14 VaR Noise

Having examined the static properties of the VaR, we now look into its dynamic behavior over time. As new positions are entered or old ones closed, and as the volatilities of the assets involved change, the VaR, recalculated every day, will change as well. Often, such VaR changes and their reasons are of more interest in risk management than the level of the VaR itself.

In order to appraise VaR changes, it is useful to first look into the behavior of our VaR measure in the special case of a constant portfolio and benign markets. It turns out that even in such a stable environment the VaR will fluctuate to some extent. This baseline of natural noise is good to keep in mind when analyzing a particular VaR change, when comparing different VaR models, or when assessing the usefulness of certain optimizations. Effects below that baseline might be inconsequential really.

To sketch this baseline, we create a pseudo-history of random normal returns (say, in Excel) for a hypothetical asset. Since we can disregard units here, these returns can directly be viewed as PnLs.

The history consists of 3 years or 750 days. For the first 2.5 years, we create standard normal returns, i.e., with standard deviation 1. For the remaining half a year, we create returns with twice that standard deviation. We proceed to take a look at the VaR behavior in the third, last year (the first 2 years only provide a full, valid input of "historical" data for that final year of interest). Since we know the underlying distribution, we know what the VaR ought to be on each day, and we can compare it to the VaR estimates of various model flavors.

We first examine the simplest one—computing the VaR from mirrored but non-rescaled returns (see Fig. 14.1). During the first 6 months, we notice how the VaR estimate from the raw returns skips between constant levels of similar magnitude. Such a skip happens whenever the return associated to the VaR scenario for a given day falls out of the historical 500-day window used for the following day.

Then, after 6 months, the VaR estimate starts to change and to converge to the newly established vola level—yet as can be seen, very slowly. It takes time for the

M. Auer, *Hands-On Value-at-Risk and Expected Shortfall*, Management for Professionals, https://doi.org/10.1007/978-3-319-72320-4_14

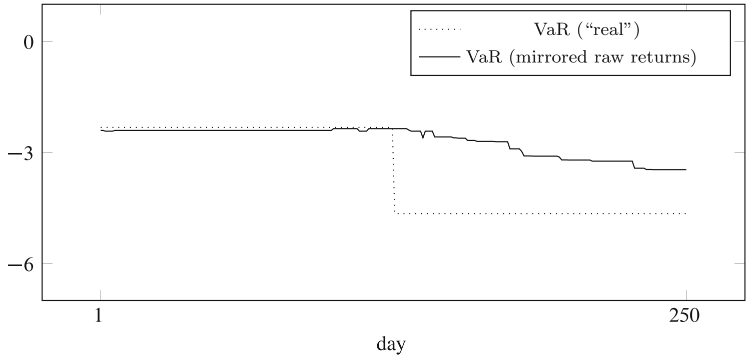

Fig. 14.1 VaR estimate from artificial raw returns; 1 asset

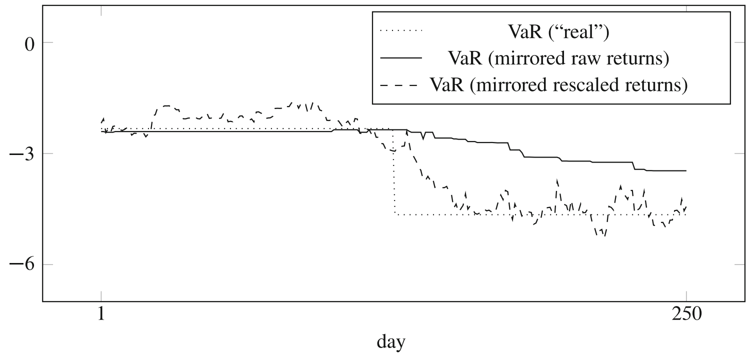

Fig. 14.2 VaR estimate from artificial rescaled returns; 1 asset

new, larger returns to influence sets of 500 returns that still mostly contain old, low-vola ones.

Let's see how our historical approach—based, again, on mirrored but now also rescaled returns—performs under the same circumstances. Figure 14.2 immediately shows that the rescaling has a major impact on the VaR estimate—it is much more volatile and overshadows any level skips (which happen underneath anyway). The reason for this VaR volatility is that each day's local vola estimate is based on only 20 returns. This small sample size hence causes the crucial target vola estimate to fluctuate more and to randomly deviate further from any "real" underlying standard deviation. The more important aspect, however, is that this VaR estimate, by design, almost immediately reacts to the mid-year vola level change and quickly converges to the newly established regime.

Now, the magnitude of our VaR's volatility is still somewhat striking. There are two ways to cope with this. The first is to notice that the deviation behavior is only an "error" because we made it so—in real life, we never know the underlying

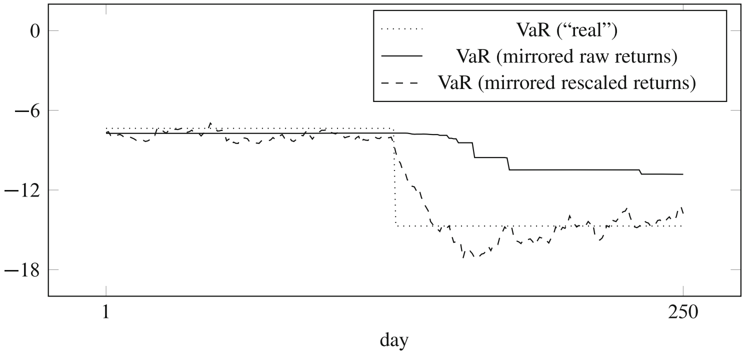

Fig. 14.3 VaR estimate from artificial rescaled returns; 10 assets

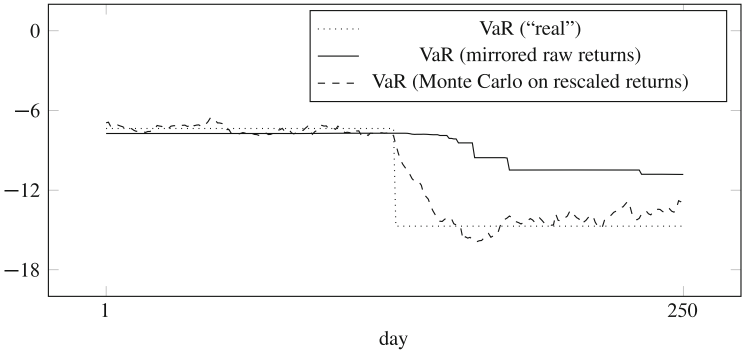

Fig. 14.4 VaR estimate from Monte Carlo on rescaled returns; 10 assets

distribution and might as well tacitly assume our estimate to be perfect (so just remove, in your mind, the dotted line of the real VaR).

Second, the effect becomes attenuated if a portfolio does not depend solely on one single asset. For a portfolio of 10 uncorrelated assets, the behavior seems less suspicious, as can be seen in Fig. 14.3. Some risk factors' target volas are overestimated, some are underestimated; combined, these errors tend to partially offset each other.

We get a very similar picture with the Monte Carlo approach on top of rescaled returns (see Fig. 14.4), again for 10 uncorrelated assets. There is one important issue to keep in mind: Monte Carlo introduces an additional random deviation or Monte Carlo error. This can be made arbitrarily small, e.g., by using a very large number of normals or by appropriately mirroring those normals as well. However, even when using an absurd 10^{10} scenarios and thus basically eliminating any Monte Carlo error, your Monte Carlo VaR estimate will only ever converge to the (dashed) line driven by the target volas (the figure, in fact, depicts this limit) and *not* to the (dotted) real VaR. So you can't Monte Carlo yourself towards the truth—it is just too elusive through our short-term vola spectacles.[1]

The area between the line of a VaR estimate and that of the real VaR is an indication of how systematically or how long an estimate is off. If it takes too long for an estimate to adjust, whole series of VaR breaches or backtesting violations may ensue. Rescaling clearly makes this area smaller, but because nothing is free, we buy this model reactivity by sacrificing some day-to-day stability of our risk measure. The proposed model mimics a hummingbird instead of a sloth.

Finally, we can take a more narrow look at the daily VaR fluctuations driven by a short-window target volatility. For one individual asset and a time series of standard normals, we can compute each day's local/target vola estimate L_i and the relative changes $L_{i+1}/L_i - 1$ over time. It turns out that the standard deviation of these changes is about 5%, par for par the standard deviation of relative VaR changes in any rescaled setup.

Like above, we can do the same exercise for 10 uncorrelated assets (we use the square-rooted sum of the local *variance* estimates for this). Here, the daily VaR changes clock in at a standard deviation of about 1.7%. Now, with 2200 risk factors we are tracking many more than just 10, but we should be aware that often just a small subset of them drives the risk, especially in sub-portfolios. Furthermore, risk factors may at times be highly correlated. This clumps them together and makes them act as if they were fewer in number, with less noise offsetting or relief.

[1]The Monte Carlo error depends mainly on the number of Monte Carlo scenarios used. You could analytically determine how far the Monte Caro estimate is likely to be off the limiting case of infinite scenarios, or you can simply try out sets of different random numbers to get an idea of this error range. You may experience, e.g., the Monte Carlo VaR with 5000 scenarios in a real-world portfolio to randomly deviate from the dashed (not dotted!) line by between $\pm 3\%$ and $\pm 5\%$ in relative terms.