

## Chapter 12

# Thoughts on Stress Testing



*Reality is the leading cause of stress.*

*(Lily Tomlin)*

There are a number of good reasons why stress-testing analysis finds itself as the final topic considered in this book: it is difficult, somewhat awkward, and requires a bird's eye view of one's portfolio and modelling techniques. This, perhaps slightly controversial, opening statement requires some additional colour and justification. Stress-testing's difficulty is easily explained. There is always a range of choice in any financial modelling exercise, but stress-testing is rather extreme. There are at least two fundamental perspectives that one might follow and, within each of these approaches, literally an infinity of possible stress scenarios that one might select, employ, and analyze. There is, in fact, basically too much choice. Creating some kind of logical order and structure, in face of this potential chaos, is one of the fundamental tasks that we will face in the following discussion.

Onto the awkwardness. The words of Rebonato [26] express this sentiment rather well, stating that stress testing

has always been the poor relation in the family of analytical techniques to control risk.

Rather harsh words, but the underlying reason stems from the fundamental structure of stress-testing. It proposes outcomes or events—which are typically extreme and lead to rather adverse portfolio effects—without an associated assessment of their probability. This stands in sharp contrast to classical risk metrics—such as VaR or expected shortfall—that structurally involve the combination of events and their associated likelihoods. Prowling in the background of both worlds, as we've seen in previous chapters, is necessarily an underlying (and unknown) loss distribution. We can envisage a stress scenario as a set (or event) in one's probability space. Unfortunately, we do not know how to assign a measure to it.<sup>1</sup> The consequence is an inherent degree of subjectivity in any stress-testing analysis.

---

<sup>1</sup> At the same time, we (rather keenly) hope that it is not a set of measure zero.

This situation creates an incongruity between standard risk measures on the one hand and stress-testing results on the other; it feels like a situation of comparing apples and pears. Aragonés et al. [2] summarize this sentiment very well with the following question:

How can we combine a probabilistic risk estimate with an estimate that such-and-such a loss will occur if such-and-such happens?

At the same time, however, there appears to be a real need for stress-testing. Quantitative analysts should not start patting one another on the back because of the logical consistency of event and probability treatment within their models. Time and again, over recent decades, probabilistic risk estimates have failed to predict a dismaying number of crises associated with realizations of extreme financial outcomes. Why is this the case? It may stem from over-optimism in one's parametrization, leading to insufficient weight on severely adverse outcomes. Part of the explanation is certainly, as argued by Taleb [31], that unknown unknowns driving financial crises structurally elude our models until it is too late.<sup>2</sup> It is also entirely possible that asking our models to accurately predict inherently unpredictable financial shocks, driven by complex human behaviour, is simply too tall an order. Whatever the reason, it is useful to (once again) underscore the fallibility of probabilistic models.

Stress-testing, when done well, is thus a complement and not a competitor to our probabilistic models. The key objective of stress-testing analysis is not necessarily the construction of extreme scenarios not fully captured in our base models—although this is certainly part of it—but rather to identify vulnerabilities in our portfolios. Where, for example, are the weak spots? What constellation of events would be particularly troublesome for our portfolio? Such information, despite its subjectivity and lack of probability assignment, is useful. When combined with our standard modelling framework, it tells a more nuanced and complete story.

The final claim—the need for a bird's eye view—is the strongest reason for treating stress-testing analysis in our final chapter. “If such-and-such happens”, to borrow again the expression from Aragonés et al. [2], we need to understand the implications for our portfolio. For our purposes, this basically amounts to estimating the impact on our firm's capital position. Such an assessment requires understanding the big picture. We need to incorporate dimensions—introduced in Chap. 1—of both capital demand and supply. The capital-demand impact will necessarily flow through our economic-capital model addressed in Chaps. 2 to 4; as a result, we need to understand how precisely this will happen and how to estimate it. Capital-supply effects arise via the profit-and-loss statement; our loan-impairment computations from Chap. 9 will play a central role in this respect. We will also recycle—and, in some ways, extend—many of the earlier ideas associated with stress-scenario generation introduced in Chaps. 7 and 8. Faced with the need to

---

<sup>2</sup> Such never-before-observed events cannot, by their very nature, be incorporated into our parameter estimates.

estimate large number of computationally intensive economic-capital values, we will again rely upon our approximation model from Chap. 5. Thus, while it might be an exaggeration to state that this final discussion relies on all of the previous chapters, it does lean heavily on most of them.

To summarize, stress-testing is not easy and requires a broader perspective. It is, however, useful and represents a welcome companion to the classical, probabilistic models presented in previous chapters. In the following sections, we will consider a range of alternative approaches—with our usual focus on the credit-risk perspective—towards the incorporation of stress scenarios into our analysis.

**Colour and Commentary 141** (THE ROLE OF STRESS-TESTING ANALYSIS): *Among quantitative analysts, stress-testing has something of a bad reputation. Its very structure—involving specification of an extreme event without an associated probability—is somewhat awkward. Such analysis can be hard to interpret and difficult to combine with traditional probabilistic models. Since the great financial crisis, however, there has been a groundswell of support for stress-testing. Failure of probabilistic models to predict numerous past crises have undermined their credibility; some even argue that such models should be entirely replaced with stress-testing analysis. These arguments, for and against stress-testing, both have some merit. Stress-testing certainly deserves more credit and attention, but it is not a replacement for more holistic financial models.<sup>a</sup> In practice, there is ample room for both approaches in the assessment of financial risk. Stress-testing, in short, is an effective complement to probabilistic models and a powerful tool for the identification of vulnerabilities in one’s portfolio.*

---

<sup>a</sup> To entirely replace probabilistic models with stress-testing would essentially amount to “throwing the baby out with the bath water.” One can certainly improve basic modelling techniques (i.e., eliminate the dirty water) while maintaining their fundamental usefulness (i.e., keeping the baby).

## 12.1 Organizing Stress-Testing

Before jumping into the technical details and considering practical examples, we first need to organize our thinking. As is often the case, it is useful to highlight the key issues as a catechism. We thus raise *three* important questions:

1. What pathways do stress-scenarios follow on their way to impact the firm’s capital position?
2. What logical approach should we take in our construction of a stress-testing framework?

3. How do we manage, in a stress-testing setting, the passage of time and our portfolio composition?

Reflecting upon these questions and providing detailed answers will help immensely in clarifying a possible, and pragmatic, way of thinking about stress-testing. Let's address each in turn.

### ***12.1.1 The Main Risk Pathway***

Our first question can basically be translated as: what does stress-testing actually even mean within our context? We are considering the firm's economic-capital position—and associated measures—from a predominately credit-risk perspective. Economic-capital, as we've indicated and discussed numerous times, is fundamentally based on a long-term, unconditional, through-the-cycle parametrization. This assumption permeates our basic credit-risk model and, as we saw in Chap. 11, underscores most of the regulatory guidance. This is logically rather problematic; to be frank, it feels like a deal-breaker. The very structure of our economic-capital computations would appear to preclude the role of stress scenarios. The long-term, through-the-cycle view presumably incorporates numerous, extreme stress outcomes, but when averaged over many, more normal events, their individual-parameter influence is strongly diluted.

If our model parameters cannot change, then what is effected by stress scenarios? A bit of reflection suggests that, although the through-the-cycle perspective needs to be maintained, stress outcomes can indeed have an important impact on our portfolio's risk profile. Since the beginning of Chap. 2, we have focused on the *three* key credit-risk variables: default probability (i.e., credit rating), loss-given-default, and exposure. All three aspects can be altered from a specific stress scenario. The most obvious, and important, relates to the credit ratings of one's obligors. An adverse event may not change the model parameters, but if it creates broad-based rating downgrade of one's credit counterparties, the impact can be quite severe.

Economic-capital is, in fact, influenced via *two* main avenues associated with the downgrade of a given credit obligor. The first is quite direct. In most cases, a downgrade will lead to an increase in the amount of economic-capital associated with a given credit counterpart.<sup>3</sup> Downgrades, all else equal, will lead to an increase in capital demand. The result is thus a reduction of the firm's capital headroom.

---

<sup>3</sup> This is not always true. Since default and migration risk often move in opposite directions, the net impact of a downgrade can sometimes lead to small decreases in economic capital. While possible, however, such outcomes are not typical and, when they occur, are generally small in magnitude.

The second avenue depends on the accounting treatment of the counterparty's exposure. If the underlying instruments are fair-valued, there will be a valuation impact associated with the downgrade.<sup>4</sup> Should the underlying instruments be held at amortized cost—as is typical for a loan portfolio—then there will be consequences for loan impairments. The expected-credit loss will move to a new (higher) default-probability term structure and, depending on the magnitude of the downgrade, may even involve stage-II lifetime treatment. Either way, the loan-impairment balance will rise. Although the magnitude of the effects will differ depending on fair-value or amortized-cost treatment, the logical consequences are the same. The resulting losses will pass through the profit-and-loss statement and lead to a reduction in capital supply and, as an unwelcome side-effect, a decrease in capital headroom.

A credit downgrade basically creates a squeeze-play for firm's capital position. Capital demand is pushed upwards, capital supply is pushed downwards, and the capital headroom is caught in the middle. A sufficient number of credit downgrades of sizable magnitude can—depending on the firm's capital position—lead to complete erosion of the firm's capital position. The result may involve a credit downgrade of its own or, in the worst case, reduced confidence in the firm, triggering a chain of events leading to default. All of this occurs within the through-the-cycle perspective without any impact on the model parameters. The entire adjustment occurs through the stress-induced deterioration of credit quality of one's portfolio.

This is, of course, a very credit-risk-centric view of stress-testing. Large, adverse shocks to market-risk factors will also—even in the through-the-cycle setting—generate negative valuation effects impacting capital supply. One might wish to revisit operational-risk capital estimates or even adjust capital buffers. These effects matter and are important—although they tend to play a lesser role in most lending institutions—but will *not* be addressed in our treatment. Our focus will—in a manner consistent with previous development—lie firmly upon the credit-risk dimension.

---

<sup>4</sup> Given increased risk, we should expect a higher discount factor, a lower value, and thus a valuation loss.

**Colour and Commentary 142** (CREDIT-RISK STRESS-TESTING PATHWAYS): *Crime dramas, investigative journalists, and political actors are fond of using the phrase “follow the money.”<sup>a</sup> As exciting as it sounds, it sadly doesn’t typically come up too often in quantitative modelling. In this case, however, we have identified a possible exception. The fundamental through-the-cycle perspective of economic-capital computations—with its stable, long-term, unconditional parameters—precludes any direct modelling impact from adverse stress scenarios. At first glance, one might reasonably conclude that stress analysis is misplaced in this context. An important indirect pathway nonetheless exists. Should a stress scenario lead to broad-based credit-obligor downgrade, two capital effects present themselves. The first is a direct increase in capital demand associated with lower portfolio credit quality. The second is—via either valuations or loan-impairments—a decrease in capital supply. Following the money through our modelling frameworks and financial statements, we can identify a consequent squeeze in one’s capital headroom. A severe stress scenario can thus, without impacting our model parameters, lead to dramatic weakening of a firm’s capital-adequacy position. Portfolio downgrade, through economic-capital, market valuation, and loan impairments, represents the (principal) risk pathway of our stress-testing analysis.*

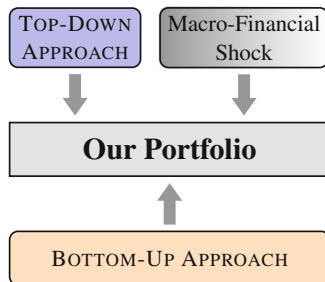
---

<sup>a</sup> Apparently, this dangerous-sounding expression originated from the 1970’s movie “All the President’s Men” surrounding the Watergate scandal starring Robert Redford and Dustin Hoffman.

### 12.1.2 Competing Approaches

Most authors on stress-testing—see, for example, Rebonato [26], Bellini [4] or Rösch and Scheule [28]—would agree that there are *two* broad-based approaches to the topic: top-down or bottom-up. Within each class, there are many variations and there are presumably arguments to be made for an alternative organization. For our purposes—that is, examining how external stress can generate portfolio downgrades and capital degradation—this structure is perfectly sufficient.

Figure 12.1 schematically illustrates these two ideas. As the name suggests, the top-down approach takes a macro perspective. The stress scenarios are defined, in a global sense, through adverse changes in economic output, inflation, monetary policy, commodities prices, and so on. This comfortably links back to our discussion of stress-scenario construction in Chap. 8.



**Fig. 12.1** *Attacking the stress-testing problem*: This schematic illustrates the two broad-based approaches to the stress-testing problem. One may approach from either a top-down or bottom-up perspective. The top-down viewpoint necessarily works through macro-financial outcomes, whereas the bottom-up method is rather less obviously defined.

The top-down approach to the stress-testing problem is, in many ways, the most natural viewpoint. It easily allows us to construct interesting, and topical, narratives. What are the implications, one might ask, of prolonged lower output and higher inflation following the resolution of the COVID-19 pandemic? What if inflationary expectations get out of hand leading to a secular increase in commodity prices and interest rates? Constructing a story for each stress-testing scenario helps to practically narrow down the, literally infinite, number of potential options.

There are other ways to identify top-down stress scenarios. It is possible, and indeed fairly common practice, to *borrow* extreme outcomes from the past to inform one's macro-financial shocks. We could go back as far as the bond-market sell-off in 1994, the Asian crisis, the dot-com bubble, the events surrounding the Lehman Brothers bankruptcy, or the most recent COVID-19 pandemic. There is no shortage of choice.<sup>5</sup> The significant advantage is that the analyst need not construct an explicit macro-financial outcome, nor does she need to worry whether it could potentially occur.<sup>6</sup> The downside is that, because it already took place, the likelihood of a repeat performance is vanishingly small.

There is also an entire industry dedicated to the collection, processing, and forecasting of key financial variables. Such firms are quite competent at shuffling through all of the complexity of economic conditions and financial markets. They not only provide baseline point forecasts and uncertainty bounds, but also stress scenarios of varying degrees of uncertainty. For those institutions with sufficient resources, this is a sensible path to scenario identification.

<sup>5</sup> Reinhart and Rogoff [27] is an excellent source of inspiration for past crises.

<sup>6</sup> This is because it has, in fact, already happened.

The top-down approach thus requires building, buying or borrowing one or more macro-financial stress scenarios.<sup>7</sup> Equipped with these events, one must then determine how precisely they impact the general credit quality (i.e., credit ratings) of one's portfolio. These downgrades then translate into capital demand and supply effects. Although the chain of logic is rather long and complex, it is a very compelling approach. Once completed and all the intermediate steps arranged, one can examine the impact of an output shock on the firm's capital position. Or, as another possible example, we could examine the sensitivity of one's current portfolio to a repeat of the 1998 Asian crisis.

The bottom-up approach attacks the stress-testing problem from the opposite direction. One asks, in a micro fashion, what would be the impact of a direct downgrade to a specific subset of counterparties. The location of downgrade could be determined by individual identity, rating level, industrial sector, region, or it might even be randomized. There is no logical link—at least, directly—to the state of the macroeconomy or financial markets. Instead, it looks to the portfolio structure and seeks to understand the vulnerabilities associated with key concentrations.

This bottom-up perspective appears rather less structured, and thereby less conceptually compelling, than the top-down approach. While entirely true, it also relies less upon the (at times) tenuous set of statistical relationships for its computation. As we've seen in Chap. 8—and will see again in later discussion—establishing a robust relationship between general credit conditions and a set of macro-financial variables is *not* an easy or unequivocal task. Although we may be rather interested in determining the impact of an output shock on our capital position, our ability to accurately measure such an effect will be limited. Rather less sexy, the bottom-up approach makes up in honesty what it lacks in conceptual appeal.

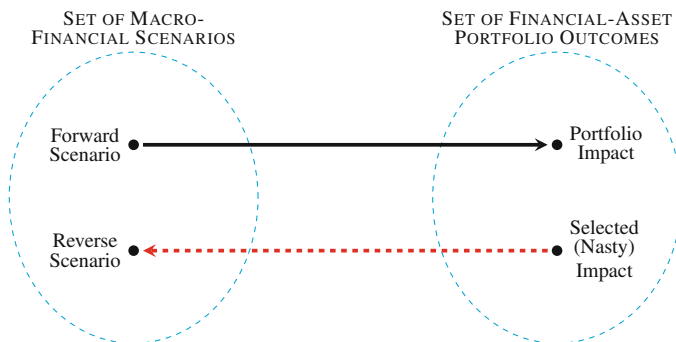
While not precisely the same thing, there is, in our case at least, a link between the bottom-up approach and the idea of reverse stress-testing. The basic idea of reverse stress testing is to identify those scenarios that create headaches for the firm. There is a relatively small, but growing quantitative literature on this topic—see, for example, Albanese et al. [1] in the market-risk setting—and it appears to be increasingly popular in regulatory circles.<sup>8</sup>

---

<sup>7</sup> Given the difficulty of identifying sensible stress scenarios, it may even be preferable to source them from multiple places.

<sup>8</sup> It finds, for example, explicit mention in BIS [6].





**Fig. 12.2** *Forward and reverse stress-testing*: This schematic illustrates, in a highly stylized fashion, the logical distinction between forward and reverse stress testing. The forward direction follows the standard approach of translating a macro-financial shock into a portfolio result. Reverse stress-testing tries to work backwards from a particularly unpleasant portfolio outcome to the original macro-financial shock. This helps to identify vulnerabilities.

Figure 12.2 provides a visualization of the notion of reverse stress testing. The forward direction follows the standard, top-down approach of translating a macro-financial shock into a portfolio result. Reverse stress-testing works backwards from a particularly unpleasant (i.e., nasty) portfolio outcome to the original macro-financial shock. Thinking a bit more mathematically, the set of macro-financial scenarios can be seen as the domain. Each stress scenario, or function argument, is passed through some (usually unknown) stress-testing function to determine the portfolio impact (i.e., the function target or image). Reverse stress-testing essentially involves selecting some portfolio loss from the image of the stress-testing function and trying to infer its pre-image.<sup>9</sup> Notice that we explicitly avoid the term *inverse*, but instead use pre-image; there is no reason to expect a unique one-to-one correspondence between portfolio outcomes and macro-financial scenarios. A bad portfolio outcome could occur through many different possible constellations of adverse macro-financial variables (and vice versa).

By focusing on adverse portfolio outcomes, the reverse stress-testing process can help to identify vulnerabilities. Once again, this is not quite equivalent to a bottom-up approach. In the forthcoming practical discussion, however, it will come tantalizingly close.

<sup>9</sup> Since we don't really know the true form of the stress-testing function, this can be a challenging task.

**Colour and Commentary 143** (STRESS-TESTING STRATEGIES): *Stress-testing analysis seeks to identify portfolio vulnerabilities associated with extreme, adverse states of the world.<sup>a</sup> There are two main ways to address this problem. The first, and most popular, is the top-down approach. Motivated from forward-looking analysis or historical crises, one identifies a collection of unpleasant macro-financial outcomes and traces out their impact on one's capital position. The logic is basically summarized as: if this macro-financial event happens, this is the implication for our portfolio. The alternative, referred to as the bottom-up approach, looks at the impact of specific adverse portfolio events. In our specific case, this would amount to the downgrade of a given set of obligors or those within a geographic region or sector. Related to the bottom-up approach, but nonetheless distinct, is the notion of reverse stress-testing. This basically takes the bottom-up method a step further by trying to identify the macro-financial shock (or shocks) that might have created it. Although the context differs, the conceptual idea of stress-testing is identical in engineering settings. When building a bridge, for example, the engineer seeks to understand the vulnerability in her design and construction to various wind shears, weight loads, and a variety of other dimensions beyond a financial analyst's expertise. Instead of physical structures, our objective is to explore and identify important vulnerabilities in our asset portfolios that might jeopardize the firm's capital adequacy position.*

---

<sup>a</sup> The creditworthiness state variable in our credit-risk economic capital model incidentally uses a similar idea, but maps out the entire credit loss distribution (including many extreme outcomes).

### 12.1.3 Managing Time

In stress-testing analysis, time can be a bit tricky. When, for example, does a scenario actually occur? We cannot simply (without some reflection, at least) snap our fingers and impose a large-scale, negative shock to our asset portfolio. Granted, some adverse situations can occur quickly, but often they unfold over multiple quarters or even years. Even if the effect is instantaneous (or close to it), the aftermath is also important for our analysis.

The consequence is that we should expect our macro-financial scenarios to unfold over time. For argument's sake, let's fix our stress horizon to one year.<sup>10</sup> This immediately raises a second question: how do we describe the portfolio? It is exceedingly unlikely that the portfolio remains fixed over a one-year period. Even

---

<sup>10</sup> This is at the lower end of the scale. It is not uncommon to see two- to five-year stress horizons.

in the unlikely case that it did, multiple aspects of the portfolio will nonetheless change. At a minimum, there will be principal repayments reducing exposures and all instruments will have a shorter tenor.<sup>11</sup> In short, the passage of time creates a number of thorny issues.

There is no single correct way to solve this problem. It basically involves a trade-off between granularity and accuracy on one hand and parsimony and manageability on the other. Looking out numerous years into the future makes it logistically almost impossible to maintain the full details of one's current portfolio structure. There are simply too many moving parts associated with the roll-down of the current portfolio and assumptions regarding the details of new assets. The only workable approach involves dimension reduction. One creates a stylized, low-dimensional view of one's portfolio and builds a model to describe its intertemporal evolution.<sup>12</sup> The top-down or bottom-up scenarios are then applied to this stylized portfolio. Corners nonetheless need to be cut and overall accuracy suffers.

The alternative, which is only tenable for reasonable short time horizons, involves assuming that the current portfolio remains fixed in its current state. As unrealistic as this might appear, it does offer some important advantages. The full granularity of the portfolio is preserved. The regional and industry identities along with credit ratings, loss-given-defaults, tenors, and exposures are known with precision. This permits continued use of our detailed economic-capital and loan-impairment frameworks in the computation. If one looks too far into the future or expects imminent portfolio changes, then this tactic will break down.

This issue basically brings us to the distinction between risk-management and strategic analysis. The risk-management viewpoint takes the portfolio as it is, operates at the highest level of detail, and restricts its attention to relatively short time periods. Strategic analysis extends over longer horizons and provides the analyst with flexibility over the portfolio structure. The ability to dial up or down risk attributes to meet strategic objectives is, after all, the main point of the exercise. To do this, one needs to stylize the portfolio construction. Generally speaking, since one cannot legitimately do both, it is necessary to pick a lane.

With these ideas in mind, we will consciously take the risk-management pathway. Preservation of full portfolio dimensionality is critical, because it allows us to borrow from the various helpful ideas presented in previous chapters. Digging into the minutiae of the portfolio, from both top-down and bottom-up perspectives, also permits us to more definitely identify vulnerabilities. There is a price. We cannot defensibly look too far into the future. A one-year time horizon is thus a fairly tight constraint for our analysis.

---

<sup>11</sup> All else equal, this roll-down effect will reduce migration risk.

<sup>12</sup> This is common in strategic portfolio-choice problems. See, for example, Bolder [7, 8] or Bolder and Deeley [9] for a practical example demonstrating just how involved such exercises can become.

### 12.1.4 Remaining Gameplan

We asked *three* fundamental questions and have provided detailed answers for each of them. Our first question enquires about the practical link between stress and portfolio outcomes. The principal risk pathway for our stress scenarios towards the firm's capital position—within economic capital's through-the-cycle goggles—is via downgrade of individual credit obligors. This provides us with a fairly clear idea of the capital demand and supply effects to be considered.

The second question is strategic: how does one actually go about stress testing? There are two fairly broadly defined alternative avenues for the investigation of stress-scenario impact on our portfolios: the top-down and bottom-up approaches. Rather than choosing a specific track, we will, as we've done throughout this entire book, consider *both* approaches to this problem. There is no compelling reason to specialize our stress-testing framework when much can be learned from exploration of both alternatives

The final question relates to management of the time dimension. We will make use of the current portfolio—in all of its detail—for our analysis. This critically assumes that the portfolio structure remains unchanged as we advance in time. To avoid this assumption turning into conceptual abuse, we must constrain our time horizon to a single year.<sup>13</sup> In making this choice, we have explicitly adopted a high-dimensional, risk-management perspective in our stress-testing analysis.<sup>14</sup>

The remaining discussion in this chapter is relatively easily described. In the context of a small (fictitious) portfolio, we will proceed to examine both the theory and implementation associated with both the top-down and bottom-up approaches to stress-testing analysis. This exercise will not only underscore the key ideas, it will also uncover the many links to the discussion from previous chapters. In this manner, it pulls the (occasionally) disparate material in this book together for the reader.

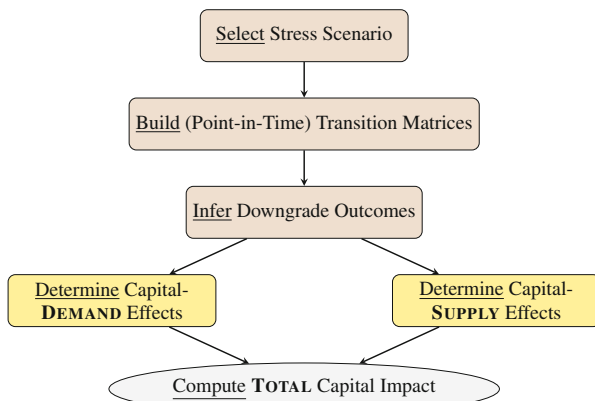
## 12.2 The Top-Down, or Macro, Approach

The top-down approach is the proper starting point. When most people think about stress-testing analysis, it is the top-down perspective that they have in mind. Figure 12.3 gets right to the point with a description of the various steps involved in our specific capital-focused, credit-risk implementation of a top-down stress analysis.

---

<sup>13</sup> This, to be perfectly frank, is already pushing things somewhat. Its defensibility will turn on whether or not one believes that the current portfolio is a reasonable estimator for the future portfolio one year hence.

<sup>14</sup> The longer-term, strategic perspective is also entirely sensible. It does, however, bring us closer to the important (strategic) question of capital forecasting and planning.



**Fig. 12.3** *Top-down sequencing*: This graphic walks through the sequence of steps—from selection of stress scenarios to computation of the total capital impact—involving the performance of a top-down stress-testing analysis. It is a relatively complex multi-step undertaking.

The sequence of steps in Fig. 12.3 is surprisingly long. Selection of one’s stress scenario—which is of utmost importance—is only the beginning. One then needs to establish a link between these adverse macro-financial outcomes and individual credit-obligor downgrades. Fortunately, in Chap. 8 we constructed a methodology for mapping stress outcomes into point-in-time default and transition probabilities. The basic loan impairment calculation—from Chap. 9—relied heavily upon the default-probability dimension. Our top-down stress-testing efforts, by contrast, will exploit the full point-in-time transition matrices to infer portfolio downgrades.<sup>15</sup>

Given the resulting portfolio downgrades, the remaining effort is rather straightforward. We simply determine the valuation, loan-impairment and economic-capital consequences, organize the results, and compute the capital impact. The lower boxes in Fig. 12.3 are common in both the top-down and bottom-up approaches. This is the good news. The bad news is that there is an infinity of possible stress scenarios that one might potentially select. We require some way to make sense of this bewildering array of choice.

Some useful guidance can be found in the field of macroeconomics. More than 40 years ago, in a path-breaking paper, Sims [29] introduced the vector auto-regression (VAR) model to the economics profession.<sup>16</sup> VAR models remain, to this day, a critical element of the economist’s toolkit. The reason is that this approach is surprisingly flexible in its ability to capture the complex interactions between a system (i.e., vector) of correlated, time-indexed random variables. Not only does

<sup>15</sup> The through-the-cycle matrix is still used to compute credit-migration economic capital given each obligor’s (shocked or current) credit state.

<sup>16</sup> Christiano [10] is a fascinating look into the history and repercussions of Sims [29] and related work.

it provide a useful statistical description, but VAR models also permit generally effective forecasts of macro-financial systems. Quite simply, the VAR model is basically custom-built for our purposes.

Vector auto-regressive models—of even moderate size—almost invariably possess a large number of estimated parameters.<sup>17</sup> It is typically fairly hopeless to try to interpret, or perform statistical inference upon, individual model coefficients. For this reason, different strategies are used to interpret and employ VAR models. One common, and powerful, technique is referred to as the impulse-response function. The core idea is to shock a single variable in one's VAR system at the current time, while simultaneously leaving the others unaffected. This shock will typically exert an instantaneous effect on the other variables in the system. The VAR model is then forecasted forward—with no other sources of uncertainty—allowing it to return to its long-term equilibrium values.<sup>18</sup> Shocking a single variable is the *impulse*, while the forecasting results represent the *response*.

While not a definitive solution, the VAR framework and its associated impulse-response functions offer a helpful strategy for organizing the vast array of possible stress scenarios. The first step is to transform our collection of macro-financial variables—we will continue to use the specific variable choices introduced in Chap. 8—into vector-auto-regressive form. We then build the appropriate impulse-response functions. Using these quantities, we may then investigate, for each of our macro-financial state variables, the joint impact of a multiple standard-deviation shock over a one-year stress horizon. Following the sequence of steps outlined in Fig. 12.3, we may then compute a separate capital impact for each macro-financial shock. Instead of a single stress scenario, we produce an logically organized collection of them. This provides some semblance of order to the confusing task of stress-scenario selection. It need not be the end of the line. One can use this information to construct more global stress scenarios, but this analysis will help identify key vulnerabilities. In this way, it lends a bit of direction to any subsequent investigation.

**Colour and Commentary 144** (THE IMPULSE-RESPONSE FUNCTION): *It is simultaneously easy and difficult to select a macro-financial stress scenario. It is practically easy, since one need only mix and match some random set of movements in one's collection of state variables. It is difficult, because there are infinitely many possible choices and it's not a priori*

(continued)

<sup>17</sup> This is not always true, since one can impose numerous clever constraints to control and guide variable interactions. In these cases, individual parameters can become quite important. The sizable parameter dimensionality, the previous point notwithstanding, typically holds.

<sup>18</sup> Although the term through-the-cycle is not actively used in macroeconomics, the long-term equilibrium of a VAR model is essentially the same idea.

**Colour and Commentary 144** (continued)

*clear how to identify a meaningful one. The field of macroeconomics offers a clue towards resolving this impasse. Introduced decades ago by Sims [29], the vector auto-regressive model admirably captures the dynamics of high-dimensional, correlated, macroeconomic and financial time series. The vector-auto-regressive framework also offers a tool—termed the impulse-response function—that describes the system-wide impact of a shock to a single variable in one’s system. Capturing the correlation effects, the impulse-response function traces out the system’s return to its long-term equilibrium. Shocking each macro-financial variable and applying the associated impulse-response function in a systematic fashion allows us to examine our portfolio’s vulnerabilities one source of uncertainty at a time. The results might represent one’s final analysis or a stepping stone to the construction of a more detailed stress scenario. In either case, this strategy lends some welcome structure to a difficult choice.*

### 12.2.1 Introducing the Vector Auto-Regressive Model

Nothing in life is free. To effectively employ impulse-response functions in our exploration of our macro-financial scenario space, we first need to fully understand how they work. This requires an associated appreciation of the vector auto-regressive model. Accomplishing this task, which will require a bit of work and patience, is the objective of this section.

### 12.2.2 The Basic Idea

Given their general popularity in economic circles, there is no shortage of superb sources on vector auto-regressive models. A central resource, for time-series analysis in general and vector auto-regressions in particular, is Hamilton [16].<sup>19</sup> We will lean heavily on his treatment and, indeed, borrow his general notation to sketch out the main elements required for our purposes.<sup>20</sup> We begin with a time-indexed, vector-valued collection of random variables,  $y_t \in \mathbb{R}^{k \times 1}$ . For our purposes, this will be our set of macro-financial variables introduced in Chap. 8. More generally, they could be anything. As the name strongly suggests, we are basically regressing the

---

<sup>19</sup> Judge et al. [18] and Lütkepohl [20] are two other excellent alternative choices.

<sup>20</sup> The reader is, of course, recommended to return to the source. Hamilton [16] is a much more complete, nuanced, and thoughtful exposition.

current values of the system on their previous values. That is, we are regressing the vector  $y_t$  on previous versions of itself. More specifically, we postulate a description of  $y_t$  as a linear function of its previous values. In general, it can be written as

$$y_t = c + \sum_{i=1}^p \Phi_i y_{t-i} + \epsilon_t, \quad (12.1)$$

where  $c \in \mathbb{R}^{\kappa \times 1}$  is a vector of constants, each  $\Phi_i \in \mathbb{R}^{\kappa \times \kappa}$  is a parameter matrix, and  $\epsilon_t \in \mathbb{R}^{\kappa \times 1}$  is an error vector.  $p \in \mathbb{N}$  is the number of lags, or time steps into the past, used to describe the current  $y_t$  outcome. This model arises very naturally. If we perform, as we did in Chap. 2, a standard discretization of a Markov process (such as a drifted geometric Brownian motion) we will recover something conceptually similar to Eq. 12.1. There is, therefore, a clear link between these econometric ideas and the study of stochastic processes.

To keep the notational clutter under control—and to specialize to our specific application—we will set  $p = 1$  leading to:<sup>21</sup>

$$y_t = c + \Phi y_{t-1} + \epsilon_t. \quad (12.2)$$

In this case, we have a single matrix,  $\Phi \in \mathbb{R}^{\kappa \times \kappa}$ , of coefficients in addition to the constant vector,  $c$ . The structure of the error term is crucially important. In particular, we assume that  $\mathbb{E}(\epsilon_s) = 0$  for all  $s$  and

$$\mathbb{E}(\epsilon_t \epsilon_s^T) = \begin{cases} \Omega & : t = s \\ 0 & : t \neq s \end{cases}, \quad (12.3)$$

where  $\Omega$  is a real-valued, positive-definite, symmetric covariance matrix. This basically means that there is a fixed error structure at each point in time, but these errors are not serially correlated throughout time.

---

<sup>21</sup> All of the following results generalize readily the  $p$ -lag case. In fact, with a bit of cleverness, any vector auto-regressive process can be written with a first-order structure. This is referred to as the so-called companion form. Again, Hamilton [16] is a good starting place for this discussion.



If we take expectations of both sides of Eq. 12.2, we can determine the long-term unconditional average value of  $y_t$ . The consequence is:

$$\begin{aligned} \mathbb{E}(y_t) &= \mathbb{E}\left(c + \Phi y_{t-1} + \epsilon_t\right), & (12.4) \\ \mathbb{E}(y_t) &= c + \Phi \mathbb{E}(y_{t-1}) + \underbrace{\mathbb{E}(\epsilon_t)}_{=0}, \\ \underbrace{\mathbb{E}(y_t)}_{\mu} - \Phi \underbrace{\mathbb{E}(y_{t-1})}_{\mu} &= c, \\ \mu &= (I - \Phi)^{-1}c, \end{aligned}$$

where  $\mu$  represents the (unconditional) mean of our vector auto-regressive process. This is the multivariate analogue of the one-dimensional average,  $\frac{c}{1-\Phi}$ . In the univariate case, it is clear that  $\Phi \in \mathbb{R}$  has to be less than unity or the long-term mean is undefined. The same intuition applies to Eq. 12.4, but the condition is a bit more complicated. We rather obviously require that the  $\kappa \times \kappa$  matrix,  $(I - \Phi)$ , be non-singular. The constraint is that all of the  $\kappa$  eigenvalues of  $\Phi$  must lie in the unit circle.<sup>22</sup> Vector auto-regressions satisfying this constraint are referred to as covariance-stationary. In practice, this is essentially a non-negotiable condition.

Our principal interest lies with the forecasting of  $y_t$  several periods into the future. This is accomplished with the use of the natural recurrence in the vector auto-regression. Let us begin, again from Eq. 12.2, with the conditional expectation of  $y_{t+1}$ ,

$$\begin{aligned} y_{t+1} &= c + \Phi y_t + \epsilon_t, & (12.5) \\ \mathbb{E}\left(y_{t+1} \middle| \sigma\{y_t\}\right) &= \mathbb{E}\left(c + \Phi y_t + \epsilon_t \middle| \sigma\{y_t\}\right), \\ &= c + \Phi \underbrace{\mathbb{E}\left(y_t \middle| \sigma\{y_t\}\right)}_{y_t} + \underbrace{\mathbb{E}\left(\epsilon_t \middle| \sigma\{y_t\}\right)}_{=0}, \\ &= c + \Phi y_t, \end{aligned}$$

where  $\sigma\{y_t\}$  denotes the  $\sigma$ -algebra—or as econometricians refer to it, the information set—generated by  $y_t$ . This result is relatively unsurprising: the conditional one-step forward expectation is a linear function of the lagged values,  $y_t$ . Trying

---

<sup>22</sup> Or, equivalently, that the eigenvalues of  $(I - \Phi)$  are all *outside* the unit circle. We refer to the unit circle because eigenvalues may occasionally take complex values. We thus consider the complex norm or modulus of the eigenvalue. If  $\lambda$  denotes the eigenvalue, then the general condition is  $|\lambda| = |r + zi| = \sqrt{r^2 + z^2} < 1$  for any  $r, z \in \mathbb{R}$  where  $i = \sqrt{-1}$ . For more on the complex norm, see Harris and Stocker [17, Chapter 17]

two steps forward—with the same information set—and using similar arguments, yields

$$\begin{aligned}
 y_{t+2} &= c + \Phi y_{t+1} + \epsilon_{t+1}, & (12.6) \\
 \mathbb{E} \left( y_{t+2} \middle| \sigma \{y_t\} \right) &= \mathbb{E} \left( c + \Phi y_{t+1} + \epsilon_{t+1} \middle| \sigma \{y_t\} \right), \\
 &= c + \Phi \underbrace{\mathbb{E} \left( y_{t+1} \middle| \sigma \{y_t\} \right)}_{\text{Eq. 12.5}} + \underbrace{\mathbb{E} \left( \epsilon_{t+1} \middle| \sigma \{y_t\} \right)}_{=0}, \\
 &= c + \Phi \left( c + \Phi y_t \right), \\
 &= c + \Phi c + \Phi^2 y_t.
 \end{aligned}$$

We can, of course, continue to play this game as long as we like. Ultimately, the general term for  $k$  steps into the future is described as

$$\begin{aligned}
 y_{t+k} &= c + \Phi c + \dots + \Phi^{k-1} c + \Phi^k y_t, & (12.7) \\
 &= (I + \Phi + \dots + \Phi^{k-1}) c + \Phi^k y_t, \\
 &= \sum_{i=0}^{k-1} \Phi^i c + \Phi^k y_t.
 \end{aligned}$$

This yields a rather elegant result. If we continue to the limit and let  $k$  tends towards infinity, then two things happen. First, given that all of the eigenvalues of  $\Phi$  are in the unit circle, for sufficiently large  $k$ ,  $\Phi^k$  tends to zero. This means that the impact of the conditioning information,  $y_t$ , disappears over time. Using the same properties of  $\Phi$ , the first term in Eq. 12.7 converges to  $(I - \Phi)^{-1} c$ , which is none other than the long-term unconditional mean from Eq. 12.4. To be very definitive, using the fact that the limit of the sum is the sum of the limits, the result is:

$$\begin{aligned}
 \lim_{k \rightarrow \infty} y_{t+k} &= \lim_{k \rightarrow \infty} \left( \sum_{i=0}^{k-1} \Phi^i c + \Phi^k y_t \right), & (12.8) \\
 &= \lim_{k \rightarrow \infty} \underbrace{\left( \sum_{i=0}^{k-1} \Phi^i c \right)}_{(I - \Phi)^{-1} c} + \underbrace{\lim_{k \rightarrow \infty} \Phi^k y_t}_{=0}, \\
 &= \mu.
 \end{aligned}$$

**Table 12.1** *Vector auto-regression details*: The underlying table highlights a number of interesting details associated with a VAR(1) implementation of our  $\kappa = 10$  macro-financial variables. We observe, by virtue of the eigenvalues of  $\Phi$ , that it is covariance stationary. The volatility of the error term and the long-term mean values are also displayed for each state variable.

Macro-financial variable	$\Phi$ Eigen-values	$\Omega$ Error volatility	Long-term mean ( $\mu$ )	Series mean
Credit spreads ( $\varphi_t$ )	0.32	0.34	2.03	2.01
Inflation ( $\pi_t$ )	0.82	0.20	0.43	0.42
Fed Funds ( $r_t$ )	0.82	0.27	-0.03	-0.04
GDP ( $\theta_t$ )	0.62	4.17	2.45	2.57
Unemployment ( $k_t$ )	0.51	0.90	5.88	5.89
(Non-Fuel) Commodities ( $w_t$ )	0.51	5.29	0.79	0.79
S&P 500 ( $m_t$ )	0.33	5.38	1.82	1.97
Oil ( $h_t$ )	0.33	0.15	0.01	0.01
Curve slope ( $s_t$ )	0.13	0.34	-0.02	-0.01
VIX ( $\sigma_t$ )	0.07	5.07	0.10	0.05

In other words, as we move into the future, our forecasted values of  $y_{t+k}$  ultimately tend towards the unconditional mean of our vector auto-regressive process. As a practical matter, this convergence can occur rather quickly, often with a few years. This is precisely the same logic—indeed, borrowed from this fundamental aspect of time-series analysis—used in Chap. 8 to motivate the notion of time decay.

At this point, let's quickly return to our concrete situation. We introduced a  $\kappa = 10$  system of macro-financial variables in Chap. 8. While perhaps somewhat USA-centric, they do represent a fairly broad view of the global macroeconomy and financial markets. It is a good idea, at this early stage, to fit a VAR(1) specification to this data-set.<sup>23</sup> Table 12.1 provides some key values for each macro-financial variable. The first point is that all of the eigenvalues of the slope coefficient matrix,  $\Phi$ , are comfortably below unity. We may thus, with some relief, conclude that a VAR(1) model of our macro-financial system is covariance stationary. Using Eq. 12.8, we also display the long-term equilibrium values for each variable. In the column immediately beside it, we find the average value for each time series across our three-decade data-set. While not precisely the same, the individual  $\mu$  values are generally quite close to their simple, long-term mean. This should underscore the conceptual link between equilibrium VAR values and the through-the-cycle perspective. The final element in Table 12.1, from the diagonal of  $\Omega$ , is an estimate of each macro-financial variable's error volatility. This provides some insight into the typical magnitude of a shock.

<sup>23</sup> Recall that our quarterly data spans almost 30 years from the early 1990s until 2021; it thus comprises about 120 observations for each macro-financial variable.

**Colour and Commentary 145** (A VECTOR AUTO-REGRESSIVE STARTING POINT): *In Chap. 8, when introducing our system of 10 macro-financial variables, we took pains to ensure that none of the individual components was non-stationary.<sup>a</sup> This was important for estimation of Yang [33]’s structural model linking through-the-cycle and point-in-time probabilities, but it also proves helpful in the stress-testing setting. Estimating a VAR(1) specification for our macro-financial system reveals a covariance-stationary result.<sup>b</sup> This basically ensures that our statistical model is stable. Moreover, when left to its own devices, it will tend back towards its well-defined long-term equilibrium values. These equilibrium outcomes, not coincidentally, are not far from the long-term unconditional mean estimates for each individual time series. The VAR model equilibrium is thus closely coupled with the idea of a through-the-cycle perspective.*

<sup>a</sup> We relied on the work of Dickey and Fuller [11] to test this condition.

<sup>b</sup> Or, in other words, all of the eigenvalues of the slope coefficient matrix,  $\Phi$ , are comfortably within the unit circle.

### 12.2.3 An Important Link

There is an equivalent representation of the vector auto-regression, which is particularly important for our purposes. It is based on a recursive argument and, to be honest, is a bit tedious to construct. It nevertheless relies on a number of the basic ideas introduced in the preceding discussion. Let’s again begin with Eq. 12.2,

$$\begin{aligned}
 & \text{Eq. 12.2} \\
 y_t &= c + \underbrace{\Phi \left( c + \underbrace{\Phi y_{t-2} + \epsilon_{t-1}}_{y_{t-1}} \right)}_{y_{t-1}} + \epsilon_t, & (12.9) \\
 &= c + \Phi c + \Phi^2 y_{t-2} + \Phi \epsilon_{t-1} + \epsilon_t, \\
 &= c + \Phi c + \Phi^2 \left( \underbrace{c + \Phi y_{t-3} + \epsilon_{t-2}}_{y_{t-2}} \right) + \Phi \epsilon_{t-1} + \epsilon_t, \\
 &= c + \Phi c + \Phi^2 c + \Phi^3 y_{t-3} + \Phi^2 \epsilon_{t-2} + \Phi \epsilon_{t-1} + \epsilon_t, \\
 &= \sum_{i=0}^2 \Phi^i c + \Phi^3 y_{t-3} + \sum_{i=0}^2 B_i \epsilon_{t-i},
 \end{aligned}$$

where  $B_i = \Phi^i$  denotes the  $i$ th moving-average coefficient matrix. After two recursions, we begin to see the general pattern. If we push this out indefinitely, we arrive at the following structure with two main components,

$$\begin{aligned} y_t &= \underbrace{\sum_{i=0}^{\infty} \Phi^i c}_{(I-\Phi)^{-1}c} + \sum_{i=0}^{\infty} B_i \epsilon_{t-i}, \\ &= \mu + \sum_{i=0}^{\infty} B_i \epsilon_{t-i}, \end{aligned} \quad (12.10)$$

which, as we saw before, is only possible if  $y_t$  is covariance stationary. This is sometimes called the Wold representation of the vector auto-regression; it is basically a vector moving-average process with an infinite number of lags. The idea is that the current value of our random vector (i.e.,  $y_t$ ) can be described as its long-term average plus the sum of all of the past  $B_i$ -weighted shocks going back as far as one can imagine. Because all of the eigenvalues of  $\Phi$  are restricted to the unit circle, however, the importance of these past errors, shocks or innovations gradually decays and, ultimately, disappears over time. Thus, the long-term unconditional average plays a central role whether we are looking forward as in Eq. 12.8 or backwards as in Eq. 12.10.

From a stress-testing perspective, our interest is in the consequences of a specific (adverse) shock to one of the key elements of  $y_t$ . Specifically, we might ask about the impact of a three standard deviation downward movement in economic output. It is very interesting to consider, on a standalone basis, what this might mean for our portfolio? Equation 12.10 shows us, rather clearly, that the mechanism to construct such a shock is through a perturbation of one or more elements in  $\epsilon_t$ . Let's consider a concrete one-period forward version of Eq. 12.10

$$y_{t+1} = \mu + \sum_{i=0}^{\infty} B_i \epsilon_{t+1-i}. \quad (12.11)$$

What if we compute the derivative of  $y_{t+1}$  with respect to  $\epsilon_t$ ? That is, the sensitivity of our vector process to a change in the innovation term. The result is

$$\begin{aligned} \frac{\partial y_{t+1}}{\partial \epsilon_t} &= \frac{\partial}{\partial \epsilon_t} \left( \mu + \sum_{i=0}^{\infty} B_i \epsilon_{t+1-i} \right), \\ &= B_1. \end{aligned} \quad (12.12)$$

This immediately allows us to describe the impact of our perturbation on  $y_{t+1}$  as

$$\underbrace{y_{t+1} - y_t}_{\Delta y_{t+1}} = \frac{\partial y_{t+1}}{\partial \epsilon_t} \epsilon_t, \quad (12.13)$$

$$= B_1 \epsilon_t.$$

Repeating this same chain of logic leads to the following general version of Eq. 12.11

$$y_{t+n} = \mu + B_n \epsilon_t, \quad (12.14)$$

for an arbitrary choice of  $n$ . This immediately suggests that we can categorize our time  $t$  shock to  $\epsilon_t$  as

$$\underbrace{y_{t+n} - y_t}_{\Delta y_{t+n}} \approx \frac{\partial y_{t+n}}{\partial \epsilon_t} \epsilon_t, \quad (12.15)$$

$$\approx B_n \epsilon_t.$$

Given that each of the moving-average coefficients is intimately related to powers of the  $\Phi$  matrix, as discussed previously, the impact of the  $\epsilon_t$  shock ultimately dissipates to zero.

It is rather interesting to think about a shock at time  $t$  of magnitude  $\epsilon_t$ , but zero for all other future periods (i.e.,  $\epsilon_{t+1}, \dots, \epsilon_{t+n} \equiv 0$ ). To make it more specific, we could give a special structure to  $\epsilon_t$ ; in particular, we might define it as a one standard-deviation shock to the  $j$ th element of  $\epsilon_t$  and zero for all other factors. Following Eq. 12.15, the system impact going forward in time is simply,

$$B_1 \epsilon_t, B_2 \epsilon_t, \dots, B_n \epsilon_t, \dots \quad (12.16)$$

The specific shock could be to output, employment, or the oil price. Indeed, we could perform a sequence of such shocks, in an organized sequential fashion, for each of our macro-financial variables. This powerful idea is called the impulse-response function. It helps us identify—within the context of a large number of variables and parameters—the impact of a shock to inflation on, for example, the oil price in three periods.

### 12.2.4 The Impulse-Response Function

There is, however, a catch. Since  $\epsilon_t$  is a correlated system of errors, it becomes very difficult to disentangle the effects of a given shock. The consequence is that Eq. 12.15—since it is not really telling us what we would like to know—is not

typically directly used for this task. Instead, econometricians use some clever tricks to orthogonalize the covariance matrix of  $\epsilon_t$ . The consequence of this adjustment is that we can interpret a modified version of Eq. 12.15 as a true partial derivative.

The general approach is to revise somewhat the Wold representation from Eq. 12.14—under the assumption of a single time- $t$  shock—as

$$\begin{aligned} y_{t+n} &= \underbrace{\mu + B_n \overbrace{H^{-1}H}^I \epsilon_t}_{\text{Eq. 12.14}}, & (12.17) \\ &= \mu + B_n H^{-1} u_t, \end{aligned}$$

where  $u_t = H\epsilon_t$  and  $H \in \mathbb{R}^{k \times k}$  is non-singular. This mathematical sleight of hand—essentially multiplying by the matrix equivalent of one—permits us to treat our revised error vector as a projection.

While this might not seem like progress, it all depends on the choice of  $H$ . The objective is to find an  $H$  such that the covariance matrix of  $u_t$  is diagonal. From Eq. 12.3, the error-covariance matrix of  $\epsilon_t$  is denoted as  $\Omega$ , which represents a fairly natural jumping-off point. In particular, given its real-valued, symmetric and positive-definite form, we may write this error covariance matrix as,

$$\Omega = P P^T, \quad (12.18)$$

where  $P \in \mathbb{R}^{k \times k}$  is a lower-triangular matrix.<sup>24</sup> The direct corollary of Eq. 12.18 is that

$$P^{-1} \Omega (P^{-1})^T = I. \quad (12.19)$$

With  $H = P^{-1}$ , or equivalently  $u_t = P^{-1}\epsilon_t$ , the variance of  $u_t$  immediately becomes

$$\begin{aligned} \text{var}(u_t) &= \text{var}\left(P^{-1}\epsilon_t\right), & (12.20) \\ &= P^{-1} \underbrace{\text{var}(\epsilon_t)}_{\Omega} (P^{-1})^T, \\ &= \underbrace{P^{-1} \Omega (P^{-1})^T}_{\text{Eq. 12.19}}, \\ &= I. \end{aligned}$$

---

<sup>24</sup> This is referred to as the Cholesky factorization or decomposition, which has already made multiple appearances in previous chapters. Loosely speaking—although, not technically quite true—it feels a bit like the square-root of a matrix. See Golub and Loan [14, Chapter 4] and Press et al. [25, Section 2.9] for much more detailed background.

This is a pretty convenient choice of  $H$ , because the covariance matrix of  $u_t$  is now the identity matrix. This amounts to an orthogonalization of the macro-financial shock dimension.

To introduce our (now independent) shocks, we let  $\delta_j$  denote a  $\kappa$ -dimensional vector with a value of 1 at element  $j$  and zero everywhere else. The value of  $\delta_1$ , for example, is simply

$$\delta_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{12.21}$$

By iterating over  $\delta_j$  for  $j = 1, \dots, 10$ , we can induce consecutive one-standard-deviation shocks for each of our individual macro-financial variables. Should we wish to generate larger shocks, we need only scale up the value of each  $\delta_j$  by the appropriate constant.

We now have all of the moving parts needed to construct our impulse-response function. For the  $i$ th step into the future, it is usefully conceptualized as the difference of two conditional expectations. The first involves a shock of  $\delta_j$  at time  $t$ , while the second is business as usual where we have no information on the shock. In both cases, we condition on the information available at time  $t - 1$ . The distance between these conditional expectations brings us to the following orthogonalized form of the impulse-response function:

$$\begin{aligned} \text{irf}(i, j) &= \underbrace{\mathbb{E} \left( y_{t+i} \mid u_t = \delta_j, \sigma\{\epsilon_{t-1}\} \right)}_{\text{Conditional expectation given shock to variable } j} - \underbrace{\mathbb{E} \left( y_{t+i} \mid \sigma\{\epsilon_{t-1}\} \right)}_{\text{Normal conditional expectation}}, \tag{12.22} \\ &= \mathbb{E} \left( \underbrace{\mu + B_i P u_t}_{\text{Eq. 12.17}} \mid u_t = \delta_j, \sigma\{\epsilon_{t-1}\} \right) - \mathbb{E} \left( \underbrace{\mu + B_i P u_t}_{\text{Eq. 12.17}} \mid \sigma\{\epsilon_{t-1}\} \right), \\ &= B_i P \delta_j - B_i P P^{-1} \underbrace{\mathbb{E} \left( \epsilon_t \mid \sigma\{\epsilon_{t-1}\} \right)}_{=0}, \\ &= B_i P \delta_j. \end{aligned}$$

for  $i = 0, \dots, n$  and  $j = 1, \dots, \kappa$  where  $B_i$  is the  $i$ th moving-average coefficient matrix and  $P$  is the lower-diagonal factorization of  $\Omega$ . Given that  $B_i, P \in \mathbb{R}^{\kappa \times \kappa}$  and  $\delta_j \in \mathbb{R}^{\kappa \times 1}$  the result of Eq. 12.21 is a  $\kappa$ -dimensional vector for each time step. Tracing out Eq. 12.22 for each  $i = 1, \dots, n$  describes the adjustment of our



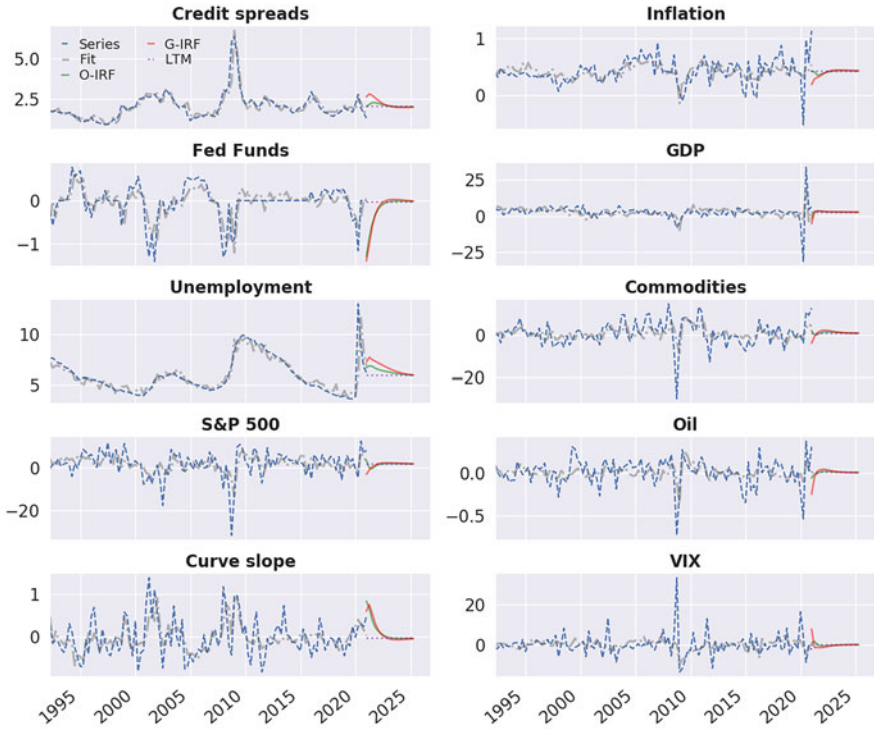
macro-financial system associated with a one-standard-deviation innovation to the  $j$ th variable.

The orthogonalized impulse-response function is a huge step forward, but it still has an important shortcoming. The order of the shocks still matters. Slightly annoying, this requires us to reflect on issues of causality and essentially increases our number of cases from  $\kappa$  to something rather larger. Koop et al. [19] and Pesaran and Shin [24]—with a significant amount of additional effort—offer a solution to this problem. Referred to as the *generalized* impulse-response function, they offer the following twist on Eq. 12.22

$$\begin{aligned}
 \text{girf}(i, j) &= \underbrace{\mathbb{E} \left( y_{t+i} \mid \epsilon_t = \delta_j \sqrt{\omega_{jj}}, \sigma \{ \epsilon_{t-1} \} \right)}_{\text{Conditional expectation given shock to variable } j} - \underbrace{\mathbb{E} \left( y_{t+i} \mid \sigma \{ \epsilon_{t-1} \} \right)}_{\text{Normal conditional expectation}}, \quad (12.23) \\
 &= \mathbb{E} \left( \underbrace{\mu + B_i \epsilon_t}_{\text{Eq. 12.14}} \mid \epsilon_t = \delta_j \sqrt{\omega_{jj}}, \sigma \{ \epsilon_{t-1} \} \right) - \mathbb{E} \left( \underbrace{\mu + B_i \epsilon_t}_{\text{Eq. 12.14}} \mid \sigma \{ \epsilon_{t-1} \} \right), \\
 &= B_i \mathbb{E} \left( \epsilon_t \mid \epsilon_t = \delta_j \sqrt{\omega_{jj}}, \sigma \{ \epsilon_{t-1} \} \right) - \underbrace{B_i \mathbb{E} \left( \epsilon_t \mid \sigma \{ \epsilon_{t-1} \} \right)}_{=0}, \\
 &= \frac{B_i \Omega \delta_j \sqrt{\omega_{jj}}}{\omega_{jj}}, \\
 &= \frac{B_i \Omega \delta_j \overbrace{\sqrt{\omega_{jj}}}^{=1}}{\underbrace{\sqrt{\omega_{jj}} \cancel{\sqrt{\omega_{jj}}}}_{\omega_{jj}}}, \\
 &= \frac{1}{\sqrt{\frac{\delta_j^T \Omega \delta_j}{\omega_{jj}}}} B_i \Omega \delta_j,
 \end{aligned}$$

which is invariant to the ordering of the macro-financial shocks for  $i = 0, \dots, n$  and  $j = 1, \dots, \kappa$ . This result relies on the fact that  $\epsilon_t$  is multivariate normally distributed and requires setting the shock to  $\epsilon_{jt}$  to the volatility of the  $j$ th element in our vector auto-regression.

We can immediately turn this, relatively abstract, discussion into practical, useable outputs. Figure 12.4, for each of our 10 macro-financial variables, illustrates the impact of a multiple standard-deviation downward shock to the Fed-funds rate over a 16-quarter (i.e., 4-year time horizon). The entire three-decade history and fitted



**Fig. 12.4** *Fed-funds impulse response functions*: These graphics illustrate the impact of a multiple standard-deviation downward shock to the Fed-funds rate variable. We observe the simultaneous impact to the other macro-financial variables and the gradual return—over a four-year period—back to their long-term values. Both the orthogonalized (O-IRF) and generalized (G-IRF) impulse-response functions are provided for comparison.

VAR(1) values are also presented. The orthogonalized and generalized impulse-response functions—from Eqs. 12.22 and 12.23, respectively—are naturally both included. Although they provide relatively similar results, there are a few notable points of deviation. In the following analysis, we will use the generalized method, from Eq. 12.23, since it is independent of the ordering.

The utility of the impulse-response function is clearly visible in Fig. 12.4. A large downward shock to the Fed-funds rate leads to instantaneous adjustment to all factors. Particularly effected are the credit spreads, unemployment, output, oil prices, and the slope of the yield curve. The magnitude, direction, and duration of the effects are interesting to examine. Figure 12.4 starkly illustrates why Sims [29]’s suggested use of VAR models found such success. It permits us to assess the plausibility of the model dynamics within an intuitive framework. It seems

**Table 12.2** *An imaginary portfolio*: To examine the various flavours of stress-testing analysis that one may perform, it is easier to conceptualize with a concrete portfolio. This entirely fictitious example, summarized in the underlying table, is hopefully small enough to see what is going on, but large enough to illustrate key trends.

#	PD	Exposure (EUR)	LGD	Tenor (yrs.)	Grace (yrs.)	Bullet Loan?	Coupon Rate	Fair-Valued?	Systemic Weight	Concentr. Index
1	3	3,500,000	0.40	6	0	True	1.50%	True	0.24	0.75
2	4	1,500,000	0.35	4	0	True	1.75%	True	0.19	0.72
3	7	2,000,000	0.45	3	1	False	1.00%	False	0.22	0.80
4	9	1,750,000	0.55	2	0	True	2.25%	False	0.30	0.85
5	13	1,500,000	0.60	5	2	False	3.50%	False	0.15	0.70
6	14	3,000,000	0.15	18	6	False	3.00%	False	0.40	0.95
7	15	1,000,000	0.35	3	1	False	4.00%	False	0.28	0.90
8	17	4,500,000	0.30	15	5	True	5.00%	False	0.22	0.75
9	18	750,000	0.40	5	2	False	6.50%	False	0.20	0.80
10	20	500,000	0.20	4	1	False	6.00%	False	0.40	0.90
Total/Mean	11.1	20,000,000	0.36	8.7	2.4	56%	3.12%	25%	0.26	0.80

rather plausible that a dramatic downward movement in interest rates would be accompanied with falling output, commodity and equity prices, a widening of credit spreads, lower inflation, higher unemployment, a steeper yield curve, and enhanced financial-market volatility. The behaviour of each individual member of our macro-financial system can be examined in this manner. For our purposes, therefore, it makes for an excellent stress-testing tool. In the coming sections, we'll investigate the portfolio and capital consequences associated with 10 separate versions of Fig. 12.4.

### 12.2.5 A Base Sample Portfolio

Equipped with all of the necessary technical elements to perform a top-down stress-testing analysis, all we need is a portfolio. For obvious reasons, the examination of the vulnerabilities of NIB's actual portfolio is a bad idea. Even if we could, the expositional and pedagogical value would be limited by the portfolio size. It is simply too hard to see precisely what is going on in the face of hundreds, or even thousands, of individual positions. Our solution, as in a few previous chapters, is to consider a fabricated example.

Table 12.2 displays a fictitious portfolio comprised of ten distinct positions with an equal number of credit obligors.<sup>25</sup> The number of positions is intended to form a happy compromise along the size dimension. Too many instruments and it is

<sup>25</sup> This abstracts from many complicating real-world factors such as multiple loans and instruments with a single credit counterparty, guarantees, and derivative contracts.

difficult to follow, but too few limits the possible dynamics and interesting cases to be examined. The total portfolio is EUR 20 million with an average PD—along our internal scale—of roughly 11, a mean tenor of roughly 9 years, and a weighted-average loss-given-default of about 0.35. The elements have an assortment of coupon rates and about half of the total notional amounts have a bullet-repayment profile. Two high-quality instruments—allocated to an internal trading portfolio—are fair-valued, while the remainder are accounted for via amortized cost. This creates some distinction between valuation and loan-impairment effects. Finally, Table 12.2 also includes (randomly assigned) systemic weights and concentration-index values to permit easy use of the economic-capital approximation model from Chap. 5.

Using Tables 12.2, 12.3 demonstrates the base, un-shocked, risk position for our imaginary portfolio. It includes both the stage-I and II through-the-cycle expected-credit loss amounts as well as the default, migration, and total economic-capital estimates. We assume that all obligors, at inception, find themselves in stage I.<sup>26</sup> The consequence is that the total loan-impairment is less than one percent of the overall portfolio.<sup>27</sup> Credit-risk economic capital, by contrast, represents almost 9% of the portfolio value. This ranges from a few percent, at the upper end of the credit scale, to about 15% at the bottom end.

Table 12.3 is our baseline. These results will be our point of comparison for all future stress-testing shocks, whether from the top-down or bottom-up approaches. A couple of simplifications need to be mentioned. First, we use the economic-capital approximation with the parameters described in Chap. 5. This is, of course, a slight abuse because these parameters apply to a rather different underlying portfolio. Hopefully the reader will excuse this discretion in the name of simplicity. The second issue, relating to the expected-credit loss, is a bit harder to justify. Chapter 9 indicated clearly that the overall expected-credit loss estimate is a (subjectively) weighted average of three forward-looking macro-financial scenarios. With plans to systematically shock each of our macro-financial variables, it is difficult to imagine the consequence for our forward-looking scenarios. With some patience and imagination, one could attempt to reconstruct a set of macro-financial shock consistent with the forward-looking point-in-time default-probability term structures. In this analysis, however, we have taken the easy way out. The expected-loss calculation is proxied with the shock-invariant through-the-cycle default curves. Given its centrality to the forward-looking construction, it will do a sensible job

---

<sup>26</sup> This implies that they have not experienced any significant decrease in credit quality since inception. See Chap. 9 for more discussion on IFRS 9 stage-allocation logic.

<sup>27</sup> The fair-valued instruments do not contribute to this overall ECL amount. These positions are excluded from the loan impairments; any valuation effects associated with credit-spread movements flow directly through the profit-and-loss statement.

**Table 12.3** *Our risk baseline:* It is centrally important to have a clear starting point, or baseline, for one’s stress analysis. This table summarizes the initial loan-impairment and economic-capital values for each instrument in our portfolio. Across all of the downgrade shocks—from both top-down and bottom-up perspectives—this will be our comparison point.

#	Trade details				Risk estimates					
	PD	EAD	LGD	Tenor	Expected-credit loss			Economic capital		
					Stage I	Stage II	Total	Default	Migration	Total
1	3	3,500,000	0.40	6	0	0	0	20,895	83,690	104,585
2	4	1,500,000	0.35	4	0	0	0	6638	14,064	20,702
3	7	2,000,000	0.45	3	920	2469	4936	41,510	13,434	54,944
4	9	1,750,000	0.55	2	2371	4936	68,130	123,618	11,630	135,248
5	13	1,500,000	0.60	5	11,187	8261	83,035	247,514	168,848	416,361
6	14	3,000,000	0.35	3	9584	26,952	364,675	467,692	60,674	528,366
7	15	1,000,000	0.30	15	84,324	27,700	66,725	98,595	2360	100,955
8	17	4,500,000	0.40	5	19,996	29,538	71,339	386,727	1,690,128	2,196,894
9	18	750,000	0.20	4	164,344	646,459	3.2%	6.5%	1.9%	8.5%
10	20	500,000	0.36	8.7	0.8%	3.2%	6.5%	1.9%	8.5%	
Total/Mean	11.1	20,000,000	0.36	8.7	164,344	646,459	3.2%	6.5%	1.9%	8.5%
Percent of portfolio					0.8%	3.2%	6.5%	1.9%	8.5%	

of capturing the downgrade and stage-allocation effects. It ignores, and this is important to stress, any macro-financial-related shock consequences.<sup>28</sup>

**Colour and Commentary 146 (A MINIMAL WORKING PORTFOLIO):**  
*Internally, as do all financial institutions, we naturally perform our stress-testing analysis using our current portfolio in all of its (gory) detail. This nonetheless creates two problems for the present discussion. The specific structure and vulnerabilities of our actual portfolio contain numerous proprietary and confidential elements. As such, it cannot really be used. Even if it was useable, its sheer size and complexity would certainly make it sub-optimal for our pedagogical discussion. To resolve this problem, we invented a small portfolio comprised of 10 positions each associated with a different credit obligor. Although entirely fictitious, it does cover a reasonable range of credit ratings, position sizes, tenors, notional repayment profiles, accounting treatments, and recovery values. It is big enough to illustrate reasonably complex effects, but small enough to actually see what is going on. As such, it strikes a workable compromise between parsimony and realism.*


## 12.2.6 From Macro Shock to Our Portfolio

Given a specific macro-financial shock, the link to an associated transition matrix is described in Chap. 8. In brief, there are three steps in the construction of a point-in-time transition matrix. The first is the impact upon the default probabilities in the final column. The second involves the so-called upgrade and downgrade ratios. These are factors, centred on unity, that scale up or down the (non-default) off-diagonal, through-the-cycle, quarterly transition matrix elements. In this way, a given macro-financial shock impacts the likelihood of default and relative probability of upgrade and downgrade. The final step is the necessary adjustment of the diagonal elements to ensure each row still sums to unity.<sup>29</sup>

Figure 12.5, to underscore the link to our transition matrices, shows these upgrade and downgrade ratios for each of our 10 macro-financial-variable shocks *four* quarters into the future. We observe rather important differences among the various shocks as well as throughout time. As we would expect, the distance from zero appears to be gradually decreasing with each time step. Eventually, as the shock dissipates, we recover our through-the-cycle transition matrix.

<sup>28</sup> Such effects can, with some additional effort and assumptions, be incorporated into one's analysis.

<sup>29</sup> This is a kind of conservation-of-(probability-)mass constraint.



**Fig. 12.5** Upgrade and downgrade multipliers: This graphic displays the evolution of the quarterly upgrade and downgrade multipliers—out to our one-year stress-testing horizon—associated with a multiple standard-deviation shock to each of our macro-financial state variables. These outcomes play a central role in the structure of the subsequent point-in-time transition matrices.

Some time-related book-keeping is required. With a quarterly model of macro-financial movements, it is also necessary to work with a quarterly transition matrix. To address this point, we need the generator matrix  $G$ , which is computed from the (annual) through-the-cycle transition matrix,  $P$ .<sup>30</sup> The quarterly through-the-cycle transition matrix is obtained through the following transformation:

$$P_Q = \exp\left(\frac{1}{4} \cdot G\right), \quad (12.24)$$

where  $\exp(\cdot)$  denotes the matrix exponential. We then combine the quarterly through-the-cycle transition matrix, our macro-financial shocks, and the previously described point-in-time transition-matrix construction logic. The result is a sequence of point-in-time transition matrices,

$$\left\{ P_Q(i, j) : i = 1, \dots, n \right\}, \quad (12.25)$$

for each  $j = 1, \dots, \kappa$  macro-financial shock. There is thus a logical correspondence between the generalized impulse-response function from Eq. 12.23 and the point-in-time transition matrices for the  $i$ th step into the future and the  $j$ th macro-financial shock.

Our interest lies with the final transition matrix after  $n$  steps into the future along our shocked macro-financial variable paths. We use the time-homogeneity feature of the transition matrix to approximate this quantity as,

$$P(n, j) = \prod_{i=1}^n P_Q(i, j) \quad (12.26)$$

with an associated generator matrix,  $G(n, j)$ . After having invested much time maligning the lack of time homogeneity in credit-rating transition matrices, it may appear somewhat suspicious to see it being used in Eq. 12.26. For relatively small  $n$ , assuming that our transition-matrices are time-homogeneous does not represent a serious crime. We can think of it as roughly equivalent to annualizing an interest rate applied over some fraction of a year.

To respect our one-year stress-testing horizon, we set  $n = 4$ . The consequence is a collection of  $\kappa$  one-year, point-in-time transition matrices: one for each macro-financial shock. The next step, in our rather lengthy chain of logic, involves inferring the associated credit rating for each instrument in Table 12.2 across each top-down shock. The transition matrix provides us with all of the requisite information. Let's begin with our through-the-cycle matrix,  $P$ . If we are told that a credit obligor is in state  $S_k(t)$  at time  $t$ , then its predicted value in one-year's time is simply,

$$S_k(t+1) = \sum_{m=1}^{21} P_{k,m} \cdot m. \quad (12.27)$$

---

<sup>30</sup> These ideas were also introduced in Chap. 8.



We exploited this trick in Chap. 5 when constructing our migration economic-capital estimator. To determine the natural progression of the portfolio, over a one-year horizon, we need only apply Eq. 12.27 for  $k = 1, \dots, 20$ . We then compare the set of distances,

$$\Upsilon_k = S_k(t + 1) - S_k(t), \quad (12.28)$$

for  $k = 1, \dots, 20$ . The sign and magnitude of  $\Upsilon_k$  determine the downgrade outcome. If  $\Upsilon_k = 0.1$ , this suggests a tendency towards credit deterioration, but it is not large enough to predict an actual downgrade. Although debatable, we set the threshold at 0.5. That is, there is a one-notch downgrade if  $\Upsilon_k > 0.5$ , a two-notch downgrade if  $\Upsilon_k > 1.5$ , and so on. Interestingly, using this criterion, the one-year, through-the-cycle transition matrix does not predict a downgrade for any credit rating.<sup>31</sup>

To concretely establish the link from macro-financial shocks to transition matrix matrices to attendant rating outcomes, we arrive at

$$\Upsilon_k(n, j) = \underbrace{\sum_{m=1}^{21} P_{k,m}(n, j) \cdot m}_{\text{Shocked rating}} - \underbrace{S_k(t + 1)}_{\text{TTC forecast}}, \quad (12.29)$$

for  $n = 4$  and  $j = 1, \dots, \kappa$ . To suit our purposes, we have slightly modified the definition of Eq. 12.28. We are comparing the one-year shocked credit-rating outcome to the one-year, through-the-cycle forecast. This ensures an apples-to-apples comparison. If, for example, the shocked rating forecast is 9.6, but the through-the-cycle downgrade expectation is 9.7, we can hardly attribute a downgrade to the macro-financial shock.

### 12.2.7 The Portfolio Consequences

Having delineated the path from macro-financial system to impulse-response function to point-in-time transition matrix to portfolio downgrades, we may now reap the benefits. Table 12.4 provides a concrete point of comparison to the base results in Table 12.3 associated with a multiple standard-deviation shock to the Fed Funds rate—as described in Fig. 12.4—over our one-year stress-testing horizon.

<sup>31</sup> Naturally, given its structure, it eventually pulls all credit counterparties to downgrade, and ultimately default, as we move sufficiently far into the future.

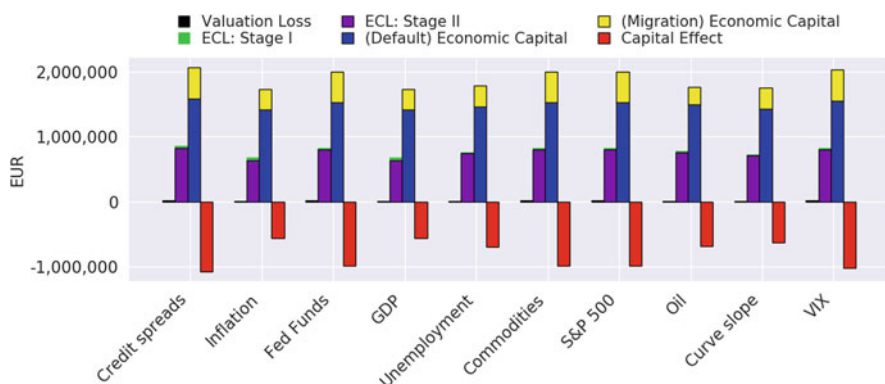
**Table 12.4** A (first) macro-financial shock analysis: The underlying table illustrates the impact—via portfolio downgrades—of a multiple standard-deviation shock to the Fed Funds rate as summarized in Fig. 12.4. The default-probability, valuation losses, expected-credit loss (including stage-allocations), and economic-capital implications are presented.

#	Trade Details					Capital Supply		Capital Demand		
	PD <sub>0</sub>	PD <sub>1</sub>	EAD	LGD	TTM	ECL	MTM	Economic Capital		
						Stage I/II	Δ-Spread	Default	Migration	Total
	1	3	4	3,500,000	0.40	6	0	10,139	27,281	57,606
2	4	5	1,500,000	0.35	4	0	3,512	8,742	13,358	22,100
3	7	7	2,000,000	0.45	3	920	0	41,510	13,434	54,944
4	9	9	1,750,000	0.55	2	2,371	0	123,618	11,630	135,248
5	13	15	1,500,000	0.60	5	104,016	0	118,496	12,373	130,868
6	14	16	3,000,000	0.15	18	103,158	0	308,689	94,721	403,410
7	15	17	1,000,000	0.35	3	47,882	0	183,635	5,040	188,676
8	17	19	4,500,000	0.30	15	470,857	0	541,849	250,467	792,315
9	18	19	750,000	0.40	5	79,747	0	104,040	16,456	120,496
10	20	20	500,000	0.20	4	19,996	0	71,339	2,676	74,016
Total/Mean	11.1	12.4	20,000,000	0.36	8.7	828,947	13,650	1,529,199	477,762	2,006,960
Percent of Portfolio						4.1%	0.1%	7.6%	2.4%	10.0%

Seven of our ten credit obligors experienced a downgrade; three of these were downgraded by a single notch, while the other four experienced a two-grade movement. The credit-risk economic capital allotment rose by about EUR 300,000 or  $1\frac{1}{2}$  percentage points of the portfolio's notional amount. Although both the migration and default dimensions were impacted, around three quarters of the increase can be attributed to default risk. One-notch downgrades to our two high-quality fair-valued securities lead to a modest valuation loss. The largest impact, by a significant margin, stems from the loan-impairment side. The expected-credit loss rise by almost EUR 700,000—about twice the economic-capital movement—amounting to more than 3% of the total portfolio. The lion's share of the loan-impairment result stems from the allocation of five positions—comprising more than one half of the portfolio's value—to IFRS 9's stage II. The corresponding lifetime expected-credit loss values represent a sizable, non-linear hit to the firm's capital position.

Combining the figures in Table 12.4 gets us to the ultimate prize: the capital impact of a given macro-financial shock. The total increase in capital demand is, as mentioned, about EUR 300,000. The corresponding capital-supply result approaches EUR 700,000. Together they squeeze the firm's capital headroom by a total of approximately EUR 1 million. This represents about 5% of the total assets and—although we have not specified the original capital position—would be a tough blow for any financial institution.

Table 12.4 is a critical milestone, but Fig. 12.6 is the end game. It illustrates—for a multiple standard-deviation shock of each macro-financial variable—the combined impact of the five capital drivers associated with credit risk: default and migration economic capital, stage I and II expected-credit losses, and valuation adjustments. The red, downward-trending, bars in Fig. 12.6 cut to the chase. They provide the overall impact to the firm's capital associated when we combined the one-two punch of capital demand and supply effects.



**Fig. 12.6** *Impulse-response-function intuition*: The preceding graphic provides—in a systematic fashion—the expected-credit loss and economic-capital outcomes associated with a multiple standard-deviation shock to each of our  $\kappa = 10$  macro-financial variables. The net capital impact, combining both capital supply and demand effects, is also presented.

Given the extreme nature of the shocks, none of our  $\kappa$  stress scenarios looks particularly appealing. The advantage of this visualization—indeed, the entire exercise—is that it allows us to classify our portfolio’s sensitivity to the various macro-financial factors. Overall, credit spreads seem to exhibit the largest capital impact, while the portfolio is least sensitive to inflation. Examining Fig. 12.6 we could, in fact, put our macro-financial factors into three groups. The highest impact group—with a capital squeeze in the neighbourhood of EUR 1 million—include credit spreads, the Fed Funds rate, commodity and equity prices, and finally financial-market volatility. A medium impact group would include unemployment, oil prices, and the yield-curve slope. The capital consequence for this group is in the neighbourhood of EUR 600,000 to 700,000. Output and inflation appear to be the final, *lower*-effect, macro-financial variables with capital reduction about half of the magnitude of the high-impact group.


While there are clearly limits to the potential interpretation of an imaginary portfolio, Table 12.4 and Fig. 12.6 illustrate the cornerstones of a larger, broader top-down analysis. It is perfectly reasonable to augment this analysis with a few forward- or backward-looking macro-financial stress scenarios. The technical steps in moving from scenario to portfolio impact remain unchanged. Most importantly, our systematic, variable-by-variable, shock-based assessment provides valuable insight into portfolio vulnerabilities and dramatically eases interpretation and communication of the results from more global stress scenarios.

**Colour and Commentary 147** (TOP-DOWN STRESS-TESTING ANALYSIS): *Stress testing is, at its heart, trying to understand how one's portfolio would fare in a financial crisis. A top-down, macro-motivated approach is consequently what most everyone has in mind. The idea of determining one's portfolio sensitivity to a specific, extreme, and adverse macro-financial scenario is inherently conceptually appealing. This is, after all, how a typical financial crisis manifests itself. Actually turning this conceptual idea into a practical outcome is not particularly easy. It involves translating a macro-financial shock into general credit conditions over some horizon, inferring the associated impact on one's portfolio composition and then computing the capital demand and supply effects. To crown it all, there are an infinity of possible stress scenarios from which to choose. Using a workhorse model from macroeconomics, we propose a systemic approach to examine macro-financial shocks—of various degrees of severity—one macro-financial variable at a time. This helps to manage the complexity of stress-scenario selection. The long and tenuous chain of logic from scenario to capital impact, however, is simply a feature of top-down analysis. It is also a weakness, which suggests a motive to intellectually diversify. For this reason, the top-down approach to the stress-testing problem is sensibly complemented with a bottom-up perspective.*

### 12.3 The Bottom-Up, or Micro, Approach

As discussed, it is possible to come at the stress-testing question from another angle. The bottom-up approach side-steps the macro-financial dimension and jumps straight to the downgrade outcomes. Figure 12.7 provides a visualization of the sequencing associated with this alternative strategy. Comparing it to the top-down logic from Fig. 12.3, we see *two* main differences. First, the macro-financial stress scenarios and point-in-time transition matrices disappear. Second, we identify portfolio downgrades rather than *infer* them from macro-financial shocks. We are basically imagining that, at the snap of our fingers, we can change the credit state of any (and every) credit obligor in our current portfolio. This is not very realistic, of course, but we can view this as the essence of the bottom-up thought experiment.

The bottom-up approach liberates us somewhat, but also creates its own challenges. The change of viewpoint does not, to be clear, resolve many of the fundamental issues associated with stress-testing analysis. Direct identification of portfolio downgrades remains absent any probability assignment and there are many—indeed, far too many—possible combinations of downgrade that one might select for consideration. The objective is nonetheless the same: to identify important



**Fig. 12.7** *Bottom-up sequencing*: This graphic describes the sequence of steps—as compared to the schematic in Fig. 12.3—involving the performance of a bottom-up stress-testing analysis. Conceptually, at least, it is rather less involved than the top-down case.

vulnerabilities in one’s portfolio. Like the top-down setting, effective use of the bottom-up approach will require some structure and organization. To the best of the author’s knowledge, there is no *correct* way to perform this task. In the remainder of this chapter, we will instead consider a number of practical variations of the bottom-up approach and, in doing so, help to complement the preceding top-down results. The first order of business is to motivate the need for a master plan.

### 12.3.1 *The Limits of Brute Force*

Although we’ve eliminated the thorny question of selecting macro-financial scenarios, we immediately face another problem.<sup>32</sup> How precisely do we identify downgrade cases? The simple, and naive, answer is *all of them*. This essentially means looking at all possible combinations of downgrades one’s portfolio might experience. While it has a certain appeal, for even a moderate number of credit obligors, this implies an enormous number of possible future combinations and permutations of different credit migrations. The number of combinations of credit ratings—across our 20-notch non-default scale—associated with several hundred different credit obligors is simply staggering.<sup>33</sup>

One might argue that the full power set of downgrade permutations is excessive.<sup>34</sup> Perhaps we might wish to consider the set of all possible one-notch


<sup>32</sup> This situation brings to mind the old proverb, which is used in Tolkien [32], about moving “out of the frying pan and into the fire.” It describes going from a bad situation to a worse one.

<sup>33</sup> Even with  $n = 10$  counterparties and  $k = 20$  credit states, there are some

$$\binom{n+k-1}{k-1} \approx 20,000,000, \quad (12.30)$$

possible combinations of portfolio credit ratings. For any meaningful number of obligors, this is an insurmountable hurdle.

<sup>34</sup> We use the term power set to represent the set of all subsets of downgrades.



**Fig. 12.8** *One-notch downgrade distribution*: The preceding histogram maps out the capital impact of the set of 1023 possible downgrade combinations for our 10-counterparty example introduced in Table 12.2. While interesting, such analysis is not possible outside of unrealistically small portfolio examples.

downgrades. With two counterparties, this is a very small set: each downgrade and both downgrades. It contains only *three* cases. If we add another counterparty, the set grows to seven elements, but it is still manageable. As a general rule, the number of outcomes is

$$n\text{-notch downgrade outcomes} = (n + 1)^{\text{Number of counterparties}} - 1, \quad (12.31)$$

using basic combinatorial logic and subtracting off the single case where no obligor downgrades. Applying Eq. 12.31 to our 10-obligor portfolio from Table 12.2, where  $n = 1$  yields 1023 different possible combinations of one-notch downgrades. Again, while not small, this seems like something we can handle. If we increase the number of obligors to a modest 100, however, the number becomes simply too large to reasonably write down.<sup>35</sup>

Figure 12.8 provides a histogram of the capital impact associated with the set of 1023 possible one-notch downgrades for our small example in Table 12.2. There is no denying that this is an interesting object. It has two modes: one around EUR 200,000 and another in the neighbourhood of EUR 550,000. Many possible combinations of downgrades can generate such reductions in capital. There are only very few constellations of one-notch downgrades, by contrast, that can generate

<sup>35</sup> Some readers might recognize this idea from the story about a clever inventor, a King, a chessboard, and many grains of rice. See Tahan [30] for more on this ancient take on Bellman [5]’s curse of dimensionality.

capital decreases of EUR 700,000. Although a bit cumbersome—given the relatively large number of cases—one could drill into Fig. 12.8 and try to identify trends and patterns within the downgrades.

The bottom line is that, however interesting it may be, Fig. 12.8 is simply *not* available. No matter how much we might wish to perform such an analysis, the dimensionality of the power set of downgrades in a real-life portfolio—even for a single notch—is just too big to tame. Another solution needs to be found.

### 12.3.2 The Extreme Cases

When faced with such a complex problem, one's first thought is typically try to reduce dimensionality. How might we start? While it might sound a bit silly at first, one option is to look at the extremes. By *extreme* we mean placing *all* of our credit obligor into the same rating category. The result is a portfolio entirely comprised of PD01 ratings, another completely in PD02 and so on out until PD20; this assignment occurs irrespective of an obligor's starting point.<sup>36</sup> We can organize this idea in the form of a matrix, which we'll call  $E$ . If we place the  $I$  instruments in one's portfolio along the horizontal axis and map the columns to the rating categories, we have  $E \in \mathbb{R}^{I \times 20}$ . Computing the capital-demand and supply effects for each element in this matrix,  $(I - 1) \times N$  instrument-level computations are involved.<sup>37</sup> For a medium-sized portfolio with 500 credit obligors, this amounts to approximately 10,000 sets of calculations. The main takeaway is that only 20 (extreme) portfolios are considered: one for each credit rating. There is much to criticize, but the huge advantage of this perspective is its invariance to the size of the portfolio.

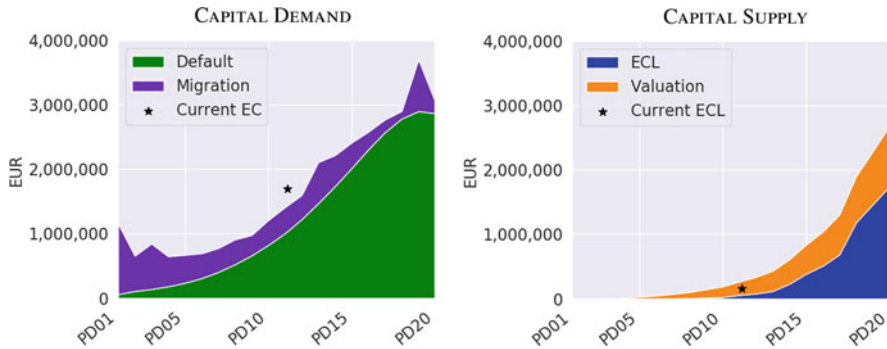
It is admittedly somewhat artificial and unnatural to assign all of the credit obligors in one's portfolio simultaneously to the *same* credit category. Information is clearly being lost, but that is the price of dimension reduction. The litmus test is the helpfulness of this approach and the associated insights it provides into one's portfolio.

For our sample portfolio where  $I = 10$ , this dimension-reduction idea involves 20 new portfolios with only 180 instrument-level calculations. As a consequence, the figures are quickly computed. As in all of the previous analysis, we employ our default and migration economic-capital approximation model. This is not without some drawbacks. Our approximations are estimated using the current portfolio with a broad mixture of rating outcomes. It is rather unfair to then ask the approximation framework to provide entirely sensible results for such extreme portfolios. It will, in particular, have difficulty with the corner cases; that is, where all obligors

---

<sup>36</sup> Each counterparty will thus experience—depending on its starting point—an upgrade, a downgrade, or no change at all.

<sup>37</sup> This is specialized to our 20-notch (non-default) credit-rating scale. It naturally generalizes to any finite number of credit ratings.



**Fig. 12.9** *An extreme portfolio perspective:* These graphics summarize the capital demand and supply effects associated with a rather extreme *what-if* exercise. All positions are simultaneously moved into the same, common credit rating.

are assigned either the highest or lowest credit quality. This situation is readily resolved by simply re-running the simulation model either for all 20 portfolios or, if computational resources are constrained, for the most extreme portfolios.

Figure 12.9 highlights the capital-demand and supply implications of—starting from our baseline—moving all positions into our 20 extreme-rating configurations. On the capital-demand side, default economic capital is a rather smooth increasing function of the credit state. Migration, however, is a bit less well-behaved. This is partly due to difficulty on the part of the approximation model and lumpiness in our small portfolio. Positions #6 and #8 with their large exposures—together representing close to 40% of the portfolio—have rather long tenors. At the corners, this can lead to odd behaviour. This could be partly resolved with proper portfolio simulations, but perhaps not completely eliminated.

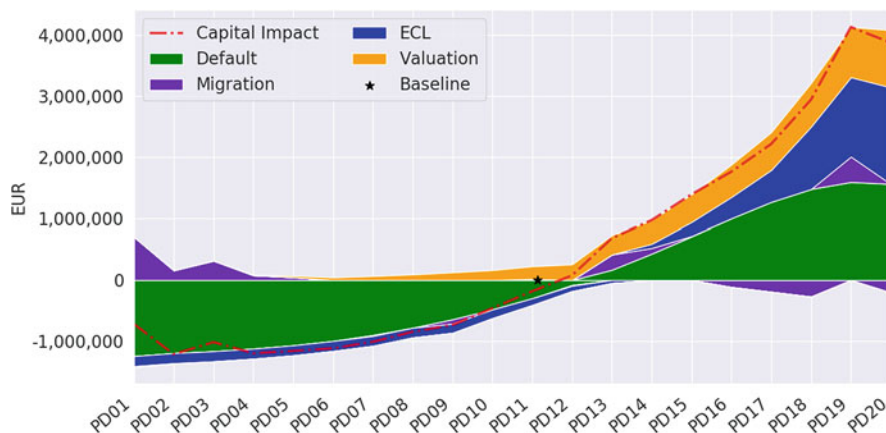
The right-hand graphic of Fig. 12.9 illustrates the capital-supply consequences for these extreme portfolios. The loan-impairment allocation ranges from effectively zero to more than 8% of the overall portfolio value.<sup>38</sup> The valuation impact, which influences only *two* securities, are surprisingly large. These are nonetheless relatively sizable, reasonably long-tenor positions currently situated at the upper echelon of the credit spectrum. Dramatic downgrade—below PD12 or so—would

<sup>38</sup> The reasonableness of this figure is easily verified by calculating:

$$\begin{aligned}
 \text{Extreme PD20 ECL} &\approx \text{Total Exposure} \cdot \text{PD20 Default Probability} \cdot \text{Average LGD}, \\
 &\approx \text{EUR } 20,000,000 \cdot 20\% \cdot 0.36, \\
 &\approx 1,440,000,
 \end{aligned}
 \tag{12.32}$$

which is slightly more than 7% of the portfolio.





**Fig. 12.10** *Capital impact at the extremes*: This graphic combines the capital demand and supply effects—across our extreme views of the portfolio—and illustrates the associated capital impact relative to our baseline setting.


involve significant spread widening with attendant valuation consequences. For small shocks of one or two credit notches, this effect is *not* visible.

Combining the capital demand and supply graphics from Fig. 12.9 and incorporating the baseline perspective, Fig. 12.10 lays out the capital impact of these extreme portfolios. It tells an interesting story. Moving south of the current portfolio—with an average credit rating of about PD11—there is an increasing reduction of capital headroom. At PD19 or PD20, this amounts to roughly EUR 4 million. At the other end of the spectrum, up to about EUR 1 million of capital is actually released. There is thus a certain asymmetry around the baseline portfolio.

Figure 12.11 provides a final perspective on our extreme-portfolio construction. At the position level—as a percentage of its notional value—the individual capital demand and supply results are presented. A heat-map format is employed to incorporate the three dimensions of position, credit rating, and capital percentage. The current baseline portfolio is represented as the shaded line running roughly through the diagonal of our heat maps. The objects in Fig. 12.11 are basically capital demand and supply versions of the previously introduced matrix,  $E$ . For lack of a better term, we refer to these as stress matrices.

In a real-world portfolio, such matrices can be quite large.<sup>39</sup> If one orders the positions from highest to lowest credit category, however, an interesting pattern emerges. Everything to the right of the shaded baseline-portfolio line represents various combinations of portfolio downgrade, whereas everything to the left of

<sup>39</sup> To manage the sheer size of these stress matrices, it is also entirely possible to collect the individual positions into meaningful sub-categories. Good examples would include regions, industries, firm-size, or initial credit rating. Indeed, the organization is limited only by the analyst's creativity (and data availability).



**Fig. 12.11** *Extreme stress matrices*: These graphics display—as a percentage of their underlying notional value—the capital demand and supply impact in grid format. These so-called stress matrices permit identification of particularly sensitive instruments. The true portfolio is highlighted in white.

the demarcation denotes upgrade. As pessimistic risk managers, we naturally focus on the right-hand side, but both provide insights into one’s portfolio. The capital-supply heat map suggests, for example, that there is little scope for capital relief associated with upgrade along this dimension; everything left of the shaded line has essentially the same colour. On the capital-demand side, the picture is a bit more subtle. Some positions—such as #4 and #7—are relatively stable for a one- or two-notch downgrade. As we move to multiple downward steps, however, their capital-demand requirements increase substantially. Moreover, something funny (and possibly incorrect) is occurring for extreme upgrades of position #6. The visualization in Fig. 12.11, while certainly imperfect, does provide a slightly different perspective on one’s portfolio.

**Colour and Commentary 148 (BRUTE-FORCE AND EXTREMES):** *A brute-force examination of the power set of one- or two-notch downgrades—despite its intellectual and practical appeal—is unfortunately off the table. The combinatorics tell the merciless story of a mind-bogglingly large set of possibilities. When the dimensionality is too large, common practice is to identify ways that it might be reduced. A simple approach, which involves a computational burden that is roughly invariant to one’s portfolio size, involves consideration of so-called extreme portfolios. This basically means placing all of one’s position sequentially into each of the rating buckets. The result is a collection of highly stylized extreme portfolios. The truth is that this is probably an excess of dimension reduction. Much information is lost and the associated portfolio outcomes are not realistic. There is, however, some insight to be gained. Our extreme-portfolio construction, while it cannot be*

(continued)

**Colour and Commentary 148** (continued)

*the sole bottom-up strategy for our stress-testing analysis, allows for rather informative visualization of our portfolio. Portfolios are complicated high-dimensional objects. Anything we can do to visualize these sensitivities and vulnerabilities is thus both helpful and entirely welcome.*

### 12.3.3 Traditional Bottom-Up Cases

While a useful starting point, the extreme-portfolio perspective should rather be viewed as exploratory analysis. We can learn a few high-level vulnerabilities of our portfolio, which might be incorporated into more realistic cases. Analogous to the top-down setting, when it comes down to it, stress-testing basically boils down to subjective selection of shocks. In this section, we will investigate a few alternative (and perhaps more typical) ways to proceed.

The term chestnut, beside its obvious meaning, is used by (typically fairly old) English-speakers to describe an over-used joke, story, or example.<sup>40</sup> Our first choice of standalone bottom-up scenario definitely falls into this definition. It involves a case of a global and simultaneous one- or two-notch downgrade to all of the positions in one's portfolio. Its conceptual shortcomings are (partially) offset by its simplicity. No deep reflection or complex logic is required; each position is treated in the same manner.

Table 12.5 summarizes the aggregate, instrument-level, capital-related results of this type of global one- and two-notch portfolio downgrade. The reduction to the firm's capital headroom is displayed in both currency and percentage terms and ordered, in descending fashion, by the two-notch downgrade percentage impact. A global one-grade downgrade leads to a roughly EUR 700,000 headroom squeeze, which corresponds to a medium-impact macro-financial impulse-response function shock. At two notches, the EUR 1.1 million headroom exceeds any of our macro-financial scenarios. This comparison of bottom-up shocks to our top-down macro-financial outcomes is not precisely reverse stress-testing, but it is rather close.

Table 12.5 also provides useful information about our individual positions that, in some cases, overlaps with the extreme-portfolio analysis from the previous section. In the final column, for example, we observe that the average capital effect—when moving from one to two notches of downgrade—is about +70%. This varies wildly across individual positions. One instrument experiences, in fact, a negative

---

<sup>40</sup> See Merriam-Webster [22, page 231].

**Table 12.5** *An analytic chestnut*: This table summarizes the (ordered) capital-headroom results of the classic (and perhaps slightly overused) trick of applying a broad one- or two-notch downgrade to one's entire portfolio. While not particularly nuanced, it does help identify the most capital-sensitive sensible instruments in both currency and percentage terms.

#	Trade Details				One Notch		Two Notches		Notch Increase Percentage
	PD	EAD	LGD	Tenor	EUR	Per cent	EUR	Per cent	
8	17	4,500,000	0.30	15	354,571	7.9%	650,482	14.5%	<b>83%</b>
9	18	750,000	0.40	5	71,587	9.5%	74,701	10.0%	<b>4%</b>
5	13	1,500,000	0.60	5	86,020	5.7%	122,121	8.1%	<b>42%</b>
7	15	1,000,000	0.35	3	43,079	4.3%	73,599	7.4%	<b>71%</b>
4	9	1,750,000	0.55	2	32,885	1.9%	79,270	4.5%	<b>141%</b>
6	14	3,000,000	0.15	18	82,577	2.8%	81,946	2.7%	<b>-1%</b>
3	7	2,000,000	0.45	3	14,018	0.7%	31,113	1.6%	<b>122%</b>
2	4	1,500,000	0.35	4	4,911	0.3%	10,587	0.7%	<b>116%</b>
1	3	3,500,000	0.40	6	-9,560	-0.3%	7,960	0.2%	<b>183%</b>
10	20	500,000	0.20	4	0	0.0%	0	0.0%	<b>0%</b>
Total/Mean	11.1	20,000,000	0.36	8.7	680,087	3.3%	1,131,779	5.0%	<b>72%</b>
Percent of Portfolio					3.4%	-	5.7%	-	-


movement.<sup>41</sup> Another exhibits no movement at all.<sup>42</sup> A handful of instruments encounter increases of around twice the portfolio average. These are, from a risk-management perspective at least, quite interesting cases.

Credit analysts work with *watch* lists of firms that appear close to downgrade or even default. Figure 12.12 runs with this idea and extends it—using the information in Table 12.5—to create a quantitatively motivated *danger* list. These firms may be in fine financial shape, in contrast to the watch list, but were they to be downgraded, they would exert the greatest negative capital impact on the financial institution. Figure 12.12 restricts this to the top-four most sensitive instruments in our small, fictitious portfolio. In a real-life portfolio, this idea can be made much more precise and refined. It nonetheless represents a pragmatic application of the chestnut one- or two-notch downgrade scenarios, which can aid in the identification of real portfolio vulnerabilities.

If a global one- or two-notch downgrade bottom-up strategy lacks finesse, then how might we do better? Again, the challenge is the sheer number of possible combinations. There is, however, an entire class of bottom-up scenarios that are constructed using portfolio-level knowledge. There are many possible examples. One might consider downgrades, of various magnitude, for a given industrial

<sup>41</sup> This is related to the vagaries of migration risk for long-tenor, high-quality securities.

<sup>42</sup> Since position #10 is already in PD20, it can no longer downgrade without moving into default. We have precluded this possibility and held it fixed. Alternatively, of course, one could move it into default and compute the loss. It is a question of preference.



**Fig. 12.12** *Creating a quantitative danger list:* This graphic uses 7th and 8th columns of Table 12.5 to construct a quantitatively motivated danger list. These are securities with the highest capital sensitivity to one- or two-notch downgrade.

sector, or firm size, or geographical region. These might be supported by portfolio concentrations, macro-financial reasoning, or a gut feeling. Typically, there is a story associated with such scenarios. The deteriorating prospects of a particular industry, as an example, may have been treated extensively in the media or identified by one’s credit analysts. A natural bottom-up stress scenario would involve a broad-based, or more complex, set of downgrades to one’s obligors in this industry.<sup>43</sup>

To illustrate and motivate this idea, we will construct a fairly artificial example for our sample portfolio. The idea is to use the concentration index from Table 12.2. These, completely invented, values shouldn’t really be taken very seriously when thinking about these ideas and interpreting the results. Instead, we should view them as a proxy for something more complicated in the underlying structure of the portfolio. In particular, we propose the following bottom-up stress scenario:

$$\text{Downgrade}_i = \begin{cases} \text{Concentration Index}_i \geq 0.9 : +3 \\ \text{Concentration Index}_i \in [0.8, 0.9) : +2 \\ \text{Concentration Index}_i < 0.8 : 0 \end{cases}, \quad (12.33)$$

<sup>43</sup> This might feel a bit like a macro-financial scenario, and to a certain extent it is, but such situations are typically too granular (or micro-focused) for treatment within a macro-financial model.

**Table 12.6** *A concentration-motivated case:* The underlying table illustrates the capital supply and demand implications of stress scenarios constructed with the portfolio’s concentrations in mind. Using the concentration index from our simple example, this is representative of a class of possible bottom-up approaches.

#	Trade Details					Capital Supply		Capital Demand		
	PD <sub>0</sub>	PD <sub>1</sub>	EAD	LGD	Conc. Index	ECL	MTM	Economic Capital		
						Stage I/II	Δ-Spread	Default	Migration	Total
1	3	3	3,500,000	0.40	0.75	0	0	20,895	83,690	104,585
2	4	4	1,500,000	0.35	0.72	0	0	6,638	14,064	20,702
3	7	9	2,000,000	0.45	0.80	5,942	0	69,383	11,652	81,035
4	9	11	1,750,000	0.55	0.85	16,398	0	185,515	14,976	200,491
5	13	13	1,500,000	0.60	0.70	11,187	0	82,369	19,207	101,576
6	14	17	3,000,000	0.15	0.95	123,291	0	334,819	67,071	401,890
7	15	18	1,000,000	0.35	0.90	62,826	0	200,007	2,993	203,000
8	17	17	4,500,000	0.30	0.75	84,324	0	467,692	60,674	528,366
9	18	20	750,000	0.40	0.80	95,166	0	103,871	4,320	108,191
10	20	20	500,000	0.20	0.90	19,996	0	71,339	2,676	74,016
Total/Mean	11.1	12.2	20,000,000	0.36	0.80	419,129	0	1,542,529	281,322	1,823,851
Percent of Portfolio						2.1%	0.0%	7.7%	1.4%	9.1%

for  $i = 1, \dots, I$ . Equation 12.33 consequently applies a three-notch downgrade to the most concentrated aspects of the portfolio, two grades for moderate concentration, and no change for everything else. In this way, specific positions or obligors are *not* targeted. Instead, some set of downgrade criteria are applied and investigated. Equation 12.33 is simply one example of a very large, probably countably infinite, class of stress scenarios. The trick is to identify a few that are meaningful and useful for one’s purposes.

Following the usual format, Table 12.6 provides the capital implications of the stress scenario from Eq. 12.33 for our example portfolio. There are three cases with a three-step downgrade; position #10, already being at the very bottom of the scale, is nonetheless left unaffected. Another three cases are downgraded by two notches, while all other positions remain unchanged. Comparing to the baseline in Table 12.3, the total capital headroom reduction is slightly less than EUR 400,000. Roughly  $\frac{2}{3}$  of the impact can be attributed to the capital-supply effects stemming from the loan-impairment calculation. Since only about 40% of the portfolio is hit by this stress scenario, the results are relatively less severe than we’ve seen in previous bottom-up and top-down cases. In a real-life portfolio, the depth of analysis and insight are potentially more significant than with our imaginary example. Equation 12.33 and Table 12.6 nonetheless provide a concrete illustration of the more traditional class of bottom-up scenarios.

**Colour and Commentary 149** (CLASSIC BOTTOM-UP STRESS SCENARIOS): *Top-down macro-financial related stress-testing scenarios, to be effective, usually involves an associated narrative. One envisions the outcome as a consequence of some forward- or backward-looking sequence of events. Bottom-up scenario construction, albeit in a different way, is similarly constrained. We can ignore this point via global  $n$ -notch portfolio downgrades.<sup>a</sup> This provides some useful insight in a more realistic manner than in our extreme-portfolio experiment. It can even help to build danger lists of instruments with particular capital sensitivity to downgrade. To really gain traction in bottom-up analysis, however, a narrative is indispensable. Classical bottom-up scenarios involve identifying downgrade stories for different aspects of one's portfolio. Related to elements such as firm size, industry, or region, they are typically inspired by intimate portfolio knowledge. A good start is the well-known risk-management question: what keeps you up at night? Practical examples are related to concentrations, known weaknesses, or specific public information too granular for inclusion in one's macro-financial model.*

---

<sup>a</sup> Usually, we set  $n = 1, 2$ , although there are no fast rules. For sufficiently large  $n$ , however, the results converge to the lower end of our extreme-portfolio analysis.

### 12.3.4 Randomization

Our final flavour of bottom-up stress-testing scenario construction is less conventional, but has a long and successful history in quantitative analysis: randomization. It has been used to solve intractable high-dimensional problems via Monte Carlo simulation,<sup>44</sup> to identify useful search directions in complex non-linear optimization problems,<sup>45</sup> and to improve the efficiency of solutions to stochastic optimal-control problems.<sup>46</sup> All of these examples have something in common; they involve trying to solve problems in the face of potentially crippling dimensionality. As we quickly learned when attacking the selection of bottom-up scenarios via brute force, we face a similar challenge in this setting. It thus stands to reason that we might also benefit from this technique.

---

<sup>44</sup> Indeed, we already investigated these ideas extensively in Chap. 4.

<sup>45</sup> See, for example, McCall [21].

<sup>46</sup> There are multiple possible references for this field, all widely outside the author's area of competence, with increasing levels of technical complexity. Ono et al. [23] is a deeply cool example—to give the reader a sense of the broad applicability of randomized strategies—involving planning trajectories of entry, descent, and landing for future Mars missions.

While randomization of downgrade outcomes might be employed in any number of ways, we will examine it in a rather specific manner. Our interest is in the significant—let’s call it catastrophic—downgrade of a small number of important credit obligors in one’s portfolio. It could also be restricted to some subset of the portfolio such as the largest exposures, those obligors with the highest loss-given-default uncertainty, or those with generally excellent credit credentials. Assigning downgrade to a handful of members of such a portfolio subset—as a specific stress scenario—is problematic. The analyst will inevitably be asked: why did you pick this obligor and not another one? The question is entirely justified. Our response is to abstract from any one (or several) obligors and select *all of them*. The difference with the brute-force setting is that we will randomly downgrade a few counterparties—by numerous credit notches—and compute the capital-headroom implications. We will then reset the portfolio and repeat the process. Performing this randomized action many times and averaging across the results will—integrating across all the possibilities—approximate the capital impact of catastrophic downgrade in a subset of one’s portfolio. It will do so, however, in a global sense without explicitly identifying any specific credit counterparties.

Let’s assume that we have  $I$  members in our portfolio and we wish to consider all possible combinations of catastrophic downgrade by  $n$  credit counterparties. This is a well known counting problem with the solution,

$$\binom{I}{n} = \frac{I!}{n!(I-n)!}. \quad (12.34)$$

For small  $I$  and  $n$ , this number of combinations is quite manageable. Imagine that you have 100 large exposures and wanted to consider all combinations of *two* multiple-step downgrades. Using Eq. 12.34 with these values generates 4950 possibilities. Randomization hardly seems necessary. We could compute the capital impact for each case and take the average. Once again, however, dimensionality becomes a problem. For  $I = 100$ , Eq. 12.34 grows very fast in  $n$ .<sup>47</sup> Setting  $n = 3$  yields more than 150,000 combinations, while  $n = 5$  already puts the total to more than 75 million. Examining all possible combinations is thus not a suitable general strategy; randomization, in such cases, can get us to good answers with much less computational effort.

We’ll use our small portfolio to make this idea a bit more tangible. Given that  $I = 10$ , a value of  $n = 5$  generates the largest number of possible combinations: 210 cases. As a consequence, a randomization strategy is not really necessary for any choice of  $n$ . Setting  $n = 2$ , however, does allow us to investigate the interplay between randomization and simply examining all possible instances of catastrophic default.

---

<sup>47</sup> Given that Pascal’s triangle is lurking in the background, Eq. 12.33 reaches a maximum when  $n = \frac{I}{2}$ —if  $I$  is an even number—and then starts to decrease again.



	1	2	3	4	5	6	7	8	9	10
1	(1, 1)	(1, 2)	(1, 3)	(1, 4)	(1, 5)	(1, 6)	(1, 7)	(1, 8)	(1, 9)	(1, 10)
2	(2, 1)	(2, 2)	(2, 3)	(2, 4)	(2, 5)	(2, 6)	(2, 7)	(2, 8)	(2, 9)	(2, 10)
3	(3, 1)	(3, 2)	(3, 3)	(3, 4)	(3, 5)	(3, 6)	(3, 7)	(3, 8)	(3, 9)	(3, 10)
4	(4, 1)	(4, 2)	(4, 3)	(4, 4)	(4, 5)	(4, 6)	(4, 7)	(4, 8)	(4, 9)	(4, 10)
5	(5, 1)	(5, 2)	(5, 3)	(5, 4)	(5, 5)	(5, 6)	(5, 7)	(5, 8)	(5, 9)	(5, 10)
6	(6, 1)	(6, 2)	(6, 3)	(6, 4)	(6, 5)	(6, 6)	(6, 7)	(6, 8)	(6, 9)	(6, 10)
7	(7, 1)	(7, 2)	(7, 3)	(7, 4)	(7, 5)	(7, 6)	(7, 7)	(7, 8)	(7, 9)	(7, 10)
8	(8, 1)	(8, 2)	(8, 3)	(8, 4)	(8, 5)	(8, 6)	(8, 7)	(8, 8)	(8, 9)	(8, 10)
9	(9, 1)	(9, 2)	(9, 3)	(9, 4)	(9, 5)	(9, 6)	(9, 7)	(9, 8)	(9, 9)	(9, 10)
10	(10, 1)	(10, 2)	(10, 3)	(10, 4)	(10, 5)	(10, 6)	(10, 7)	(10, 8)	(10, 9)	(10, 10)

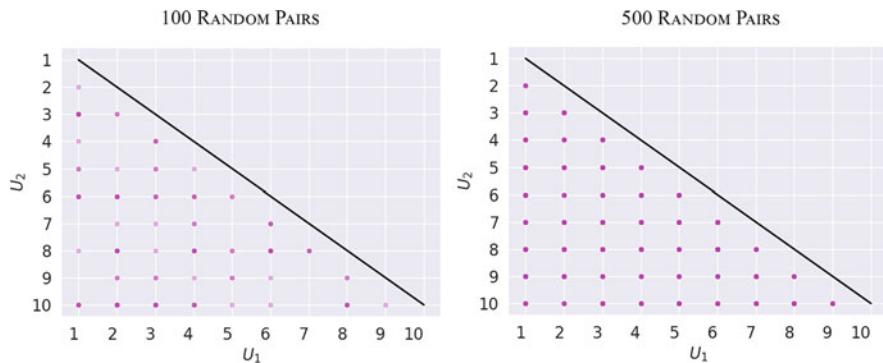
**Fig. 12.13** *Partitioning the unit square*: To motivate the randomization strategy among our  $I = 10$  exposures, this square allows us to visualize the set of all possible combinations of  $n = 2$  catastrophic downgrades. By symmetry, the values above and below the diagonal have an identical portfolio impact.

We will also use a rather simple method for random selection of our two downgrade candidates. There are other, rather more clever ways to do this, but the presented method helps us see clearly what is going on. Our plan is to choose, in a uniformly random manner, two of  $I = 10$  obligors.<sup>48</sup> This can be accomplished by partitioning the unit square into  $I \times I$  sub-squares, where each square is described by two integer coordinates.<sup>49</sup> This object is presented in Fig. 12.13. There are  $K \times K = 100$  sub-squares. Due to the inherent symmetry of our problem—the order of the pair has no importance—and a need to avoid the diagonal, there are only  $\frac{I^2 - I}{2} = 45$  unique combinations.<sup>50</sup>

<sup>48</sup> The technique we present is conceptually related to the more general stratified sampling method termed Latin-hypercube sampling. See, for example, Glasserman [13, Section 4.3], Fishman [12, Section 4.3], and Asmussen and Glynn [3, Section V.7]. The basic idea is to divide the one’s sample space into a disjoint set of sub-regions—typically referred to as *strata*—and to ensure that simulated random variates fall into each of these strata. This partitioning of the sample space ultimately leads to reduction of variance in one’s simulation estimators without biasing its convergence.

<sup>49</sup> For  $n = 3$ , this becomes the unit cube, while for larger values it generalizes to the unit hypercube.

<sup>50</sup> This happily coincides, as it must, with evaluation of  $\binom{10}{2} = 45$  from Eq. 12.34.



**Fig. 12.14** *Unit-square coverage:* By uniform random selection of downgrade pairs, it takes a rather large number of draws to cover the lower-diagonal of our partitioned unit square. 100 draws is not enough, while 500 seems to do the job. Considering there are only 45 combinations, this is not very efficient. As we add dimensions to our hypercube and increase  $I$ , however, this becomes a viable strategy.

Randomly selecting a pair of obligor downgrades thus reduces to picking a single sub-square. This is operationalized by selecting uniformly distributed random variates,  $U_1, U_2 \sim \mathcal{U}[0, 1]$ . The corresponding square is given by

$$\left( \underbrace{[\lceil I \cdot U_1 \rceil]}_{\text{Row}}, \underbrace{[\lceil I \cdot U_2 \rceil]}_{\text{Column}} \right) \equiv \left( \underbrace{[\lceil I \cdot U_2 \rceil]}_{\text{Row}}, \underbrace{[\lceil I \cdot U_1 \rceil]}_{\text{Column}} \right). \tag{12.35}$$

where  $\lceil \cdot \rceil$  denotes the ceiling operator to ensure that we obtain integer coordinates.<sup>51</sup> The only constraint is that we need to ensure that no draw falls along the diagonal, since this amounts to having a single counterpart downgrading twice.<sup>52</sup>

How well does repeated draws from Eq. 12.34 fill in the unit square? Not particularly well. Figure 12.14 displays—for 100 and 500 randomly drawn pairs of downgrades—the coverage of the lower-diagonal of our partitioned unit square from Fig. 12.13. 100 draws is not enough, while 500 seems enough to do the job. Considering there are only 45 combinations, this is fairly inefficient. For our (very small) sample portfolio, this is clearly overkill. As we add dimensions to our hypercube (i.e.,  $n > 4$  or so) and increase the number of counterparties (i.e.,  $I > 50$ ), however, this becomes an entirely (indeed, perhaps the only) viable strategy.

<sup>51</sup> This is readily generalized for arbitrary  $n$  by selecting  $n$  independent uniform variates and adjusting Eq. 12.35 accordingly.

<sup>52</sup> Practically, this means we toss out any randomized draws from the diagonal. This is admittedly not very elegant and a bit computationally wasteful.



**Fig. 12.15** *Randomized convergence*: After a few hundred randomly drawn downgrade pairs, the average capital impact converges rather closely to the grid-based solution involving the (known) 45 possible combinations. Once again, as we increase both  $I$  and  $n$ , randomization strategies represent our only hope of solving this problem.

Figure 12.15 submits the (rather sad) results of comparing various numbers of random pairs of downgrades to the average across all of the known combinations for our simple 10-security portfolio. It takes a few hundred random samples to get near to the final figure, which is just south of about EUR 800,000. This is rather typical. In low dimensions, randomization is a poor, inefficient strategy. Numerical approaches using grids or working with the power set routinely dominate in terms of accuracy and speed. A randomized approach, for all its shortcomings, does have an important advantage. It is roughly invariant to dimensionality; that is, its (slow) convergence rate does not change (much) for very large problems.

At this point in our discussion, we should take a moment to underscore the importance of Chap. 5. While the expected-credit-loss and credit-spread-valuation computations are fast, estimating economic capital consequences are not. If we had to turn to our simulation engine to estimate the implications of every downgrade, the majority of the ideas we've considered in our bottom-up analysis would be infeasible. Our economic-capital approximation models thus open up broad avenues of investigation, which were previously inaccessible. This is a huge analytic advantage. The approximation does, of course, have some shortcomings. In particular, for extreme changes to our portfolio, we should be somewhat cautious in interpreting the result.<sup>53</sup> Indeed, in such cases, there is certainly value in turning back to the simulation engine to *verify* the approximation-model results.

<sup>53</sup> Handling the downgrade of a few positions within the overall portfolio, by contrast, is an entirely natural application of the approximation model.

**Table 12.7** *Randomized two-obligor catastrophic downgrade*: The underlying table, in the usual way, outlines the average capital demand and supply repercussions of a randomized pair of catastrophic downgrades. To obtain these results, we employed 500 random downgrades of two obligor's current credit ratings down to PD20.

#	Trade Details					Capital Supply		Capital Demand		
	PD <sub>0</sub>	PD <sub>1</sub>	EAD	LGD	Conc. Index	ECL	MTM	Economic Capital		
						Stage I/II	Δ-Spread	Default	Migration	Total
1	3	6.4	3,500,000	0.40	0.75	0	144,538	103,273	72,440	175,713
2	4	7.3	1,500,000	0.35	0.72	0	43,217	29,968	12,809	42,777
3	7	9.1	2,000,000	0.45	0.80	42,765	0	90,146	12,586	102,733
4	9	11.8	1,750,000	0.55	0.85	67,623	0	220,890	10,296	231,185
5	13	14.4	1,500,000	0.60	0.70	66,505	0	98,620	16,927	115,547
6	14	15.1	3,000,000	0.15	0.95	40,289	0	269,501	147,059	416,560
7	15	16.1	1,000,000	0.35	0.90	28,698	0	158,317	8,981	167,298
8	17	17.6	4,500,000	0.30	0.75	178,757	0	483,600	61,142	544,742
9	18	18.4	750,000	0.40	0.80	40,384	0	99,587	2,728	102,315
10	20	20.0	500,000	0.20	0.90	19,996	0	71,339	2,676	74,016
Total/Mean	11.1	12.9	20,000,000	0.36	0.80	485,017	187,755	1,625,242	347,645	1,972,888
Percent of Portfolio						2.4%	0.9%	8.1%	1.7%	9.9%

Table 12.7 outlines, in the now familiar format, the detail capital demand and supply results associated with many randomized pairs of catastrophic downgrades. As in Fig. 12.15, catastrophic implies an obligor being downgraded from its current credit rating all the way down to PD20.<sup>54</sup> The average increase in capital demand is a bit over EUR 250,000, while the loan-impairment impact comes to around EUR 300,000. Interestingly, the valuation repercussions total almost EUR 200,000. This stems from the current high quality of the fair-valued securities in our portfolio. Overall, as we saw in Fig. 12.15, the total capital squeeze amounts to just shy of EUR 800,000. As always in a stress-testing analysis, we cannot assign a probability to this outcome. We can, however, state that it is twice our concentration-motivated, bottom-up scenario from Table 12.6, it is roughly equivalent to an extreme portfolio with all obligors assigned about PD13, and it lies at the upper end of our medium-impact impulse-response function macro-financial shocks.

<sup>54</sup> As before, for credit counterparties already at PD20, we do not force default.

**Colour and Commentary 150** (THE POWER OF RANDOMIZATION): *Traditional bottom-up stress-testing, in contrast to the top-down approach, places higher importance on the idiosyncratic dimension. These are risks specific to the financial institution's portfolio. They consider downgrades—in terms of concentrations, industries, and regions—that may really do harm to one's capital position. Taking idiosyncratic risk to its logical limit, we should also consider a small number of individual obligors experiencing severe deterioration in their credit quality. Independent of macro-financial shocks and portfolio structure, this basically captures bad luck. While tempting, it is not a good idea to simply pick a handful of names from one's portfolio, downgrade them by multiple notches, and call it misfortune. You will naturally be asked: how did you come to choose those particular names? It also represents a single, extremely low probability, case. An alternative solution is to use a randomization strategy. One thus computes the average capital demand and supply outcomes over many randomly selected sets of hard-luck (or catastrophic) downgrades.<sup>a</sup> The results shine a different light on (more global) idiosyncratic vulnerabilities in one's portfolio.*

---

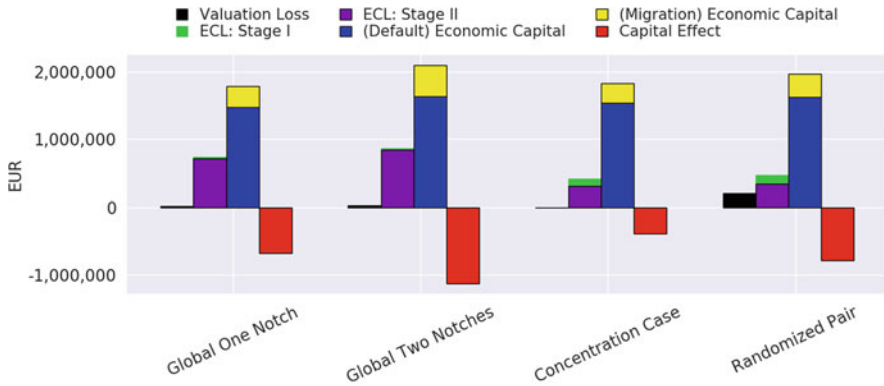
<sup>a</sup> One can make this even more meaningful by focusing on subsets of one's portfolio. Selecting a small number of obligors from the firm's 100 largest exposures or a particularly important industry or region.

### 12.3.5 Collecting Our Bottom-Up Alternatives

The various alternatives examined in the previous sections should *not* be viewed as an exhaustive review of the entire universe of bottom-up stress scenarios. Instead, they represent some practical ideas for getting a handle on the complexity associated with scenario selection. Whatever strategy one follows to identify meaningful bottom-up stress scenarios, however, it is important to impose some structure and organization. Failure to do so does not immediately imply disaster, but does substantially increase the chances of losing the forest for the trees.

To close out our bottom-up stress-testing analysis, Fig. 12.16 presents *four* of the main scenarios examined in this section. Designed analogously to the macro-financial results presented in Fig. 12.6, it allows comparison of capital demand and supply effects in a single glance. Our simple portfolio example does not permit incredibly complex dynamics, but there is nonetheless a surprising amount of intra-scenario deviation among the individual components. The loan-impairment stage allocation and valuation effects are rather more nuanced in the concentration-index and randomized-downgrade scenarios.

Stress-testing, as the preceding discussion hopefully makes clear, is hard work. There are no short cuts; it is inherently a bit messy, hard to manage, and high



**Fig. 12.16** *Bottom-up intuition*: This graphic illustrates, in a manner analogous to the macro-financial results from Fig. 12.6, the individual capital demand and supply consequences for *four* of the bottom-up scenarios considered in this section. It is enlightening to inspect the capital impact and its various sources across the different stress-tests.

dimensional. It is precisely such multifaceted analysis of alternative bottom-up—and, of course, also top-down—scenarios that contributes to a deeper understanding of the vulnerabilities in one’s portfolio. No one scenario provides the full picture. It is rather through the judicious selection and comparison of various scenarios, constructed following different strategies, that stress-testing analysis adds value.

## 12.4 Wrapping Up

In the banking world over the last few decades, much (entirely warranted) emphasis has been placed on the notion of *know-your-customer* to help minimize money laundering activities and other financial crimes.<sup>55</sup> The core idea is that more knowledge about a phenomenon leads to improved understanding and management of associated risks. This excellent and uncontroversial point is readily extended to our stress-testing discussion in the form of the snappy, but entirely pertinent, catchphrase: *know your portfolio*.

Our risk-management models, as we’ve seen in previous chapters, provide us with a wealth of information. Much of it, however, is embedded in the through-the-cycle perspective with the assumption of a constant portfolio. Stress-testing expands our horizon by examining—from a rich array of perspectives—the repercussions of downgrade-related portfolio changes. Although broadly defined and difficult to organize, the principal contribution of stress-testing is portfolio knowledge. Better comprehension and classification of portfolio weaknesses and vulnerabilities

<sup>55</sup> There are many sources on this area, but Graham [15] provides an interesting introduction.

**Table 12.8** *A stress-testing menu*: The underlying table, like an *à la carte* menu, illustrates the seven alternative top-down and bottom-up stress-testing techniques treated in this chapter. While certainly non-exhaustive, hopefully it can help readers design and execute their own menus.

I. TOP-DOWN	I.1 Forward-looking adverse scenarios
	I.2 Historical backward-looking crisis outcomes
	I.3 Structured impulse-response-function motivated analysis
II. BOTTOM-UP	II.1 The power set of all downgrades
	II.2 Extreme (all in one credit class) portfolios
	II.3 The $n$ -notch global downgrade chestnut
	II.4 Classic portfolio-knowledge motivated scenarios
	II.5 Randomized catastrophic defaults

complement traditional probabilistic models and help stakeholders take better decisions.

Table 12.8 is given the last word. Organized like an *à la carte* menu from a fancy restaurant, it summarizes the *seven* alternative top-down and bottom-up stress-testing techniques examined in the previous sections. While certainly not an exhaustive list, it does provide a reasonable amount of choice. Hopefully, when combined with the ideas presented in preceding chapters, it can help the reader design and execute her own stress-testing menu.

## References

1. C. Albanese, S. Crépey, and S. Iabichino. Reverse Stress Testing. Social Science Research Network, February 2021.
2. J. R. Aragonés, C. Blanco, and K. Dowd. Incorporating stress tests into market risk modeling. *Institutional Investor*, pages 44–49, Spring 2001.
3. S. Asmussen and P. W. Glynn. *Stochastic Simulation: Algorithms and Analysis*. Springer-Verlag, Heidelberg, Germany, first edition, 2007. Stochastic Modelling and Applied Probability.
4. T. Bellini. *Stress Testing and Risk Integration in Banks: A Statistical Framework and Practical Software Guide*. Elsevier Inc., London, UK, 2017.
5. R. Bellman. *Adaptive Control Processes: A Guided Tour*. Princeton University Press, 1961.
6. BIS. Stress testing principles. Technical report, Bank for International Settlements, October 2018.
7. D. J. Bolder. A Proposed Stochastic Simulation Framework for the Government of Canada's Debt Strategy Problem. Bank of Canada: Working Paper 2003-10, April 2003.
8. D. J. Bolder. The Canadian Debt Strategy Model. Bank of Canada Review: Summer 2008, October 2008.
9. D. J. Bolder and S. Deeley. The Canadian Debt-Strategy Model: An Overview of the Principal Elements. Bank of Canada: Discussion Paper 2011-3, May 2011.
10. L. J. Christiano. Christopher A. Sims and Vector Autoregressions. *Scandinavian Journal of Economics*, 114(4):1082–1104, 2012.
11. D. A. Dickey and W. A. Fuller. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366):427–431, 1979.

12. G. S. Fishman. *Monte Carlo: Concepts, Algorithms, and Applications*. Springer-Verlag, 175 Fifth Avenue, New York, NY, 1995. Springer Series in Operations Research.
13. P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, Berlin, first edition, 2004.
14. G. H. Golub and C. F. V. Loan. *Matrix Computations*. The John Hopkins University Press, Baltimore, Maryland, 2012.
15. N. M. Graham. Know your customer: Good for compliance, great for client service. *Journal of Investment Compliance*, 4(3), 2003.
16. J. D. Hamilton. *Time Series Analysis*. Princeton University Press, Princeton, New Jersey, 1994.
17. J. W. Harris and H. Stocker. *Handbook of Mathematics and Computational Science*. Springer Verlag, New York, New York, 1998.
18. G. G. Judge, W. Griffiths, R. C. Hill, H. Lütkepohl, and T.-C. Lee. *The Theory and Practice of Econometrics*. John Wiley and Sons, New York, New York, second edition, 1985.
19. G. Koop, M. Pesaran, and S. M. Potter. Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1):119–147, 1996.
20. H. Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer, Heidelberg, Germany, 2005.
21. J. McCall. Genetic algorithms for modelling and optimisation. *Journal of Computational and Applied Mathematics*, 184(1):205–222, 2005.
22. Merriam-Webster. *Webster's Ninth New Collegiate Dictionary*. Merriam-Webster, 9th edition, 1991.
23. M. Ono, M. E. Chamie, M. Pavone, and B. Açikmese. Mixed strategy for constrained stochastic optimal control. *CoRR*, abs/1607.01478, 2016.
24. M. Pesaran and Y. Shin. Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58(1):17–29, 1998.
25. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Trumpington Street, Cambridge, second edition, 1992.
26. R. Rebonato. *Coherent Stress Testing: A Bayesian Approach to the Analysis of Financial Stress*. John Wiley & Sons, West Sussex, UK, 2010.
27. C. M. Reinhart and K. S. Rogoff. *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press, Princeton, New Jersey, 2009.
28. D. Rösch and H. Scheule. *Stress Testing for Financial Institutions: Applications, Regulations, and Techniques*. Risk Books, London, UK, 2008.
29. C. A. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980.
30. M. Tahan. *The Man Who Counted: A Collection of Mathematical Adventures*. W.W. Norton & Co., New York, NY, 1993.
31. N. N. Taleb. *The Black Swan: The Impact of the Highly Improbable*. Random House, New York, NY, 2007.
32. J. Tolkien. *The Hobbit*. Harper Collins, Hammersmith, London, 1937.
33. B. H. Yang. Point-in-time PD term structure models for multi-period scenario loss projection: Methodologies and implementations for IFRS 9 ECL and CCAR stress testing. MPRA Paper 76271, University Library of Munich, Germany, January 2017.