Chapter 10
Measuring Derivative Exposure

If a man tells you he knows a thing exactly, then you can be safe in inferring that you are speaking to an inexact man.

(Bertrand Russell)

One of the central inputs in the computation of credit-risk economic capital, which is often somewhat taken for granted, is a position's expected exposure at default. In full generality, this is a random quantity following some (potentially complicated) probability law. For standard linear instruments—such as loans, deposits, or bonds—it is common practice to ignore their stochastic nature and assume deterministic exposures. This is (mostly) defensible, because the uncertainty surrounding future exposure outcomes is usually sufficiently small that the extra effort (and incremental complexity) is unwarranted. For some financial instruments, this argument is less compelling. Derivative contracts are the classic counterexample.

There is a fundamentally large amount of variability in the exposure associated with a future default event involving (interest-rate, cross-currency, or default) swaps, options, or forward contracts. Swaps and forward contracts are linear, but at any given point in time they may represent an asset or a liability to their holders; more importantly, their status can change quickly with market movements. The same cannot (typically) be said for option contracts, but their situation is often complicated by non-linear pay-off structures.[1] To be blunt, assuming deterministic exposure-at-default values for derivative contracts is a singularly bad idea.

This important realization, it turns out, opens something of a Pandora's box of incremental questions for consideration. Indeed, it is so involved that it has spawned a separate, albeit related, field of analysis: counterparty credit risk. The bilateral *counterparty* aspect is the deciding factor. Very often, collateral is exchanged to offset and mitigate risks associated with the swings in derivative contract values. Collateral exchange, of course, is governed by legal agreements and operates across multiple derivative contracts with a single counterparty. This fact, combined with the

[1] One can, of course, have swap contracts with embedded optionality. In this way, one has to juggle both sign change in exposures *and* non-linearities with respect to underlying risk factors.

D. Bolder, *Modelling Economic Capital*, Contributions to Finance and Accounting,
https://doi.org/10.1007/978-3-030-95096-5_10

multiplicity of derivative contracts and their associated features, gives rise to many special cases and detailed exceptions. To summarize, the world of counterparty credit risk is *not* simple.

It is nonetheless necessary to wade into the realm of counterparty credit risk to understand the techniques used to approximate (random) exposure-at-default amounts associated with one's derivative contracts. This chapter attempts to describe the basic challenge, briefly sketch out the *correct* way to address it, review a widely used (and surprisingly useful) regulatory derivative-exposure approximation, and thereby motivate NIB's approach.

10.1 The Big Picture

As the name strongly suggests, counterparty credit risk turns around the idea of a counterparty. There are many other possible filters through which one can examine this risk dimension, but this represents the highest level of aggregation. Exposure at default is defined by each individual counterparty. One might, however, have hundreds or even thousands of individual positions of various types, sizes, tenors, and currencies. Moreover, it is entirely possible that multiple legal agreements might be in place—for different sub-entities or instrument types—with a single counterparty. This immediately raises a second critical point: how should these various positions be aggregated? This falls under the category of questions that are notoriously easy to ask, but rather more difficult to answer.

A useful way to think about aggregation is to introduce the notion of a netting set. Duffie and Canabarro [11] refer to this helpful idea as a netting node, but the idea is identical. Closely related, although not necessarily identical, is the so-called margin set or node. Duffie and Canabarro [11] describe this as:

> a collection of trades whose values should be added in order to determine the collateral to be posted or received.

The central idea behind netting is to permit, within a collection of trades, market values to offset one another when aggregating credit exposures. If, for example, you have two trades—one negative and the other positive—netting permits the negative exposure to (at least, partly) nullify the positive component.[2] As clearly highlighted in Gregory [13], such an arrangement has value under *two* main conditions. First, some (or all) of the trades must potentially take negative values. A bit of reflection reveals that no netting benefit is possible if *all* market values are strictly positive (or negative). The second, perhaps less obvious, point touches upon the dependence between market-value outcomes. In the event of perfect (positive) correlation between current and future market values, there will be limited scope for

[2] When you think about it, this is far from obvious, since these represent two separate contracts. Salonen [20] provides an interesting and comprehensive discussion of how this notion developed, which is far outside of the author's area of expertise.

any difference in sign between individual positions; the consequence would, once again, be no netting benefit. Naturally, perfect correlation is somewhat exaggerated and rarely observed in practice, but it represents a limiting case. In general, the stronger the positive correlation between the underlying risk factors driving the valuation of the derivative trades in one's netting set, the lower the potential netting benefit.

Although the most natural form of aggregation occurs along netting sets, the margin set determines the collateral. Collateral levels, in turn, have important implications for one's exposure computations. A central divide, therefore, is between margined and unmargined netting sets. Unmargined netting sets are not, it should be stressed, completely absent of any collateral. Typically, some initial margin or independent amount is posted at the inception of these trades, which can often be adjusted in the event of a credit-migration event by either counterparty.[3] The distinction between margined and unmargined netting sets comes from the role of variation margin; these are amounts that are exchanged (usually on a daily basis) in the event of changes in the netting set's market value.

In straightforward cases—which happily coincides with NIB's situation—a counterparty can be mapped to a single, common netting and margin set. It is possible to have a margin set comprise multiple netting sets; that is, there may be a one-to-many mapping between margin and netting sets.[4] For the purposes of this discussion, we will assume a one-to-one mapping between netting sets and counterparties as well as between netting and margin sets. To be very explicit, all derivative trades (or positions) with a single counterparty are netted and follow a common margin policy. Deviations from this rule are practically essential to capture in one's portfolio—and most definitely should not be assumed away—but they do not add much to the overall narrative. They result in more dimensionality and logistical complexity, but their management mostly just requires careful book-keeping.

Figure 10.1 tries to conceptualize our simplified view of the interaction between the counterparty, netting, margining, and our set of derivative trades. Practically, as indicated, there might be multiple hedging sets within each margin set and many margin sets for a given counterparty. The viewpoint in Fig. 10.1 will, however, be sufficient to address the main concepts associated with the determination of credit counterparty exposures.

[3] Initial margin and independent amount are conceptually the same thing, but the specific choice of term appears to depend on the situation and instrument type. See Gregory [13] for more background on basic collateral concepts.

[4] The reverse does not seem to be the case. That is, one does not (to the best of the author's knowledge) observe multiple margin sets within a single netting set.

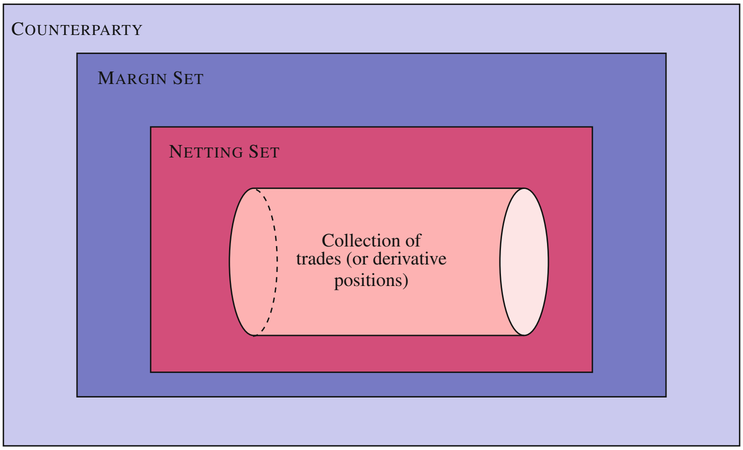

Fig. 10.1 *A simplified hierarchy*: This schematic illustrates, in a high-level conceptual manner, the interactions between counterparty, margin and netting sets, and the collection of derivative trades. This is a simplified representation; in practice, there are often many netting and margin sets associated with a given counterparty.

Colour and Commentary 115 (A KEY TASK IN CREDIT COUNTERPARTY RISK)*: Aggregation of exposures for plain-vanilla instruments—such as loans, bonds, or deposits—is dead simple: you just need to add them up. Moreover, their relative stability (typically) precludes the need to treat such exposures as stochastic. The exposure picture for derivative contracts is rather murkier due to their ability to quickly move from asset to liability, imperfect correlation structures, and pay-off non-linearities. Derivative exposures, as a consequence, cannot responsibly be treated in a deterministic fashion. The quantitative risk analyst's task is to determine a (defensible) derivative exposure-at-default amount for each of her institution's credit counterparties at a given point in time. This necessitates the aggregation of individual trade (or position) market values along netting relationships. It further involves the incorporation of collateral over each margin set. The entire point of collateral, of course, is to mitigate credit losses in the event of a counterparty default. There will be different amounts of collateral depending on the margining policy and this needs to be taken into consideration. The final, and not least complicated, issue involves making some adjustment for the fact that individual netting set market values can actually move against us. Despite the presence of collateral, market forces can (and do) lead to increased exposure and consequently larger potential default losses. Building a comprehensive framework to capture all of these moving parts is, in terms of a big picture, not particularly easy. We will have our hands full in the following discussion.*

10.2 Some Important Definitions

There are a number of fundamental definitions that need to be introduced to make any useful headway in this area. Almost as important as the definitions themselves is the establishment of a meaningful and descriptive notation. Given the numerous possible levels of aggregation, keeping track of the details will prove essential to success in this venture. There are many excellent sources on these basic definitions, but this treatment draws principally from Gregory [13] and Duffie and Canabarro [11]. Our notation has a few unique elements, but it is certainly inspired from these and other helpful sources.

We begin at the trade or position level. Although we use the terms trade and position interchangeably in this discussion, they are not precisely identical. One might buy a futures contract in a number of separate transactions, where each one would be a trade. The collection of these identical trades represents a position. Systems operate at either level, and occasionally both, so exactly how this is performed can vary. Whichever one uses, we can think of these as being (close to) the most atomistic level of derivative exposure definition. We define $V_{kt}(i)$ as the market value of the kth derivative position with the ith counterparty at time t.[5] The notation thus makes reference to the position, the counterparty, and the current point in time. Assuming a single netting and margin set for the ith counterparty, we can compute the netting-set market value as,

$$V_t(i) = \sum_{k=1}^{K_i} V_{kt}(i), \tag{10.1}$$

where K_i represents the number of individual derivative trades in the ith counterparty's netting set.[6]

The notion of exposure comes in a variety of flavours in the field of counterparty credit risk. Counterparty *exposure* is typically written as,

$$\begin{aligned} E_t(i) &= \max\Big(V_t(i), 0\Big), \\ &\equiv \Big(V_t(i)\Big)^+. \end{aligned} \tag{10.2}$$

In plain English, the exposure is the positive part of the netting set's market value. The sole focus on positive outcomes is natural since, in the event of default, no

[5] Unless otherwise specified, as elsewhere in this book, we use t to denote the current point in time.

[6] In the case of multiple netting sets, then one has multiple versions of Eq. 10.1.

loss is incurred in the face of negative market values.[7] Sometimes the term *current exposure* is used to refer to Eq. 10.2, because it denotes the time t exposure.

The origin of most of the complexity in this area stems from precisely the fact that derivative transactions—and by extension derivative books—can take both positive and negative values. That one's current exposure is zero, by virtue of a negative netted market valuation, is *no* assurance that it will remain so tomorrow, the day after, or next week. Not being able to peek into the future, we cannot know what will happen. As prudent risk managers, however, we can try to prepare for the worst. This leads to the idea of future exposure, which we could denote as $E_{t+T}(i)$. This would be the (unknown) exposure to counterparty i at time T or $T-t$ units of time into the future.

The best we can do, absent a functioning crystal ball, is to treat $E_{t+T}(i)$ as a random variable. It is thus interesting to consider the mathematical expectation of this value as,

$$\mathbb{E}\left(E_{t+T}(i)\middle|\mathcal{F}_t\right) = \mathbb{E}\left(\underbrace{\max\left(V_{t+T}(i), 0\right)}_{\text{Eq. 10.2}}\middle|\mathcal{F}_t\right), \tag{10.3}$$

$$\mathbb{E}_t\left(E_{t+T}(i)\right) = \mathbb{E}_t\left(\left(V_{t+T}(i)\right)^+\right),$$

where $\mathcal{F}_t$ is the information set (or σ-algebra) at time t. Probabilistic details aside, the $\mathcal{F}_t$ explicitly captures the idea that such computations can only be performed with the information available at time t; anything else would be essentially cheating. To avoid the clutter of the conditioning set, however, we define the *expected exposure* (or EE) as $\mathbb{E}_t\left(E_{t+T}(i)\right)$. To actually evaluate Eq. 10.3, it is necessary to make some sort of distributional assumption about the market value of our counterparty's netting set, $V_{t+T}(i)$.

Equation 10.3 illustrates the expected exposure of counterparty i for a given point of time in the future, T. Practically, for example, we might set $T-t$ to be one-year in the future. It is also interesting to consider the full time profile of expected exposures for $\tau \in (t, T]$. To boil this down to a single number, it is common to compute the average expected exposure over such a time interval. Mathematically, this would reduce to integrating Eq. 10.3 over the time dimension and dividing the result by

[7] It would be helpful to describe the quantity in Eq. 10.2 as *positive* exposure to avoid confusion. Sadly, this is not the case.

$T - t$. This yields

$$\mathbb{E}_t\left(E(i,t,T)\right) = \frac{1}{T-t}\int_t^T \mathbb{E}_t\left(\left(V_{t+\tau}(i)\right)^+\right)d\tau. \tag{10.4}$$

This quantity is generally referred to as the *expected positive exposure* or EPE. We opted to write it as $\mathbb{E}_t\left(E(i,t,T)\right)$ to illustrate the explicit dependence on the counterparty, the information set, and the time interval.

Another twist on counterparty exposure is the so-called *potential future exposure* or PFE measure. It may be defined as,

$$\text{PFE}_\alpha(i,t,T) = \inf_{x\in\mathbb{R}}\left(x \,\middle|\, \mathbb{P}\left(V_{t+T}(i) \geq x\right) \geq 1-\alpha\right). \tag{10.5}$$

where α is a predefined confidence level. Those familiar with the Value-at-Risk (or VaR) metric—or have already reviewed Chaps. 2 or 5 or even Chap. 7—will see a clear similarity. The key difference between PFE and VaR is which tail of the value distribution they focus upon. VaR is concerned with downside losses, whereas PFE attempts to describe upside derivative-valuation gains. The reason, of course, is that the larger the positive value of one's netted derivative exposure with a given counterparty, the larger the scope for losses in the event of default. As we will see in the next section, however, Eq. 10.5 might be the typical definition of PFE, but it is not the only way this quantity might be computed.

Expected and potential-future exposure (i.e., EE and PFE) are thus looking at the same underlying distribution; the former examines the expectation, while the latter concerns itself with a given quantile. The choice of exposure measure along with the specification of α and, not least, the determination of T all depend on the specific application and one's desired degree of conservatism. For the reporting of one's exposure, it might make sense to use the expected exposure measure. Economic-capital—with its focus on worst-case events—might usefully employ the PFE perspective. Whatever one's selection, it requires some reflection and justification.

Figure 10.2 helps to turn the previous equations into visualizations by illustrating the interaction between netting set valuations, (positive) exposure, expected exposure, and potential future exposure. The left-hand side considers one-step into the future over the interval, $(t, T]$; here we clearly see the current, expected, and potential-future exposures. The right-hand graphic, conversely, takes *five* steps to span the time interval, $(t, T]$; many more steps, of course, are possible. This allows for the introduction of the expected positive exposure measure.

There are a few final concepts that are worth introducing to avoid confusion down the road. These ideas stem principally from the regulatory setting.[8] They

[8] See BIS [3] for more detailed background.

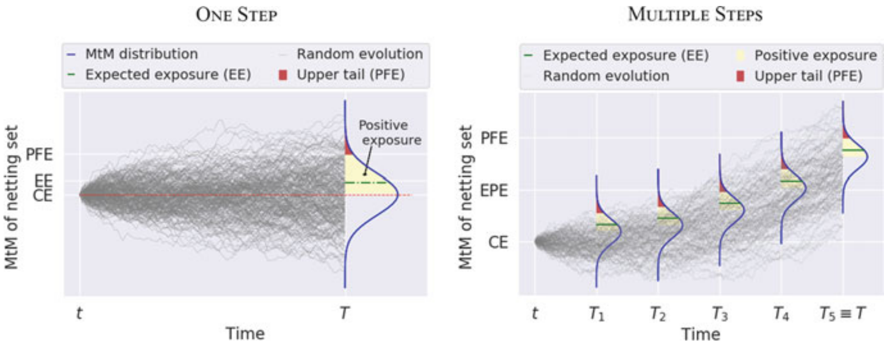

Fig. 10.2 *Some useful visualizations*: The graphics above illustrate the interaction between netting set valuations, (positive) exposure, expected exposure, and potential future exposure. The two graphics examine both one- and multiple-step perspectives, which allow for the introduction of expected positive exposure.

all relate to the passage of time. Either the plain expected exposure (i.e., EE) or expected positive exposure (i.e., EPE) can be computed for a range of possible future points in time. As we move further out the horizon, there will be some maturities in one's derivative netting set. This will naturally lead to a gradual decrease in both the EE and EPE measures. Although, in practice, these positions are likely to be either rolled-over or otherwise replaced, such assumptions are complicated and typically not included in one's analysis. The regulatory solution was to introduce the idea of *effective* EE and EPE measures. The effective means that the EE and EPE profiles are computed, but then are adjusted so that they are (weakly) monotonically increasing (i.e., non-decreasing).[9] Imagine, for example, that you calculated the effective exposure in monthly time steps out to a year. If the highest EE estimate occurred during the seventh month, then the effective EE would fix all subsequent (lower) values to the maximal amount.

Colour and Commentary 116 (Derivative Exposure Definitions)*: Counterparty credit risk can, and should, be viewed as a risk-management sub-field. Like all areas of endeavour, it has its own central ideas and vernacular. The preceding definitions represent essential groundwork for the computation of exposure-at-default values for one's derivative portfolios; while not entirely comprehensive, they do touch upon the core of counterparty credit-risk measurement. A few important commonalities arise. First, attention is strongly on the positive part of the future netting-set valuations.*

(continued)

[9] This is a fairly unfortunate choice of word, since *effective* does not usually, in the author's understanding of general English-language parlance, imply weak positive monotonicity.

Colour and Commentary 116 (continued)
This makes logical sense since this is where the risk lies. The second point is that there are both cross-sectional and time dimensions at play. There is a collection of derivative contracts for a given point in time (i.e., the cross section), but we are also interested in how the value of these contracts behaves over various periods (i.e., the time dimension). We will have to manage randomness of the underlying risk drivers associated with our derivative positions along both dimensions; this will necessarily involve implicitly or explicitly introducing stochastic processes. Given that netting effectiveness depends importantly on correlation assumptions, the joint distribution of the various marginal risk-factor dynamics will also play an important role.

10.3 An Important Choice

Financial institutions with large derivative portfolios inevitably find themselves in the valuation business. To permit proper exchange of collateral, a market-consistent current value for all derivative contracts in one's portfolio is required. This necessitates implementation of pricing formulae (or numerical algorithms) linking the individual cash-flows of each instrument to its associated underlying risk factors: interest rates, exchange rates, commodity prices, and so on. Often embedded within these pricing approaches—or determined for relatively modest additional cost—are associated hedging ratios and risk-factor sensitivities. These quantities can help predict changes in the value of individual instruments—and by extension one's portfolio—for given movements in key risk factors. Development, construction, and maintenance of such valuation frameworks involve much work and significant resources. It is, however, basically a cost of being a player in derivative markets.

Linking this back to our previous definitions, we can see how this valuation infrastructure feeds into counterparty credit risk. Without proper and timely valuations, we cannot hope to reasonably exchange collateral or compute current exposure levels. Equations 10.1 and 10.2 are only knowable with the appropriate valuation machinery. This is good news, since part of what we need is, literally by construction, already available. As we move to Eq. 10.3, however, the situation changes. No matter how sophisticated one's valuation framework, it cannot tell you the value of individual derivative instruments in the future. These outcomes are unknown and depend, of course, on the stochastic evolution of multiple market-related random variables.

The *correct* solution to this problem involves the creation of a second forward-looking valuation framework, which lies on top of the base pricing infrastructure. It is an analytically and computationally intensive undertaking. Generally, it relies on simulation methods. This is partly due to the complexity of the underlying instruments, but principally due to the non-analytic nature of collateral exchange and the management of netting and margin sets. It begins with the calibration and simulation of a collection of correlated stochastic processes, over some pre-defined analysis horizon, describing the set of risk drivers for one's derivative portfolio.

Stepping forward into the future—for each sample path and along time grids of varying granularity—one revalues each individual instrument, revises cash-flow patterns as necessary, exchanges collateral (or not), and determines the time profile of the netting set for each credit counterparty.[10] This process helps to trace out the (simulated) evolution of derivative exposure for each counterparty in one's portfolio. With such an engine, (positive) exposure, expected exposure, expected positive exposure, potential future exposures, and many other possible variations are readily numerically approximated.

The ability to describe the future evolution and distribution of one's derivative portfolio—by individual instrument and counterparty—offers value beyond exposure computation. Following the great financial recession, there have been fairly dramatic changes in how credit risk is incorporated into derivative pricing.[11] This leads to the idea of the credit valuation adjustment (CVA) and its many related flavours, which is generally referred to as XVA.[12] Such simulation engines also play a central role in this derivative-pricing activity. As a consequence, it is common for financial institutions heavily involved with derivative instruments to have not only a valuation engine, but also a complex apparatus for PFE and XVA computations.

At first glance, the solution to our problem seems obvious. Given an army of quantitative analysts already computing exposure profiles for individual instruments and counterparties, we should simply borrow their results for our purposes. It is always pleasant to find someone competent to perform a difficult task for you. Ultimately, despite a few clear advantages, we have nonetheless opted *not* to follow this route. Instead, we make use of a regulatory approximation. There are *three* main reasons. First, and perhaps most importantly, our focus is on the risk-management dimension. Most forward-looking simulation engines are firmed focused on pricing; they work with different time horizons and probability measures. A second point relates to conditionality. We seek a through-the-cycle perspective in our economic-capital computations to avoid—as mentioned numerous times in previous chapters—pro-cyclicality. Any sensible forward-looking simulation model is calibrated to current market conditions—otherwise it cannot reproduce current valuations—and is thus, by its very construction, working under the point-in-time viewpoint. The final point relates back to our conceptual axioms introduced in the preface. It is a difficult, full-time job to maintain and run a forward-looking

[10] For some computations, particularly those over short periods, one may ignore the collateral exchange to understand just how bad one's exposure position could become with a given counterparty.

[11] The basic issue was that, prior to the 2008 crisis, most derivative pricing and hedging (more or less) structurally ignored default risk. See Brigo [10] for a much more eloquent and detailed description of this point.

[12] See Ruiz [19] for an excellent jumping off point into this world.

simulation platform; such an engine seeks to perform a complex array of tasks. Were we to simply borrow these results, there is a very real danger of it becoming a black box for risk-management staff. It seems more defensible to use a simpler, more easily understandable approach for our purposes.

Colour and Commentary 117 (PICKING AN ANALYTIC LANE)*: Treatment of derivative exposure as a random variable, for use in our economic-capital model, complicates our life. It turns out, however, that we are not the only group of quantitative analysts taking this perspective. Subsequent to the 2008 global financial crisis, a major rethink occurred in the area of derivative pricing. One of the (many) consequences has been the construction of forward-looking derivative exposure simulation engines to support the revised pricing algorithms. This infrastructure would, at first blush, seem to be tailor-made for our purposes. Sadly, it is not. Forward-looking engines have a pricing, and not risk-management, focus. They are also, by absolute necessity, firmed rooted in the point-in-time perspective. Finally, they embed a level of complexity rather far beyond what is required for our purposes. Such exposure estimates could, rather easily, become a black-box input into our economic-capital computations. From an economic-capital perspective, we require short-term,* $\mathbb{P}$*-measure, through-the-cycle, and relatively easily interpretable derivative-exposure estimates. With some reluctance, therefore, we have firmly elected to follow the simpler (pseudo-)analytic lane offered by the regulatory authorities.*[a]

[a] This choice, depending on one's organization and resources, certainly need not apply to all financial institutions.

10.4 A General, But Simplified Structure

To make any further progress, it is necessary to make some concrete assumptions about the underlying distributional dynamics of the market value of one's netting sets. As just discussed, the most direct way would be to begin by describing the joint behaviour of the underlying collection of (correlated) market risk factors—such as interest rates, exchange rates, commodity prices, or volatilities—that drive derivative valuations. While this approach is the most defensible, it also involves the largest degree of complexity. We have opted for an alternative strategy.

Understanding our requirements, there is a strong incentive to simplify. A useful approach involves working with the overall netting-set market value as a *single*, aggregated random variable. In other words, we exploit the fact that a netting set is a kind of sum of many other random variables. How might this work? Imagine that we define the following stochastic process on $(\Omega, \mathcal{F}, \mathbb{P})$ describing the intertemporal dynamics of the ith counterparty's netting set:

$$dV_t(i) = \mu_i dt + \sigma_i dW_t, \tag{10.6}$$

where $\mu_i \in \mathbb{R}$, $\sigma_i \in \mathbb{R}_+$ and $\{W_t, \mathcal{F}_t\}$ is a standard, scalar Wiener process on $(\Omega, \mathcal{F}, \mathbb{P})$. This is a drifted Brownian motion. If we would like to solve Eq. 10.6 for $V_t(i)$, then we need to make use of the underlying, rather famous, theorem from the stochastic calculus.[13]

Theorem 10.1 (Itô) *Let X_t be a continuous semi-martingale taking values in an open subset $U \subset \mathbb{R}$. Then, for any twice continuously differentiable function $f : U \to \mathbb{R}$, $f(X_t)$ is a semi-martingale and,*

$$f(X_T) - f(X_t) = \int_t^T \frac{\partial f}{\partial X_u} dX_u + \frac{1}{2}\int_t^T \frac{\partial^2 f}{\partial X_u^2} d\langle X_u \rangle, \tag{10.7}$$

where $d\langle X_t \rangle$ denotes the quadratic variation of the process X_t.

This is a fairly trivial application, where we need only select the identity function $f(V_t(i)) \equiv V_t(i)$ to apply Theorem 10.1. This leads to:

$$\begin{aligned} V_T(i) - V_t(i) &= \int_t^T \frac{\partial f(V_u(i))}{\partial V_u(i)} dV_u(i) + \frac{1}{2}\int_t^T \frac{\partial^2 f(V_u(i))}{\partial V_u(i)^2} d\langle V_u(i) \rangle u, \quad (10.8)\\ &= \int_t^T 1\,(\mu_i du + \sigma_i dW_u) - \cancel{\frac{1}{2}\int_t^T 0 \cdot d\langle \mu_i du + \sigma_i dW_u \rangle}, \\ &= \mu_i \int_t^T du + \sigma_i \int_t^T dW_u, \\ V_T(i) &= V_t(i) + \mu_i (T - t) + \sigma_i (W_T - W_t), \end{aligned}$$

where $W_T - W_t \sim \mathcal{N}(0, T - t)$ is a Gaussian-distributed, independent increment of the Wiener process. The consequence is that for any future time period, $T \geq t$, the market value of our netting set is distributed as

$$V_{t+T}(i) \equiv V_T(i) \sim \mathcal{N}\left(\underbrace{V_t(i) + \mu_i(T - t)}_{\mu_i(t,T)}, \underbrace{\sigma_i^2 (T - t)}_{\sigma_i(t,T)^2} \right), \tag{10.9}$$

for the ith credit counterparty where the functions $\mu_i(t, T)$ and $\sigma_i(t, T)$ are introduced to keep the notation moderately under control. In other words, starting

[13] This is only a very quick (and non-rigorous) peek into this vast array of mathematical endeavour. See Karatzas and Shreve [15], Oksendal [18], and Heunis [14] for much more information, rigour, and intuition.

from the current value of $V_t(i)$, it will grow at the rate of μ_i in a manner proportional to the passage of time. The volatility around this value is σ_i scaled by the square-root of the time interval.

The consequence of this mathematical detour is a cohesive approach to the intertemporal dynamics of the total market value of the individual derivative positions with counterparty i.[14] Equation 10.9 basically tells us how we can describe the value of a given netting set for any time $\tau \in (t, T]$. All of the intricacy of underlying risk factors is embedded in the value of μ_i and σ_i. Information is certainly lost, but this foundational assumption allows us to put more structure—and specificity—to the exposure definitions introduced in the previous section.

10.4.1 Expected Exposure

Using Eq. 10.3 and the final result from Eq. 10.9, we may now proceed to put a face to the notion of expected exposure. At first glance, it might seem difficult to actually determine the expectation of only the positive part of $V_{t+T}(i)$. It turns out, however, to be relatively straightforward. Since $V_{t+T}(i)$ is a Gaussian random variable with support on $\mathbb{R}$, the positive expectation is determined by simply integrating over the positive part of this domain, $\mathbb{R}_+$. Practically, this reduces to:

$$\begin{aligned}\mathbb{E}_t\bigg(E_{t+T}(i)\bigg) &= \underbrace{\mathbb{E}_t\left(\bigg(V_{t+T}(i)\bigg)^+\right)}_{\text{Eq. 10.3}}, \\ &= \int_0^\infty v f_V(v)dv, \\ &= \int_0^\infty \frac{v}{\sigma_i(t,T)\sqrt{2\pi}} e^{\frac{-(v-\mu_i(t,T))^2}{2\sigma_i(t,T)^2}} dv.\end{aligned} \tag{10.10}$$

Although this involves a wearisome bit of calculus, it is instructive to actually solve this integral.[15] It does, however, necessitate a few substitutions. Let us first define $g = v - \mu_i(t, T)$. This immediately implies that $v = g + \mu_i(t, T)$ and consequently

[14] Basically this approach paves over all of the underlying market-risk factors—such as interest rates, foreign-exchange, and key spreads—and operates immediately at the aggregate netting-set level.

[15] This is the type of exercise that is typically left for the reader, which always feels slightly unfair.

$dv = dg$ leading to:

$$\mathbb{E}_t\left(E_{t+T}(i)\right) = \int_0^\infty \frac{g+\mu_i(t,T)}{\sigma_i(t,T)\sqrt{2\pi}} e^{\frac{-g^2}{2\sigma_i(t,T)^2}} dg, \quad (10.11)$$
$$= \int_0^\infty \frac{g}{\sigma_i(t,T)\sqrt{2\pi}} e^{\frac{-g^2}{2\sigma_i(t,T)^2}} dg + \mu_i(t,T)\int_0^\infty \frac{1}{\sigma_i(t,T)\sqrt{2\pi}} e^{\frac{-g^2}{2\sigma_i(t,T)^2}} dg.$$

Resolution of the first integral in Eq. 10.11 requires a second substitution. Setting $u = \frac{-g^2}{2\sigma_i(t,T)^2}$ gives us $du = \frac{-gdg}{\sigma_i(t,T)^2}$ where $gdg = -\sigma_i(t,T)^2 du$. Plugging this back into Eq. 10.11 and using some basic properties of the normal distribution, we may simplify as:

$$\mathbb{E}_t\left(E_{t+T}(i)\right) = \frac{-\sigma_i(t,T)^{\not 2}}{\not\sigma_i(t,T)\sqrt{2\pi}}\int_0^\infty e^u du$$
$$+\mu_i(t,T)\int_0^\infty \frac{1}{\sigma_i(t,T)\sqrt{2\pi}} e^{\frac{-(v-\mu_i(t,T))^2}{2\sigma_i(t,T)^2}} dv, \quad (10.12)$$
$$= \frac{-\sigma_i(t,T)}{\sqrt{2\pi}}\Big[e^u\Big]_0^\infty + \mu_i(t,T)\underbrace{\int_{\frac{\mu_i(t,T)}{\sigma_i(t,T)}}^{-\infty} \frac{1}{\sqrt{2\pi}} e^{\frac{-v^2}{2}} dv}_{\text{To standard normal}},$$
$$= \frac{-\sigma_i(t,T)}{\sqrt{2\pi}}\left[e^{\frac{-(v-\mu_i(t,T))^2}{2\sigma_i(t,T)^2}}\right]_0^{-\infty} + \mu_i(t,T)\underbrace{\int_{-\infty}^{\frac{\mu_i(t,T)}{\sigma_i(t,T)}} \frac{1}{\sqrt{2\pi}} e^{\frac{-v^2}{2}} dv}_{\text{By symmetry}},$$
$$= \frac{-\sigma_i(t,T)}{\sqrt{2\pi}}\left(\underbrace{\lim_{v\to\infty} e^{\frac{-(v-\mu_i(t,T))^2}{2\sigma_i(t,T)^2}}}_{=0} - e^{\frac{-\mu_i(t,T)^2}{2\sigma_i(t,T)^2}}\right)$$
$$+\mu_i(t,T)\underbrace{\Phi\left(\frac{\mu_i(t,T)}{\sigma_i(t,T)}\right)}_{\text{By definition}},$$
$$= \frac{\sigma_i(t,T)e^{\frac{-\mu_i(t,T)^2}{2\sigma_i(t,T)^2}}}{\sqrt{2\pi}} + \mu_i(t,T)\Phi\left(\frac{\mu_i(t,T)}{\sigma_i(t,T)}\right),$$
$$= \sigma_i(t,T)\phi\left(\frac{\mu_i(t,T)}{\sigma_i(t,T)}\right) + \mu_i(t,T)\Phi\left(\frac{\mu_i(t,T)}{\sigma_i(t,T)}\right),$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal probability density and cumulative distribution functions, respectively.

If we recall our definitions of $\mu_i(t, T)$ and $\sigma_i(t, T)$, the expected exposure (or EE) becomes, in its full glory,

$$\mathbb{E}_t\left(E_{t+T}(i)\right) = \sigma_i\sqrt{T-t}\phi\left(\frac{V_t(i)+\mu_i(T-t)}{\sigma_i\sqrt{T-t}}\right) + \left(V_t(i)+\mu_i(T-t)\right)\Phi\left(\frac{V_t(i)+\mu_i(T-t)}{\sigma_i\sqrt{T-t}}\right), \quad (10.13)$$

for the ith counterparty and any arbitrary choice of $T \in (t, \infty)$. This may not be the most pleasant looking or intuitive formula, but it is a concrete analytic consequence of assuming one's netting set's market value follows a drifted Brownian motion process and a firm focus on positive value outcomes. This expression also lies, as we will see in following discussion, at the heart of the suggested regulatory approach towards computation of derivative exposure-at-default estimates.[16]

10.4.2 Expected Positive Exposure

Given the form of Eq. 10.13, it would be convenient to derive another formula for the expected positive exposure. Unfortunately, it is not so straightforward. The cumulative distribution function, $\Phi(\cdot)$, is not integrable. Adding another time-related integral over $(t, T]$ does not help matters. In practical settings, it is common practice to analytically compute the expected exposure using Eq. 10.13 over a number of steps along a time partition of $(t, T]$; we then approximate the EPE as the average over this discrete partition of the time domain. This is not particularly satisfying, but it is reasonably manageable. If one desires more precision, then it is always entirely possible to numerically integrate Eq. 10.13 over any individual's time interval of interest.

Figure 10.3 attempts to provide a bit more colour by displaying the analytic expected exposure over a given time period. This example was (arbitrarily) computed using the formula in Eq. 10.13 with $V_t(i) = 5$, $\mu_i = 0.05$, $\sigma_i = 1.75$, and $T - t = 10$. The expected positive exposure is also computed in *two* ways: as a simple average of the expected exposure observations and via numerical integration.[17] There is little or no (practical) difference between these two approaches.

---

[16] Equation 10.13 plays, roughly speaking, the same role for regulatory counterparty risk exposure calculations as Gordy [12]'s ASRF model in the Basel IRB methodology. See Bolder [9, Chapter 6] for an introduction to these concepts. We'll also return to this point in Chap. 11 when we examine the broader regulatory capital perspective.

[17] For the average computation, *four* discrete steps are taken for each year over $T - t$ for a total of 40 expected-exposure evaluations. As we increase the granularity of the time grid, of course, these two approaches converge.

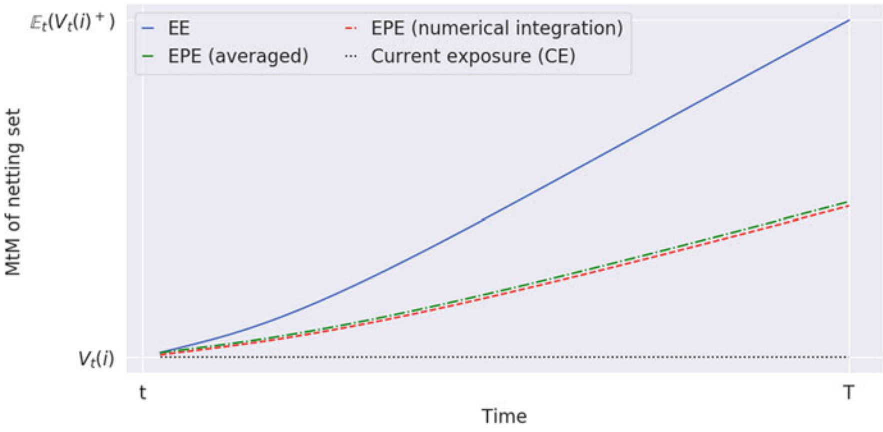

Fig. 10.3 *Analytic exposure*: This graphic displays the analytic expected exposure over a given time period. The EPE is also computed as a simple average of the EE observations and via numerical integration; there is little or no economic difference. This example was (arbitrarily) computed with $V_t(i) = 5$, $\mu_i = 0.05$, $\sigma_i = 1.75$, and $T - t = 10$.

If one is willing to make an additional assumption, then it is possible to obtain an analytic expression for the expected positive exposure. We need to assume that $\mu_i(t, T) = V_t(i) + \mu_i(T - t) \equiv 0$. Practically, this means that we must eliminate the drift element from our Brownian motion in Eq. 10.6. This transforms Eq. 10.13 into

$$\mathbb{E}_t\left(E_{t+T}(i)\right) = \underbrace{\sigma_i\sqrt{T-t}\phi(0) + (0)\cdot\Phi\left(\frac{0}{\sigma_i\sqrt{T-t}}\right)}_{\text{Eq. 10.13}}, \tag{10.14}$$

$$= \frac{\sigma_i\sqrt{T-t}}{\sqrt{2\pi}}.$$

This eliminates the annoying $\Phi(\cdot)$ term and easily allows us to determine an expected positive exposure expression as,

$$\begin{aligned}
\frac{1}{T-t}\int_t^T \mathbb{E}_t\left(E_{t+\tau}(i)\right)d\tau &= \frac{1}{T-t}\int_t^T \frac{\sigma_i\sqrt{\tau-t}}{\sqrt{2\pi}}d\tau, \qquad (10.15)\\
&= \frac{\sigma_i}{\sqrt{2\pi}(T-t)}\int_t^T \sqrt{\tau-t}\,d\tau,\\
&= \frac{\sigma_i}{\sqrt{2\pi}(T-t)}\left[\frac{2}{3}(\tau-t)^{\frac{3}{2}}\right]_t^T,
\end{aligned}$$

$$= \frac{2\sigma_i}{3\sqrt{2\pi}\cancel{(T-t)}}(\tau - t)^{\cancel{\frac{3}{2}}},$$

$$= \frac{2\sigma_i\sqrt{\tau - t}}{3\sqrt{2\pi}}.$$

The conclusion is that if the netting-set market value follows a Brownian motion with a zero drift and time-homogeneous variance, then both the EE and EPE measures are simply multiples of the volatility and the passage of time. Admittedly, these are perhaps not the most reasonable possible assumptions one might take. They do, however, permit (mostly) analytic formulae and a point of comparison.

10.4.3 Potential Future Exposure

Morgan/Reuters [17] originally suggested the VaR measure in the context of portfolio loss. Equation 10.5 provides a twist on this idea. Let's repeat it here—in a slightly different form—to examine it in more detail, help consider alternatives, and understand under what conditions it might be determined analytically. The α-VaR of the ith counterparty's netting-set exposure over the horizon $T - t$ can be written as,

$$q_\alpha(i, t, T) = \inf_{x \in \mathbb{R}} \left(x \,\middle|\, \mathbb{P}\Big(V_{t+T}(i) \geq x \Big) \geq 1 - \alpha \right). \tag{10.16}$$

The function depends on *two* parameters: α, which a threshold for the probability on the right-hand side of Eq. 10.16 and the time horizon, $T - t$. Imagine that we set α to 0.95 and $T - t = \frac{1}{52}$ or one week. With these parameter choices, the VaR describes the smallest exposure outcome that exceeds 95% of possible valuations, but is less than 5% of them. In other words, $q_\alpha(i, t, T)$ is the upper $(1 - \alpha)$-quantile of the netting set's exposure distribution.[18] By common convention, the VaR is denoted as $q_\alpha(i, t, T)$ to describe how far, for the ith counterparty, out into the future and the distribution's tail one wishes to go.

A second, increasingly common risk measure, is defined as the following conditional expectation,

$$\mathbb{E}\left(V_{t+T}(i) \,\middle|\, V_{t+T}(i) \geq q_\alpha(i, t, T) \right) = \frac{1}{1-\alpha} \int_{q_\alpha(i,t,T)}^{\infty} v f_{V_{t+T}(i)}(v) dv. \tag{10.17}$$

In words, Eq. 10.17 essentially describes the exposure given that one find's oneself at or beyond the $(1 - \alpha)$-quantile, or $q_\alpha(i, t, T)$, level. This quantity is, for this

[18] Quantiles are points in a distribution pertaining to the rank order of the distribution's values. It is a generalization of the notion of a percentile, which is defined on a scale of 0 to 100.

reason, often termed the conditional Value-at-Risk, the tail VaR, or the expected shortfall.[19] We will use the latter term and denote it as $\mathcal{E}_\alpha(i, t, T)$ to explicitly include the desired quantile defining the tail of the return distribution (as well as the time horizon and credit counterparty). Either Eq. 10.16 or 10.17 is a reasonable candidate for the potential future exposure.

Both of these quantities are fairly abstract. With a bit of additional effort and the assumption of Gaussianity, we may derive analytic expressions for both measures. From Eq. 10.16, the $(1-\alpha)$-quantile satisfies the following (by-now-quite-familiar) relation for a standard normal variate,

$$\begin{aligned} \int_{q_\alpha(i,t,T)}^{\infty} f_{V_{t+T}(i)}(v)dv &= 1-\alpha, \\ 1 - \int_{-\infty}^{q_\alpha(i,t,T)} f_{V_{t+T}(i)}(v)dv &= 1-\alpha, \\ \int_{-\infty}^{q_\alpha(i,t,T)} f_{V_{t+T}(i)}(v)dv &= \alpha, \\ \mathbb{P}\bigg(V_{t+T}(i) \leq q_\alpha(i,t,T)\bigg) &= \alpha, \\ \Phi\bigg(q_\alpha(i,t,T)\bigg) &= \alpha, \\ q_\alpha(i,t,T) &= \Phi^{-1}(\alpha). \end{aligned} \tag{10.18}$$

The standard inverse normal cumulative distribution function, $\Phi^{-1}(\cdot)$, does not have a closed-form solution, but through a variety of approximations, it is found in virtually any software package.[20]

Equation 10.18 summarizes the results for a standard normal random variable. All that remains is to scale the outcome to the moments of our portfolio-return distribution. Recall that if $Z \sim \mathcal{N}(0,1)$, then $X = a + \sqrt{b}Z \sim \mathcal{N}(a, b)$. Using this logic, it follows that

$$\begin{aligned} q_p(\alpha, T-t) &= \mu_i(t,T) + q_\alpha(i,t,T)\sigma_i(t,T), \\ &= V_t(i) + \mu_i(T-t) + \underbrace{\Phi^{-1}(\alpha)}_{\text{Eq. 10.18}} \sigma_i\sqrt{T-t}. \end{aligned} \tag{10.19}$$

[19] The expected shortfall measure—a central metric for our economic-capital model—also fulfills all of the criteria for a co-called *coherent* risk measure. See Artzner et al. [1] for an introduction to this theoretically important notion.

[20] An α of 0.95, for example, leads to a constant value $\Phi^{-1}(0.95) \approx 1.65$, whereas an α of 0.99 leads to $\Phi^{-1}(0.99) \approx 2.33$.

Given an α of 0.99, $V_t(i) = 5$, $\mu_i = 0.05$, $\sigma_i = 1.75$, and $T - t = \frac{1}{52}$, a possible VaR-based PFE estimate takes the value,

$$
\begin{aligned}
q_\alpha(i,t,T) &= \underbrace{5 + 0.05 \cdot \left(\frac{1}{52}\right)}_{\mu_i(t,T)} + \Phi^{-1}(0.99) \cdot \underbrace{1.75 \cdot \sqrt{\frac{1}{52}}}_{\sigma_i(t,T)}, \qquad (10.20)\\
&= 5.10 + 2.33 \cdot 0.24,\\
&= 5.66.
\end{aligned}
$$

In this setting, the 99% VaR would be 5.66 units of currency. Alternatively, we estimate that the worst-case increase in exposure over the next week, with a 99% degree of confidence, is $q_\alpha(i,t,T) - V_t(i) = 0.66$ units of currency. The computation itself is trivial. All of the effort is found in the identification of robust and sensible expected return and portfolio volatility estimates.

To find a workable expression for an expected-shortfall-based PFE, we again begin with a standard normal random variate, $V_{t+T}(i) \sim \mathcal{N}(0,1)$. The average value in the tail beyond the $(1-\alpha)$-quantile is defined as follows,

$$
\begin{aligned}
\mathbb{E}\left(V_{t+T}(i) \middle| V_{t+T}(i) \geq q_\alpha(i,t,T)\right) &= \frac{1}{1-\alpha}\int_{q_\alpha(i,t,T)}^{\infty} v f_{V_{t+T}(i)}(v)dv, \qquad (10.21)\\
&= \frac{1}{1-\alpha}\int_{\Phi^{-1}(\alpha)}^{\infty} v\phi(v)dv,\\
&= \frac{1}{1-\alpha}\int_{\Phi^{-1}(\alpha)}^{\infty} \frac{v}{\sqrt{2\pi}} e^{-\frac{v^2}{2}}dv.
\end{aligned}
$$

This is not, at first glance, a terribly easy integral to solve. If one computes $\phi'(v)$, the solution presents itself directly. In particular,

$$
\begin{aligned}
\phi(v) &= \frac{1}{\sqrt{2\pi}} e^{-\frac{v^2}{2}}, \qquad (10.22)\\
\phi'(v) &= \left(\frac{-2v}{2}\right)\frac{1}{\sqrt{2\pi}} e^{-\frac{v^2}{2}},\\
&= \frac{-v}{\sqrt{2\pi}} e^{-\frac{v^2}{2}},\\
&= -v\phi(v),
\end{aligned}
$$

where $\phi(\cdot)$ denotes the standard normal density function. Inserting Eq. 10.22 into 10.21, we have

$$\begin{aligned}\mathbb{E}\left(V_{t+T}(i)\middle| V_{t+T}(i) \geq q_\alpha(i,t,T)\right) &= \frac{1}{1-\alpha}\int_{\Phi^{-1}(\alpha)}^{\infty} \underbrace{\frac{v}{\sqrt{2\pi}}e^{-\frac{v^2}{2}}}_{\text{Eq. 10.22}} dv, \\ &= \frac{1}{1-\alpha}\int_{\Phi^{-1}(\alpha)}^{\infty} -\phi'(v)dv, \\ &= \frac{1}{1-\alpha}\Big[-\phi(v)\Big]_{\Phi^{-1}(\alpha)}^{\infty}, \\ &= \frac{1}{1-\alpha}\Big[-\underbrace{\lim_{v\to\infty}\phi(v)}_{=0} + \phi(\Phi^{-1}(\alpha))\Big], \\ &= \frac{\phi\left(\Phi^{-1}(\alpha)\right)}{1-\alpha}.\end{aligned} \tag{10.23}$$

Once again, this need only be scaled using our netting-set moments.

$$\begin{aligned}\mathcal{E}_\alpha(i,t,T) &\equiv \mathbb{E}\left(V_{t+T}(i)\middle| V_{t+T}(i) \geq q_\alpha(i,t,T)\right), \\ &= \mu_i(t,T) + \underbrace{\frac{\phi\left(\Phi^{-1}(\alpha)\right)}{1-\alpha}}_{\text{Eq. 10.23}}\sigma_i(t,T), \\ &= V_t(i) + \mu_i(T-t) + \frac{\phi\left(\Phi^{-1}(\alpha)\right)}{1-\alpha}\sigma_i\sqrt{T-t}.\end{aligned} \tag{10.24}$$

Returning to our previous example with $\alpha = 0.99$ and the same set of parameter values used earlier in the VaR setting, we may compute

$$\begin{aligned}\mathcal{E}_{0.99}(i,t,T) &= \underbrace{5 + 0.05\cdot\left(\frac{1}{52}\right)}_{\mu_i(t,T)} + \frac{\phi(\Phi^{-1}(0.99))}{1-0.99}\cdot\underbrace{1.75\cdot\sqrt{\frac{1}{52}}}_{\sigma_i(t,T)}, \\ &= 5.10 + 2.67\cdot 0.24, \\ &= 5.74.\end{aligned} \tag{10.25}$$

It makes logical sense that the expected shortfall should exceed the VaR for the same value of α. Indeed, since the expected shortfall is the average of all values beyond the VaR values, it must logically be greater than (or, at the very least, equal to) the VaR computation. Conceptually, since we are already quite far out in the tail

Table 10.1 *Simple example details*: This table summarizes, at a glance, the numerical details of various exposure measures in the context of a rather simple example.

Characteristic	Notation	Value
Starting point	$V_t(i)$	5.00
Drift (or trend)	μ_i	0.05
Diffusion (or volatility)	σ_i	1.75
Time horizon	$T - t$	0.02
Confidence level	α	0.99
VaR-based PFE	$q_\alpha(i, t, T)$	5.66
Expected-shortfall-based PFE	$\mathcal{E}_\alpha(i, t, T)$	5.74

and the Gaussian distribution has relatively small amounts of probability mass in its extremes, it also makes sense that the expected shortfall is not dramatically larger than the VaR values—in this case, it exceeds it by about $\mathcal{E}_\alpha(i, t, T) - q_\alpha(i, t, T) = 0.08$ units of currency.[21]

Table 10.1 summarizes the results of this simple numerical exercise. Either the VaR measure from Eq. 10.19 or the expected shortfall metric in Eq. 10.24 is a reasonable choice for the ith counterparty's potential future exposure. The expected shortfall is, by construction, somewhat more conservative since it explicitly delves into the tail of the netting set's value distribution.[22]

Colour and Commentary 118 (CREDIT COUNTERPARTY RISK ANALYTIC FORMULAE)*: This section has illustrated how far we can go, in terms of analytic expressions for important counterparty credit risk quantities, on the back of the assumption of drifted Brownian motion for one's netting set dynamics. Expected exposure, expected positive exposure (in the special case of* $\mu_i(t, T) \equiv 0$)*, and potential future exposure all boil down to relatively concise formulae. While useful and insightful, it is important to keep our enthusiasm somewhat in check. The central assumption is, if we are honest with ourselves, quite likely flawed. Extensive empirical evidence indicates that few, if any, important financial risk factors follow a Gaussian distribution. The basic form and symmetry are consistent, but actual risk-factor and market-value distributions are significantly more heavy tailed than the assumption of normality would suggest. Perhaps more problematic is the underlying complexity of the true netting-set valuation process. Given the broad range of potential pay-off patterns and underlying risk factors, treating*

(continued)

[21] In the this Gaussian case, the parametric expected shortfall estimator simply amounts to a change in the magnitude of the multiplier constant. The VaR multiplier is, for the 99% confidence level, equal to about 87% of the expected-shortfall value.

[22] A similar expression can be derived for different assumptions regarding the stochastic dynamics of the netting-set process. Chapter 10, for example, provides similar computations under the t distribution.

Colour and Commentary 118 (continued)
the overall netting-set as a single monolithic object is hard to defend. For this reason, as we've already highlighted, these various derivative exposures definitions are typically computed using rather complicated simulation engines. The previous investigation, however, remains pertinent since these formulae represent the foundation for our chosen path: simplified regulatory approximations. It is also useful as an easy-to-understand diagnostic and comparator for more complex derivative-exposure calculations performed in other parts of one's organization.

10.5 The Regulatory Approach

Virtually every piece of regulatory guidance has a history; it can often prove surprisingly helpful in appreciating the general context and what the regulators seek to accomplish. It is nevertheless often difficult, for any given regulatory standard, to know how far in time one should go back. In the area of counterparty credit risk, perhaps the first reference is BIS [2] where the idea of a netted exposure add-on, to account for potential value increases, was first introduced. BIS [2] is a revision to the original 1988 Basel Accord. While interesting, this is probably a bit too far back into the past. A better starting point would be BIS [3, Part I], which was written about ten years after the first revision. This document provides a thorough description of the so-called current-exposure and standardized methods; referred to understandably as CEM and SM, respectively. These models are very useful background into the regulator's mindset with regard to counterparty credit risk.

The directives found in BIS [3] were nonetheless superseded about another decade later with BIS [5]. This paper contains an updated exposure-at-default measurement methodology, which has been catchily dubbed the *standard approach for counterparty credit risk*. This being rather a mouthful has led to the broad-based use of the unfortunate acronym, SA-CCR. Jumping into BIS [5] without any preparation can be a bit overwhelming; perhaps appreciating this fact, a helpful companion paper—see BIS [6]—was produced a few months later. BIS [7] provides further responses to frequently asked questions. This section seeks to provide a workable description of the Basel Committee on Banking Supervision's (BCBS) SA-CCR methodology. It makes ample use of BIS [2, 3, 5, 6, 7]—along with a number of other references mentioned along the way—to accomplish this task. It is precisely this so-called SA-CCR approach that we will use to describe our (non-deterministic) derivative exposures for the purposes of economic-capital computations.[23]

Like most regulatory guidance, the SA-CCR is highly prescriptive and stylized. This is, given regulatory objectives, both natural and understandable. It nonetheless consists of a sequence of formulae that require some unpacking and deciphering. Let's not lose sight of our principal objective: a reasonably realistic and parsimonious description of the current exposure-at-default (i.e., EAD) associated with an arbitrary counterparty i. The SA-CCR high-level, entry-point formula for this object has the following form:

$$\text{EAD}_t(i) = \alpha \cdot \bigg(\text{RC}_t(i) + \text{PFE}_{t+T}(i)\bigg), \tag{10.26}$$

where $\text{RC}_t(i)$ denotes the replacement cost of the ith credit counterparty at time t and $\alpha \in \mathbb{R}_+$. We use the Greek letter α in this context to be consistent with the regulatory documentation. One should be careful to avoid confusion with the confidence level, which uses the same symbol.

Replacement cost is intended to represent the loss incurred in the event of *immediate* default. There is no lapse of time and, as a consequence, no possibility of upward movement in the market value of one's netting set. The PFE term, as you might expect, looks $T - t$ units of time into the future to capture possible value increases. The sum of these current and forward-looking effects are modified by a constant multiplier, α. As of the writing of this document, α was set to a value of 1.4.[24] The rationale for α is, it would appear, to provide a cushion for increased conservatism. The replacement-cost and PFE computations are distinct pieces with different perspectives. Although obviously not unrelated, it is convenient to address them separately. We'll follow this strategy in the coming sections and then pull them back together to complete the picture.

10.5.1 Replacement Cost

Of the two quantities in Eq. 10.26, the replacement cost is the least complicated to compute. The reason is simple; it only depends on the current point in time. In its basic form, the replacement cost is the sum of the market value of the individual trades in the ith countparty's netting set *less* the existing collateral. If the total market value is EUR 10 million and one holds EUR 8 million of collateral, then the replacement cost is EUR 2 million. Conversely, if you have a negative market

[23] Derivative exposure—as we'll soon see—also typically makes an appearance in leverage calculations.

[24] The constant α may take different values for different applications of the resulting exposure estimate. When, for example, computing leverage—see BIS [4]—the value of α is fixed at unity.

value of EUR −20 million and have posted EUR 15 million of collateral, then you have a EUR −5 million replacement cost. This logic is approximately correct, but also a bit faulty. We are not terribly interested in negative replacement costs and we need to be a bit more conservative in our treatment of collateral.

To address these issues, BIS [5] defines the replacement cost as

$$\mathrm{RC}_t(i) = \begin{cases} \max\left(V_t(i) - \tilde{C}_{t+T}(i), 0\right) & : \text{No margin} \\ \max\left(V_t(i) - \tilde{C}_{t+T}(i), 0, \mathrm{TH}_i + \mathrm{MTA}_i - \mathrm{NICA}_t(i)\right) & : \text{Margin} \end{cases} \tag{10.27}$$

This is fairly detailed definition that requires a bit of digestion. Indeed, Eq. 10.27 contains a number of acronyms that are desperately in need of explanation. The first point is that the replacement-cost definition differs by margin policy. For unmargined netting sets, the replacement cost is simply the positive part of the difference between the current market value and collateral holdings. Since there is always the danger that non-cash collateral can lose value—due to adverse market movements—$\tilde{C}_t(i)$ denotes the haircut value of any net collateral position.[25] It is, in a technical sense, described as,

$$\tilde{C}_{t+T}(i) = \begin{cases} C_t(i) \cdot \left(1 - h(t,T)\right) & : \text{Holding collateral where } C_t(i) > 0 \\ C_t(i) \cdot \left(1 + h(t,T)\right) & : \text{Posting collateral where } C_t(i) < 0 \end{cases}, \tag{10.28}$$

where $h(t, T)$ is the haircut value over the time interval $T - t$ and $C_t(i)$ represents the market value of the ith counterparty's net collateral position at time t.[26] The introduction of T, as an argument, is necessary to reflect the role of the time-horizon in the collateral haircut computation and our broader analysis.

The notion of *net* collateral will help us make sense of the second part of the replacement-cost definition for margined netting sets in Eq. 10.27. The initial margin or independent amount, given their similarity and the potential for confusion, are referred to as independent collateral amount (or ICA) in the regulatory setting. Unlike variation margin, the ICA can be simultaneously posted and received. This immediately leads to the idea of net ICA, or NICA, which represents the difference between received and posted initial collateral. The NICA value might be either positive or negative; it really depends on the magnitude of the ICA values associated with both counterparties. In the non-margin setting, therefore, the $\tilde{C}_{t+T}(i)$ value

[25] A haircut is a rather colourful, colloquial term used to describe any adjustment to the valuation of a collateral position to account for possible unfavourable market movements.

[26] This is a very simplified set of haircut rules; it can, in practice, be much more nuanced.

basically refers to the NICA amount with the appropriate application of haircuts as described in Eq. 10.28.

In the margined cases, the current collateral amount, $\tilde{C}_{t+T}(i)$, is related to variation margin. The second line of Eq. 10.28, however, has three arguments in the max operator. The third argument, $\text{TH}_i + \text{MTA}_i - \text{NICA}_t(i)$, includes the net independent collateral amount at the current time, t. This maintains a separation between the variation and initial margins. The TH and MTA variables are determined through the legal counterparty contracts governing the management of margins. TH is an acronym for the collateral threshold; this is a lower bound, below which collateral is not exchanged. For example, if there is a threshold of EUR 100,000 and the market value of the netting set is EUR 50,000, then no collateral would be received. The point of the threshold is to reduce the operational costs associated with calling and returning collateral. Naturally, this creates a trade-off between risk mitigation and logistical costs.[27]

MTA represents the minimum transfer amount, which unsurprisingly is the lowest amount of collateral that may be transferred between two entities. Much like the threshold, the intention is to reduce operational overhead. It probably doesn't make logical sense to exchange EUR 10 of collateral for a minuscule change in market value.[28] The threshold and minimum transfer amount should *not*, however, be considered separately. As clearly stated in Gregory [13]:

> The minimum transfer amount and threshold are additive in the sense that the exposure must exceed the sum of the two before any collateral can be called.

This succinctly explains why these two values enter jointly into the margined netting-set replacement cost in Eq. 10.27. This extra condition covers the situation where there is positive exposure, but it has not broken through the lower bound required for receipt of collateral. The net ICA is subtracted from this amount to account for the offsetting impact of the initial margin. In short, Eq. 10.27 is a rather conservative view of the current losses associated with immediate counterparty default.

10.5.2 The Add-On

Conservatism of its construction aside, the replacement cost does not tell the whole story. It completely ignores the possibility of future value increases in the counterparty's netting set. This is precisely the role of the second term in the high-level SA-CCR expression presented in Eq. 10.26. It has the following form,

$$\text{PFE}_t(i) = \mathcal{M}_i \cdot \mathbb{A}_i^{\Sigma}, \tag{10.29}$$

[27] One may, of course, set the threshold to zero to move completely towards the risk-mitigation direction.

[28] Such a value should be somehow related—logically, at least, if not practically—to the confidence interval around our valuation estimates.

where $\mathcal{M}_i \in (0.05, 1]$ is a multiplier with a reasonably complex form and $\mathbb{A}_i^{\Sigma}$ is referred to as the aggregate *add-on* for the ith credit counterparty. Putting the multiplier aside for the moment, a central point is that what SA-CCR refers to as PFE is actually something else. In fact, BIS [6] indicates that it is intended to represent a conservative estimate of the netting set's effective expected positive exposure. To repeat, it does *not* technically describe potential future exposure.

The aggregate add-on, as the name strongly suggests, is some kind of function of the underlying trades within the netting set designed to capture the risk of future derivative-value increases. This naturally involves a number of levels and, of course, a logical hierarchy to determine the proper order of aggregation. It is, quite frankly, a bit messy. Figure 10.4 provides, in an attempt to help matters, a visualization of the *five* levels associated with the SA-CCR methodology. Individual trades are first organized into *five* distinct asset classes. These are essentially flavours of derivative contracts: interest rate, currency, credit, equity, and commodity. This immediately raises an important question: what is done with derivatives that touch upon multiple asset classes?[29] BIS [5]'s answer reads:

> The designation should be made according to the nature of the primary risk factor. [...] For more complex trades, where it is difficult to determine a single primary risk factor, bank supervisors may require that trades be allocated to more than one asset class.

For the purposes of this discussion, we will assume that each derivative trade can be reasonably and defensibly assigned to a single asset class. The aggregate add-on is thus,

$$\mathbb{A}_i^{\Sigma} = \sum_{a \in \mathcal{A}_i} \mathbb{A}_i^{(a)}, \tag{10.30}$$

where $\mathcal{A}_i$ denotes the collection of asset classes associated with netting set i. One thus simply sums over the asset classes to get to the aggregate level.[30]

Within an asset class, the trades are further allocated into so-called hedging sets. Again BIS [5], who introduce the concept, offer the best definition:

> a hedging set is the largest collection of trades of a given asset class within a netting set for which netting benefits are recognized in the PFE add-on of the SA-CCR.

The corollary, of course, is that hedging sets within an asset class are not permitted to offset one another. If we write the collection of hedging sets within asset class a and netting set i as $\mathcal{H}_i^{(a)}$, then these hedging sets are aggregated as

$$\mathbb{A}_i^{(a)} = \sum_{h \in \mathcal{H}_i^{(a)}} \left| \mathbb{A}_i^{(a,h)} \right|, \tag{10.31}$$

[29] A cross-currency swap, for example, has both currency and interest rate exposure.

[30] To be crystal clear with this invented notation, if a netting set includes N assets classes then $\mathcal{A}_i = \left\{ \mathcal{A}_i^{(1)}, \mathcal{A}_i^{(2)}, \cdots, \mathcal{A}_i^{(N)} \right\}$.

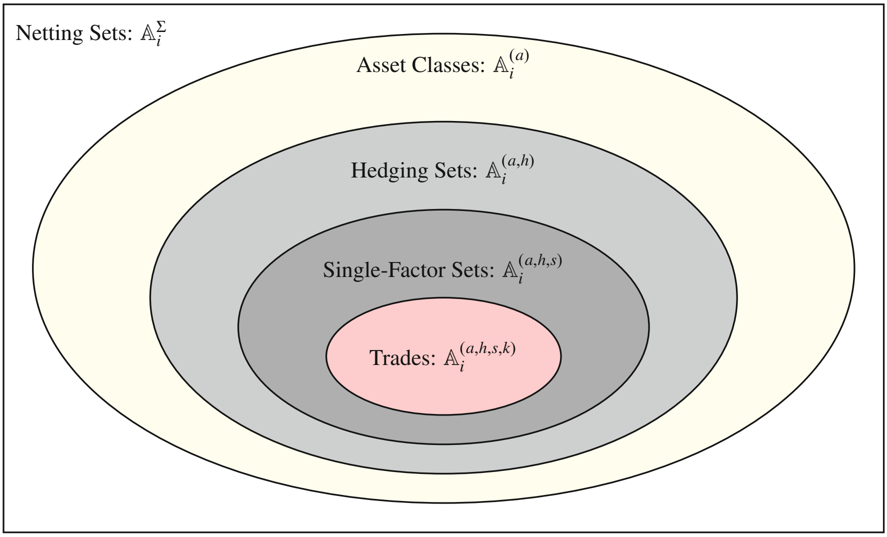

Fig. 10.4 *The add-on hierarchy*: The schematic above outlines the five levels of aggregation embedded in the SA-CCR methodology. In addition to visualizing how the aggregate add-on is computed, it also introduces the notation used to describe the various components.

where $\mathbb{A}_i^{(a,h)}$ represents the calculated add-on for the hth hedging set of asset class a associated with netting set i. $\mathcal{H}_i^{(a)}$ is the related collection of hedging sets. The notation is admittedly dreadful, but given the complexity of the hierarchy in Fig. 10.4, it will only get worse. The absolute value operator in Eq. 10.31 precludes any possible netting between these hedging sets.

The hedging set, as the name indicates, permits some form of diversification. This is captured through non-perfect positive correlation between the individual single factor sets comprising the hedging set. We further define the collection of hedging sets in asset class a and netting set i as $\mathcal{S}_i^{(a,h)}$. This allows us to write the aggregation of each hedging set as,

$$\mathbb{A}_i^{(a,h)} = \sqrt{\sum_{s\in\mathcal{S}_i^{(a,h)}} \sum_{r\in\mathcal{S}_i^{(a,h)}} \rho(s,r)\mathbb{A}_i^{(a,h,s)}\mathbb{A}_i^{(a,h,r)}}, \tag{10.32}$$

where $\mathbb{A}_i^{(a,h,s)}$ represents the sth single-factor set of hedging set h in asset class a and $\rho(r,s)$ is the (regulator determined) correlation between single factor sets r and s. These latter values are explicitly provided by the regulatory authorities. Each

$\mathbb{A}_i^{(a,h,s)}$ is referred to as a single-factor set. Equation 10.32 should be recognized as the standard deviation of a group of correlated random variables.[31] For it to work, however, we must be able to interpret the individual $\mathbb{A}_i^{(a,h,s)}$ terms as standard deviations.

The single-factor set add-ons, in turn, are the direct sum of the trades falling into that asset class, hedging set, and single-factor set. Calling this collection $\mathcal{T}_i^{(a,h,s)}$, the final level of aggregation is given as,

$$\mathbb{A}_i^{(a,h,s)} = \sum_{k \in \mathcal{T}_i^{(a,h,s)}} \mathbb{A}_i^{(a,h,s,k)}. \tag{10.34}$$

Practically, a single factor set is a sub-group of interest-rate tenors for a given interest curve, a currency, or a credit single entity. The driving idea is that the correlation within a single factor set is approximately perfectly positive, so they can be directly summed.

Each individual derivative trade follows a clear hierarchy up to the netting set level. The underlying logical chain illustrates how each trade, single-risk-factor, hedging-set, and asset-class collection is nested,

$$\mathcal{T}_i^{(a,h,s)} \subseteq \mathcal{S}_i^{(a,h)} \subseteq \mathcal{H}_i^{(a)} \subseteq \mathcal{A}_i, \tag{10.35}$$

for the ith netting set.[32]

Just for fun, let's now combine all of these elements together to create a single, abstract representation of the aggregate add-on associated with the ith counterparty.

[31] If $X_1, ..., X_N$ are a collection of N correlated random variables, then the variance of their sum can be represented as

$$\text{var}\left(X_1 + \cdots + X_N\right) = \underbrace{\sum_{i=1}^{N}\sum_{i=j}^{N} \text{corr}(X_i, X_j)\sqrt{\text{var}(X_i)}\sqrt{\text{var}(X_j)}}_{\sum_{i=1}^{N}\sum_{i=j}^{N}\text{cov}(X_i, X_j)}, \tag{10.33}$$

where we recover Eq. 10.32 by taking the square root of both sides.

[32] Or, if you prefer a word equation, we might describe the idea in Eq. 10.35 as

$$\underbrace{\text{Collection of Trades}}_{k \in \mathcal{T}_i^{(a,h,s)}} \subseteq \underbrace{\text{Single-Factor Sets}}_{s \in \mathcal{S}_i^{(a,h)}} \subseteq \underbrace{\text{Hedging Sets}}_{h \in \mathcal{H}_i^{(a)}} \subseteq \underbrace{\text{Asset Class}}_{a \in \mathcal{A}_i}. \tag{10.36}$$

Starting from Eq. 10.30, we have

$$
\begin{aligned}
\mathbb{A}_i^{\Sigma} &= \underbrace{\sum_{a\in\mathcal{A}_i} \mathbb{A}_i^{(a)}}_{\text{Eq. 10.30}}, \qquad (10.37)\\
&= \sum_{a\in\mathcal{A}_i} \underbrace{\sum_{h\in\mathcal{H}_i^{(a)}} \left|\mathbb{A}_i^{(a,h)}\right|}_{\text{Eq. 10.31}},\\
&= \sum_{a\in\mathcal{A}_i} \sum_{h\in\mathcal{H}_i^{(a)}} \left| \underbrace{\sqrt{\sum_{s\in\mathcal{S}_i^{(a,h)}} \sum_{r\in\mathcal{S}_i^{(a,h)}} \rho(s,r)\mathbb{A}_i^{(a,h,s)}\mathbb{A}_i^{(a,h,r)}}}_{\text{Eq. 10.32}} \right|,\\
&= \sum_{a\in\mathcal{A}_i} \sum_{h\in\mathcal{H}_i^{(a)}} \left| \sqrt{\sum_{s\in\mathcal{S}_i^{(a,h)}} \sum_{r\in\mathcal{S}_i^{(a,h)}} \rho(s,r) \underbrace{\sum_{k_s\in\mathcal{T}_i^{(a,h,s)}} \mathbb{A}_i^{(a,h,s,k_s)}}_{\text{Eq. 10.34}} \underbrace{\sum_{k_r\in\mathcal{T}_i^{(a,h,s)}} \mathbb{A}_i^{(a,h,r,k_r)}}_{\text{Eq. 10.34}}} \right|.
\end{aligned}
$$

With its almost frightening array of sums, Eq. 10.37, is not particularly easy to look at. It does, however, clearly chart out the path from individual trade to single-factor set to hedging set to asset class and, ultimately, to netting set as described in Fig. 10.4. Although slightly masochistic and certainly not the typical treatment of this material, this final expression does provide some useful intuition to a complex aggregation.

Colour and Commentary 119 (SA-CCR AGGREGATION LOGIC)*: As anyone with a sock drawer can tell you, getting all of the right socks paired up correctly with their partner can be a lengthy and annoying undertaking. The SA-CCR regulatory guidance add-on calculations involve a five-level hierarchy with three alternative aggregation approaches. The aggregation logic is, to be blunt, ugly and rather hard to follow. Like a particularly untidy and convoluted sock drawer, a bit of organization can go a long way. The preceding discussion consequently introduces some (imperfect) notation and logic to help keep some order in this hierarchy. A netting set is comprised of a number of distinct asset classes. Within each asset class, trades are*

(continued)

Colour and Commentary 119 (continued)
allocated into hedging sets. The various single factor sets, within a hedging set, combine together in a manner similar to variance. Between hedging sets, no offsetting is permitted; aggregation occurs with absolute values. Trades within each single factor set, by virtue of their perfect positive correlation, can be simply summed. Perhaps somewhat ironically, the need to create a rather straightforward formulaic exposure calculation is the main driver of this relatively messy structure. In a classic potential-future-exposure computation, the complexity is embedded in the dependence structure underlying the risk factors and their mapping to derivative valuations. Netting set aggregation is, in this setting, comparatively trivial.[a] *In the SA-CCR, the risk-factor structure is simplified. The cost is somewhat unwieldy aggregation logic.*

[a] The heavy-lifting has to be performed somewhere. In this case, it is performed by the forward-looking exposure engine.

10.5.3 The Trade Level

The effective expected positive exposure form of the add-on needs to meet a central objective. As we saw in the previous section, to use the variance-motivated aggregation at the hedging-set level as in Eq. 10.32, it is necessary to view the individual (hedging-set) add-ons as a standard deviation. To accomplish this, a number of assumptions regarding the market value of each individual trade are required. These include—and are also outlined in both BIS [5, 6]—a zero current market value, an absence of collateral, no cash-flows over the next one-year time horizon, and undrifted Brownian-motion dynamics.[33]

We now reap the benefits of having worked through the details of the Brownian-motion case. Let us employ k to denote an arbitrary trade with the i credit counterparty. The market value of the ith trade at time $t+T$, following from the preceding assumptions, is given as

$$V_{t+T}(i,k) \sim \mathcal{N}\left(0, \underbrace{\mathbb{I}_{\{m_k \geq T\}}\sigma_k\sqrt{T-t}}_{\sigma_{ik}(t,T)}\right), \tag{10.38}$$

[33] These are fairly strong assumptions and, to a certain extent, they are relaxed through the multiplier. We will address this point in forthcoming discussion.

where m_k is the remaining maturity. This is entirely consistent with the undrifted Brownian-motion case. The only real twist is that all trades maturing before time $t + T$ are excluded.

Equipped with the definition of an arbitrary trade's market-value, we observe that we are in the general, but simplified, structure described in the preceding sections. From Eqs. 10.13 and 10.14, the expected exposure associated with the kth trade is simply

$$\mathbb{E}_t\left(V_{t+T}(i,k)^+\right) = \frac{\mathbb{I}_{\{m_k \geq T\}}\sigma_k\sqrt{T-t}}{\sqrt{2\pi}}. \tag{10.39}$$

The actual value of T will vary depending on whether or not the netting set is margined or non-margined. In the non-margined case, $T - t = 1$ year. For margined netting sets, one needs to define a margin period of risk, which we will refer to as $t + \tau_{M_i}$ to reflect its dependence on the netting set. As discussed earlier, we can compute the expected exposure profile by evaluating (and averaging) Eq. 10.39 for a number of choices of $T \in (t, t+1]$ or $T \in (t, t+\tau_{M_i}]$ for non-margined and margined netting sets, respectively.

The authors of the SA-CCR (understandably) wanted to avoid—as highlighted in BIS [6]—the need to partition the time interval and repeatedly compute (and average) Eq. 10.39. This leads to a restatement of our maturity indicator variable as

$$\mathbb{I}_{\{m_k \geq t\}} \equiv 1. \tag{10.40}$$

This is equivalent to forcing all trades in one's portfolio to have a minimum maturity of—depending on the margin policy—either one-year or the margin period of risk.

This assumption allows us, following directly from Eq. 10.15, to represent the expected positive exposure as

$$\frac{1}{T-t}\int_t^T \mathbb{E}_t\left(V_{t+\tau}(i,k)^+\right)d\tau = \frac{2\sigma_k\sqrt{T-t}}{3\sqrt{2\pi}}, \tag{10.41}$$

where $T = 1$ or τ_{M_i}. Two features of Eq. 10.39 also make it the *effective* expected positive exposure: over $(0, T]$ there are neither cash-flows nor maturities and $\sigma_k \in \mathbb{R}$. The presence of these conditions ensures that the expression in Eq. 10.41 is (at least) weakly positively monotonic in T. Under this monotonicity condition the expected positive exposure (EPE) and effective expected positive exposure (EEPE) coincide.

Pulling this all together, the trade-level add-on has the following theoretical form,

$$\mathbb{A}_i^{(a,h,s,k)} = \underbrace{\frac{2\sigma_k\sqrt{T-t}}{3\sqrt{2\pi}}}_{\text{Eq. 10.41}}. \tag{10.42}$$

It is important to stress that this is not the final form of the trade-level add-on. Ultimately, as we'll see shortly, it has a rather different practical form. Nonetheless, following from BIS [6], Eq. 10.42 describes its origins.

Colour and Commentary 120 (SA-CCR PFE COMPONENT)*: It should be clear by this point that what the SA-CCR refers to as the PFE component is not actually, in the typical definition of the term, potential future exposure. It is, to be clear,* not *a tail-based measure of the worst-case upside movement in a trade or netting set's market value. Instead, it is defined with regard to the expectation of the positive part of the trade or netting set's future exposure; in the non-margined case, this becomes effective positive exposure, whereas this reduces to expected exposure in the margined setting. In both cases, numerous assumptions are necessary to preserve analytic formulae for these quantities. This choice is perfectly within the purview of the regulatory authorities and, given their objectives, makes logical sense. Their characterization of PFE does, after all, respect the fundamental idea of incorporating some additional amount to accommodate the possibility of future value increases. It nonetheless bears stressing that the add-on element should* not *be strictly interpreted—relying on its name—as a quantile-based characterization of the future counterparty level value distribution. Understanding this fact may very well help to avoid unnecessary confusion.*

The next step is to adjust Eq. 10.42 so that it naturally handles both margin and non-margined netting sets. In principle, one could simply change the value of T depending on the margin policy. As a practical matter, however, the EEPE form from Eq. 10.42 only really applies to non-margined netting sets. For margined netting sets, the margin period of risk is typically quite short: something between 10 to 20 working days. Averaging the expected positive exposure of this period would, according to the formula, reduce the add-on by one third. Apparently, this was judged to be insufficiently conservative. The SA-CCR consequently restates the trade-level add-on as,

$$\mathbb{A}_i^{(a,h,s,k)} = \begin{cases} \dfrac{2\sigma_k}{3\sqrt{2\pi}} & : \text{No margin (EEPE)} \\ \dfrac{\sigma_k\sqrt{\tau_{M_i} - t}}{\sqrt{2\pi}} & : \text{Margin (EE)} \end{cases}. \tag{10.43}$$

Again, this is not precisely how it is represented in the SA-CCR documentation. Instead, they express the trade-level add-on as,

$$\begin{aligned} \mathbb{A}_i^{(a,h,s,k)} &= \frac{2\sigma_k\sqrt{\text{1-year}}}{3\sqrt{2\pi}}\text{MF}_k, \qquad (10.44) \\ &= \frac{2\sigma_k}{3\sqrt{2\pi}}\text{MF}_k, \end{aligned}$$

where MF_k is called the maturity factor and it is defined as,

$$\text{MF}_k = \begin{cases} 1 & : \text{No margin (EEPE)} \\ \frac{3\sqrt{\tau_{M_i} - t}}{2} & : \text{Margin (EE)} \end{cases}. \tag{10.45}$$

Rather obviously, Eqs. 10.44 and 10.45 are entirely equivalent to the representation in Eq. 10.43. To avoid confusion in the application and discussion of the SA-CCR computation, it is probably best to make use of the maturity factor.

Colour and Commentary 121 (THE SA-CCR MATURITY FACTOR)*: The maturity factor is intended to simplify the SA-CCR implementation by permitting a common form irrespective of a netting set's margin policy. It is a rather small point, but role of the maturity factor in the trade-level add-on obscures a fundamental modelling choice. Non-margined netting set add-on values are based on the effective expected positive exposure measure. Margined netting sets, conversely, use straight-up expected exposure. Given the dramatic difference between the time interval involved, this is an understandable and conservative regulatory choice. It is still useful to* see through *the maturity factor form and to be aware of this choice. If one is comparing SA-CCR results to a model-based implementation, this could potentially lead to some confusion.*

BIS [6] indicates clearly that "the SA-CCR does not directly operate with trade volatilities." The actual form of the trade-level add-on is given as,

$$\mathbb{A}_i^{(a,h,s,k)} = \delta_k \cdot d_k^{(a)} \cdot \text{SF}_k^{(a)} \cdot \text{MF}_k, \tag{10.46}$$

where δ_k is the directional delta, $d_k^{(a)}$ is the adjusted notional, and $\text{SF}_k^{(a)}$ is the supervisory factor. The a subscript arises, because as we'll see shortly, some of these quantities are handled differently depending on their asset-class membership. δ_k essentially denotes the sign of the add-on, while δ_k describes its magnitude.[34] The supervisory factor provides the link back to the volatility structure found in Eq. 10.44. Indeed, if we equate Eqs. 10.44 and the absolute value of 10.46 and solve for σ_k we arrive at

$$\underbrace{\frac{2\sigma_k}{3\sqrt{2\pi}} \cancel{\text{MF}_k}}_{\text{Eq. 10.44}} = \left| \underbrace{\delta_k \cdot d_k^{(a)} \cdot \text{SF}_k^{(a)} \cdot \cancel{\text{MF}_k}}_{\text{Eq. 10.46}} \right|, \tag{10.47}$$

[34] For option contracts, the directional delta captures the sign, but also tells us something about moneyness.

$$\sigma_k = \left(\frac{3\sqrt{\pi}\mathrm{SF}_k^{(a)}}{\sqrt{2}}\right)\left|\delta_k\right| \cdot d_k^{(a)},$$
$$\sigma_k \approx \left(2.66 \cdot \mathrm{SF}_k^{(a)}\right)\left|\delta_k\right| \cdot d_k^{(a)}.$$

The punchline is that the volatility of the trade is simply a scaled version of the supervisory factor. What is particularly useful to know is that the term $\left(2.66 \cdot \mathrm{SF}_k^{(a)}\right)$ should be interpreted, according to BIS [6], as

> the standard deviation of the primary risk factor at the one-year horizon.

The individual analyst does not determine these values; they are provided by one's friendly neighbourhood regulator. These supervisory factors are presumably determined by individual asset class and single risk-factor definitions as a combination of calibration to historical outcomes and regulatory judgement.

Once again, pulling together the high-level aggregate add-on viewpoint from Eq. 10.37 and incorporating Eq. 10.46, we arrive at:

$$\mathbb{A}_i^{\Sigma} = \underbrace{\sum_{a\in\mathcal{A}_i}\sum_{h\in\mathcal{H}_i^{(a)}}\left|\sqrt{\sum_{s\in\mathcal{S}_i^{(a,h)}}\sum_{r\in\mathcal{S}_i^{(a,h)}}\rho(s,r)\sum_{k_s\in\mathcal{T}_i^{(a,h,s)}}\mathbb{A}_i^{(a,h,s,k_s)} \times \sum_{k_r\in\mathcal{T}_i^{(a,h,s)}}\mathbb{A}_i^{(a,h,r,k_r)}}\right|}_{\text{Eq. 10.37}}, \tag{10.48}$$

$$= \sum_{a\in\mathcal{A}_i}\sum_{h\in\mathcal{H}_i^{(a)}}\left|\sqrt{\sum_{s\in\mathcal{S}_i^{(a,h)}}\sum_{r\in\mathcal{S}_i^{(a,h)}}\rho(s,r)\sum_{k_s\in\mathcal{T}_i^{(a,h,s)}}\underbrace{\delta_{k_s}\cdot d_{k_s}^{(a,h,s)}\cdot \mathrm{SF}_{k_s}^{(a,h,s)}\cdot \mathrm{MF}_{k_s}}_{\text{Eq. 10.46}} \times \sum_{k_r\in\mathcal{T}_i^{(a,h,r)}}\underbrace{\delta_{k_r}\cdot d_{k_r}^{(a,h,s)}\cdot \mathrm{SF}_{k_r}^{(a,h,r)}\cdot \mathrm{MF}_{k_r}}_{\text{Eq. 10.46}}}\right|.$$

Admittedly somewhat overwhelming, this representation nevertheless indicates how the various pieces fall into place. The rather puzzling application of the absolute-value operator to the result of a square root suggests differential treatment among certain asset classes. We will turn to this important point after completing the overall add-on computation.

10.5.4 The Multiplier

The multiplier is an attempt to attenuate the rather strong assumptions involved in the construction of the add-on. The most violent choices, from an add-on perspective at least, entail ignoring any future collateral and forcing the current market value to zero. Unlike the individual add-ons, however, the multiplier is defined at the netting-set level. A bit of reflection suggests that the starting point, of any given netting set, could be significantly negative for one (or a combination) of two reasons: the overall netting-set market value is large and negative and/or there is over-collateralization. Both situations can, and do, happen in practice. In the event that the starting point is significantly negative, it seems reasonable to adjust the aggregate add-on downwards to reflect this fact. To not do so, would be to unfairly overestimate the SA-CCR aggregate add-on component.

Quite simply, the multipler is a netting-set level value in the unit interval—that is, $\mathcal{M}_i \in [0, 1]$—that modifies the aggregate add-on for a negative starting position. As with many regulatory quantities, it is a bit funky, but it also follows a certain logic. The actual construction takes a few steps, but relies on elements already introduced and discussed in previous sections. The first step involves a distributional assumption regarding the value of the ith netting set. Similar in spirit, but different in level of aggregation, to Eq. 10.38 we have

$$V_{t+T}(i) \sim \mathcal{N}\left(\overbrace{V_t(i) - \tilde{C}_{t+T}(i)}^{\mu_i(t,T)}, \overbrace{\underbrace{\mathbb{I}_{\{m_k \geq t\}}}_{=1} \sigma_i \sqrt{T-t}}^{\sigma_i(t,T)}\right), \tag{10.49}$$

Basically the mean starting-level has been added back. Immediately, again with recourse to Eq. 10.13, we can write the expected exposure for the ith netting set as,

$$\begin{aligned}\mathbb{E}_t\left(E_{t+T}(i)\right) &= \sigma_i\sqrt{T-t}\phi\left(\frac{V_t(i)-\tilde{C}_{t+T}(i)}{\sigma_i\sqrt{T-t}}\right)\\ &+\left(V_t(i)-\tilde{C}_{t+T}(i)\right)\Phi\left(\frac{V_t(i)-\tilde{C}_{t+T}(i)}{\sigma_i\sqrt{T-t}}\right).\end{aligned} \tag{10.50}$$

This form is not immediately helpful, since we do not actually work with the volatilities, but rather with add-ons. Equation 10.43 illustrates the theoretical link between the trade-level add-on and the volatility. Conceptually, for a margined netting set, this reduces to:

$$\begin{aligned}\mathbb{A}_i^{\Sigma} &= \frac{\sigma_i\sqrt{T-t}}{\sqrt{2\pi}}, \\ \sigma_i &= \frac{\mathbb{A}_i^{\Sigma}\sqrt{2\pi}}{\sqrt{T-t}}.\end{aligned} \tag{10.51}$$

We are, in this case, using the expected-exposure perspective add-on assumption associated with a margined netting set. This is an important point, because the expected positive exposure approach used for non-margined sets will not permit an analytic solution. We saw this result in an earlier section.

If we plug Eq. 10.51 back into Eq. 10.50 and simplify things a little bit, we arrive at

$$\mathbb{E}_t\left(E_{t+T}(i)\right) = \underbrace{\begin{array}{c}\sigma_i\sqrt{T-t}\phi\left(\frac{V_t(i)-\tilde{C}_{t+T}(i)}{\sigma_i\sqrt{T-t}}\right) + \left(V_t(i)-\tilde{C}_{t+T}(i)\right) \\ \times\Phi\left(\frac{V_t(i)-\tilde{C}_{t+T}(i)}{\sigma_i\sqrt{T-t}}\right)\end{array}}_{\text{Eq. 10.13}}, \tag{10.52}$$

$$= \underbrace{\left(\frac{\mathbb{A}_i^{\Sigma}\sqrt{2\pi}}{\cancel{\sqrt{T-t}}}\right)}_{\text{Eq. 10.51}}\cancel{\sqrt{T-t}}\phi\left(\frac{V_t(i)-\tilde{C}_{t+T}(i)}{\underbrace{\left(\frac{\mathbb{A}_i^{\Sigma}\sqrt{2\pi}}{\cancel{\sqrt{T-t}}}\right)}_{\text{Eq. 10.51}}\cancel{\sqrt{T-t}}}\right)$$

$$+\left(V_t(i)-\tilde{C}_{t+T}(i)\right)\Phi\left(\frac{V_t(i)-\tilde{C}_{t+T}(i)}{\underbrace{\left(\frac{\mathbb{A}_i^{\Sigma}\sqrt{2\pi}}{\cancel{\sqrt{T-t}}}\right)}_{\text{Eq. 10.51}}\cancel{\sqrt{T-t}}}\right),$$

$$= \mathbb{A}_i^{\Sigma}\sqrt{2\pi}\phi\left(\frac{V_t(i)-\tilde{C}_{t+T}(i)}{\mathbb{A}_i^{\Sigma}\sqrt{2\pi}}\right)$$

$$+\left(V_t(i)-\tilde{C}_{t+T}(i)\right)\Phi\left(\frac{V_t(i)-\tilde{C}_{t+T}(i)}{\mathbb{A}_i^{\Sigma}\sqrt{2\pi}}\right),$$

$$= \mathbb{A}_i^{\Sigma}\left(\sqrt{2\pi}\cdot\phi\left(\frac{\upsilon}{\sqrt{2\pi}}\right)+\upsilon\cdot\Phi\left(\frac{\upsilon}{\sqrt{2\pi}}\right)\right),$$

where

$$\upsilon = \frac{V_t(i)-\tilde{C}_{t+T}(i)}{\mathbb{A}_i^{\Sigma}}. \tag{10.53}$$

The consequence is a (fairly) concise analytic representation of the expected exposure associated with the ith netting set written in terms of known quantities: the aggregate add-on, the current netting-set valuation, and the collateral position.

The central part, or kernel, of the SA-CCR multiplier is now defined as the quotient of the expected exposure representation in Eq. 10.52 and the aggregate add-on. In particular, this gives

$$\frac{\mathbb{E}_t\left(E_{t+T}(i)\right)}{\mathbb{A}_i^{\Sigma}} = \frac{\overbrace{\cancel{\mathbb{A}_i^{\Sigma}}\left(\sqrt{2\pi}\cdot\phi\left(\frac{v}{\sqrt{2\pi}}\right)+v\cdot\Phi\left(\frac{v}{\sqrt{2\pi}}\right)\right)}^{\text{Eq. 10.52}}}{\cancel{\mathbb{A}_i^{\Sigma}}}, \qquad (10.54)$$
$$= \sqrt{2\pi}\cdot\phi\left(\frac{v}{\sqrt{2\pi}}\right)+v\cdot\Phi\left(\frac{v}{\sqrt{2\pi}}\right).$$

We can think of this as a form of standardization. If the expected exposure—computed with the starting point and collateral in mind—exceeds the add-on (which is determined assuming a zero replacement cost), then the ratio in Eq. 10.54 will exceed one. This means the expected evolution of the netting set value is even more positive than the aggregate add-on would suggest. This is not the situation the regulators are interested in. The multiplier is intended to cover the opposite case where $\mathbb{E}_t\left(E_{t+T}(i)\right) < \mathbb{A}_i^{\Sigma}$ (or equivalently, where Eq. 10.54 is less than one). In this case, the aggregate add-on assumptions are too conservative and the multiplier's job is to mitigate this somewhat.

With this in mind, the *theoretical* form of the multiplier is expressed as,

$$\underbrace{\tilde{f}_i(v)}_{\tilde{\mathcal{M}}_i} = \min\left(1, \frac{\mathbb{E}_t\left(E_{t+T}(i)\right)}{\mathbb{A}_i^{\Sigma}}\right), \qquad (10.55)$$
$$\tilde{f}_i(v) = \min\left(1, \underbrace{\sqrt{2\pi}\cdot\phi\left(\frac{v}{\sqrt{2\pi}}\right)+v\cdot\Phi\left(\frac{v}{\sqrt{2\pi}}\right)}_{\text{Eq. 10.54}}\right),$$

where we can think of the multiplier as a function of v, denoted $\tilde{f}_i$. The role of the $\min(\cdot)$ operator is to ensure that the multiplier only considers those cases where the ratio is in the unit interval. A bit of reflection reveals that this occurs only when the numerator of v—the initial replacement cost $V_t(i) - \tilde{C}_{t+T}(i)$—is less than zero. This, in turn, only happens when the initial market value is out-of-the-money or overcollaterized.

At this point, the construction of the multiplier starts to get a bit weird. BIS [6] claims that:

> MTM values of real nettings sets are likely to exhibit heavier tail behaviour than the one of the normal distribution.

This is likely to be empirically true, but no other justification or motivation is provided. The proposed solution—which comes a bit out of thin air—is to select the following alternative multiplier function:

$$\underbrace{\breve{f}_i(v)}_{\breve{\mathcal{M}}_i} = \min\left(1, e^{\gamma v}\right), \tag{10.56}$$

where $\gamma \in \mathbb{R}_+$ is a parameter that requires specification. BIS [6] further indicates that γ is determined by equating the derivatives of the theoretical and proposed multiplier ratios evaluated at zero. This is readily verified

$$\begin{aligned}
\frac{d}{dv}\underbrace{\left(e^{\gamma v}\right)}_{\substack{\text{Equation}\\ 10.56}}\Bigg|_{v=0} &= \frac{d}{dv}\underbrace{\left(\sqrt{2\pi}\cdot\phi\left(\frac{v}{\sqrt{2\pi}}\right) + v\cdot\Phi\left(\frac{v}{\sqrt{2\pi}}\right)\right)}_{\text{Eq. }10.52}\Bigg|_{v=0}, \\
\left(\gamma e^{\gamma v}\right)\Big|_{v=0} &= \left(\frac{\cancel{\sqrt{2\pi}}}{\cancel{\sqrt{2\pi}}}\cdot\phi'\left(\frac{v}{\sqrt{2\pi}}\right) + \Phi\left(\frac{v}{\sqrt{2\pi}}\right) + \frac{v}{\sqrt{2\pi}}\cdot\Phi'\left(\frac{v}{\sqrt{2\pi}}\right)\right)\Bigg|_{v=0}, \\
\gamma &= \phi'(0) + \Phi(0), \\
\gamma &= \underbrace{\left(\frac{-x}{\sqrt{2\pi}}e^{-\frac{x^2}{2}}\right)\Bigg|_{x=0}}_{=0} + \frac{1}{2}, \\
\gamma &= \frac{1}{2}.
\end{aligned} \tag{10.57}$$

This leads to updating the proposed multiplier function to

$$\breve{\mathcal{M}}_i \equiv \breve{f}_i(v) = \min\left(1, e^{\frac{v}{2}}\right). \tag{10.58}$$

This is *not*, confusingly, the end of the story. BIS [6], making reference to the fact that Eq. 10.58 "would still approach zero with infinite collateralization", impose a floored version with the following *final* form:

$$\underbrace{f_i(v)}_{\mathcal{M}_i} = \min\left(1, \vartheta + (1-\vartheta)\cdot e^{\frac{v}{2\cdot(1-\vartheta)}}\right), \tag{10.59}$$

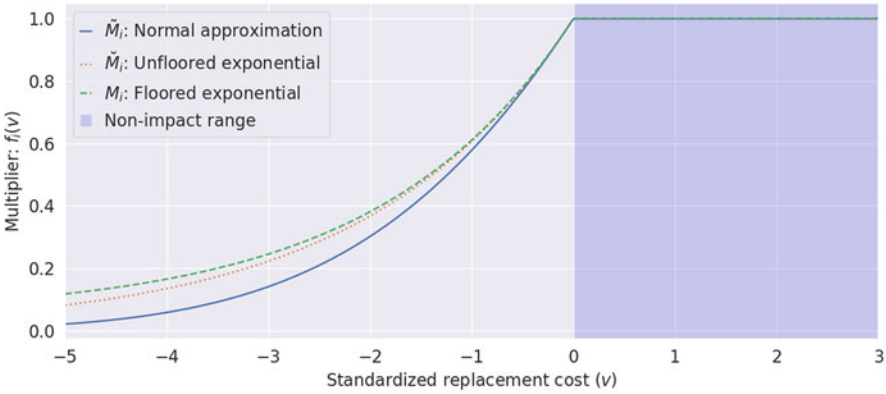

Fig. 10.5 *The add-on multiplier*: This graphic illustrates three alternative forms of the add-on multiplier: the (theoretical) normal approximation from Eq. 10.55, the raw exponential form in Eq. 10.58, and the official floored exponential version from Eq. 10.59. The only real difference among these three choices is their speed of decrease in $-v$.

where, currently, $\vartheta = 0.05$. The idea of Eq. 10.59 is that as v gets very large and negative, the multiplier tends to ϑ rather than zero. This appears to be a mathematical adjustment, rather than predicated on any theory.

Figure 10.5 illustrates *three* alternative forms of the add-on multiplier over the interval $[-5, 3]$. It includes the *theoretical* normal approximation from Eq. 10.55, the heuristic exponential form in Eq. 10.58, and the official floored exponential version from Eq. 10.59. The only practical difference among these three choices is their speed of decrease in $-v$. It is also quite clear, from Fig. 10.5, that none of the multipliers plays an role for positive realizations of Eq. 10.53.

As a final point, the only distinction between the multiplier in margined and non-margined netting sets relates to the treatment of the collateral in the definition of v. That is,

$$v = \begin{cases} \dfrac{V_t(i) - \tilde{C}_{t+1}(i)}{\mathbb{A}_i^{\Sigma}} & : \text{No margin} \\ \dfrac{V_t(i) - \tilde{C}_{t+\tau_{M_i}}(i)}{\mathbb{A}_i^{\Sigma}} & : \text{Margin} \end{cases}, \tag{10.60}$$

where the haircut period applies to one-year and the margin period of risk for non-margined, and margined netting sets, respectively. Although it is a bit difficult to see after the multiple interventions, both computations also share the same expected exposure foundation.

Colour and Commentary 122 (THE SA-CCR ADD-ON MULTIPLIER)*: The role of the multiplier is simple: it attempts to reduce the aggregate netting-set add-on for excess collateral or large negative market values. Its construction is, conversely,* not *at all simple. One could even claim that its derivation follows a relatively weird, possibly over-engineered, or at least convoluted, path. It starts out as the ratio of expected exposure to the aggregate add-on. The expected-exposure numerator is computed by incorporating a non-zero starting value and collateral under the typical assumption of Brownian-motion dynamics. So far, so good. This yields a fairly intuitive expression. Then the basic form is replaced with an exponential function, which is vaguely calibrated to the original result. This can be accepted with a slight "suspension of disbelief." Finally, a floor is imposed for "infinite overcollateralization." The ultimate result is practical, workable, and appears to meet the key regulatory objectives. The weirdness, therefore, has no lasting consequences. One simply needs to invest some additional effort in understanding the origins and justification of the add-on multiplier. Without this background, examination of the final form might lead to a bit of head scratching.*

10.5.5 Bringing It All Together

The preceding sections have been, quite frankly, jumping around somewhat. The idea has been to derive and motivate the various elements of the SA-CCR exposure methodology, which forms the basis of our derivative-exposure model. It is nonetheless admittedly hard to follow the thread connecting these diverse formulae. Figure 10.6 attempts to (at least partially) rectify this situation by offering a high-level overview of the main SA-CCR points. It begins from the derivative exposure (i.e., EAD) equation and then splits into *two* streams: the replacement cost and the add-on. Attempting to incorporate the margin and non-margin perspectives, it also chronicles the various levels of the add-on hierarchy and a number of miscellaneous definitions. It can be viewed as something of a cheat-sheet for anyone looking to implement (or understand) this important calculation.

The description in Fig. 10.6 does, however, remain somewhat abstract. It does not, for example, delve into the important instrument level details of the directional delta, adjusted notional and supervisory factors. To actually use the SA-CCR approach, it is necessary to grapple with these concepts. For this reason, these components, determined at the individual asset-class level, are the next topic of consideration.

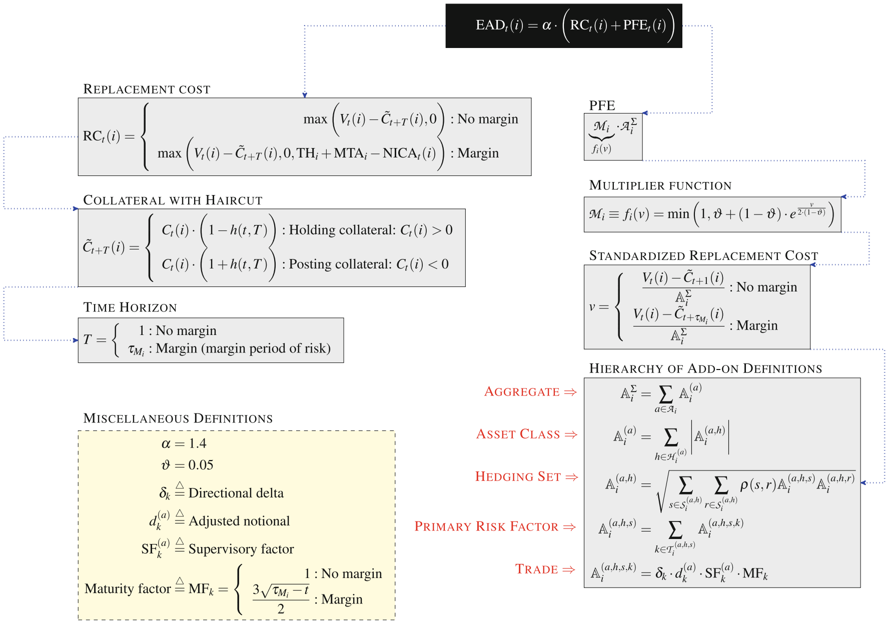

Fig. 10.6 *SA-CCR cheat-sheet*: This schematic outlines, at a single glance, the principal formulae involved in the SA-CCR exposure calculation. It also attempts to highlight the most important relationships between the key elements.

10.6 The Asset-Class Perspective

The SA-CCR methodology decomposes the universe of derivative contracts into *five* distinct asset classes: interest rates, currency, credit, equity, and commodities. This is a challenging task. Much like peeling an onion, one could easily construct many other additional asset classes. There are also many sub-layers within a given asset class. The notions of primary risk factors and hedging sets are, in fact, an attempt to capture some of the complexity of these sub-layers.

To actually implement the SA-CCR methodology, Figs. 10.4 and 10.6, while hopefully useful, are not enough. Additional colour is required—regarding the directional deltas, adjusted notionals, and supervisory factors—to evaluate Eq. 10.46. Detailed direction, and helpful examples, are found in BIS [5, 7]. The following sections will walk through the key details in the area of interest-rate and currency derivatives; these are the two areas of principal interest to the NIB. They also coincide with the currently most popular and widely used OTC derivatives found in financial markets: interest-rate and cross-currency swaps.

10.6.1 Interest Rates

For most lending institutions using interest-rate swaps to match their assets and liabilities—as discussed in detail in Chap. 6—this will likely be the most important category of asset class. This makes it the most sensible entry point. There are basically *three* additional aspects—in addition to the overall formulae—required for each instrument within each asset class: its sign, initial magnitude, and risk scaling. We'll begin with the sign, which is captured by the general definition of the *directional* delta and is given as

$$\delta_k = \begin{cases} 1 : \text{Long} \\ -1 : \text{Short} \end{cases}, \tag{10.61}$$

in the primary risk factor. This definition is sufficient for linear derivative instruments, but more complexity is required for non-linear contracts.[35]

Equation 10.61 remains fairly abstract. Let's consider the most popular case: interest-rate swaps. The floating swap rate is, for the purposes of these instruments, viewed as the primary risk factor. This implies that a receiver swap—where one receives the fixed rate and pays floating—is a long position leading to $\delta_k = 1$. By the reverse logic, a payer swap is short in the primary risk factor leading to $\delta_k = -1$. Proper classification of one's swaps can thus help tremendously in determination of the directional delta. Basis swaps, where both legs represent alternative floating rates, are less obvious. According to BIS [7], each floating pair needs to be treated as a separate hedging set. In our implementation, we assume that the shortest reset frequency is the primary-risk factor. In those (very rare) cases where the frequencies of the floating components are equal, then the order of the indices determines the sign.[36]

The *adjusted* notional amount of an interest-rate derivative with a lifetime over the interval $[S_k, E_k]$ is given as,

$$d_k^{(\text{IR})} = N_k \cdot \text{SD}_k, \tag{10.62}$$

where N_k denotes the *average* notional position over the instrument's lifetime and SD_k represents the so-called supervisory duration. S_k and E_k are, as one might expect, the start and end dates of the derivative contract. Using t, as usual, to represent the current point in time, the remaining term to maturity (or tenor) is simply $E_k - t$.

[35] If the position possesses optionality, then the ± 1 is replaced with the appropriately signed option deltas. The model that one must use to determine these quantities—and the key input parameters, such as implied volatility—is provided in the regulatory guidance. See BIS [5, 6] for more details.

[36] In this case, the assignment of a sign is arbitrary. It just needs to be consistent across all ordered pairs.

As usual for the SA-CCR model parameters, the supervisory duration is determined via regulator-provided formulae. It also applies to both interest-rate and credit derivatives. According to BIS [6] it stems from

$$
\begin{aligned}
\text{SD}_k &= \int_{\max(S_k,t)}^{E_k} e^{-r(\tau-t)} d\tau, \qquad (10.63)\\
&= \left[-\frac{e^{-r(\tau-t)}}{r} \right]_{\max(S_k,t)}^{E_k},\\
&= \frac{e^{-r\cdot(\max(S_k,t)-t)} - e^{-r(E_k-t)}}{r},
\end{aligned}
$$

where $\max(S_k, t)$ is intended to handle ongoing and forward-start derivative contracts. In the current implementation of SA-CCR, the value of r has been arbitrarily—and rather non-conservatively given current interest-rate levels—set to 5%.

Given that we typically think of the modified duration of a fixed-income instrument as a normalized partial derivative of its value function with respect to its yield, it is somewhat surprising to see the integral form in Eq. 10.63.[37] It is not, in fact, a classical estimate of duration, but rather the cumulative discount factor over the interval, $\left[\max(S_k, t), E_k \right]$.

Colour and Commentary 123 (SUPERVISORY DURATION)*: The term* supervisory duration*, in the context of SA-CCR, is not easily interpreted. As a (normalized) sensitivity to an interest-rate movement, duration is typically defined as a derivative taken with respect to some dimension of the yield curve. This could be a single rate, the instrument's yield, or a spread. Equation 10.63 illustrates the supervisory duration—rather uniquely in the author's experience—as the sum (or integral) of a continuously compounded discount factor—at a fixed rate of 5%—over the trade's lifetime. The consequence is a rather high "duration" estimate; for* $E_k - S_k = 10$*, for example, the result is approximately 7.9. Interestingly, as* $r \rightarrow 0$*, the supervisory duration estimate tends towards* $E_k - S_k$*. This leads to the conclusion that supervisory duration is, loosely speaking, constructed under the (rather conservative) treatment of a swap as something similar to the cumulative discount factor of a bond with tenor* $E_k - S_k$ *and a yield of 5%. Once again, as we've already seen on numerous occasions, there is a certain logic to this approach, but the rationale is not entirely clear from the outset.*

[37] See Bolder [8, Chapter 2] for rather more on the foundations of modified duration and related concepts.

An interest-rate hedging set encapsulates all trades in a given currency. As the tenor of fixed-income securities is quite important—short-term instruments can act rather differently, for example, than their long-term equivalents—this dimension is used to construct primary risk factors. In particular, the notion of a maturity bucket is introduced by BIS [5] and defined as,

$$\text{Maturity Bucket} \equiv \text{MB} = \begin{cases} \text{Short } (S) : E_k - t \in (0, 1] \\ \text{Medium } (M) : E_k - t \in (1, 5] \\ \text{Long } (L) : E_k - t \in (5, \infty) \end{cases}. \tag{10.64}$$

This partition of yield-curve space is not particularly granular, but it dramatically eases the implementation.

Although it is a bit tedious and repetitive, we will combine all of these elements together and explicitly write out the aggregate relations from the trade to the hedging-set level. This begins with the following trade description

$$\mathbb{A}_i^{(\text{IR},\text{CCY},\text{MB},k)} = \pm N_k \underbrace{\overbrace{\left(\frac{e^{-r\cdot(\max(S_k,t)-t)} - e^{-r(E_k-t)}}{r}\right)}^{\text{Eq. 10.62}}}_{\text{Eq. 10.63}} \cdot \text{MF}_k, \tag{10.65}$$

where the sign is determined by Eq. 10.61. The add-on by primary factor set, or rather maturity bucket in this case, is now

$$\mathbb{A}_i^{(\text{IR},\text{CCY},\text{MB})} = \sum_{k\in \text{MB}} \underbrace{\pm N_k \left(\frac{e^{-r\cdot(\max(S_k,t)-t)} - e^{-r(E_k-t)}}{r}\right)\cdot \text{MF}_k}_{\text{Eq. 10.65}}. \tag{10.66}$$

Each interest-rate hedging set is aggregated following the logic in Eq. 10.32 as

$$\mathbb{A}_i^{(\text{IR},\text{CCY})} = \underbrace{0.05\%}_{\text{SF}_k} \sqrt{\sum_{s\in\{S,M,L\}} \sum_{r\in\{S,M,L\}} \rho(s,r) \underbrace{\mathbb{A}_i^{(\text{IR},\text{CCY},s)} \mathbb{A}_i^{(\text{IR},\text{CCY},r)}}_{\text{Eq. 10.66}}}, \tag{10.67}$$

where CCY denotes a given currency defined hedging set and the supervisory factor, SF_k is given as 0.05% in BIS [5, Table 2].[38] Practically, the correlation coefficients are provided by the regulators and set to $\rho(S, M) = \rho(M, L) = 0.7$ and $\rho(S, L) = 0.3$. Although the absolute levels are not easily motivated across all currencies, the choice makes sense given that adjacent maturity categories tend to be significantly

[38] The previously discussed basis swaps, where each pair of floating indices are treated as separate hedging sets, receive a supervisory factor of 0.025%. Cutting the base interest-rate swap amount in half is explicitly indicated in BIS [7].

more highly correlated than more distant nodes along the yield curve. Finally, the sum over the currency hedging sets requires no application of the absolute-value operator, because the hedging-set aggregation in Eq. 10.67 always yields positive values.[39]

Computation of the interest-rate add-on, therefore, requires only a limited amount of information for each instrument: the average notional amount, position currency, start and end dates, and whether it is a short or long position. In line with typical regulatory calculations, this imposes a limited operational and data burden on the calculating agent.

10.6.2 Currencies

The currency asset class is a bit larger than one might initially imagine. This is because cross currency swaps are allocated to this category. This might seem somewhat odd, but given the significantly higher volatility of most currencies relative to interest rates, this seems like a judicious choice.[40] The ability of a proper forward-looking simulation engine to simultaneously manage all underlying risk factors associated with each position—by comparison—is definitely an advantage. This is one of the costs of parsimony.

In contrast to interest rates, the currency asset class is relatively straightforward. The adjusted notional amount is expressed as,

$$d_k^{(\text{CCY})} = N_k^{\text{LC}}, \tag{10.68}$$

where LC denotes the local-currency value of the foreign-currency claim. If both legs of the trade are in foreign currency, then BIS [5] recommends computing both in local currency terms and using the larger of the two estimates.[41] The directional delta has the same basic form as found in Eq. 10.61, while the supervisory duration is implicitly $\text{SD}_k \equiv 1$. The corollary of this choice is that the trade-level add-on is independent of the instrument's maturity.

In the currency asset class, a hedging set is defined as a currency pair; we will denote this by CP. The idea, it seems, is that each member of the hedging set's primary risk factor is the spot exchange rate associated with its currency pair. The key practical task is to get the correct directionality of the currency pairs. A EUR-

[39] If it did not, with the presence of the square-root operator, we might find ourselves with a difficult-to-explain, complex-valued interest-rate add-on.

[40] The regulators could, of course, have assigned them to both asset classes. This would, however, lead to more complexity and immediately break the general rule of assigning each instrument to a single asset class.

[41] Our current implementation, incidentally, always conservatively uses the larger of the two legs in local-currency terms.

AUD pair is, of course, equivalent to an AUD-EUR position. What is important is to identify the set of unique currency pairs and determine long and short positions with respect to that combination.[42]

There is no primary risk factor sub-category in the aggregation, since it is already defined at the trade level by the choice of currency pair. The trade-level formula is given as,

$$\mathbb{A}_i^{(\text{CCY},\text{CP},k)} = \pm N_k^{\text{LC}} \cdot \underbrace{4.0\%}_{\text{SF}_k} \cdot \text{MF}_k, \tag{10.69}$$

where the supervisory factor is significantly (and understandably) much larger than in the interest-rate setting. Summing over all of the trades within a given currency pair yields,

$$\mathbb{A}_i^{(\text{CCY},\text{CP})} = \underbrace{\sum_{k\in\text{CP}} \pm N_k^{\text{LC}} \cdot \underbrace{4.0\%}_{\text{SF}_k} \cdot \text{MF}_k}_{\text{Eq. 10.69}}. \tag{10.70}$$

The currency asset class add-on then reduces to the following final aggregation over all of the individual currency pairs,

$$\mathbb{A}_i^{(\text{CCY})} = \sum_{c\in\text{CCY}} \underbrace{\left|\mathbb{A}_i^{(\text{CCY},c)}\right|}_{\text{Eq. 10.70}}. \tag{10.71}$$

The uniqueness of the currency asset class stems from the overlap between the trade and primary-risk factor levels. Equation 10.71 also reveals the reason for the presence of the absolute-value operator in Eqs. 10.37 and 10.48. If it were not present, then positive and negative exposures from different currencies would (rather indefensibly) be able to offset one another.

10.7 A Pair of Practical Applications

As usual, we would like to conclude with some concrete discussion to help solidify these ideas. Unfortunately, there are fundamental reasons that make it challenging in this case. As should be clear from the previous sections, the individual SA-CCR calculations are rather data-heavy and multi-layered; there is a real danger of spend-

[42] Ultimately, the actual sign assigned to long and short trades does not matter since absolute values will be imposed in a latter step. The key is consistency of sign treatment among short and long positions with each unique currency pair.

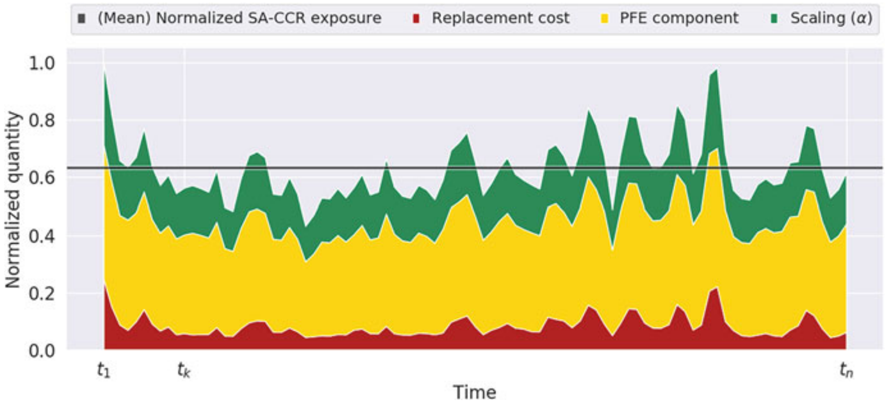

Fig. 10.7 *Normalized portfolio-level derivative exposure*: This graphic traces our normalized portfolio-level SA-CCR derivative exposure estimate over a roughly six-month period in 2021. It should help to visualize the relative roles of replacement cost, the model-based PFE component, and the additional regulatory conservatism.

ing more time introducing the data than actually discussing the results. It is also impossible to directly use NIB-related inputs, because these derivative obligations represent rather sensitive internal data. Even if we could, the dimensionality of the data would be a bit overwhelming.

To counteract these shortcomings, we will look at a pair of, rather stylized, examples. The first examines a normalized view of our overall portfolio-level derivative exposure. These results generically describe the key relationships between the various elements of the SA-CCR computation. Our second example investigates a practical application—important for all financial institutions—of derivative exposure estimates: the so-called leverage ratio.

10.7.1 Normalized Derivative Exposures

Derivative exposures are computed at the counterparty level; this is an unavoidable consequence of counterparty credit risk. When the exposures of all individual derivative counterparts are summed up, however, we arrive at a surprisingly interesting perspective on portfolio-level exposure.

Figure 10.7 provides, for an arbitrarily selected six-month time interval during 2021, the normalized view of our portfolio-level derivative exposure.[43] This vantage point allows us to establish a few basic facts about the SA-CCR computation.

[43] The normalization involves imposing unit volatility and scaling by the maximum observation. This allows us to abstract from the actual values and focus on trends and broad-based insights.

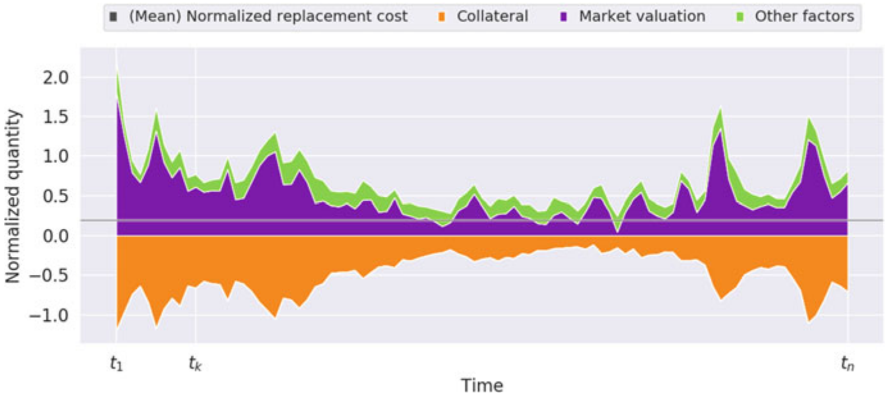

Fig. 10.8 *Normalized portfolio-level replacement cost*: The graphic above traces the our normalized portfolio-level replacement cost over a roughly six-month period in 2021. This should be viewed as drilling deeper into the replacement-cost component presented in Fig. 10.7.

The portfolio-level values—which are a function of *three* main components—are relatively stable across time. Moreover, the replacement cost represents a fairly modest part of the overall picture. This is essentially the point of the SA-CCR; the add-on-related PFE component is an adjustment for risk placed on top of the collateralized current exposure. This PFE element, which is determined by the structure of the underlying derivative instruments, is naturally quite persistent.[44] The final component, stemming from the constant in Eq. 10.26, is a bit of additional conservatism provided by our regulators.

Figure 10.8 drills deeper into the replacement-cost component. As we learned in the previous discussion, the key replacement-cost drivers are market value and collateral. Again using a normalized, portfolio-level perspective, we see clearly that collateral and derivative valuations are essentially mirror images of one another.[45] This is, in fact, the entire point of the collateral process. As market values move in the financial institution's favour, collateral is received to offset the increasing counterparty credit risk.[46] The component called *Other factors* in Fig. 10.8 relates to the net initial collateral, collateral threshold, and minimum transfer amounts introduced in Eq. 10.27. These elements—which are simply the difference between the SA-CCR replacement cost and the residual market value after collateral—are uniformly positive. They represent an additional element of caution in the SA-CCR construction.

[44] A big part of the reason is the fixed regulatory parameters, which are not recalibrated on a daily basis as is the case with most forward-looking exposure engines.

[45] The offset is not perfect, but the correlation coefficient between these two series over this period is approximately −0.9.

[46] Over this period, NIB was a net receiver of collateral at the portfolio level. An institution's collateral position tends to ebb and flow over time with market movements and portfolio composition.

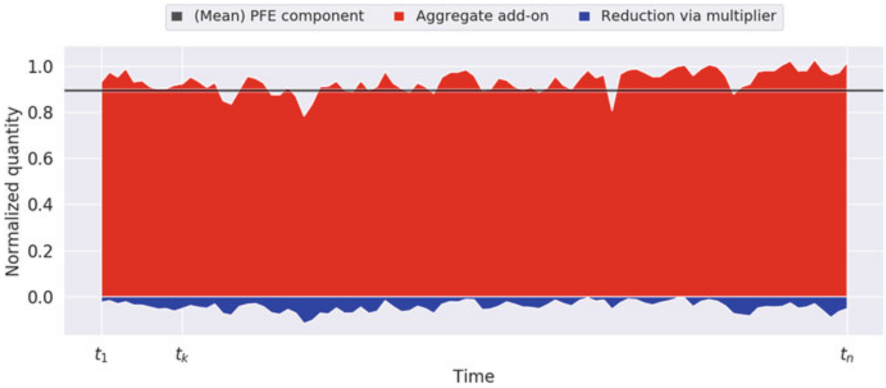

Fig. 10.9 *Normalized portfolio-level SA-CCR PFE*: This graphic traces our normalized portfolio-level SA-CCR PFE over a roughly six-month period in 2021. This should, once again, be visualized as zooming in on the PFE component presented in Fig. 10.7.

Figure 10.9 concludes our high-level examination by zooming in on the PFE component. This normalized, portfolio-level risk-adjustment is decomposed into two sub-components: the aggregate add-on and the multiplier reduction. Recalling from Eq. 10.59, the multiplier acts to reduce the overall PFE component in cases of negative market values or overcollateralization. Over this period, the multiplier plays a comparatively limited role in the determination of the PFE contribution to overall portfolio derivative exposure. This is both typical and consistent with the regulatory guidance; the multiplier is an adjustment to handle a special case.

Colour and Commentary 124 (INTERPRETING DERIVATIVE EXPOSURES)*: Much can be learned from a high-level analysis of one's own derivative portfolio. Using a normalized daily portfolio-level view, we consider values over a six-month period during 2021. This bird's eye perspective yields a number of interesting points. The total portfolio-level derivative exposure is generally quite stable over time; this is consistent with our desired through-the-cycle exposure estimate.*[a] *Replacement cost is, in general, only a relatively small fraction of the overall SA-CCR estimate; the risk-related PFE component represents the dominant contribution. Collateral and derivative market values, as one would expect, largely offset one another. This drives the stability and (relatively small) magnitude observed in the replacement cost. The add-on multiplier, while conceptually important, appears to play a fairly limited role. Finally, the presence of other (collateral-related) factors in the replacement cost and the scaling constant in the overall computation lend additional conservatism to the final result.*

[a] This is a direct consequence of fixed regulatory risk parameters.

10.7.2 Defining and Measuring Leverage

Leverage, at its heart, seeks to measure the amount of debt employed to finance a firm's assets. All else equal, less leverage is preferred to more, because the greater the debt burden, the more difficult it is for a firm to meet its claims. Excessive amounts of debt can become particularly problematic in the event of financial stress. Debt in a firm's capital structure is, of course, not without its benefits. In particular, tax treatment typically makes debt financing relatively cost effective.[47] As is so often the case, the firm must find a balance between the benefits and risk associated with the use of leverage.

Given the accounting identity that firm assets are the sum of debt and equity, there are a variety of (broadly equivalent) ways to estimate leverage. Sometimes it is represented as a debt-to-equity ratio, while in other cases one examines the ratio of debt to total assets. In the financial-institution setting, the leverage ratio is computed as

$$\text{Leverage Ratio} = \frac{\text{Available Financial Resources (Equity)}}{\text{Total Asset Exposure}}. \tag{10.72}$$

This particular definition is consistent with European Union banking regulation.[48] The benefit of the unit-less ratio format is that it is easy to interpret, enforce, and compare with other financial institutions.

It is natural to compare Eq. 10.72 to the concept of capital headroom discussed on numerous occasions in previous chapters. Both involve, after all, the firm's capital supply. When inspecting Eq. 10.72, we find that it is silent on the riskiness of the underlying assets. It depends solely on their overall magnitude or, as it is commonly referred to, volume. This is simultaneously a strength and a weakness. On the positive side, it is relatively easily described and computed. There is little scope for analytical freedom in determination of Eq. 10.72. More negatively, asset-risk composition matters. Two financial institutions might have the save leverage ratio, but vastly different asset-risk profiles. The leverage ratio provides no information on this dimension. A solid compromise is to avoid arguing about the relative merits of economic capital and leverage; instead, one can easily use both and thereby gain a broader complementary perspective.

The leverage ratio, while vastly simpler to compute than economic capital, still has a few tricky parts. While the numerator of the leverage-ratio definition in Eq. 10.72 is fairly clearly related to some variation on internal capital supply, the denominator remains somewhat vague. The regulation is, however, quite specific. In particular, the guidance stipulates that it must include:

---

[47] This fact has been appreciated for decades and is a key driver of capital-structure decisions; see, for example, Modigliani and Miller [16].

[48] Further background, from the regulatory side, is found in BIS [4].

1. an accounting representation of firm assets *excluding* derivative contracts and associated collateral;
2. an allocation for non-derivative-related off-balance sheet items; and
3. an allowance for derivative contracts incorporating **an add-on for counterparty credit risk**.

It is this final point that, in the context of this chapter, catches our interest. Pulling these directions together, we arrive at the following revised, and more precise, description of the leverage ratio:

$$\text{Leverage Ratio} = \frac{\text{Tier I Capital Position}}{\begin{array}{l}\text{Accounting Assets ex Derivatives}\\ \quad + \text{Off-Balance Sheet Items} + \text{Derivatives}\end{array}}. \tag{10.73}$$

Practically, most of the ingredients are relatively easily procured. Accounting asset values are sourced from internal systems. Off balance-sheet items relate exclusively to agreed, but not-yet-disbursed loans. Committed agreed, not disbursed loans are transformed with a 50% credit-conversion factor into loan equivalents; only 10% of uncommitted loans are included.[49] The wild-card in this calculation, however, relates to derivative exposure.

The regulatory term "allowance for derivative contracts" is a bit vague. Assuming a financial institution can convince its regulator of the reasonableness of its computation, then presumably it has a fairly broad range of alternatives. Common practice, and in this respect NIB is no exception, is to simply employ the SA-CCR methodology for this purpose. This completes Eq. 10.73 and readily permits its computation.[50]

Colour and Commentary 125 (MEASURING LEVERAGE)*: Leverage seeks to describe the amount of debt employed to finance one's assets.*[a] *Unlike the risk-based economic-capital measure, which sits at the heart of this book, it focuses exclusively on volume. It can nonetheless provide a very useful complement to risk-based metrics. While leverage can be measured in many ways, it is common to construct ratios involving a firm's assets and equity (i.e., capital supply). The majority of the inputs into any leverage assessment are not particularly controversial and easily obtained from a firm's financial*

(continued)

[49] The credit-conversion factor is a regulatory invention, which we have already encountered in previous chapters. It is a value between zero and one that, when multiplied by the instrument's notional value, transforms an off-balance-sheet item into a balance-sheet equivalent.

[50] Somewhat ironically, given that it typically only represents a small fraction of a financial institution's total assets, much of the effort associated with the leverage-ratio computation relates directly to the derivative component.

Colour and Commentary 125 (continued)
statements. Derivatives, which may be either an asset or a liability depending on their mood, are the exception. Regulatory guidance on the measurement of leverage counsel excluding all derivative-related balance-sheet inputs—both collateral and valuations—and replacing them with an allowance for derivative contracts. *Those wishing to measure leverage thus find themselves in the business of estimating derivative exposure. The SA-CCR computation jumps into this breach and offers a tailor-made solution to this problem.*

[a] Assets, in this context, go somewhat beyond the accounting view to consider off-balance sheet items.

10.8 Wrapping Up

There are many ways, of varying degrees of complexity, to compute forward-looking measures of derivative exposure. The most conceptually correct methods involve the construction of a high-dimensional instrument-level, simulation-based valuation engine that also incorporates time-varying collateral behaviour. For quantitative analysts involved in derivative pricing, complex simulation engines are a non-optional prerequisite. For economic-capital purposes, however, such complexity is neither necessary nor desired. The more stable, through-the-cycle, $\mathbb{P}$-measure-based regulatory capital calculations make for a more natural choice. They provide defensible risk-based derivative exposure estimates and, as a side benefit, represent a helpful input into the measurement of one's leverage position.

References

1. P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, July 2001.
2. BIS. Treatment of potential exposure for off-balance-sheet items. Technical report, Bank for International Settlements, April 1995.
3. BIS. The application of Basel II to trading activities and the treatment of double default effects. Technical report, Bank for International Settlements, July 2005.
4. BIS. Basel III leverage ratio framework and disclosure requirements. Technical report, Bank for International Settlements, January 2014a.
5. BIS. The standardised approach for measuring counterparty credit exposures. Technical report, Bank for International Settlements, March (revised April 2014) 2014b.
6. BIS. Foundations of the standardised approach for measuring counterparty credit risk exposures. Technical report, Bank for International Settlements, August (revised June 2017) 2014c.

7. BIS. Calculation of RWA for credit risk: Standarized approach to counterparty credit risk. Technical report, Bank for International Settlements, December 2019.
8. D. J. Bolder. *Fixed Income Portfolio Analytics: A Practical Guide to Implementing, Monitoring and Understanding Fixed-Income Portfolios*. Springer, Heidelberg, Germany, 2015.
9. D. J. Bolder. *Credit-Risk Modelling: Theoretical Foundations, Diagnostic Tools, Practical Examples, and Numerical Recipes in Python*. Springer-Nature, Heidelberg, Germany, 2018.
10. D. Brigo. Counterparty risk faq; credit var, pfe, cva, dva, closeout, netting, collateral, re-hypothecation, wwr, basel, funding, CCDS and margin lending. In D. Brigo, editor, *Counterparty Credit Risk, Collateral and Funding*. Wiley Finance, 2012.
11. D. Duffie and E. Canabarro. *Measuring and marking counterparty risk*, chapter 9, pages 122–134. Euromoney Books, 2003.
12. M. B. Gordy. A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation*, 12:199–232, 2003.
13. J. Gregory. *Counterparty Credit Risk: The New Challenge for Global Financial Markets*. John Wiley & Sons Ltd., Chicester, UK, 2010.
14. A. J. Heunis. Notes on stochastic calculus. University of Waterloo, Winter 2011.
15. I. Karatzas and S. E. Shreve. *Methods of Mathematical Finance*. Springer-Verlag, New York, NY, 1998.
16. F. Modigliani and M. H. Miller. The cost of capital, corporation finance and the theory of investment. *The American Economic Review*, 48(3):261–297, 1958.
17. J. Morgan/Reuters. *RiskMetrics — Technical Document*. Morgan Guaranty Trust Company, New York, 1996.
18. B. K. Oksendal. *Stochastic Differential Equations*. Springer Verlag, Berlin, fourth edition, 1995.
19. I. Ruiz. *XVA Desks: A New Era for Risk Management*. Palgrave Macmillan, New York, NY, 2015.
20. A. P. Salonen. *Transnational Law of the Over-the-Counter Derivatives Market: A Study of the Interactions of Law and Finance*. PhD thesis, Faculty of Law, University of Helsinki, 2019. Unpublished Thesis.