In order to be properly compensated for CVA then, as pointed out by [Arvanitis and Gregory (2001)], the following must hold:

$$\mathrm{V_D}(c) = \widetilde{\mathrm{V_D}}(c+\epsilon), \tag{2}$$

Since the risk-free (without counterparty risk) value of the considered derivatives (or portfolio of netted derivatives) is linear with respect to c, then we obtain:

$$\epsilon\,\mathrm{DV01} = \mathrm{CVA}_D(c+\epsilon) \tag{3}$$

where DV01 stands for the present value (risk-free) of the running payment of value 1[2]. In other words, the running spread must act as compensation for the CVA: ϵ is such that it creates a P&L profit at inception exactly equal to the CVA on the realized trade, that is settled at fixed rate $c+\epsilon$. This typically requires iterations, as discussed in Section 3. This above equation corresponds, for example, to the procedure implemented in Bloomberg for the calculation of upfront CVA on a plain vanilla IRS[3].

The above definition can also encapsulate the effect of other trades under a netting agreement (and anything else of relevance, such as a collateral arrangement). In the case of netting, we can obtain via a similar logic to the stand-alone case:

$$\epsilon\,\mathrm{DV01} = \mathrm{CVA}_{NS+D}(c+\epsilon) - \mathrm{CVA}_{NS}(.) \tag{4}$$

where $\mathrm{CVA}_{NS+D}(c+\epsilon)$ represents the CVA of the netting set including the new trade and $\mathrm{CVA}_{NS}(.)$ is the CVA of the netting set before the new trade was added. The difference between these terms is often known as the *incremental CVA*. Examples of both stand-alone and incremental trades will be analysed later. We can now consider how to convert an up-front premium to a equivalent running CVA.

2.2 Risk-free duration

The *upfront-to-running* conversion is most easily achieved by dividing the upfront amount by a duration. There are two possible durations measures we can choose: either the *risk-free* or the *risky* duration. The risk-free duration is the DV01 of the above leg. Therefore, the running CVA is here defined as

$$\boxed{\epsilon_A \doteq \mathrm{CVA}_D(c)/\,\mathrm{DV01} \quad \text{(Method A)}}$$

The DV01 can be readily computed from the interest rate term structure and payment calendar and should be understood as being that of the specific leg that will be adjusted to account for CVA. Although this method is appealing due to its simplicity, the premium ϵ_A is considered to be *not at risk* and we therefore expect the result to be too low.

[2]For non-linear products, $V(c) + \epsilon\,\mathrm{DV01}$ is just the first order approximation of $V(c+\epsilon)$

[3]We are grateful to Harvey Stein, from Bloomberg L.P., for discussion on the aforementioned implementation

2.3 Risky duration

A more realistic point of view is to assume that the premium will not be paid during the whole life of the trade which accounts for the fact that the counterparty is not default-free. We therefore divide by the risky duration, $\widetilde{\mathrm{DV01}}$, which can be computed similarly to the risk-free duration using additionally the risk-neutral survival probability curve [4]. We therefore assume:

$$\boxed{\epsilon_B \doteq \mathrm{CVA}_D(c)/\,\widetilde{\mathrm{DV01}} \quad \text{(Method B)}}$$

The above method reflects the (static) *cost of hedging*. Indeed, the running CVA will finance the buying of credit protection on the counterparty. Whether or not we can simply include this spread in the ad-hoc leg of the underlying transaction is another question, that will be addressed below.

3 Discussion

Both methods A and B are simple and have a concrete interpretation. Method A is intuitive as we simply convert the upfront CVA into a running payment by using the risk-free duration. This is contradictory, however, as CVA precisely results from the non-zero probability that all the payments might not be made as a result of counterparty risk. Method B, by contrast, features the idea that payments are at risk. However, it does not recognise the change in CVA when the contractual rate is itself changed, and this approximation will be too large as will be shown in more detail below. The fact that the CVA that needs to be taken into consideration is that on the adjusted trade (that is for a receiver IRS, with fixed rate $c+\epsilon$, $\epsilon \geq 0$, instead of c) is a key point. Because none of the above method incorporate this reality, *both Method A and Method B will introduce a P&L jump*. Indeed, suppose we book a receiver trade with fixed coupon modified for CVA. Method A will lead to a P&L loss when a trade is booked whilst Method B will lead to a P&L gain.

By contrast, there exists a third method, which gives the correct result. The latter consists in iterating over the value of the running spread, each time re-computing the CVA to satisfy Eq. (3) or Eq. (4) depending on the type of CVA we are interested in; we refer this approach to as Method C. Unfortunately, in a realistic CVA framework where computation resource is probably a major concern, such iteration scheme should be avoided.

Consequently, we investigate whether there is another possible approach, that would be of similar complexity than Method A or Method B, but that would limit the aforementioned P&L hit. In other words, we seek for a competitive (but non-iterative) alternative to Method C.

4 Analytical proxy of the running CVA (stand-alone)

We start by considering stand-alone CVA and assume we are in a position where we will receive a higher contractual rate in order to compensate us for CVA. This would be the case, for example, in a receiver swap. Adopting this convention amounts at seeking a positive running CVA spread only. We note that the opposite situation (such as a payer swap) or, more generally, as a result of netting effects when using incremental CVA follows naturally from this. The upper and lower bounds derived below may simply need to be swapped, as illustrated in the Example section.

4.1 Closed-form bounds on the running CVA

In this section, we derive bounds to the correct running CVA, ϵ. We are considering the case where the CVA of a derivative, D, is an increasing function of the fixed rate, that is $\mathrm{CVA}_D(c) \leq \mathrm{CVA}_D(c+\epsilon)$ where c is the fixed rate and $\epsilon \geq 0$ is a positive spread. On the other hand, we know that $\mathrm{CVA}_D(c+\epsilon)$ is upper-bounded

[4] In the case of quarterly payment, it is actually very close to the DV01 of the fee leg of the CDS contracted on the counterparty to a risk-free third party

(due to netting effects) by the sum of the CVA and the CVA on the running credit premiums only. We note that latter quantity by $\mathrm{CVA}_P(\epsilon)$ and write:

$$\mathrm{CVA}_D(c) \leq \mathrm{CVA}_D(c+\epsilon) \leq \mathrm{CVA}_D(c) + \mathrm{CVA}_P(\epsilon) \tag{5}$$

Observe that $\mathrm{CVA}_P(\epsilon) = \epsilon\,\mathrm{CVA}_P(1)$. We will denote $\mathrm{CVA01} = (1-R)(\mathrm{DV01} - \widetilde{\mathrm{DV01}})$ for $\mathrm{CVA}_P(1)$ so that we can drop the D subscript in the swap CVA, to simplify the notations (R stands for the counterparty's recovery rate). The detailed calculation of CVA01 is given in the Appendix. The above inequalities imply that there exists α_ϵ in $[0,1]$ such that

$$\mathrm{CVA}_D(c+\epsilon) = \mathrm{CVA}(c) + \epsilon\alpha_\epsilon\,\mathrm{CVA01} \tag{6}$$

The ϵ subscript appearing in the α_ϵ factor suggests that the function $\mathrm{CVA}(c+\epsilon)$ is not linear with respect to the spread ϵ. Using Eq. (3), we get

$$\epsilon = \frac{\mathrm{CVA}(c)}{\mathrm{DV01} - \alpha_\epsilon\,\mathrm{CVA01}} \tag{7}$$

Further defining

$$\epsilon_\alpha \doteq \frac{\mathrm{CVA}(c)}{\mathrm{DV01} - \alpha\,\mathrm{CVA01}} \tag{8}$$

for $\alpha \in [0,1]$, we get $\epsilon_{\alpha_\epsilon} = \epsilon$. Notice that the risk-free duration proxy ϵ_A given by Method A corresponds to the $\alpha = 0$ case: $\epsilon_A = \epsilon_0$. The running CVA ϵ_B yielded by the risky-duration approach (Method B) is proven to be lower-bounded by ϵ_1 and is therefore of little interest in terms of upper bound to the correct spread; from that perspective, ϵ_1 should always be preferred to ϵ_B [5].

The previous equations suggest that one can "proxy" $\alpha_\epsilon \approx \alpha$ for some $\alpha \in [0,1]$, and that ϵ_α is thus a proxy of the CVA spread parametrized by the "portfolio effect factor" α. The truth is somewhere in between the extreme cases, $\epsilon = \epsilon_{\alpha_\epsilon}$, but α_ϵ being unknown, iterations are required; this is Method C.

4.2 Parametric Proxy Running CVA

As explained above, equations Eq. (6) and Eq. (7) are of little help in absence of the precise definition of the function α_ϵ. By contrast, Eq. (8) defining ϵ_α is interesting in that it yields a proxy value for ϵ for a given α. Without further information, one could choose a reasonable value for α with the knowledge that the higher the α, the more conservative we are. One could simply decide to take, for example, $\alpha = 1/2$. This has the advantage of requiring only one evaluation of the CVA (rather than the six or seven evaluations that would be required in a good root searching algorithm). We could be slightly more sophisticated and obtain an estimate of α_ϵ via a "proxy" $\alpha_\partial \doteq \left.\frac{\partial\,\mathrm{CVA}(x)}{\partial x}\right|_{x=c} / \,\mathrm{CVA01}$, as a result of a first order extension of $\mathrm{CVA}(c+\epsilon)$ with respect to ϵ. This would require two evaluations of the CVA which would still be preferable to a full iteration.

The nice thing is that, by looking at the range using Method A and Method B, one has an idea of the degree of uncertainty we have on the quality of the approximation. The bounds and the proxy can be intuitively understood by looking at figures Fig. 1(a) and Fig. 1(b). They can be compared with the correct running CVA.

5 Extension to incremental CVA and DVA

As mentioned at the beginning of Section 4.1, the sign of inequalities can be reversed if we are receiving (the positive) or paying (the negative) CVA spread. The same applies in the case where the CVA is negative. This can occur when quantifying incremental CVA, as a result of netting effects (adding a trade can reduce

[5] In the case where both $\mathrm{CVA}(c) \neq 0$ and $\mathrm{CVA01} \neq 0$, the equality case in $\epsilon_1 \leq \epsilon_B$ is obtained if and only if $R = 0$.

the total CVA of all trades within the "netting set" by compensating an existing risk, hence leading to a negative CVA contribution of the new trade. For instance, suppose we have a netting set built from a single payer IRS EUR 5Y with fixed rate 2%. The incremental CVA on the symmetric receiver swap is an increasing function of the swap rate, but it can be substantially negative for sufficiently low swap rates. For example, suppose for some reason that the receiver rate is 1%, so that the trade is deeply OTM, generating almost no exposure. On the other hand, because of netting effects, the negative PV of this trade will help reducing the total CVA, so that the inclusion of that receiver 1% contributes negatively to the total CVA. As the associated fixed rate increases to 2%, the incremental CVA tends to the opposite of the stand alone CVA of the payer swap, and we get a zero total CVA. Further increasing the fixed rate of the receiver swap will increase the total CVA, and the incremental CVA associated to the receiver swap will now increase without bound with the fixed rate.[6]

As for the stand-alone case, we illustrate the procedure on two simple examples. They are receiver fixed-floating cross-currency swaps (CCS) that will be added to a portfolio which consists in a payer CCS swap. Figure Fig. 1(c) adresses the case where the profolio is built of the mirror trade where no CVA spread is taken into account (this may seem not realistic, but it is an interesting case to study; suppose the sales decided not to charge the CVA spread for the first deal to get a new client, and that the market did not change). Because the floating leg will always compensate, we are only exposed to the fixed leg. As long as the running spread is negative (we are receiver), we shall always pay more than what we will receive (floating cashflows compensates, and deterministic fixed cashflows are always negative) and so our exposure decreases by including this OTM receiver trade (negative incremental CVA). The portfolio CVA is zero in all these cases, and the incremental CVA is just the opposite of the payer CCS CVA. In other words, in this case, the CVA spread will be negative, and we will have $\mathrm{CVA}(c+\epsilon) = \mathrm{CVA}(c) = 0$ and so $\epsilon = \epsilon_0$.

If we consider now a partial risk compensation (eg by including a receiver CCS to a payer CCS with different maturity for example), the scenario is slightly different, as exposed in Fig. 1(d).

Remark: We have considered the CVA spread to be included on the receive leg. Therefore, both the DV01 and the CVA01 are positive. In the case of incremental CVA, $\mathrm{CVA}(c)$ is negative, and the running CVA will be negative as well. In that case, the lower and upperbounds should be swapped.

6 Examples

In this section, we provide some examples to assess the quality of the proxy CVA. In all cases we consider a flat credit curve of 500 bps and no CSA (collateral posting). We have considered the following examples. :

1. A 5-year receiver interest-rate swap (Rec IRS).
2. A 5-year payer interest-rate swap (Pay IRS).
3. As 2 but including DVA with our own credit curve flat at 250 bps (Pay IRS DVA).
4. Cross-currency swap (CCS).
5. Cross-currency swap incrementally with an existing netting set of 17 trades (CCS incremental).

The results are given in Table 1. Table 2 gives the corresponding relative error. We have deliberately chosen cases where the conversion will be most challenging (long dated trades with a relatively risky counterparty).

In general, we can see that the simple approximation ($\alpha = 1/2$) gives a better estimate that using simple duration approximations for no extra computation effort and is within 2% of the exact result in all cases. The approximation of α_∂ requires only limited additional computation time and gives a result within 1% of the actual result in all cases. Hence, requiring a full iterative solution is not really necessary in light of many of the other uncertainties involved in CVA computation.

[6] The inclusion of DVA (own default risk) can have a similar impact as has been well-documented see, for example, [Gregory (2009)]

α	Pay IRS	Rec IRS	Rec IRS DVA	CCS	CCS incremental
0	3.44	-5.35	-3.49	-44.76	-8.02
1	4.29	-6.37	-4.70	-53.01	-9.52
1/2	3.64	-5.67	-3.78	-48,54	-8,71
α_∂	3.64	-5.64	-3.78	-49.06	-8.7
α_ϵ	**3.65**	**-5.68**	**-3.80**	**-49.5**	**-8.67**

Table 1: Comparison of running CVA (in bps) versus α. The CVA derivative is estimated by bumping the rate c by 10 bps. Last row (bolded values) gives the exact results.

α	Pay IRS	Rec IRS	Rec IRS DVA	CCS	CCS incremental
0	-5.7	-5.8	-8.1	-9.6	-7.5
1	15.7	15.6	23.7	7.1	9.8
1/2	-0.2	-0.2	-0.4	-1.9	0.5
α_∂	-0.2	-0.8	-0.4	-0.9	0.3

Table 2: Comparison of running CVA errors (in percentage) for the various approaches.

7 Conclusion

In this paper, we have focused on the conversion of the upfront CVA to a running CVA spread. We have reviewed three intuitive methods, and explained why only the third one is theoretically sound. The problem, however, is that it requires iterating over the running spread, hence leading to material calculation time when the upfront CVA is not calculated with an analytical method. This is typically the case when one is interested in the incremental CVA, or in the CVA of non-standard swaps; this is the reason why Monte Carlo simulations are used in standard CVA systems. The CVA calculation time is too expensive in this case to allow for iterations (just think about the case where we are interested in the incremental CVA on a large portfolio). For all these reasons, it is interesting to work out proxies to the running CVA spread, which do not require these iterations.

We have derived some closed-form bounds on the running CVA based on a parametric function, which can be used to derive a closed-form proxy of the running CVA, together with some confidence level we can have in the above proxy. The proxy is shown to perform very well on quite different kind of swaps, both for stand-alone and incremental CVA, giving results with no more than 2% error with no additional computation effort. The proxy and its bounds can be easily computed without iteration and can be graphically understood based on simple figures. This means that CVA desks can calculate running premiums with no additional computation effort, limiting a CVA hit in the P&L and preventing trades running away from the relevant trading desks.

8 Disclaimer

This paper expresses the view of the authors and does not necessarily reflect the opinion of their respective employers.

9 Appendix: CVA01

Approximating the continuous integral by a discrete sum, the quantity CVA01 can be obtained as follows:

$$\mathrm{CVA01} = (1-R)\sum_{i=1}^{N}(S_{i-1}-S_i)DF(0,t_i)\sum_{j=i}^{N}\mathrm{Not}_j\, DF(t_i,t_j)\tau_j = (1-R)\sum_{i=1}^{N}\mathrm{Not}_i(1-S_i)DF(0,t_i)\tau_i \quad (9)$$

$$= (1-R)\sum_{i=1}^{N}\mathrm{Not}_i\, DF(0,t_i)\tau_i - (1-R)\sum_{i=1}^{N}\mathrm{Not}_i\, S_i DF(0,t_i)\tau_i = (1-R)(\mathrm{DV01} - \widetilde{\mathrm{DV01}}) \qquad (10)$$

where $DF(t_1,t_2) = DF(0,t_2)/DF(0,t_1)$ is the ratio of discount factors, τ_j is the coupon accrual for the j-th period, Not_i is the associated notional amount and S_i is the risk-neutral survival probability curve of the counterparty up to time t_i. The sum ranges over the payment dates of the fixed coupons, and we had $\mathrm{DV01} = \sum_{i=1}^{N} DF(0,t_i)\tau_i$.

Clearly, because α belongs to then $[0,1]$ interval,

$$\alpha\,\mathrm{CVA01} \leq \mathrm{CVA01} \leq \mathrm{DV01} \qquad (11)$$

and so for an ATM swap, $\epsilon_\alpha \geq 0$ for a receiver, as expected.

References

[Arvanitis and Gregory (2001)] Arvanitis., A., and Gregory, J. *Credit: The Complete Guide to Pricing, Hedging and Risk Management*, Risk Publications, 2001.

[Cesari (2009)] Cesari, G. et al.*Modelling, Pricing and Hedging Counterparty Credit Exposure*, Springer Finance, 2009.

[Canabarro (2009)] Canabarro, E. (Ed) *Counterparty Credit Risk*, Risk Books, 2009.

[Pykhtin (2005)] Pykhtin, M. (Ed) *Counterparty Credit Risk Modelling: Risk Management, Pricing and Regulation*, Risk Books, 2005.

[Gregory (2010)] Gregory, J. *Counterparty Credit Risk*, Wiley, 2010.

[Gregory (2009)] Gregory, J. *Being Two Faced Over Counterparty Credit Risk*, Risk, 22(2), pp 86-90.

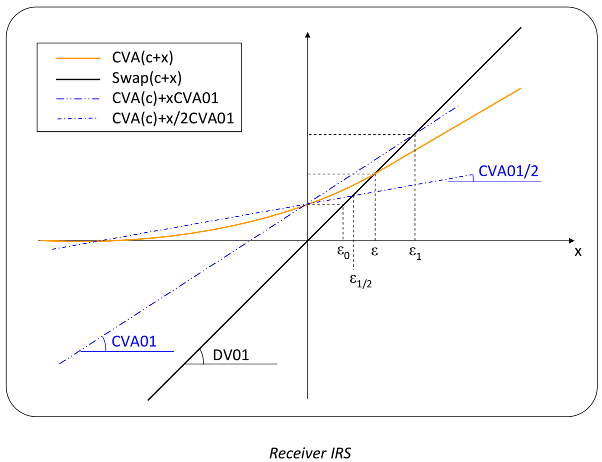

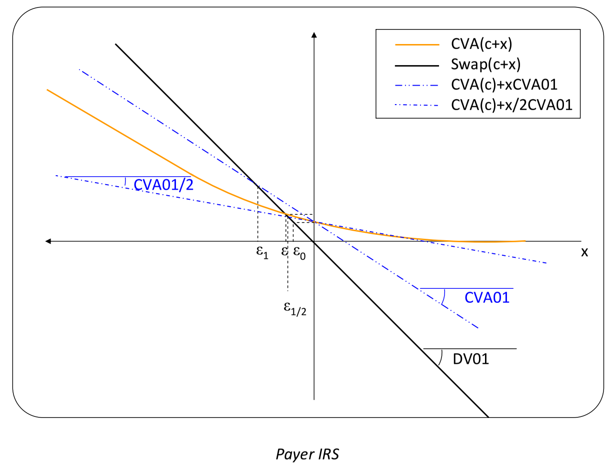

(a) Bounds, proxy and exact running CVA in the case of a IRS payer as a function of the margin x added to the swap rate.

(b) Bounds, proxy and exact running CVA in the case of a IRS receiver as a function of the margin x added to the swap rate.

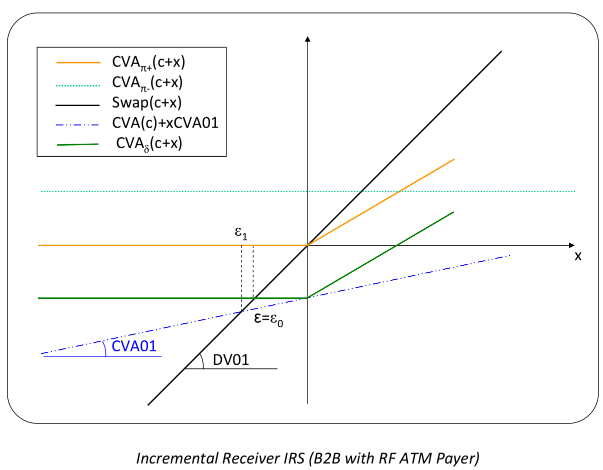

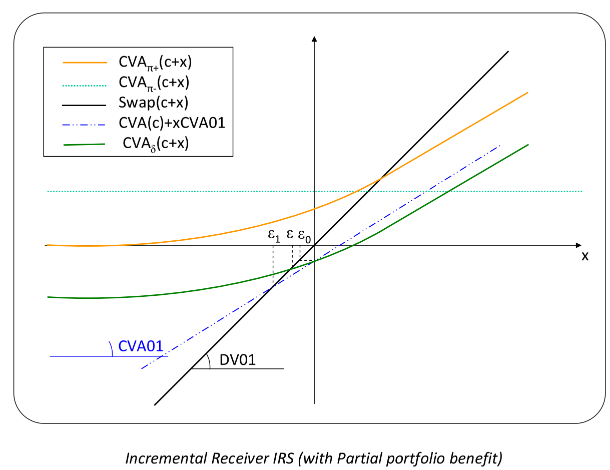

(c) Bounds, proxy and exact running CVA in the case of a IRS payer as a function of the margin x added to the swap rate.

(d) Bounds, proxy and exact running CVA in the case of a IRS receiver as a function of the margin x added to the swap rate.

Figure 1: Illustrations of the running CVA and the proposed proxy estimates.