

## Economic Scenario Generators and Solvency II

---

### Abstract of the Discussion

[Institute of Actuaries, 23 November 2009]

#### Contact

Elliot Varnell, MSc., F.I.A., E-mail: [elliott.varnell@kpmg.co.uk](mailto:elliott.varnell@kpmg.co.uk)

This abstract relates to the following paper:

Varnell, E.M. Economic Scenario Generators and Solvency II. *British Actuarial Journal*, doi:10.1017/S1357321711000079

**Mr E. M. Varnell, F.I.A.** (introducing the paper): First, I address the need for the paper; then I will talk about market consistent valuation, which is a very hot topic at the moment. I will discuss where ESGs fit in with solvency capital requirements and, finally, I touch on ESG governance.

I have been involved in helping insurers with ESG models for a number of years. ESG models are widely used for many different applications in the UK insurance sector, especially in the life insurance sector. There are various commercial solutions available on the market.

However, one of the points I make in the paper is about the failure of ESG models to pick up the micro-market features of markets, such as the premia for illiquidity or tax effects.

A practical issue with ESG models is dealing with fitting the ESG model to the company's view. Discussions with companies have focussed around who in the company owns the economic assumptions and to what extent these differ from those embedded in the ESG.

Where views differ, making the ESG model do what you want it to do can be challenging. This is particularly the case where a company may be using a different model for assumption setting to that of the ESG provider.

Another key area is getting the ESG model embedded and understood by the company, including getting it embedded within its governance structure.

Solvency II is a very significant area where ESG models are going to be used, not just in the UK, but more widely in Europe and beyond. Indeed, in South Africa ESGs are used for market consistent valuation and regulatory capital requirements under the PGN-110 actuarial guidance.

ESG models have also been widely used in financial reporting, such as European Embedded Value (EEV) and Market Consistent Embedded Value (MCEV), asset liability management and dynamic hedging.

ESG models are used in product design. Product design was one of the first places that insurance companies started using ESG models. They were used for designing guaranteed products and trying

to understand what the value of the guarantees was that they were writing. ESG models have been used to a lesser extent for product communication, that is trying to explain some of the risks and rewards associated with insurance products to the retail market.

ESG models cover two different types of models: risk-neutral models and real-world models. Risk-neutral models are used for market consistent valuations. Their sole objective is to re-create observed market prices. The idea is that if you can re-create market prices you can use the same model to create a market consistent value of your liabilities.

There are also real-world ESG models. Their key objective is to capture the true dynamics of market prices and are used to project the market consistent balance sheet and help insurers understand how much economic capital they need.

One of the sections in the paper is about the validity of market consistent valuation. The reason I put this section in was to address recent criticism of market consistent valuation.

I have cited frequent misconceptions about market consistent valuations. These are the sort of questions that I have been posed over the last five years. The first is whether an arbitrage-free ESG model will, by itself, give a market consistent valuation. It will not. If you were to take an arbitrage-free model from a textbook and code it up and use it to value a liability it would not of itself give you what I consider to be a market consistent valuation. This leads to the second misconception in the paper: that an ESG model calibrated to deep and liquid market data will of itself not give a market consistent valuation. We have to combine a model with arbitrage-free dynamics and then calibrate it to market data to give a market consistent valuation.

It is possible to construct models that are not arbitrage-free and calibrate those models to market data. It is not always obvious that the resulting model is not market consistent and therefore care is needed in how the model is used. It can be easy to convince yourself that you have an arbitrage-free model when in fact you have not.

The third point is that market consistent valuation does not necessarily give the right valuation. I put this in because there is sometimes a perception that a market consistent valuation is a holy grail with a definitive truth. This is not the case. A market consistent valuation has to be caveated with all the assumptions that have gone into that model. Those are assumptions about the model that has been used and the calibration of the parameters of the model.

Another misconception is that market consistent valuation gives the amount a third party will pay for a business. We find in practice that intangible assets and intangible liabilities are often significant to the valuation of a transaction. The way in which these are calculated is typically much less sophisticated and more subjective than the way an ESG model would value liabilities. So it is important to be aware that a transaction price would not necessarily be the same as a market consistent valuation of the business.

Finally, having criticised market consistent valuation in the last four points, I am certainly not calling for a return to discounted cash-flow style calculations previously used. Neither do I subscribe to the view that market consistent valuation is no more objective than traditional discount cash-flow technique using long-term subjective rates of return.

I make the argument that despite its flaws, which we need to recognise, market consistent valuation builds up a valuation using very clear assumptions about what economic theory is being applied and what assumptions and data have been used. In that sense, market consistent valuation using an ESG is a better way of coming up with a valuation than some of the techniques that we might have used in the past.

A number of criticisms have been levelled at market consistent valuations. One of the strongest criticisms has been that it creates a pro-cyclical capital regime. What we mean by that is that, in a crisis, asset values may fall at the same time as liability values rise sharply – for example, as interest rates fall and as equity implied volatilities rise. The result is that an insurer's own funds get squeezed from both sides and it ends up being in a weak capital position relatively quickly.

A second criticism is that what we are actually doing when we are doing market consistent valuation is importing marginal traded prices on distressed assets straight onto the insurer's balance sheet.

A third criticism is that the squeezing of insurers' own funds leads to asset fire sales. So insurers would need to get rid of risky assets in the market at much the same time and further depress prices through a negative feedback loop.

Another criticism of market consistent valuation using ESG models is that they do not reflect liquidity premiums and other micro-market features that we might like to see captured. When building an ESG model, we make a lot of assumptions, such as markets are deep and liquid, transaction costs are zero, and often taxes are zero.

A point worth noting is that when we want to include liquidity premiums in an ESG model we cannot use the ESG model itself to determine what that liquidity premium should be. We use some other model external to our ESG model in order to infer what liquidity premium should apply.

Volkswagen share price illustrates the issue with marginal pricing. About 18 months ago, for a short period of time, Volkswagen became the most valuable company in the world. This event acutely illustrates that the quoted market price is just a marginal price.

This was an extreme example of marginal pricing but the point is to illustrate what we are doing when using marginal prices to calibrate ESG models. We are taking prices from the market and we need to remember that those are marginal prices and that there can be technical irregularities in the market, as indeed there were in the Volkswagen case. We need to remember that the marginal prices can give us a misleading impression of the value of our liabilities, of our assets, and therefore of our company.

Consider also a volatility index over time. A graph of, say, the average of the short term implied volatility derived from option prices on the S&P 500 in the United States since 1990 has spikes in volatility when markets are in crisis. There is a spike in 1999, and another large spike during the crisis at the tail end of 2008. Insurers will start to consider if they should be importing these sorts of volatility spikes into their balance sheet assessments.

If you happen to have a volatility spike like this at the end of the financial year, it could lead to an explosion of the value of options and guarantees on an insurer's balance sheet. We should consider

whether there is some illiquidity associated with the implied volatility, and therefore whether we should perhaps ignore these very sharp spikes in implied volatility.

Another topic I cover in the paper is why we use ESGs at all in deriving market valuations as there are also closed-form solutions available. There are a couple of references in the paper to sources of closed-form solutions. The formulae available that are analytically tractable are quite simple and often too simple to capture all the dynamics and complexity of life insurance liabilities.

The underlying dynamics of the models can also be too simple. We often have to assume geometric Brownian motion in order to make the mathematics tractable.

It is worth mentioning replicating portfolios as well. A lot of you are probably investigating whether you should be using them. Generally these are not used for what I would call primary market consistent valuations, that is coming up with the balance sheet valuations at time zero. What they are very good for is for coming up with fast recalibrations, for example, understanding how the balance sheet valuation may change over the next 12 months in lots of different scenarios.

That brings us back to stochastic ESGs and when they are best used. Stochastic ESG models are really useful when the value of a policy depends on how a market has moved over the course of a policy rather than where the market ended up at the end of the policy. This sort of policy has path-dependency and for this sort of valuation stochastic ESG models come into their own.

Stochastic ESG models are also particularly useful when a policy depends on lots of underlying variables. If we have a policy whose payouts are just based on an equity index, it is easier to come up with a closed-form solution. If we have a policy based on lots of equity indices, lots of different bond prices and perhaps property indices too, we have a high-dimensional valuation to which stochastic ESG models are well adapted. Stochastic ESG models are also particularly useful when there are feedback loops, such as policyholder actions or management actions.

For example, management actions in response to the solvency of the company lend themselves much better to the stochastic approach that ESGs use to generate a market consistent valuation.

Finally for market consistent valuation, I will talk about the market consistent balance sheet and where we tend to apply ESGs and where we do not. ESGs tend to lend themselves naturally to the valuation of with-profits policies as with-profits policies meet all of the criteria set out above: high dimensionality, path-dependency and feedback loops. Continental participating products and variable annuity products are also well disposed to stochastic valuation.

ESGs lend themselves less to some of the other products that insurers sell, such as pension products, unit-linked products, general insurance products and standard protection products like term assurance.

Asset liability coherence is not something that is discussed but it is something we are going to have to think more about with Solvency II. When you calculate a liability using an ESG model, there are a set of assumptions about the model and its parameterisation which underpin the valuation. If we were to look at the asset side of the balance sheet we might have a derivative and typically we are going to get that price directly from the market. We would not use the ESG to value the derivative. If we were to use the ESG to value the derivative we might find a different price to the market

price simply because of approximations that are present in the ESG or in the assumptions that have been used in the calibration.

It is not obvious whether we set out the asset side of a balance sheet to be consistent with the ESG or whether we make it exactly consistent with the market. Either way we are going to have some sort of discrepancy in the market value balance sheet.

I also wanted to highlight that CEIOPS Paper CP40 uses AAA government debt as the risk-free rate and therefore potentially introduces a mis-match between the way in which we value the liabilities under Solvency II and the way in which we value the sort of assets that would be used to hedge those liabilities on the asset side of the balance sheet. These assets are valued using swap rates.

The next aspect to consider is solvency capital requirements and where the ESGs tend to get used in these calculations. When we talk about the solvency capital requirement, we are mainly considering the use of real-world ESG models. Real-world ESGs for projecting economic balance sheets will almost always have risk premiums. They will also have slightly different volatilities to the market consistent ESG models. The volatilities will be realistic as opposed to implied volatilities.

Real-world models are also the models where real world dynamic features such as tail correlations are modelled. By tail correlation I mean the situation where assets tend to fall in sympathy to each other – in other words are more correlated – when there is a sharp down move in markets. Real-world models also tend to capture fat tailed distributions.

The places where we might use an ESG for calculating capital requirements are, for example, in the standard formulae. We might have to build a calibration of our market consistent ESG under univariate stresses. That in itself can create quite a few challenges. For those of you who have tried this it will probably not have escaped your notice that when you change one part of an ESG calibration, another part of the calibration changes. It is not untypical that a stress to the swaption implied volatility leads to a significant change in the correlation between equities and interest rates. Moves of the order of 20% are not untypical. This is just one of the things to bear in mind when considering how to document your ESG under Solvency II.

There are a variety of different techniques for calculating sensitivities of market consistent balance sheets to the underlying risk drivers: examples include replicating portfolios, nested stochastic approaches and curve fitting techniques.

When using a partial internal model it is possible to use a common ESG across various business units. One of the interesting things to consider is how the correlations at the very top of the structure would need to change.

One of the main aspects of running internal models is how to meet the seven internal model tests: the use test, the documentation test, the statistical quality test, the calibration test, the validation test, the external models test and the profit and loss attribution test.

Firstly, let us consider the use test. It could be said that the entire paper was really about applying the use test to an ESG. There are two particular aspects of the use test.

The first is the potential trade-off you have between the use test and the statistical quality test. On one hand the use test challenges the insurer to ensure that the ESG and the internal model are understood within the company and therefore used for decision-making. That, to some extent, necessitates something that is not too complicated in order that there are enough people in the organisation who are able to understand the model and trust the results of the model.

On the other hand, you have a statistical quality test to ensure that the model gives an accurate picture of the risks faced. For an ESG model that potentially means building in additional features to the model. So there is going to be a natural trade-off with insurers needing to strike a balance between having a model which is simple enough to understand in business and to use for decision-making, but is accurate enough to satisfy the regulator that they are capturing the dynamics of markets.

Perhaps the important point to make is that the use test requires that there should be constant pressure to improve the model. One could imagine a starting position might be the current ESG but that a process would be required around the ESG to constantly challenge whether or not it is giving sufficiently accurate results and to gradually improve the ESG model or its calibration each year.

An aspect of the validation test worth considering is the reverse stress test. This is a scenario (or set of scenarios) that would lead the insurer to ruin. Some insurers have approached this by picking a single scenario using a one time step approach. In other words, they derive a single combined stress test including stresses on interest rates, equity prices, property values, etc that would lead the insurer to the point of ruin.

An ESG can, on the other hand, provide a richer approach to reverse stress testing by helping insurers understand the path to ruin rather than the stress to ruin. One of the important lessons of the banking crisis of 2008 was that risks build on each other and that a stress in one risk factor can lead to stress in another risk factor later. The only way of picking up that from a model is to have a stochastic approach which projects a variety of risk factors and allows exploration of the many different paths to ruin. This helps understanding of why these scenarios led to ruin.

An important aspect of statistical quality is data. The three tests around data are as follows: Is it accurate? Is it complete? Is it appropriate?

A typical consideration in ESG calibration is whether to use all the information available or some subset of the data. If we consider a market consistent ESG not all of the option prices available can be used because it might require quite a complex ESG to use all the available prices. There is an interesting decision as to what would be considered complete data.

The validation test introduces the idea of back-testing. As an ESG modeller, I have frequently been asked whether the ESG models would have predicted the events of 2008. The answer is typically that the model did not but this begs the question of how one should back-test an ESG model I discuss this in the paper. For internal models in general, back-testing in the same way that banks do back-testing would involve going back each year, running a projection of the distribution of excess own funds, and working out where the actual excess own funds at the end of the year appear in the distribution. This would be done for a large enough sample that it could be statistically inferred whether or not the distribution was reasonable.

Banks can do this because they have a relatively short risk horizon – typically a few days. Therefore they can get a lot of data and come up with a robust estimate of the quality of their VaR models. For insurers, who need to use a one-year time horizon, there are not many independent one-year observations to use. Typically data systems would allow 10–20 observations at best – not enough to make a reliable statistical estimate of the 99.5<sup>th</sup> percentile.

Back-testing for ESG models can also be quite challenging. You cannot go back many years before one starts running out of some of the key data items that we use in our ESG calibrations.

The paper also discusses the calibration test. If you happen to use a risk measure which is not the standard Solvency II risk measure it is not easy to demonstrate how this equates to a one-year 99.5% calibration standard. However an ESG model can (at least for the risk factors included within it) provide a reasonable way of demonstrating equivalence between the two different risk measures.

CEIOP consultation paper CP69 is interesting because it is the only place in the Level 2 guidance where a stochastic stress test is applied. There is a helpful graph in CP69.

The intention is to provide a dampener on the equity market stress test. The green line in the graph in CP69 is a reverse engineered example of how the adjusted equity stress should work. What is illustrated is that, despite giving a low stress test at the start of the financial crisis in September 2008, between June and August 2009 the stress test rose rapidly from –35% to –55%.

The reason for highlighting this is that it raises interesting questions about the calibration of an internal model. You might have a process in place to derive a 1 in 200 stress test. It would be interesting to ask if your internal model stress test for the equity index (perhaps derived from an ESG) will be benchmarked against this dynamic stress test which is going to be used for the SCR standard formula.

In the paper I have listed the steps that need to be undertaken by way of ESG documentation. I have tried to break the approach into various components of documentation.

Firstly, the methodology component: there is a need to document the mathematical basis on which the ESG is based. This would be a document explaining why the theory underlying the ESG model is valid. This might include proof that the models are arbitrage-free. It may be that this documentation is largely taken from the textbooks that most ESG models are based on.

The second component is more difficult and is the empirical basis of the ESG model. This would include evidence that the way the model has been designed is actually capturing the true dynamics of the market that you are trying to reproduce.

The third component is assumption setting. This would include the key assumptions used in the ESG calibration, for example, correlation, volatilities and risk premia.

The fourth point is the application of expert judgment. Expert judgment involves documenting all the areas where subjectivity (through the use of expert judgement) has been applied. It also involves demonstrating the judgment was made by an appropriately qualified individual. Perhaps most important of all is documenting where the model does not work. This is probably the area

where most work has to be undertaken. Such a document might include the economic circumstances under which you believe that the ESG model would fail. This allows the insurer to know in advance when the model doesn't work so that action can be taken or allowances made.

The documentation should also cover the formulae implemented, the parameters used, the methods for estimating parameters, the data policy and the source code.

Finally, I have included a section about ESG model governance and how the governance in the ESG could be constructed. CEIOPS paper CP33 talks about governance from the regulator's point of view and sets out what CEIOPS see as governance roles within an insurer.

CP33 sets out an actuarial function, a risk management function, internal audit and internal control. Solvency II sees the actuarial risk function as being a role covering the technical provisions. From my perspective it is difficult to draw a line between the actuarial and risk management functions in many insurers – especially from an ESG perspective. Unlike Solvency II, I see a role for a distinct economist function which determines the economic basis that the company uses.

I also see a lot of governance connections with the finance function and the sales and marketing function. I see these functions being key consumers of what an ESG model produces.

I divide functions into producer functions and consumer functions. On the producer side there is an economist function which produces the economic basis. This sets out, for example, the company's view of future volatility, of interest rates and equity markets. There is the actuarial and risk management function which is typically the owner of the ESG model. This function contains the expertise on how the model is operated and its design.

I also include the IT function as well because many of the companies I have spoken to are struggling with how they are going to cope with the significant requirements of Solvency II. It appears to me that getting the IT side right is going to be of critical importance to the success of Solvency II.

Regarding the consumer functions, we first consider the senior management and the board who will be consumers of (high level) information on the ESG model. The finance function will typically be a consumer of the ESG calibration report for its financial reporting. The sales and marketing function is where the risk is manufactured but also where the value of the insurance company is often created.

It might be helpful when considering the key relationships of an ESG governance structure to begin with financial risk reporting. The idea is that we have a common function which will come up with this economic basis, the economist function. The economic basis is agreed in discussion with the senior management. Ultimately it has got to be the senior management or the board which sign off what the insurance company believes are the economic prospects for the markets in which the company operates.

The actuarial and risk management function gets involved because in trying to impose the economists view of the future on the model they will likely come up against limitations in the fitting of the model. Therefore there needs to be a feedback loop between the actuarial and risk function and the economist function to find a common position which the economist function are happy with and which can be fitted to the ESG model.



The second relationship I want to highlight is the risk and reward manufacturing process. In the scenario we have, the economist function produces the economic basis and agrees this with the senior management. Meanwhile the sales and marketing function, perhaps using their own models, would decide what they could manufacture and what would be profit-making for the company.

I would expect the sales and the marketing function to have an interaction with the risk and actuarial function to check whether or not what they are proposing to produce would actually be profitable (after risk-adjusting). The actuarial and risk management function would report their findings on the sales and marketing proposals to the senior management. We could expect that there would be an interaction between senior management and sales and marketing as well. In particular there would be discussion over what sales and marketing see happening in the market and what sort of products would sell.

**Mr I. C. Marshall, F.I.A.** (opening the discussion): With all the things that are going on with Solvency II at the moment, it is timely to look in detail at a very important part of the model so thank you for that.

The first thing I want to cover is validation. The author made a point of contrasting validation and the use test. Validation goes hand-in-hand with the use test. The reason I say this is that the main purpose of validation is really to get comfort that the model you are going to be using is right. No senior management, no board, would use their model if it has not been validated properly. One way of looking at the use test is that it is a part of validation; only if the company actually uses the model does it gain comfort in the results it is producing.

Another point about validation is around back-testing. This is a more general point than that addressed in the paper. There has been quite a lot of comment that there is not always enough data to properly do back-testing. If you are looking at a 1 in 200 year event then you obviously need lots of data to see whether your model is right. Back-testing is one tool used within validating. Back-testing does not provide enough comfort that the model is indeed correct, so we should be looking at using additional measures or additional tools alongside back-testing.

The final point is one that was also touched on by the author and is that there is a lot of expert judgement within ESGs. Even on the market-consistent side, there is still expert judgement relating to the fact that market data is not actually always available, for example, data on equity volatility over a long period of time. In addition, the exact details of the models used is substantially determined by expert judgement.

**Mr O. J. Lockwood, F.I.A.:** I want to focus on some issues that have a bearing on how ESG models are used currently to meet the internal model requirements for Solvency II.

My first point relates to statistical quality. In my experience it is usual for an ESG model run for 1,000 or more scenarios for 40 or more years to produce a not insignificant number of extreme scenarios, with inflation and interest rates running into hundreds of percentages. My question would be whether the assumptions of the stochastic model that utilises these economic scenarios are consistent with the way the business would be managed in such a situation, and also with how policyholders would behave in such a situation. For example, smoothing rules stating that with-profits bonus rates can change by no more than, say, 15% over a six-month period clearly do not make sense in an environment with interest rates of hundreds of percentages.

Turning to documentation: I recently set up a spreadsheet within my company aimed at reproducing the calculations of an externally provided ESG model. The extent and nature of the difficulties I had doing this reveal something about how close the documentation of the methodology supplied by the provider was to what will be required for an internal model under Solvency II. In general, I found that each individual component of the model was clearly documented, but the documents were often difficult to find as they were situated in many different locations on the provider's website, and I had to ask the provider a significant number of questions. It is clear from this that the work required to meet the Solvency II documentation standards should not be underestimated.

Turning now to the validation and to the use test. In my experience, the areas where effective validation procedures for ESG models tend to be in place are in reproducing the market prices of the underlying assets and simple options on them, and ensuring that sufficient scenarios have been run to reduce sampling error to an acceptable level. However, a key area where more validation work could usefully be done is in reconciling stochastic with deterministic models. A company is likely to have difficulty meeting the use test if those elements of information which are supplied to decision-makers on the basis of deterministic models are pointing in a different direction from those elements supplied on the basis of stochastic models and no explanation is given of the differences. It could be that the differences are genuinely due to the time value of options embedded in the business. However, my experience suggests that a company is unlikely to be able to demonstrate that the differences are genuinely time value. In particular, it is likely that deterministic and stochastic modelling work will be carried out on different systems, with little understanding of how the deterministic model differs from a deterministic run of the ESG that feeds the stochastic model. It is also likely that system constraints will mean it is not possible to perform a deterministic run of the ESG – for example, if your ESG has a discrete random variable representing how many of a portfolio of corporate bonds default in each time period, then that random number of defaults would need to be replaced by its expected value in a deterministic calculation and the ESG model might not have the functionality to calculate that expected value. Resolving this issue will therefore require work from both insurance companies and external ESG providers.

**Mr J. A. Jenkins, F.I.A.:** I agree with a lot in the paper, such as an ESG is essential for valuing business with guarantees. The paper mentions replicating portfolios, but they are not sufficiently well-developed as a technique to be used for most companies – certainly not in the timeframe that we require.

On governance, it is still very difficult to get boards to engage on ESGs. I have attended many board meetings and audit committee meetings, sometimes as an adviser and sometimes as an auditor, and observed this engagement difficulty.

It is an extremely technical matter and the view tends to be taken that, “The actuary says that is what is recommended. That is what we should do.” There are very few real discussions about the pros and cons of ESGs. It is not helped by the fact that there are just four providers in the market for ESG software, including two main ones and one very new one.

A particular concern I have in relation to the ESG providers is that we have noticed in the last few years that they have started to have to make some key assumptions themselves. For example, ultra-long term interest rates when there are not any such observable interest rates in the market, e.g. a fifty-year interest rate, which is what you need if you have a deferred annuity with a long time to go to retirement and a guaranteed annuity rate which might extend the term for, say, another 20 years.

These decisions really are the judgements of the board but we do find companies lapsing almost by default into using the assumption that the ESG provided, and this is a problem.

My last point relates back to what the author said in his introduction. Market consistency is all very well, but it does imply that if the market goes crazy your results are going to go crazy. Companies do not like their results going crazy. This is an issue that should be addressed. The author mentioned that one of the research papers seemed to treat this point. So this is something which should be properly addressed prior to the introduction of Solvency II.

**Mr N. C. Dexter, F.I.A.:** Section 7 on governance is important and is going to be one of the biggest issues, as implied by the last speaker. Boards do not often understand this type of thing. I have likened it in the past to our friends in the tax departments of companies.

A particular bone of contention with tax people is that the way they do their numbers is to do a balance sheet at the beginning of the year and a balance sheet at the end of the year and calculate the tax on profit as the difference between the two. I am sure you have tried getting your tax team to explain how the tax numbers have moved. I think ESGs can have similar problems. I have tried to get clients to explain how and why the result of the ESG calibration they have at the beginning of the year have changed to those at the mid-year and those at the end of the year. We need to develop how to validate these numbers and explain why they moved the way that they have.

One thing which is particularly important, because of the importance of the profit and loss attribution, is understanding whether or not profits or losses have arisen because the model has gone wrong. It is sometimes exceedingly difficult to prove whether an ESG is giving an odd answer because the market is a bit odd or whether it is because the calibration has actually got an underlying problem.

Quite often the reason why the result has come out the way it has is not easily explained. For example, you may have to go to your board to explain why the value of a particular product has not turned out the way that the product designers thought it would, as happened for a lot of products last year, or why the numbers have gone whacky. Last year end a lot of senior management time was taken up by those adopting MCEV in trying to make decisions about what to do about adjusting for the state of the market, and explaining what the implications are. It is important for us to devise a way to give the information to senior management that they need so they can properly understand the basis on which they are making their decisions.

When reviewing a new economic capital framework for one of my clients, it was evident that there was one particular scenario which was driving all its economic capital requirements. The same scenario drove all the capital assessments in all its global subsidiaries. No one, when we were reviewing it, could explain what that scenario actually represented in the real world (e.g. an economic downturn driving high inflation followed by high interest rates) and why it gave rise to such amounts of economic capital being required. There is an onus on us in investigating scenarios to understand back-testing. It is also important to validate the underlying management actions and/or policyholder actions that are assumed and really understand whether or not those decisions would be taken in those scenarios.

**Mrs K. A. Morgan, F.I.A.:** I liked the detail in the paper on the misconceptions around market consistent valuations. It is important to understand these problems as they lead to common misunderstandings about Solvency II.

I also liked that you highlighted the partial internal model consultation paper. I recommend to any company that is thinking of applying to use an internal model to read Consultation Paper 65, as most internal models are actually going to be partial internal models rather than full ones, and there are extra requirements for approval.

I should also like to add to what the author said about the different tests and standards that CEIOPS is recommending for internal models. The whole theme of the advice that we have given to the European Commission is that these internal models are the one used in firms. The idea is not to have a regulatory capital model, but to have an economic capital model that is used in the firm to calculate their own capital requirements based on their own risk appetites, and that the regulatory capital is just one output from that model.

If you are using that model for your own decision-making, there are tests and standards which are described in CEIOPS papers which are basic commonsense. You would expect models to be validated, documented and have quality data. The tests and standards are good practice and what I am sure firms should be doing anyway if they are using the economic capital models.

In terms of linking use and documentation, my view is that it is possible to have a complex model which is explained simply, something which the actuarial profession has been doing for a long time. It is necessary to have a simple economic scenario generator to fit in with the fact that the board have to understand it. It is worth remembering that CEIOPS gave some advice to the European Commission about proportionality about 18 months ago. We were very clear that proportionality is a two-way street. If risks are complex, you would expect more complex models to quantify those risks.

Proportionality is not linked to the size of the insurer, but to the nature, scale and complexity of the risks. You can have a small insurer with complex risks. You would expect them to be modelled in a complex way.

The points made about quality standards are really good. The big issue is how they can link the documentation of what has actually been done into their own documentation so that we, as supervisors, can review it. Not that we want to read all the documentation; we want to see evidence that it exists.

A question was asked earlier about whether we would be benchmarking to a standard formula when we are assessing internal models. No, we will not. The point of the internal model is that it reflects an organisation's own risk profile. An internal model might be one of the things to look at but we will not be benchmarking against that.

I also liked the mention of IT in the paper. Solvency II is often seen as an actuarial project but it is a multi-functional project and IT is really key and is perhaps overlooked.

For Solvency II, it is best to involve IT people as early as possible. The FSA IT team is reading through the CEIOPS consultation papers to work out the effect it is going to have on the FSA's own IT systems.

Finally, on the tests and standards for internal model approval, one of the reasons why they are all covered together is the need to look at them altogether. This point is made effectively in the paper by reviewing the use test and validation together. You cannot look at each test individually; you are going to have to pass all of them in order to get your model approved.

**Mr A. N. Hitchcox, F.I.A.:** I am not from life insurance but from general insurance. I am a Chief Risk Officer and my business is much more non-executive than executive. I view this paper with great interest because I help my board challenge complex modelling issues in insurance, and I find it very interesting to pick up lessons from other areas. I read the paper as a non-life person, as an actuary, and I understand the issues, not the technical issues, but the issues I needed to understand for governance purposes.

I have two questions which I should like to put to the audience to see whether anybody has any answers.

We are designing a training programme on the internal model for the board and the big debate we are having at the moment is how far to take them down the technical road. We have to recognise that they are required to sign off on these models, my internal model, your ESGs. Our risk management team have devised a programme in-house, which covers what we think the board need to understand about the internal model. We have come up with probably 30 to 50 topics that need to be understood. If for each of these topics you assign half an hour's explanation, that is a lot of time for a board.

So when I initially took that to the CEO, I knew I was going to be pushed back. He said I am allowed a maximum of two whole days of the board's time to train them on an internal model before they sign it off. So my question for your life ESG is about your experience. Your ESGs must produce, say, 50% of your total capital? Therefore it is important for your board to sign off on this. How much time are you asking your board to give you to let you train them on this subject?

My second question relates to applying benchmark portfolios in order to test model results. The investment analysts. are looking at the relative valuations of the largest companies in the UK and Europe. They are bound to ask: "I need enough information about the way you have all used your ESGs, so that I can rank your own investment compared to all your peers, for example, I need to say whether your stock valuations are higher or lower than those of your competitors."

I like the option table approach described in section 7.11.2. I will try to apply this idea back in my own industry. The biggest risk I deal with is natural catastrophe exposures; I work for a firm in the London Market and 50% of my capital comes from natural catastrophe exposures. I get a question frequently from stockmarket and investment analysts. They say, "Tell me, Andrew, what percentage load do you apply to RMS (a software vendor) to value your capital and your future earnings flows?"

They want to know if I am using exactly 100% of the published model, or I am taking 90%, or being conservative and actually using the standard output times a loading factor.

My question is: how do you explain to an investment analyst how your ESG valuation indices compare to those of your competitors?

**Mrs K. A. Morgan, F.I.A.:** To expand on the first question of Mr Hitchcox, it is probably worth spelling out the responsibilities of the board in respect of internal models.

The Pillar V recommendation of CEIOPS is the general principle that there is a joint understanding – a joint responsibility – for governance. But for internal models, CEIOPS' advice is slightly different.

What we are looking for is an individual understanding, so each board member needs to understand the internal model to some extent.

If they are using it in decision making and are executive directors, they need to understand the internal model in the areas where they are using it.

**Mr J. A. Jenkins, F.I.A.:** Mr Hitchcox's second question, about analysts, is not that difficult to answer. There is already a table in FSA returns where you have to work out what answer your ESG produces for certain standard derivative instruments. This enables analysts to make comparisons between companies' bases.

On the board time issue, in our practice we do board training for various things, including Solvency II, and getting time allocated from a board on any subject, including internal models, is quite a challenge. Companies do not allocate such amounts of board time without very good reason. So something up to a week would be the maximum for most boards altogether in a room for a training session.

On Mrs Morgan's point, this is quite difficult. You can write down on one side of a sheet of A4 paper the key aspects of the internal model, and it is therefore a question of how much detail they need to understand. You can do it at a very high level or you could go right down to the full detailed level. My concern is that some board members tend to switch off at a fairly early stage and will not be willing to spend the amount of time that it might take to achieve what the FSA requires.

We seem to have a mismatch between what the regulator is saying and what most board members believe is the right amount of time and detail for them.

**Mr R. Frankland, F.I.A.:** In the interests of transparency, I ought to state that the author and I have been members of a working party for a couple of years now, the Extreme Events working party. Some of our results have been published and there are a few more due out at the Life Convention later this week.

This wonderful paper does the opposite to what we did on the Extreme Events working party in the sense that we were taking data and essentially analysing it to come up with estimates of risk levels and corresponding movements in market indicators. Here the subject is more about synthesising those same sort of extreme events to generate a model which is capable of estimating the capital that is needed for Solvency II or indeed for any other value-at-risk based capital environment for insurance.

The paper explores the importance of the use test within Solvency II, and there was a similar principle applied in the ICA environment. It was intended that internal models should be those which were used in relation to the business. I do have concerns about this concept.

I assume that internal models can be quite variable in nature, in the extent to which they ascribe different levels of risk to different events. We are working at levels of probability which are very, very low. It is not surprising one model may give a significantly higher risk of one event happening whilst another model might have higher risks for other events happening.

I want to compare some simple facets of the two hypothetical models. These two hypothetical models differ only in one respect: for the two risks, which I have labelled 1 and 2, model A ascribes

risk 1 a one in 150 year probability and ascribes a one in 250 year probability to risk 2; the other model, model B, conveniently has the two risks the other way around, so it ascribes risk 1 as a one in 250 year probability and risk 2 as a one in 150 year event probability.

For the proposed application of a use test, the company which uses model A will be more concerned at the risk which has the one in 150 year probability, that is to say, event 1, and will take action either through reinsurance for through changing its asset strategy or through other methods of managing the business to try to reduce that risk, or possibly to eliminate it entirely.

The company that uses model B will similarly be led to focus on the second of the two risks which it sees as having higher probability, assuming it is using its ESG and its internal model for its company management. There is nothing actually wrong with that approach. It works fairly well, and different companies should focus on the risks for which they are most concerned.

It is also important to note that relatively the company using model A will set aside more capital in relation to event 1 than it would if it used model B and less in relation to event 2 than if it used model B.

My concern is, if you use the same model for both the management of the company and for setting capital standards, there is a tendency, unintentionally, to focus management effort on those risks which the model dictates would tend to require over-capitalisation, relative to other models, and to reduce effort in relation to those risks which would be relatively under-capitalised.

My question to the author is: "Is this a risk with the Solvency II regime with the requirement to adopt the use test or to support the use of an internal model through these tests, or have I missed something somewhere in the analysis?"

**Mr E. M. Varnell, F.I.A.:** We should probably let Mrs Morgan and Mr Marshall talk about what their intentions are in the use of internal models. From my point of view, it is beneficial to stick with one particular model, year-in and year-out, but to use different models to investigate the model error.

In an ideal world I would want an insurer I worked in to have an ESG model function receiving an economic basis from an economist function.

Then I would like to have ten different model combinations which would be calibrated against that same economic basis.

I would try to test my results against each one of those ten different model combinations. Where I saw differences in the results, I would learn about weakness in the model configuration I used in my financial reporting.

When I think about the internal model I do not just think about one particular model. Mrs Morgan will probably correct me, but it seems to me that one could use several models. An internal model could perhaps be some sort of average of the results calculated using a selection on underlying ESG model configurations. We should probably ask Mrs Morgan and Mr Marshall whether that would pass the internal model tests.

**Mr I. Marshall, F.I.A.:** There are quite a few questions to answer.

The first query concerns the internal model. Are we talking about the internal model from a regulatory perspective or the internal model which is used in the business? If we are talking about the internal model used for the business you could have your own risk appetite as the calibration standard, which does not have to be the final measure. There could be a different risk measure that you actually look at.

What actually comes up in the statistical quality standards is the quality of the internal model and, if you look at the way it is described in the advice given by CEIOPS, it is a probability distribution function. What we are keen on promoting within CEIOPS and the FSA is to look at the whole distribution. So, if you look at the one in 200 basis, what will happen at different probability intervals as well? I hope that addresses your question to some extent.

The one in 200 test is actually for the whole distribution. It requires that we can answer: what is the one in ten event that could happen? It is a bit of a challenge in terms of getting the whole probability distribution function. The ICA techniques and models that have been used are fairly comprehensive. You actually have to derive the whole distribution, whereas for other risks, such as underwriting risks, they may be looking at specific points. This is one of the key challenges for the life industry.

The CEIOPS advice does have quite a bit of information about using key points on the distribution and how you can look at further validating.

The other point which the author raised was the aspect of having different models. It is clear that aspects to consider cover the different bits and pieces of that model: that is fine so long as you define the model and you define the model in such a way that you do not all of a sudden start picking and choosing, using different bits and pieces for this and that. It all needs to be included – a lot of pieces – as long as you define what it is you have.

To some extent some of the definitions given by the author of all these different techniques of calibration form part of the validation of the results actually coming out of the model.

I hope that answers a lot of the questions.

**Mr C. I. Chappell, F.I.A.** (closing the discussion):

First, Mr Varnell, thank you for the paper.

Given the advantages we have had in the UK with the ICA review, it does make me wonder how easily other countries will find it to make the significant transition that the Solvency II frame requires. In particular, I have some sympathy with the regulators as they deal with the difficulties this creates trying to maintain that level playing field.

The requirements of the tests for internal models remain somewhat nebulous, and although I believe companies are working with best endeavours to achieve what we believe will be compliance, section 3 of the paper draws out some of the difficulties in getting to the calibration standard or areas where expert judgement will remain crucial as objective bases remain illusive.

The consultation papers from CEIOPS seem to be articulating the ideal standard, which may not be wholly achievable, for example, permanently deep and liquid markets and talk about market



consistency in real-world scenarios. It would be interesting to hear views on whether the market consistent approach is inherently pro-cyclical and where the market consistent valuation could be modified in the light of the recent financial crisis.

That leads to further questions as to whether our current ESGs are good enough or whether further improvement is needed. With that in mind, the comments about how the ESG scenarios best be disclosed and whether the FSA tables used presently are sufficient are very relevant.

Section 4 details alternative approaches that may be useful for valuing liabilities. Whilst the more simplified modelling approaches, such as closed-form solutions, may have limitations, they may be necessary for projecting balance sheets to achieve reasonable and timely management information. It is evident that this will add increased pressure on our communications skills and I do start to feel sorry for the non-executives of 2012 bombarded with explanations of the ESG used to value our balance sheet day 1 and how we project balance sheets using closed-form solutions which are sufficiently robust.

Mr Jenkins also adds to this by saying the ESG providers are starting to make assumptions such as those for long term interest rates. Some of these assumptions are those that should be set by the board but inadvertently the boards are defaulting to ESG providers. I would be interested in views of the implications to members of the board, and maybe the dangers of getting constrained by technical accuracy and the potential for the disappearance of some elements of common sense.

We appreciate the list of documentation that you think will be necessary. Interestingly, the providers may be preparing significant elements of these, but they may not be to the quality appropriate for users and extra documentation requirements may still remain for companies. We should not underestimate what this will bring.

Additionally, I would be interested in views on the change control process of ESGs. We do not want to stifle technical enhancements, but do need to ensure control and stability. However, we need to recognise that sometimes it is in moments of commercial pressure that enhancements are most easily justified for development expenditure when competing for resources.

**Mr Varnell, F.I.A.** (responding): Thank you very much for the kind comments on the paper. It is much appreciated.

I have made some notes based on the comments that were made. I noted that Mr Marshall spoke first and he did not say anything with which I would disagree.

Mr Lockwood spoke about deterministic scenarios, matching those up against stochastic scenarios and how this would work. My take on that is that, in the past, people have taken decisions based on deterministic scenarios. This has got to change if we are moving into a world where you need to consider a stochastic model for calculating the capital requirement and this has to be signed off as being something that the business uses to make decisions. This necessitates that senior management need to start using stochastic models in order to make decisions in their businesses.

My view is that there needs to be an education process for management and we as actuaries can play a key part in helping senior management with that process.

Mr Jenkins talked about the governance issues, particularly around the difficulty in getting engagement with board of directors. The fact is that you might well send them to sleep if you try talking to them about ESG models in too much technical depth. Another thing that crops up when you are trying to talk to non-executives about ESG models is that they can get hooked on a particular technical issue and you spend too much time on a technical cul-de-sac, such as the choice of seed in the pseudo random number generator.

The point I am seeking to make is that mathematics is something you need to keep well away from boards of directors, unless you have got someone who is particularly keen and has the right background. This is where I feel the economist function really comes into its own.

One thing that I would hope that non-executive directors and boards are ready to get their heads around is some of the key economic assumptions that we are making in ESG models. For example, what do we expect equity markets to return over the next few years? How uncertain do we think equity markets are going to be? Where do we expect interest rates to go?

That feels like it could be made into a real discussion about what is happening in the wider economy, what is happening in the news, and what the Bank of England are doing. I hope that the non-executive directors of a financial institution have enough interest in what is happening in the wider financial markets and in the economy to have views about what the economic basis looks like going forward.

Mr Dexter talked about the governance issue too. He explained that there is an issue around disclosure and helping people to understand what is inside the ESG box. What I thought would really address this well is the option table that was discussed. I think Mr Hitchcox may have attributed the idea of an option table to me; I cannot however claim any responsibility for that, much as I would have liked to.

I found the option table that was put in place as part of the realistic balance sheet regime of the FSA to be particularly useful as an ESG practitioner and I would be disappointed if that were to disappear in Solvency II. It could benefit from some modification but it seems to me an extremely valuable way of benchmarking different ESG models between companies or over time.

These tables address the point Mr Dexter was making about assessing how an ESG calibration has moved over the course of a year. An option table is a relatively concise, clear way of articulating what has changed in the model. Some things, which are missing from the current FSA table, that I would like to see in Solvency II are more correlations and a better implied volatility surface. You can infer some correlations from the FSA's current option table. However there is a lot more that one could have inferred with a redesign that would be very helpful for Solvency II.

Mrs Morgan raised a lot of very good points and talked about how to interpret the advice that CEIOPS has produced. I was very pleased to hear her say that internal model calibrations would not be benchmarked against the standard formula stress tests.

I would like to ask her if she is 100% sure on this because I do wonder if a 99.5% percentile stress test is available from the standard formula and a company designed their own equity stress test would the regulator feel comfortable with an internal model firm using 35% while standard formula companies were using 55%.

Mr Franklin talked about the problem that you might have a variety of models which you could apply. Using a variety of models would, in some sense, help define the model error incurred by using one particular model.

Mr Chappell raised a range of issues which I should talk through in detail. Is market consistent approach inherently pro-cyclical? I have to say that it is. If you set yourself a capital requirement target at a fixed level of 1 in 200 one-year VaR and value the balance sheet using a market consistent balance sheet approach you are inherently setting yourself up for something that is pro-cyclical when markets crash, as almost all the model assumptions or inputs are going to go the wrong way on a balance sheet at the same time.

My view of this is having an equity stress like that described in CP69 is quite narrow and that a wider dampener based on all asset classes would be quite useful for Solvency II. The CRO Forum, in particular, suggested a shadow SCR which would give you some way of coming back up after a stress event. There is some merit in looking at an approach like that.

Mr Chappell also asked if ESG models are good enough. A lot of the ESGs are good enough in the sense that the models which have been built so far have been relatively well-implemented. What is missing is the ability to easily calibrate the different model configurations of a particular dataset. Providers have tended to provide a particular set of calibrations which most of their ESG clients tend to use, and many insurers tend to stick to those default calibrations from quarter to quarter.

There are not really enough people trying out different models and seeing if the answers they get differ markedly. In the early days of ESGs there was a bit of that type of experimentation going on. Some providers had different models and the people using ESG models at that time in the insurers would try them out. It becomes quite interesting when you start to compare different model combinations, such as sophisticated equity models with different interest rate models. Then one starts to find interesting interactions between the models.

The use of more sophisticated combinations of models all calibrated to the same economic basis gives insights into model error.

This is something banks do already. If you talk to trading desks they will not use one model, but a range of models, and, in this way, decide where the true valuation is taking account of the limitations of the models used.

In terms of the other points that were made, I refer back to my point that having a well designed option table is of particular help in communicating the economic basis of an ESG.

**The Chair (Mr C. A. Cowling, F.I.A.):** It remains for me to express my thanks and your thanks to the author, our closer, and all those who spoke this evening. Thank you very much for your contributions to a very interesting and important discussion. I suspect there are a number of issues here which are going to feature very highly over the next few years as Solvency II gets closer and closer.

May I therefore ask you all to thank Elliot and the speakers in the usual way.