

Here are a few of the glorious blunders I committed that will remain on my head. I first thought it wise to output sorted PnL aggregation vectors (like destroying information is a good thing). I also proposed to add an additional column of zero to the PnL results for some “future use” as special indicator flag—a use that never materialized and now requires constant nudging when we have to explain our file formats. I unthinkingly used the VaR prediction date as label for our daily results, which is just plain stupid—everybody will always refer to the previous, close-of-business date. Each email now requires a qualifier, and each new hire will at least once perform an analysis of the wrong data. (The latter two imbecilities are actually still in place, ingrained and phlegmatic like a glacier.) Hopefully, this account helps you avoid similar pitfalls.

I have tried my best to outline a minimal VaR system that works. I believe that the model can and should be kept simple, that data rules in all its imperfection, and that it takes a band of many with different backgrounds to nourish our handle on risk every day. Here’s then to clarity, humility, and concert.

Do not hesitate to contact me at martin@value-at-risk.com with your questions and feedback and my bugs and typos.

There’s nothing left for me but to wish you goespeed with *your* endeavors.