11 Support Measures

We have selectively presented a few risk measures in the preceding chapters that, in our experience, cover many relevant aspects and tasks in a real-world market risk setup. We propose to mainly use the volatility-rescaled historical VaR$[\Omega]$ for daily risk management. It is especially well-suited to capturing "tomorrow's PnL," as it reacts fast to changes in volatility levels. The concurrent use of the sensitivity-based analytical VaR$(\mathbf{s}^{\Omega})$ serves as a sanity check and provides an additive decomposition to VaR-contributions of the risk factors, which is a handy analysis tool because it appropriately weighs risk factors by both their sensitivity and volatility. Finally, the most helpful measure we take away from the expected shortfall world is the position-wise conditional expected shortfall cES$[\alpha|\Omega]$, which provides a useful complementary breakdown of risk to positions.

The main difficulty in analyzing VaR figures and thus the need for additional support measures arise because the VaR is generally not additive:

$$\text{VaR}[\alpha + \beta] \neq \text{VaR}[\alpha] + \text{VaR}[\beta].$$

More background on this and actual use cases of our measures will be given in Part II. Before that, the current chapter will mention additional helper measures you should be able to reference. Some of them are useful, others less so; either way, certain ones might be mandated by the regulator.

First, let's address an apparent gap in our measures presented earlier. As mentioned in Chap. 8, the analytical VaR approach immediately translates to an analytical ES approach (we only need to tweak the final multiplier of the standard deviation ever so slightly). But can, reversely, the concept of the conditional expected shortfall be translated to the VaR as well? Can we find an additive decomposition of the VaR to positions?

The answer is we sure can, and easily too, but we should nevertheless steer well clear of it. First, how would we do it? Well, in analogy to the cES, we could determine the index of the portfolio PnL vector's VaR scenario; the corresponding

M. Auer, *Hands-On Value-at-Risk and Expected Shortfall*, Management for Professionals, https://doi.org/10.1007/978-3-319-72320-4_11

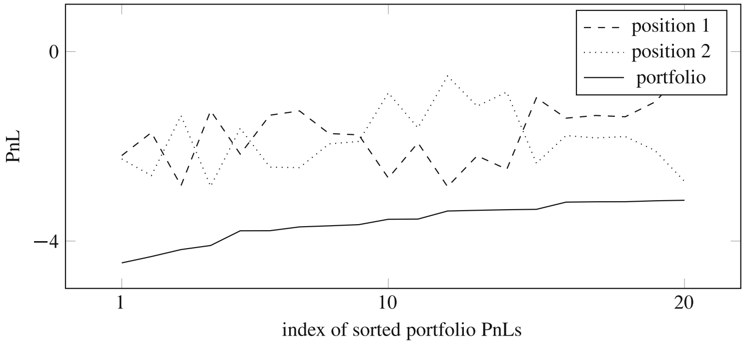

Fig. 11.1 Individual and combined PnLs near the VaR scenario

entries in the positions' PnL vectors could then be considered the individual "cVaR" values for each position (of course summing up to the portfolio PnL in that scenario, i.e., the VaR). We did the exact same for the cES—we just remembered all the original indices of the 25 most negative portfolio PnLs.

Why should we avoid doing this? The short answer is that this measure would be far too unstable, essentially yielding almost random numbers (the averaging involved in the cES calculation, on the other hand, provides for stability). One way to illustrate this unreliability of a cVaR is to artificially create, e.g., in Excel, 1000 pairs of (random normal) returns/PnLs of uncorrelated positions, along with their sums, the hypothetical portfolio PnLs. Figure 11.1 depicts—for a random example—the subset of those 20 return pairs with the largest portfolio losses. The tenth return pair from the left (call it pair *A*) can lay claim to represent our VaR and proclaim two cVaR values.

A first hint at the fickleness of all this is that the return pairs neighboring the VaR one (e.g., pair *B*, the 11th from the left) exhibit quite different cVaR candidates. Now assume the return sets to change slightly (after all, as time progresses, new returns materialize and old ones disappear). Say, the fifth most negative PnL above is dropped because its corresponding return pair vanishes. Suddenly, pair B sits at the tenth location and provides apparently very different cVaR values.[1] Workarounds have been proposed to modify such a cVaR and to make it more stable, but they essentially perform some sort of averaging over several entries around the portfolio's VaR scenario index and therefore basically converge to the cES behavior, with the added baggage of custom heuristics.

Keep in mind that the desire for an additive decomposition of the VaR to positions is comprehensible. It would allow us, for example, to cleanly assign parts of the risk as described by the overall portfolio VaR to positions and sub-portfolios. Along with this, the risk costs arising from the capital requirements could be allocated to units

[1]This issue gets even more pronounced for negatively correlated positions, where more return pairs yield close VaR contenders.

and departments and desks and individuals. A cVaR is too fickle, as we have seen, but even the cES is not ideally suited: while additive and stable, its values can be both positive or negative, offering no obvious "weight" interpretation akin to, e.g., some always non-negative probabilities. Such same-sign additive decompositions are, alas, not available.[2]

Let's now examine some other helper measures.

Individual VaR We are mostly interested in the VaR of a portfolio of positions, i.e., some VaR[Ω]. Of course we could compute each individual position's VaR as well, as in VaR[α], maybe with the aim to detect outlier positions or track down suspicious changes or jumps in the overall VaR at the position level.

In practice, such an individual VaR is not that useful. A position with a conspicuously extreme individual VaR might, in the end, not affect the portfolio VaR by much (it depends on how the position is correlated to the remaining portfolio). As another example, two positions that hedge each other could signal two extreme individual VaR values but actually have, combined, no effect on the portfolio VaR at all. If such hedges were to involve non-linear positions, we could potentially get one extreme and one moderate individual VaR value, muddying the analysis waters further (those positions' combined influence on the VaR is, again, zero).

In some instances, sub-portfolio VaRs can of course be helpful, as they help restrict the search space to position subsets when VaR changes need to be pinned down. Calculating and storing each position's VaR, however, can usually be avoided.

Incremental VaR When adding a new position α' to an existing portfolio $\Omega = \alpha + \beta + \dots$, the portfolio VaR changes. By how much mainly depends on the new position's size and its correlation to the portfolio's PnL behavior. The resulting new portfolio VaR can range from 0 (if the deal mirrors the portfolio exactly) to arbitrarily negative values (if the deal is dominant). Naturally, we'd like to know in advance how the VaR (and thus our costs, i.e., capital requirements) would change if we entered a new position. Computing such an impact is often referred to as performing a *pre-deal inquiry*.

Plainly, we just compute the new VaR and relate it to the current one. The new position's impact is called its *incremental* VaR:

$$\mathrm{iVaR}[\alpha'|\Omega] = \mathrm{VaR}[\Omega + \alpha'] - \mathrm{VaR}[\Omega].$$

[2]The concepts involved here are more deeply connected, as one could consider each position to represent a separate asset. While we tried to delineate those views on risk decomposition, they really represent two sides of the same coin. The terms marginal VaR and component VaR are commonly used in this context. The names used here I chose by sympathy and memorability; marginal VaR often also refers to what we called incremental VaR, while component VaR sounds a bit like poet laureate or Astronomer Royal.

We can do the same with a portfolio's current positions. For a deal α already contained in our portfolio, we can determine the VaR impact of its *removal* from the portfolio[3]:

$$\mathrm{VaR}[\Omega - \alpha] - \mathrm{VaR}[\Omega].$$

This corresponds to the incremental VaR of the deal's hedge, $-\alpha$, which compensates for or cancels the original position's impact:

$$\mathrm{iVaR}[-\alpha|\Omega] = \mathrm{VaR}[\Omega - \alpha] - \mathrm{VaR}[\Omega].$$

This expression should help drive home one particular point. When calculating the incremental VaR, we usually have already calculated the portfolio VaR and thus have at our disposal the PnL vectors of all positions and of the portfolio. Determining the VaR impact is then cheap:

- When adding a new position, we only have to compute its PnL vector $\Delta\mathbf{p}^{\alpha'}$ and *add* it to the known portfolio PnL vector $\Delta\mathbf{p}^{\Omega}$ before the subsequent sort and lookup steps.
- When removing an existing position, we *subtract* the known $\Delta\mathbf{p}^{\alpha}$ from the known $\Delta\mathbf{p}^{\Omega}$.

For existing positions, this measure has similar drawbacks as the individual, position-level VaR.

Partial VaR We usually generate scenarios on all risk factors. We can reduce this scope to subsets of risk-factors, e.g., to only foreign exchange (FX) or to only interest rate (IR) risk factors, which yields *partial* VaRs denoted as $\mathrm{VaR}^{\mathrm{FX}}$ or $\mathrm{VaR}^{\mathrm{IR}}$. This is done by creating new, distinct scenario sets where all but those risk factors we're interested in are kept constant.

Partial VaR figures facilitate locating possible sources of overall VaR changes. A VaR jump in only one of the tracked partial VaRs expeditiously narrows down the set of positions or risk factors we have to examine further. All major risk factor classes should therefore be tracked this way. Theoretically, we could break this down to individual risk factors, but this might become computationally too expensive.

Note that partial VaRs also do not add up to the overall portfolio VaR, the same way individual VaRs don't. Chapter 12 will illustrate this further.

There is a shortcut for actually implementing partial VaRs that avoids creating separate scenario sets with (somewhat redundant) constant rows. The pricing step can simply rely on the original scenarios and, before performing the computation, force the appropriate scenarios to be constant on the fly (see Sect. 19.5).

[3]We must refrain from dubbing this excremental VaR.

Synthetic Marginals Computing the partial VaR of a single risk factor (e.g., $\text{VaR}^{\text{IR-USD-Y10}}$) would allow us to track the model performance with respect to just that one risk factor and to its individual or *marginal* distribution. This would require performing a whole VaR calculation over a million positions 2200 times over, which is often unfeasible. A quicker alternative is to create artificial or *synthetic* test positions that are only sensitive to individual risk factors or small sets thereof.

To keep the set of tracked risk factors small and manageable, we can use the risk factors' VaR-contributions to determine the main and thus most interesting risk drivers and only set up and track synthetic positions for them.

Analyzing the behavior of those synthetic portfolios can then either show that that model performs well with respect to major risk factors or help expose problematic risk factors that might otherwise remain hidden in the joint distribution.

Stressed VaR Our VaR setup relates the recent two-year period of market activity to tomorrow's PnL behavior. Alternatively, one might ask which historic precedent of a previously observed market period would, now and for our current positions, indicate a high degree of risk—after all, such past periods might conceivably occur in similar form again.

To answer this, we need to find the historical period[4] of returns that projects the most extreme VaR for our current positions. So for each past day, we take its corresponding return window, create new scenarios based on it and on today's market scenario $\mathbf{S}_0$, and compute a VaR. One such return window will yield the most extreme or *stressed* VaR.

Finding such a worst-case past period is computationally very expensive, as it entails running a full VaR evaluation for each day in our history of thousands of days. This step is therefore typically only performed once a year in a separate calibration exercise.[5] It should also be explicitly triggered whenever the characteristics of our portfolio composition change drastically (e.g., when a new trading strategy is put in place or when different risk factors start to dominate the portfolio's risk). The hereby settled stress period returns can then be used each day anew to compute the current portfolio's stressed VaR, off of the current market scenario. (Note that for the stressed VaR the volatility rescaling step is omitted, for it would effectively mean selecting the worst-case 20-day return period.)

A main reason for using this measure is simply that you might have to, as regulators increasingly rely on it. As a mathematical instrument, though, it is not very elegant. The stressed VaR has no direct relation to PnLs that are actually observed, and it is thus nigh impossible to plausibilize (except that it should exceed the VaR in all but rare instances). And its warning-signaling power may

[4]The regulator prescribes the window size to be used for this purpose; one-year periods are typically used for the stressed VaR.

[5]The fast analytical approach can provide valuable support for speeding up this procedure. It can, for example, run a first tentative selection, thus limiting the number of full VaR evaluations required.

be overstated: benign market conditions, like happy families, are all alike; every market crisis is probably messed up in its own way.

The individual, incremental, partial, and stressed expected shortfall are computed along the very same lines as their VaR cousins.