

Cambridge University Press

978-1-107-04811-9 - Portfolio Management Under Stress: A Bayesian-Net Approach to Coherent Asset Allocation

Riccardo Rebonato and Alexander Denev

Excerpt

[More information](#)

Part I

Our approach in its context

Cambridge University Press

978-1-107-04811-9 - Portfolio Management Under Stress: A Bayesian-Net Approach to Coherent Asset Allocation

Riccardo Rebonato and Alexander Denev

Excerpt

[More information](#)

Cambridge University Press

978-1-107-04811-9 - Portfolio Management Under Stress: A Bayesian-Net Approach to Coherent Asset Allocation

Riccardo Rebonato and Alexander Denev

Excerpt

[More information](#)

In the first part of this book we explain why taking ‘stress’ events into account is so important in portfolio construction. We argue that the possibility of non-normal events should always be taken into account in the construction of the portfolio, and not only in periods of crisis. We also provide a conceptual framework for dealing with exceptional events, based on causation rather than association.

Cambridge University Press

978-1-107-04811-9 - Portfolio Management Under Stress: A Bayesian-Net Approach to Coherent Asset Allocation

Riccardo Rebonato and Alexander Denev

Excerpt

[More information](#)

Cambridge University Press

978-1-107-04811-9 - Portfolio Management Under Stress: A Bayesian-Net Approach to Coherent Asset Allocation

Riccardo Rebonato and Alexander Denev

Excerpt

[More information](#)

1 How this book came about

[Under uncertainty] there is no scientific basis on which to form any calculable probability whatever. We simply don't know. Nevertheless, the necessity for action and for decision compels us as practical men to do our best to overlook this awkward fact and to behave exactly as we should if we had behind us . . . a series of prospective advantages and disadvantages, each multiplied by its appropriate probability waiting to be summed. *J M Keynes, 1937*

This book deals with asset allocation in the presence of stress events or user-specified scenarios. To arrive at the optimal allocation, we employ a classic optimization procedure – albeit one which is adapted to our needs. The tools employed to deal consistently and coherently with stress events and scenario analysis are Bayesian nets.

The idea of applying the Bayesian-net technology, recently introduced in Rebonato (2010a, b), Rebonato and Denev (2012) and Denev (2013) in the context of stress testing and asset allocation, seems a very straightforward one. So straightforward, indeed, that one may well wonder whether a 500+ page book is truly needed, especially given that two thirty-page articles are already available on the topic.

We decided that this book was indeed needed when we began using this technique in earnest in practical asset-allocation situations. We soon discovered that many slips are possible between the cup of a promising idea and the lips of real-life applications, and that only a thorough understanding of these intermediate steps can turn a promising idea into something really useful and practical.

The steps we refer to are not purely computational or, in a wider sense, ‘practical’ (although in the book we deal with quite a few of these). We have found that the praxis of actually obtaining asset allocation using the Bayesian-net methodology has also enriched our theoretical understanding of how it works, and has given us a better appreciation of what it can deliver – and of what it can’t. It has also shed light on, and enriched our understanding of, broader aspects of asset allocation in general.

In short, we would like in this book to share with the investment community a way of dealing with the problem of choosing an optimal portfolio composition when the probability of occurrence of stress events is non-negligible (that is to say, always). And we like to think that the way we propose to do this is novel, practical, theoretically justifiable, intuitive and, last but not least, rewarding to work with.

We do not consider our offering as *the* solution of the asset-allocation problem, but we hope that it can give a tool to be used alongside, say, the Black–Litterman prescription or the Michaud technique(s). Above all, we feel that the greatest virtue of the approach we propose does not lie in any one formula or numerical ‘trick’. Rather, we concur with General Eisenhower that, yes, the best-formulated plans will prove useless in the heat of the battle, but that the preparation itself will turn out to be priceless.¹

1.1 An outline of our approach

Diversification is at the heart of modern portfolio theory. Before Markowitz, asset allocation was a rather straightforward affair. As late as 1938, the then-standard reference work on financial investments by Williams (1938) implicitly recommended that investors should place all their wealth in the investment with the maximum expected return. This recommendation was not made in the get-rich-quick investment column of an evening paper, but in a book on investing published by no less august an institution than Harvard University Press. So, in the late 1930s, putting all your eggs in the most promising basket was still ‘best industry practice’.

The work by Markowitz (e.g. 1991 [1959], 1987) fundamentally changed all that because of his insights about diversification and risk-reward trade-offs. But the reach of his ideas spread further, as the concept of diversifiable and undiversifiable risk was to pave the way to building the Capital Asset Pricing Model (CAPM), and the subsequent developments in asset pricing. Diversification – and, more generally, the idea that to build a desirable portfolio it is essential to understand how asset prices move together, not just in isolation – has transformed finance for good.

Our work takes this for granted. If anything, we put diversification at the very heart of our approach. We are acutely aware of a problem, however. Our concerns are well expressed by the words of Mohamed A. El-Erian, PIMCO’s co-CEO, written in the wake of the 2007–9 financial crisis:

[Recent] developments will serve to further highlight the danger of . . . being overly committed to an historical policy portfolio whose rigid backward-looking characterization no longer corresponds to the realities of today and tomorrow . . .

[W]ith its conventional (or, to be more precise, reduced-form) analytical foundation now subject to some motion, it will become even more difficult to rely just on a traditional portfolio diversification as both necessary and sufficient to deliver high returns and mitigate risks. Diversification will remain necessary, *but a lot more attention will be devoted to the appropriate specification of tail scenarios and tail hedges . . .*²

These words clearly suggest that the investment community requires something more than a minor tinkering at the edges of the usual statistically based assessment

¹ ‘In preparing for battle I have always found that plans are useless, but planning is indispensable.’ Dwight D. Eisenhower, 34th President of the US.

² El-Erian (2010, p. 4, our emphasis).

1 How this book came about

7

diversification techniques. It requires something different. And something different is indeed what we offer in this book.

In a nutshell, this is our ‘philosophy’.

First, we believe that, at least ‘locally’, correlations are relatively stable (or, at least, statistically predictable) during normal market conditions.³ These local and relatively stable correlations (covariances, really) can therefore be profitably extracted using traditional statistical (frequentist) techniques.

Second, we believe that in conditions of market turmoil the codependence among changes in asset prices is radically altered. Along with many econophysicists, we do believe that extreme events are ‘in a class of their own’.

Third we have found little empirical support for the belief that extreme events leave a regular, invariant ‘signature’ in the distribution of joint returns. We believe instead that each crisis unfolds according to its own idiosyncratic dynamics.⁴ So, for instance, we do not think that the way asset prices moved together during, say, the Mexican peso crisis of 1994 carries a lot of information about how asset prices moved together in the occasion of the subprime crisis. More generally, we do not believe that what happened during any one particular crisis can *ex ante* be relied upon to tell us much about how prices will move together in the next.⁵

Fourth, we recognize that establishing these crisis-specific codependencies is very difficult. But we also believe that we can go a long way towards building a sound understanding of how asset prices may move in stress conditions by injecting *causal* information about how the world works. This causal information comes from our understanding, imperfect as it may be, of what causes what, and of which events should *today* be at the forefront of our attention. We know that water causes mud, and that mud does not cause water – or, in a financial context, we know that a fall in equity prices would cause equity-implied volatilities to spike, and not the other way around. When we adopt a purely correlation-based description we relinquish this precious information. This cannot be good. We feel that the task faced by asset managers is difficult enough without throwing away any useful piece of information they may have. Every nugget of information, however limited, should be cherished and ‘squeezed’ as much as possible, exactly *because* the task at hand is so difficult. Bayesian nets will be the tool we employ for the task.

³ What we really mean here is that in normal market conditions correlations can remain stable over periods of time of comparable length to the rebalancing horizon of most asset managers. See the discussion in Chapters 20 and 21.

⁴ We stress that, if the reader does not agree with our views on this matter, our approach is still applicable. We discuss this point in Chapter 18.

⁵ Of course, we do not deny that one can detect some crisis-specific regularities: asset prices, for instance, move ‘more’ (their volatilities increase) during crises. And, yes, some common factors of returns such as liquidity, which in normal market conditions are almost latent, take on a much greater role during periods of turmoil, thereby inducing a greater degree of codependence among assets which are normally weakly correlated. But, as we argue in what follows, even mantra-like statements like ‘in crises all correlations go to 1’ need strong qualifications and, taken at face value, offer little help in tail hedging and robust portfolio construction.

We discuss the regularity or otherwise of crisis patterns in Chapter 4 (see Section 4.2 in particular).

Fifth, we combine the statistical information that pertains to the normal market conditions with the information that we have extracted from our understanding of how the world may work *today*. Once this conditional joint distribution of returns has been obtained, we can move to the traditional asset-allocation procedures based, implicitly or explicitly, on utility maximization. Our approach can be seen as a natural extension of the research programme ushered in by Black and Litterman (1992): Black and Litterman ask the asset manager to provide subjective estimates of the first moments of the return distributions; later developments asked the asset manager for subjective inputs about the covariance matrix (see, e.g., Rebonato and Jaeckel 1999); what we require in our approach is for the portfolio manager to add her subjective assessment about the joint tail behaviour of asset prices.⁶

We realize that doing so directly is very difficult. This is why we ‘distil’ this complex information from simpler building blocks, linked together and made consistent via the Bayesian-net technology. We can summarize our approach by saying that we *elicit* ‘simple’ probabilities and *construct* ‘complex’ ones.

Sixth, there are two overriding principles, arguably the most important of all, that inform our work. The first is that we want our results to be intuitively understandable and challengeable by the intelligent-but-not-mathematically-versed professional. We believe that the time of black-boxes is rapidly running out for asset allocation systems – and, for that matter, for risk management in general. Yes, in our approach we employ relatively sophisticated mathematical techniques. However, if an asset manager cannot interrogate the output of her asset-allocation system, if she cannot carry out simple and transparent sensitivity analyses, if she cannot reconstruct in an intuitive way why the suggested allocation turns out to be like this rather than like that, then we really recommend that she should not trust the output of the model she uses. She should move to something simpler – something, that is, that she can *understand* and *trust*.

Our second overriding principle is the robustness of the output. We are ambitious but modest. We know that we are trying something very difficult, and we therefore want our final results to be forgiving, in the sense of not displaying a high sensitivity to small changes in the necessarily approximate inputs. To use a hackneyed expression, we want to be approximately right, not precisely wrong.

If the reader ‘buys into’ our philosophy, we believe that we can offer such an intuitively understandable, auditable and robust approach to asset allocation. Our method has the added advantage of treating ‘normal’ and ‘exceptional’ events on a consistent footing, and of allocating assets in a coherent manner.

Of course, asset managers have always been aware of what ‘lurks beneath’. They have typically tried to ‘defend’ their portfolios with the safety nets and parachutes provided by tail hedges. However, these have typically been added as afterthoughts to the allocation suggested by their stress-naive model. Doing so clearly destroys whatever optimality there was in the original allocation, as the parachutes were not included at

⁶ We have not defined what we mean by ‘tails’. We refer the reader to Chapters 4 and 5, where we deal with exceptional returns in general.

1 How this book came about

9

the outset in the universe of allocation assets. We propose instead a *coherent* way to combine all assets (including the protection hedges) in one single allocation framework.

1.2 Portfolio management as a process

This book is mainly about portfolio management and asset allocation. What is needed for the task? Should a manager rely on her intuition, or use a formalized method – of which Markowitz’s mean-variance optimization is an early and still prime example, and of which the suggestions in this book are just another and modest variation? Modifying the taxonomy in Markowitz (2010), one can list the following alternatives. The portfolio manager could

1. use the tried-and-tested mean-variance approach pioneered by Markowitz, possibly with some of the many ‘improvements’ that have been suggested;
2. use other measures of risk or return in a risk-return analysis;
3. determine the investor’s utility function explicitly and maximize its expected value;
4. use constraints and guidelines instead of optimizing;
5. proceed intuitively.

From the point of view of logical coherence, the third alternative is arguably the best. However, as Markowitz (2010) points out, it is not without its drawbacks: an inappropriate utility function may be chosen for analytical tractability; assuming time invariance of preferences may have serious drawbacks; the process may end up being too much of a black-box; it may be difficult to model the nexus of agency relationships in a real investment process (think of the interactions between the sponsor, the trustees, and the present and future retirees in the case of a pension fund).⁷

As for the question of whether it is better to use constraints on their own or as part of a risk-return analysis, Markowitz, as usual, makes the case very clearly.

The answer is that rarely, if ever, do the intuitively desired constraints completely specify one and only one portfolio. (If the portfolio analyst tried to do so, he or she would probably overspecify the portfolio and end up with a system with no feasible solution.) So then, how is the analyst to choose among the many portfolios that meet the constraints? . . . If the choice is made using some kind of decision rule (such as, ‘Try to get as close to equal weighting as possible, subject to the constraints’), either the procedure employed would be ‘closet’ risk-return (e.g., equal weighting subject to constraints is the same as mean-variance with all assets having the same means, variances, and covariances), or it is probably a procedure unrelated to the maximization of expected utility . . .

Concerning the [fifth] alternative to proceed intuitively, the behavioral finance literature is full of ways that seat-of-the-pants investing leads to irrational choices. For example, Barber & Odean (2000) show that investors who fancy that they can outperform the market by active trading typically make a lot of money for their brokers and little, if any, for themselves.

So, what is the asset or portfolio manager to do? Judging from our experience, and from the insights we have gathered by talking to professionals who do manage

⁷ See in this respect the discussion in Part VIII of this book.

investment funds (and who very often do use variations of the Markowitz mean-variance approach) what they need is a framework that allows them to bring their financial understanding into play, but does so in a coherent, organized and ‘auditable’ manner. What do we mean by this?

Even when these asset managers use mean-variance analysis, or the Black–Litterman approach, or any of the many variants that have been proposed, they don’t really ‘believe in’ mean-variance, or Black–Litterman, or the asset allocation approach of choice in the same way that a physicist ‘believes in’, say, quantum mechanics. Rather, they have the need for a coherent (as opposed to haphazard) way to organize their intuition, make the result ‘inspectable’ and to point at a well-defined procedure to go from their ‘hunches’ to investment actions. They need (and cherish) intuition, but they also need a process. Quoting Markowitz again, ‘[d]uring a conference discussion session, Peter Bernstein described the haphazard way in which institutional portfolios were formed in the 1950s and earlier. He concluded with the statement, “Now you have a process.” [Markowitz’s] thought at the time was, “Now I understand what I started.”’

This is a simple but powerful insight: the asset management profession needs intuitive tools that can be ‘understood’ even leaving the maths aside (mean-variance is, in this respect, probably unbeatable); but they also need a logical framework by means of which this intuition can be organized into an auditable process. This is what Markowitz started. Arguably his legacy in this area is probably more important, and may end up being more enduring, than his legacy for the specific contribution to asset allocation – impressive as this is.

The work presented in this book tries (with modesty) to offer a complementary approach. The intellectual origin of the work is still firmly in the Markowitz camp: diversification, trade-off between risk and return, probabilistic (although not necessarily frequentist) approach: all these building blocks are still there. It proposes some (important) variations on the theme, but tries to keep in mind that what the asset or portfolio manager needs is not ‘the truth’, but a logically sound, intuitively appealing and logically ‘inspectable’ investment process.

1.3 Plan of the book

The book has been organized in a modular fashion, with each module reasonably autonomous and self-contained. Like every author, we hope that the reader will go through every single page of our book, from the acknowledgements to the last entry of the index. However, we have structured our work in such a way that several shortcuts can be taken. In this section we therefore outline the structure of the book first, and then propose several suggested itineraries, from the scenic historic route to the equivalent of the high-speed option of ‘taking the highway’.

The structure of the book is as follows.

Part I spells out our philosophy, and places it in its context. In particular we stress the difference between correlation (which is an association-related concept) and causation. We place causal models at the heart of our approach.