SOCIETY OF ACTUARIES

Investment Section

INVESTMENT FALLACIES 2014

Mathematical modelling of social phenomena

By Nicholas John Macleod

Introduction

The application of mathematical models to social situations is seen by some as a natural extension of their use in physical science. But there are crucial differences. For example, the lack of controlled and repeatable experimentation stands in sharp contrast to the situation in physics, where theories make precise predictions that can be tested and falsified. A physical theory whose predictions cannot be verified by experiment is normally discarded, but this process of selection is largely absent from social science.

In part that's because the objects of interest in social science are not inanimate particles whose behavior can be expressed in terms of simple laws. They are people - with individual thoughts and the freedom to adopt individual actions.

The normal approach to individuality is to fall back on statistical aggregation. In the classic example where country fair-goers are asked to guess the weight of a pig, the average guess comes out strikingly close to the animal's actual weight. That is a triumph of statistics, but the situation is atypical in that the fair-goers are given no information with respect to earlier guesses. In most real social situations the information flow between the participants is a critical determinant of system behavior.

As might be expected, modelling information flow takes us to a new level of complexity, and comes with its own problems. As models become more complex, it becomes increasingly difficult to distinguish between genuine properties of the system, and properties that arise from the particular assumptions of the model.

Does this mean that mathematical modelling in a social context is a waste of time? I don't think it does, but I do believe that more attention needs to be given to the fundamental differences between social and scientific modelling. In particular, the failure to reject social models whose outcomes differ significantly from their intentions or predictions undermines any claim to scientific method and greatly impedes progress. While we cannot set such strict criteria for acceptance as we do in proper science, that doesn't mean that we should set none.

This is too large a subject to be dealt with in a short essay, so I will try just to highlight some general issues by means of an example drawn from finance. My approach is to set up a straw man – something I've called the Old Model – and to demonstrate that a simple extension of its underlying assumptions leads to distinctly different prescriptions. I don't insist on the details of the Old Model; some people may see it as a caricature, although most will recognize at least some of the elements of conventional investment theory. My point has to do with the need to test the robustness of a model's prescriptions. If a modest change in what are necessarily uncertain or approximate assumptions leads to a radical change in indicated action, the assumptions must be carefully reviewed for dependability.

Old Model

The conventional approach to asset allocation is based on the following ideas[1]:

- There is essentially one state of the world;
- Return variation is fluctuation that reflects new information that by definition cannot be predicted;
- Volatility (the amplitude of fluctuation) measures risk;
- Reduction of volatility depends on low correlation

1 Here, and in what follows, I'm tacitly assuming that the assets are supporting a funding program.

so finding uncorrelated assets is the key to effective diversification;
- There is a long run positive relationship between expected return and risk;
- You can't time the market.

This model supports the general practice of holding fairly fixed allocations across as wide a range of asset classes as possible. Since return variation is intrinsically unpredictable, there is no point trying to time the market. Instead you should combine the positive risk-return relationship with the volatility-reducing properties of low correlation to maximize the likelihood of achieving your required return. The process can be thought of as starting with the asset with the highest expected return, and progressively shrinking portfolio volatility by adding imperfectly-correlated assets. The volatility reduction comes at a cost in expected return, and relatively few additions account for most of the reduction, so there is a natural point at which the portfolio has the maximum chance of achieving the required return. That's your optimal portfolio and it will remain pretty stable throughout the life of your funding program.

Variation around the long term average return of the portfolio also declines with time, in exactly the same way that it declines with the addition of uncorrelated assets, except that here there's no associated give-up of expected return.

As you expand the range of assets and extend the length of the investment period, return converges around a favorable long run average and the likelihood of achieving the target return over a normal funding period increases to near-certainty.

The idea of trading expected return for reduced risk in order to increase the likelihood of achieving the target return seems to make sense, and the statistical elements of the approach appear to be sound, so why hasn't it worked in practice?

It's helpful here to contrast the statistical modelling used for liability estimation with the use of statistical models on the asset side of the equation.

- Mortality is a natural process that conforms to regular biological and statistical laws, so making allowance for increases in longevity, the statistics of past mortality tend to be reliable indicators of the incidence of future mortality. There are other variables that require estimation (for example, for a final earnings pension scheme, the level of benefits will depend upon the recipients' final salaries, which are not known today) but here again, past experience generally provides a fairly dependable basis for the estimation of averages.

- Finance, on the other hand, is a social activity with occasional regularities, but no fundamental laws. Where there are no well-established laws, models must be justified by their consistency with real world experience. But the asset allocation methodology described above is based on theoretical assumptions about the way markets *should* work, rather than on experience of how they *do* work. Since the prescriptions of the Old Model have not led to the anticipated outcomes, it is falsified by application.

<u>What do I mean by "falsified"? All models are simplistic, and therefore false.</u>

What does it mean to say that one model is better than another when the underlying reality is infinitely more complex than either of them? I think the answer has to do with the models' qualitative prescriptions. The Old Model essentially recommends fixed allocations to a "diversified" (= non-correlated) set of assets. But what if we extend the model to a two-state system? Let's say there are bear markets and bull markets, and that each of those states is sticky, or persistent to some degree. In other words, suppose that market behavior is more like the weather, in that things can change, and when they do, they tend to remain in the new state for some time[2].

This model is only one level more complex than the Old Model, but its implications are entirely different. For example:

- Shifts from bull to bear markets pose a much greater risk than local fluctuation.
- Assets that are statistically uncorrelated may well respond in the same way to environmental shift. As a result, they do not diversify each other.
- It makes no sense to stick with fixed allocations in the face of market change any more than it does to stick with the same set of clothes in all weather conditions. Persistence means that when the world changes, you have to do something about it; risk management should be dynamic, not passive.

Once we agree that market change is real and persistent, the pillars of the Old Model collapse. Risk is not simply volatility, low correlation does not guarantee effective diversification, and fixed allocations are not the best way to fund liabilities. Going from a single state model to a two-state model changes everything. It also explains some of the mysteries of the Old Model; why do outsized losses among assets that were not previously correlated often occur at the same time, for example?[3]

When we generalize the two-state model to cover multiple states, there is no *qualitative* difference in its implications and recommendations; all of the points above still apply. So while a two-state model is obviously far too simple to describe the real behavior of markets in any detail, its qualitative prescriptions are not the by-products of over-simplification.

It's also easy to see how a two-state model can be extended to a multi-state model without introducing any new concepts: a multi-state model is just a string of inter-connected two-state models, so, while it's more complex, it's not fundamentally different. That isn't the case in going from a single-state model to a two-state model, where we have to bring in new mechanisms like **transition** and **persistence** that don't appear in the one-state model. And it's those mechanisms that explain the correlation dynamics and other things that the Old Model can't.

Statistical analysis of recent market activity[4] suggests very strongly that market conditions are persistent. Who can doubt that 2000–2003 was a very different environment from 2003–2007, and that 2008 was different again?

---

2 This is not to argue that the markets are as simple as the weather. We know what causes the seasons and we understand fluid dynamics, but weather is still difficult to predict. We have a much more limited understanding of financial markets.

3 This is the "volatility spiking and correlations going to 1" phenomenon, that's unexplained in the Old Model, but perfectly natural in a multi-state model.

4 And perhaps more important, professional experience and common sense.

The evidence indicates that, far from being as simple as the Old Model suggests, reality is better represented by a highly dynamic, but persistent, multi-state model in which the states are not fixed, there is potentially an unlimited number of them, the degree of persistence – i.e., the stability of market conditions - varies irregularly, and so on. But even the two-state model—the simplest possible multi-state model—is significantly more complex than the Old Model, and qualitatively quite different. It explains phenomena[5] (outsized losses, coincident losses among uncorrelated assets, etc.) that are not just theoretical mysteries under the Old Model, but real-world events that can damage or destroy a funding program. Following the prescriptions of a model that doesn't even recognize their existence is not what's normally thought of as prudent.

Nicholas John Macleod, ASA, C.Math., FIMA, is a Principal, Bespoke Financial Modelling, Jersey, Channel Islands. He can be reached at *nicholasjohnmacleod@gmail.com*.

[5] It might be more accurate to say that these phenomena arise from dynamics that are built into the structure of a multi-state model.