MILLIMAN RESEARCH REPORT

Optimisation

Improving Problem Formulation and Human Interaction

May 2017

Neil Cantle, MA, FIA, MIoD, ASA, CERA
Corey Grigg, ASA
Dr. Lucy Allan

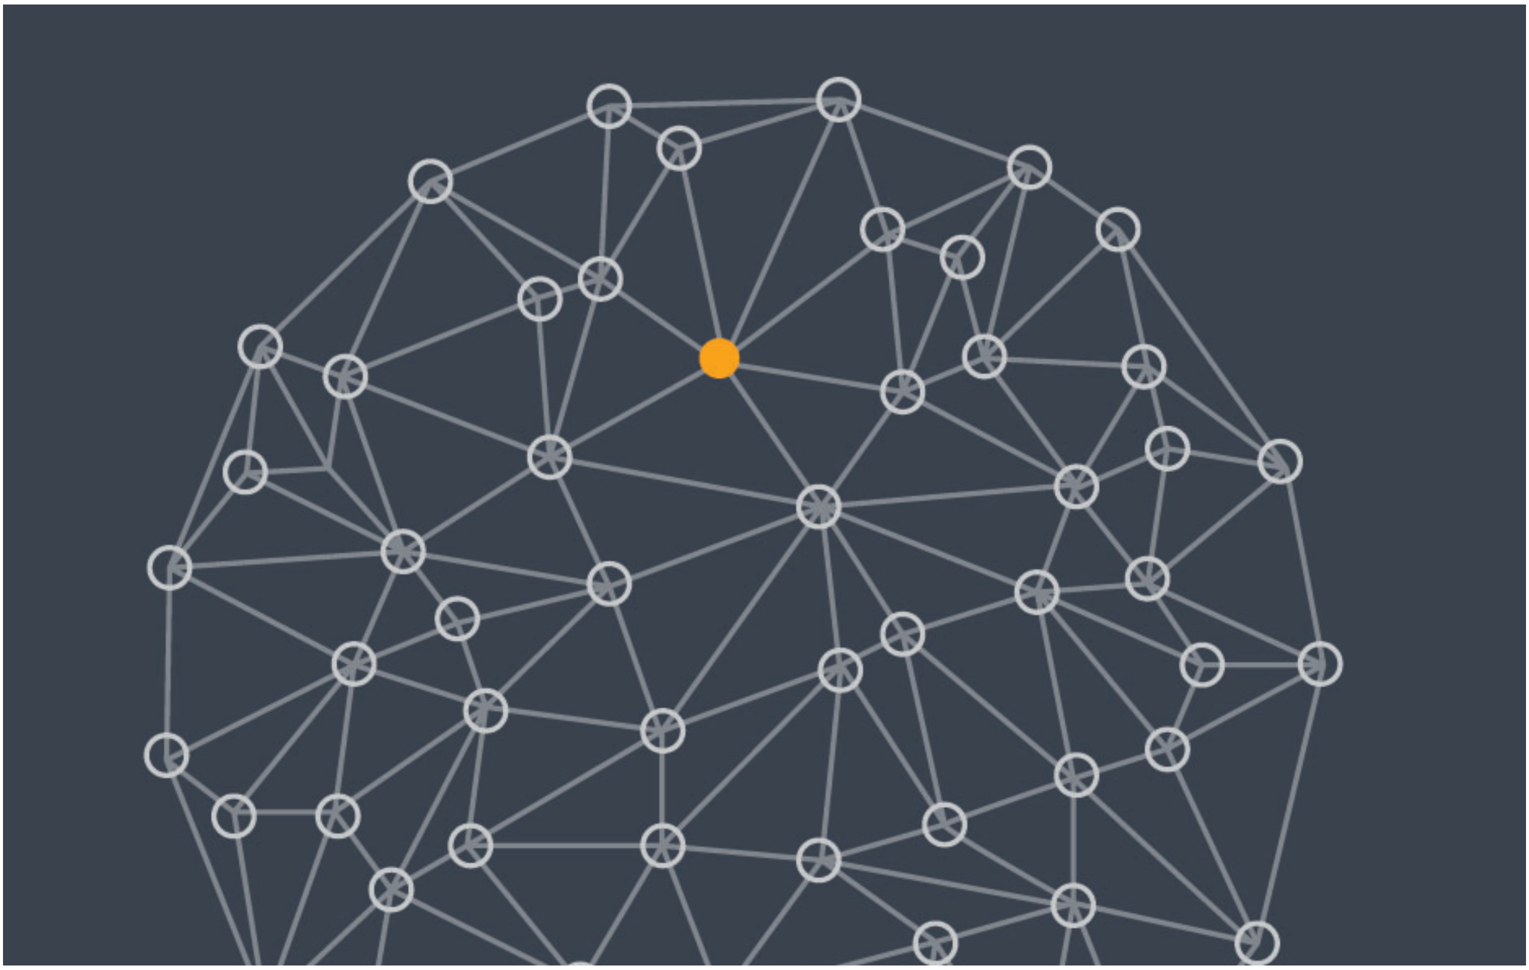

Milliman

Milliman

Table of Contents

1. INTRODUCTION

Optimisation problems have traditionally been formulated as single objective and solved with the use of gradient-based or direct search methods. Most practical real-world problems involve multiple, often conflicting objectives, and also highly complex search spaces. Competing goals and objectives necessarily give rise to a set of compromise options and solutions.

To counteract some of these difficulties, multiple-criteria decision-making (MCDM), which is a sub-discipline of operations research and which explicitly considers multiple criteria in decision-making environments, is brought together with evolutionary multi-objective optimisation (EMO). Evolutionary algorithms possess several characteristics that are desirable for this kind of problem, helping to direct the decision-maker to regions of preferred solutions.

In the literature there are three main types of decision-maker interventions reflecting ‘how’ and ‘when’ to express or articulate EMO preferences with respect to the individual criteria:

- *A priori preference articulation*: Denotes the process of introducing and incorporating the preferences before the search process
- *A posteriori preference articulation*: Denotes the process of introducing and incorporating the preferences at the end of the search process
- *Progressive preference articulation*: Denotes the process of introducing, incorporating, and modifying the preferences in an interactive and progressive way at any time during the search process

Given equally adequate solvers, with the equal ability to identify promising solution areas in the search space, the decision-maker is seen as having the capability to control the optimisation process by focusing on criteria thresholds to extract viable solutions. This approach has been useful, however it involves mostly unidirectional workflows with the assumption of known and definable problems with known and definable constraints. The challenge in these circumstances is posed as one of interpreting and deciding upon trade-offs and when and how to best express them.

In practice there are few, if any, real occasions when there is not feedback between the problem definition and solution requirements, and this is especially true when dealing with dynamic multi-criteria dimensionality. The solver of choice plays an important role and is required to efficiently evaluate the solution space under consideration. However, this is only one element in a more challenging process. For example, wider interaction, which optimises methods and scope for thinking about the problem, is also required. In essence, the potential solution space to be searched is not just the conceived landscape, but also those yet to be conceived. And, conventionally, while little heed is paid to these hidden opportunities and threats lurking elsewhere, they represent a chance to identify superior solutions. Problem re-formulations and bi-directional workflows can be used to bring forward richer problem formulations. The search then is realised as not just for optimal or robust solutions in the given landscape, but for the best solution in all conceivable landscapes—a significantly different opportunity. What is discovered during the search of solution spaces affects the understanding of the problem. This establishes a learning link that can bring new landscapes into the analysis that will likely contain better solutions. Metaphorically, why continue to look for diamonds in the dust when the dust tells us there may be richer pickings elsewhere?

In contemporary optimisation techniques, there is little if any mention of bringing together the problem, domain, expert, solutions, preferences, and an open discovery process into one unifying framework. By providing a means for users to engage more easily, collaboratively, and intuitively with their business and modelling knowledge and with heuristic insights, richer discussions can be brought forward and previously hidden or obscured solutions can be revealed.

In attempting to shift towards a problem-centric perspective in the optimisation processes, an open discussion of problems and objectives, as well as an interactive exploration of all alternatives, is necessary. This approach will provide a better understanding of the problem and reveal the impact of selecting one alternative solution over another. With improved human interaction, usually best applied through visual exploration techniques, more real-world experience can be brought to the situation, enabling more appropriate solutions for the business.

The methods proposed provide the following advantages:

- Better use of 'know-how' through collaboration and wider input of pooled experience
- The combination of quantitative and qualitative information for gaining understanding
- More robust decision outcomes
- Deeper understanding of critical relationships in larger sets of data and complex models

In this report, we illustrate how some of these techniques can be applied to a common problem in the life insurance industry. Specifically, we describe how to develop insights into drivers of economic capital within an internal model framework, and how to use these insights to make risk management decisions. For this case study, we use the ECSight™ system[1] to project the balance sheet of a hypothetical company over a set of 100,000 real-world scenarios, and we use DACORD[2] to analyse the data produced. The aim is to identify key drivers of outcomes from within the complete range of solutions. The exploration of the results using parallel coordinates illustrates one stage in the decision process, where a range of outcomes is selected from thousands or millions of alternatives.

For this case study, we model a hypothetical insurance company with three lines of business:

- Fixed deferred annuities with minimum crediting guarantees, backed by a portfolio of fixed-rate bonds
- Life insurance, including term, whole life, and traditional mortality insurance products, backed by a portfolio of fixed-rate bonds
- Variable annuities, including products with guaranteed benefit riders, backed by a portfolio of fixed-rate bonds and a static hedge portfolio including equity futures and interest-rate swaps

The one-year, real-world scenarios used consider changes in both market conditions and liability assumptions, and encompass the following categories of risks:

- Equity level
- Equity volatility
- Credit spreads
- Credit migration & default loss
- Interest rates
- Lapse
- Mortality

Please see the appendix for a more detailed description of the underlying model.

[1] For information about ECSight, please visit http://us.milliman.com/Solutions/Products/Resources/ECSight/ECSight-Advantage/.
[2] For information about DACORD please visit http://www.dacord.co.uk.

2. BACKGROUND

MULTI CRITERIA DECISION-MAKING

Multi-criteria decision-making (MCDM) involves considering trade-offs between multiple criteria or objectives in order to prioritise or rank alternatives. Based on work on Value Theory and Utility Theory in Operational Research, Zoints[3] has published work on decision making to select from discrete alternatives. This is based on using mathematics to solve quantifiable variables and other variables that are difficult to quantify. In previous papers he has addressed methods to interactively select alternatives. Some attributes will be constrained to take on certain values, whereas for others the ranking of the outcomes will be used as a basis for decision-making.

FIGURE 1: THE PROCESS OF MULTI CRITERIA DECISION-MAKING

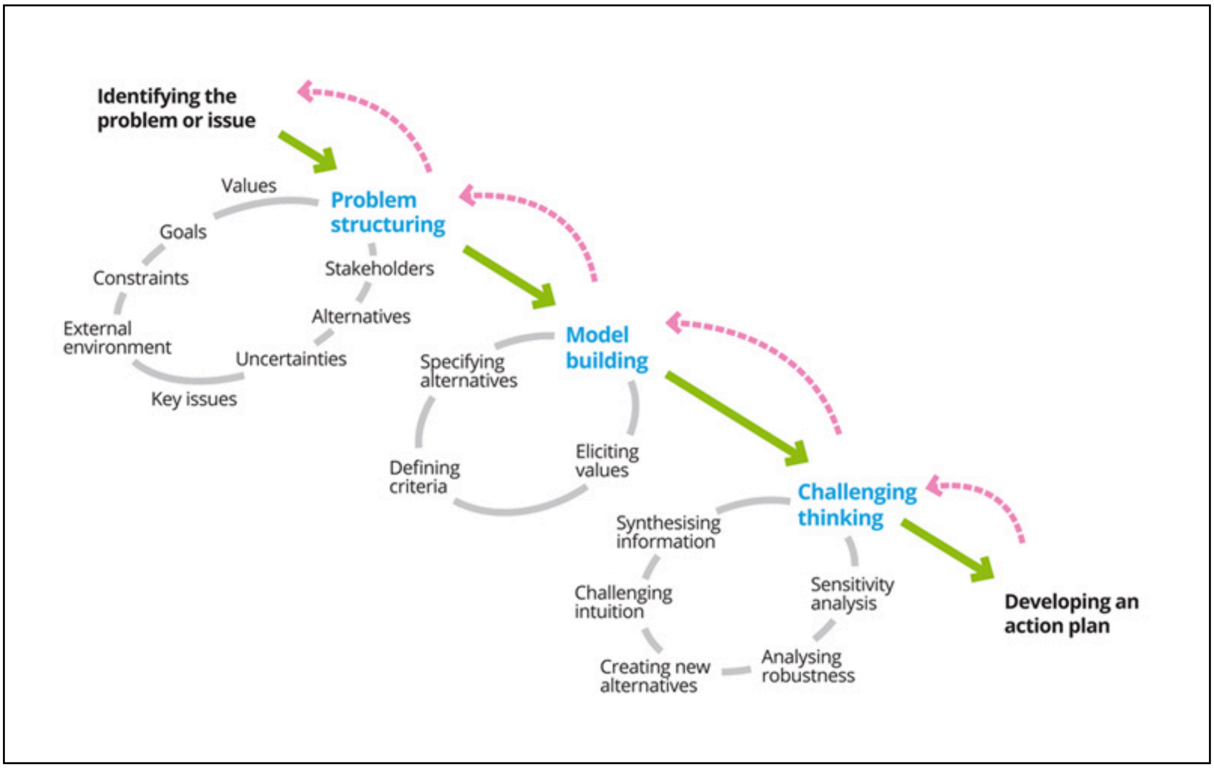

When using multi-criteria in decision-making the following definitions apply:

Alternative: An alternative is an outcome or a result that can be selected. Alternatives can be produced by simulation (e.g. stochastic modelling) or already exist in a discrete set of data.

Preferences: Preferences form the choices that are used to create the alternatives. This is typically phase 1 of the selection process. The preference is used to control simulation iterations to steer the algorithm and reduce the time it takes to calculate alternatives. Therefore, the algorithm spends more time producing alternatives that might be considered, than producing alternatives that will never be considered.

Constraints: Constraints limiting the range of an input variable to specified values. A complex constraint maybe have dependency rules, such that one variable must be greater than another, or a variable will be proportional to the value of another variable.

Objectives: Objectives guide the range of input or output variables to target values. Alternatives produced that use values outside the 'objective' range may still be valid alternatives. The objectives can be thought of as soft constraints.

[3] S. Zoints, A multiple criteria method for choosing among discrete alternatives, European Journal of Operational Research 7 (1981) 143–147. And, S. Zoints, Having and Using Alternatives in Negotiation, Plenary at MCDM2011, https://www.jyu.fi/en/congress/mcdm2011/Prog/plenarists/talk1.

PARALLEL CO-ORDINATES

Parallel coordinate (PC) visualisation provides an easily interpreted overview of large-scale multi-variable datasets. It is especially useful for fast detection of outliers and the general distribution of solutions. In MCDM, constraints can be applied to input and output variables to narrow down the choice of solutions. Additionally, objectives can used to direct simulation runs.

> *The strength of parallel coordinates isn't in their ability to communicate some truth in the data to others, but rather in their ability to bring meaningful multivariate patterns and comparisons to light when used interactively for analysis.* [4]
>
> *Rotations in Cartesian coordinates become translations in parallel coordinates and vice versa. Perhaps more interesting from a statistical point of view is that points of inflection in Cartesian space become cusps in parallel coordinate space and vice versa. Thus the relatively hard-to-detect inflection point property of a function becomes the notably more easy to detect cusp in the coordinate representation. Inselberg (1985) discusses these properties in detail.* [5]

In Figure 2, we show an example parallel coordinates display, for the case study discussed in the introduction. Each of the parallel coordinates represents a risk driver within the economic capital model. Each scenario is represented via a series of line segments between risk driver values for that scenario.

FIGURE 2: PARALLEL COORDINATES DISPLAY OF ECSIGHT SIMULATION RESULTS

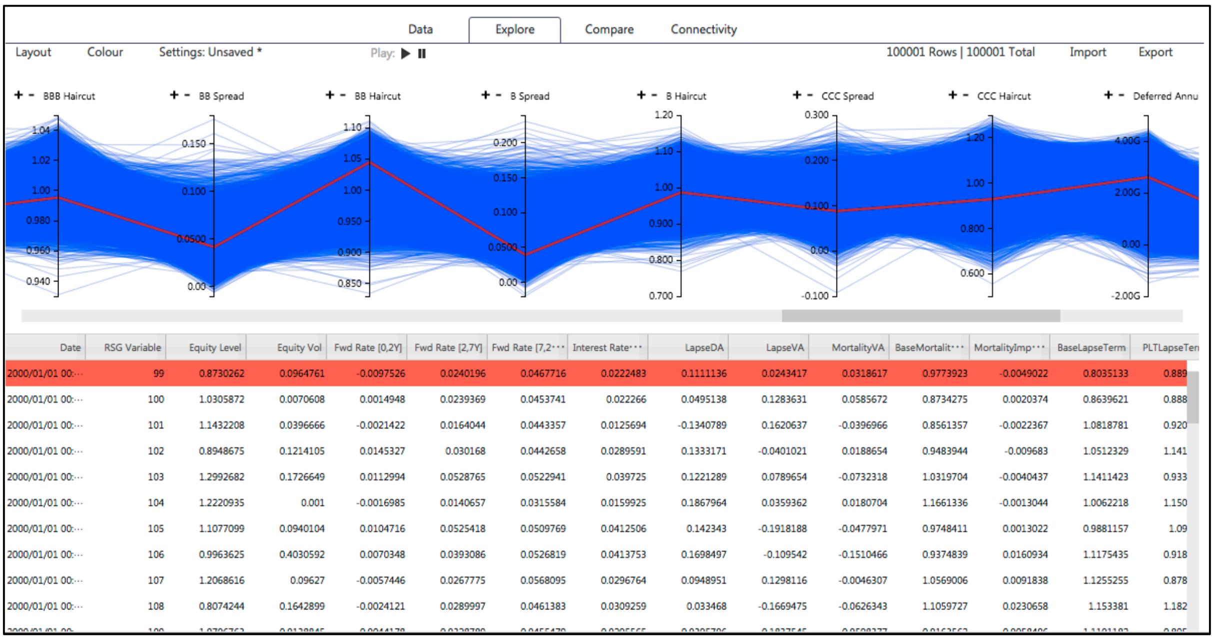

The main conclusions from Inselberg's discovery process in the Multidimensional Detective are:

- Do not let the picture intimidate you.
- Understand the objectives and use them to obtain 'visual cues'.

Inselberg investigated the mathematical relationship between variable pairs as displayed on the Parallel Coordinates graph. The shapes and crossing points can be determined mathematically. However, for MCDM the main issue is relating the range of solutions. Specifically, when using parallel coordinate plots in optimisation the relationships are stated as:[6]

[4] S. Few. Multivariate Analysis Using Parallel Coordinates. Originally published September 12, 2006, http://www.b-eye-network.com/view/3355.

[5] E. Wegman and J. Solka. On Some Mathematics for Visualizing High Dimensional Data, http://binf.gmu.edu/jsolka/PAPERS/MathVisRevision.pdf. And, Inselberg, A. (1985), "The plane with parallel coordinates," - The Visual Computer, 1, 6991.

[6] P.J. Fleming, R.C. Purshouse and R.J. Lygoe. 2004. Many-Objective Optimization: An Engineering Design Perspective.

'…crossing lines' indicates conflict between the two adjacent objectives. The degree of conflict is demonstrated by the intensity, or degree to which, the lines cross. Conversely, lines that do not cross demonstrate objectives which are in relative harmony with one another.

The advantage of viewing multi-variables on the same plot and interaction has its strength in multi-criteria problems. Gettinger et al.[7] stated:

Many decision problems involve multiple, conflicting, and incommensurate criteria. Methods of multi-criteria decision analysis aim at supporting decision-makers (DMs) in such tasks. In discrete decision problems the number of solutions is finite, but may comprise hundreds, if not thousands, of alternatives. Portfolio selection problems in which collections of items (e.g. projects) are evaluated according to several properties, may serve as a prominent example.

Gettinger et al. discuss a two-phase approach to decision-making. In the first phase, a selection of efficient alternatives is determined. In the second phase, DMs interactively explore this set in order to identify their most preferred solution. The first phase can be eliminated if the alternatives are created in an interactive manner, by steering simulations, which is also demonstrated in two Operational Research papers, Korhonen[8] and Zoints.[9]

Among the advantages of parallel coordinates are that there is no loss of data in the representation, which in turn ensures that there is a unique representation for each unique set of data. It also has a low representational complexity, O(n), where n *is the number of variables modelled, allowing the technique to scale well to large numbers of variables. Weaknesses of this visualisation method, however, are (i) that it requires multiple views (different orderings of objectives) to see different trade-offs, and (ii) that it can be hard to see what is going on when many vectors are represented.*[10]

COGNITIVE MAPPING

Cognitive mapping is a technique which visualises the way in which concepts are related to each other. Such a map is a useful way to capture theories about complex phenomena. Sharing Filter Settings can be achieved by exporting Filter Settings to a Map. In this context, a Filter Setting is considered as a set of individual variable filters, with lower and upper limit values. The Filter Settings represents a number of individual constraints to apply to the data when exploring the range of alternatives.

FIGURE 3: USE OF MAPS FOR DISCUSSION OF OPTIMISATION CRITERIA

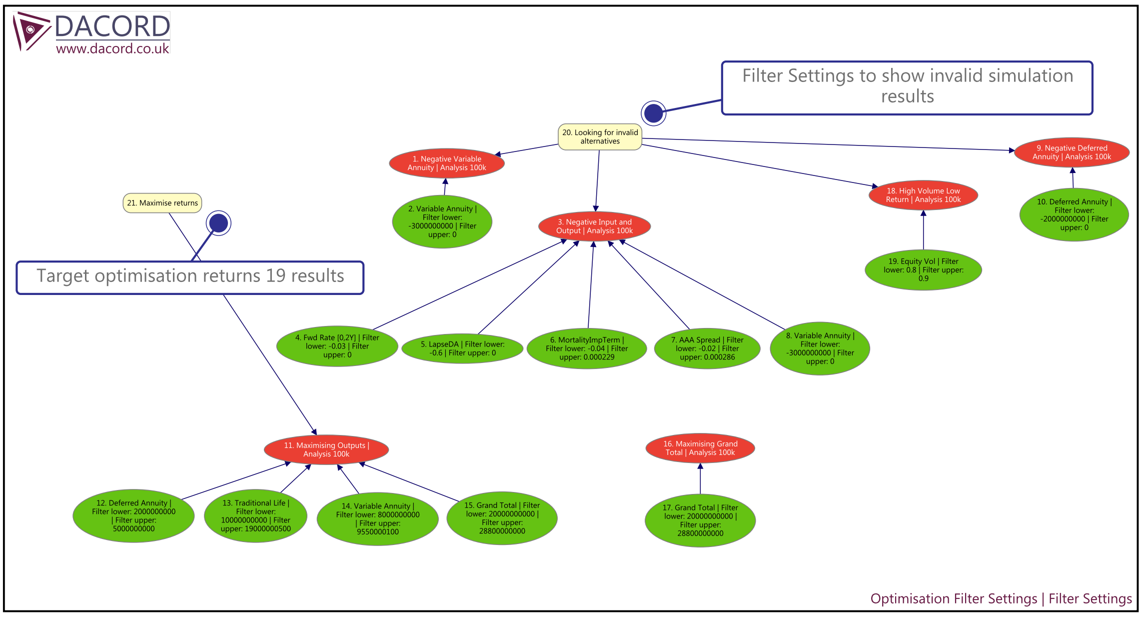

[7] J. Gettinger, E. Kiesling, C. Stummer, R. Vetschera. 2013. A comparison of representations for discrete multi-criteria decision problems, http://www.sciencedirect.com/science/article/pii/S0167923612002783.
[8] P. Korhonen. A visual reference direction approach to solving discrete multiple criteria problems. European Journal of Operational Research 34 (1988) 152–159.
[9] S. Zoints. A multiple criteria method for choosing among discrete alternatives. European Journal of Operational Research 7 (1981) 143–147.
[10] Ibid.

In the Map the Constraints definition is captured to enable discussion about the reasons for individual filter ranges and the impact of the combined Filter Settings. Additionally, after exporting a Filter Setting to a Map, the Maps can be merged to allow commentary over all the settings.

Example reasons for annotated filter limits:

- Valid range for a variable, for example, only positive values for prices or costs to constraint to show only realistic results and exclude unrealistic result.
- Desired range for a variable, for example, areas of most revenue.
- Highlighting all invalid values. Filter Settings are only inclusive between the lower and upper limits, therefore a filter or a combination of filters can be used to show the set of alternatives that are out of range and unacceptable solutions.

Sharing the commentary of Filter Settings using Maps is useful for a discussion on criteria settings and in turn their linkage to the problem. Tool mechanics make it easy to modify the values in the Node text for an individual filter and then to import the changes. When in DACORD Analysis, importing from Maps will find all the Maps with compatible Filter Settings and present a list. The list contains the Filter Setting name, how many individual variable filters it contains, and from which Map. Importing with the same name as an existing filter will replace it locally in the Analysis, and is optionally then saved into the Analysis.

FIGURE 4: SHARING FILTER SETTINGS AND MAPPING THE PROBLEM DOMAIN

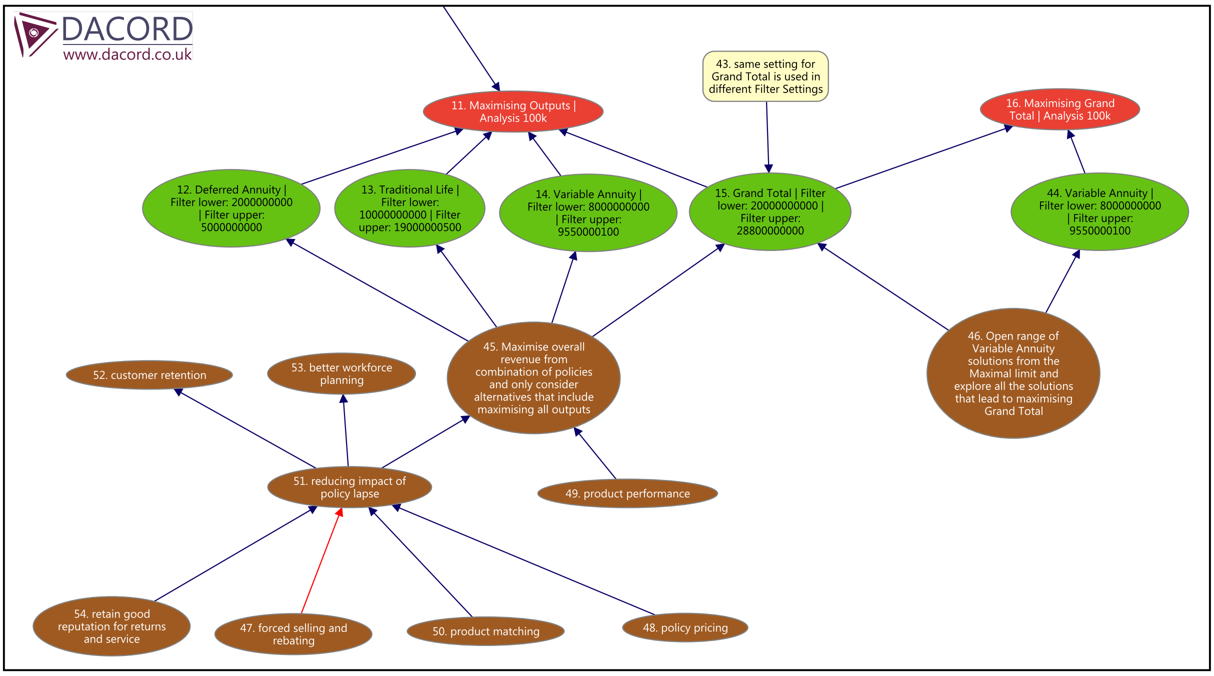

Importing the Filter Settings from a Map is a very powerful way of sharing and applying settings. Figure 4 contains an illustration of two Filter Settings (red Nodes), the individual Filters (green Nodes) and the commentary Nodes (in brown and yellow). This illustration shows some contributing factors and impacts of Policy Lapse, where the Filter Settings are used to find the alternatives (product mixes and investment strategies) that maximise returns.

The benefits of using Maps can be summarised as:

- Share Map between different people for collaboration on filter settings
- Import Map to different Analysis that uses the same Dataset
- Apply the Filter Settings from a Map to a different Dataset with the same Variable Names
- Share Individual Variable Filters with other Filter Settings
- Provide commentary on the Filter Settings and how they apply to the problem domain

OPTIMISATION OF OBJECTIVES AND CONSTRAINTS

The main assumption is that the decision-maker is presented with a number of discrete alternatives from which to choose and each alternative has a value for every parameter. The alternatives are produced from simulation modelling with a range of input parameter values and the resulting output parameter values. The overall aim is to choose one or more optimal solutions from the total population of alternatives.

In Figure 5, the information contained in Figure 2 has been filtered to enable one to focus on a specific subset of the risk factor movements over a year. More details about the filtering are contained in Section 4.

FIGURE 5: EXAMPLE OF CONSTRAINT SETTING

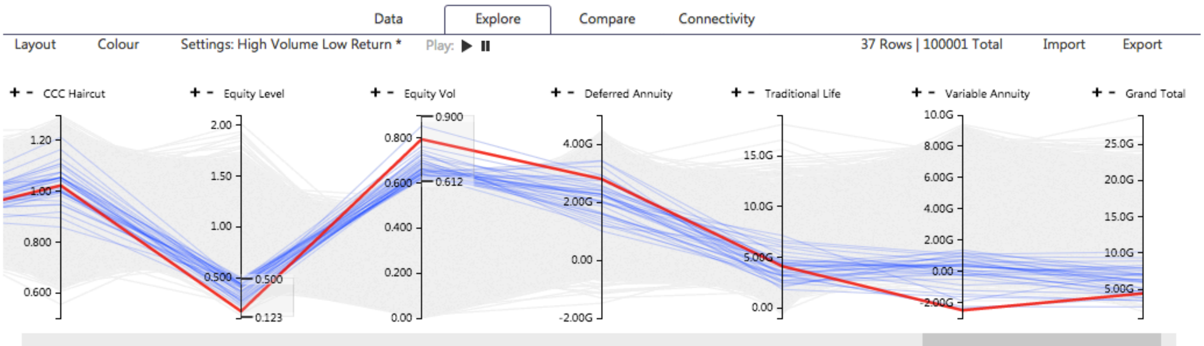

B Haircut	BB Spread	BB Haircut	B Spread	B Haircut	CCC Spread	CCC Haircut	Equity Level	Equity Vol	Deferred Ann···	Traditional Life	Variable Ann···	Grand Total
1.0125879	0.0409151	1.0519071	0.053121	1.0634511	0.1039925	1.0544703	0.28142	0.6585898	1476683799	4759928443	80256756.88	6316868999
0.9980856	0.0222919	1.0175336	0.0728013	1.0202936	0.0731658	1.0227667	0.1669431	0.7927158	2793135602	4110158826	-2498715704	4404578725
1.0091928	0.0461994	1.0093382	0.0962743	1.0261264	0.1312276	1.0543158	0.2689771	0.6508409	2259435403	2638515371	-508363710.2	4389587064
1.0048756	0.0458244	1.0236869	0.062383	1.0520623	0.0853949	1.0114516	0.4006769	0.6246251	1864407811	2059725887	1180855088	5104988786
1.0188437	0.0665244	1.0702012	0.0788061	1.0573207	0.1177027	1.0753103	0.4196269	0.7226588	2392419174	2491498269	-1486709215	3397208227
1.0016748	0.0244501	1.0096714	0.0339248	0.9921678	0.0890556	1.0543337	0.4409802	0.7447347	2004951599	4332085264	-281834342.5	6055202521
0.9850207	0.0448519	0.9477826	0.1250777	0.9584478	0.1201862	0.8942099	0.2903452	0.7007923	1309654874	2689145924	-205966099.7	3792834698

Choosing an optimal alternative requires consideration of many objectives that are often in conflict or related by complex trade-offs. The benefits of improving one objective needs to be evaluated against a potential negative decrease of another objective. For example, given the parameters of capital investment, exposure to risk and return on investment, the objective would be to minimise capital investment and exposure to risk whilst maximising return on investment. Increased investment would increase return, but it would be in conflict with other objectives.

Constraints are limitations of the parameters in which, outside of these limits, it would be not possible or too risky to operate. Examples of such limitations might be the maximum available capital to invest, the maximum exposure to risk as a calculation of worst case outcomes, or the minimum return on investment that would allow an organisation to remain in business.

Using the facilities within the DACORD platform, an interactive method has been developed for setting objectives and constraints, whilst exploring the population of alternatives and understanding the impact of conflicting objectives and trade-offs.

In Figure 5, the variables Equity Level and Equity Vol have constraints applied to limit the number of alternatives, represented by the blue lines. The resulting alternatives are also shown in the table and the number of results shown at the top (in this example 37 alternatives out of the 100,000 simulations are consistent with the constraints). The pre-filtered alternatives are shown as a grey shadow. In this example the constraining variables have been moved nearer to the output variables so the effect of interacting with the constraint values can be seen on the important output variables, such as Grand Total. The red line and red row in the table is used to highlight an individual solution and the resultant set of solutions can be exported.

Live interaction with the data is important to gain knowledge about the shape of the data and the relationship between variables. The knowledge about the total range of alternatives can be used to reinforce or challenge the knowledge about the scenarios the alternatives represent. Exploration of the data using the constraints can be used to challenge assumptions about simulation models that produce the set of alternatives.

3. OPTIMISATION WORKFLOW

FIGURE 6: MULTI-CRITERIA DECISION WORKFLOW FOR SELECTING OPTIMAL RESULTS

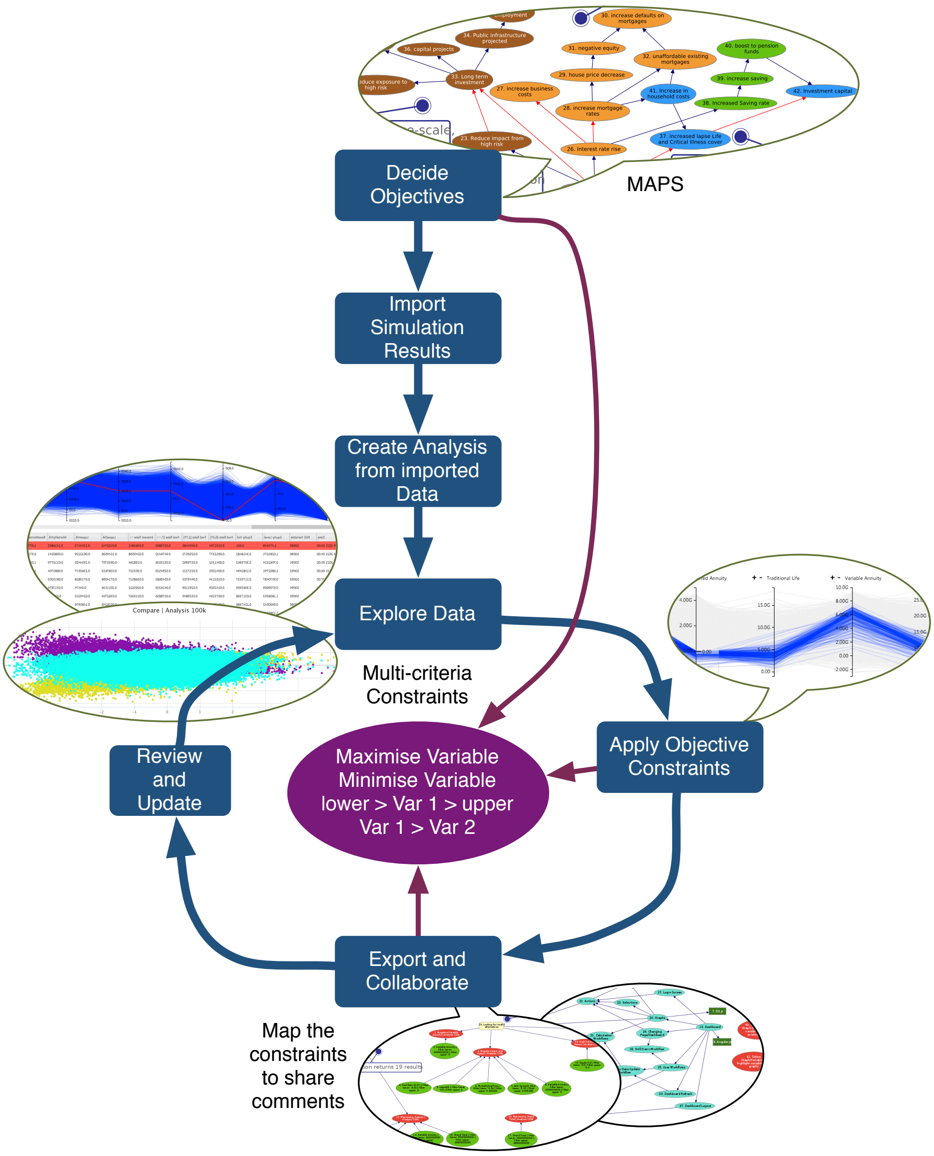

DECIDE MAIN OBJECTIVES FOR OPTIMISATION AND GATHER RESULTS DATA

The objectives describe the desired outcomes from modelling different scenario conditions. The outcomes would be determined by the values of the outputs of the simulation compared to the inputs. An example of this might be the likely return on investment using a defined investment strategy and evaluating the outcome in value to represent spread of results and exposure to risk.

The objectives will ultimately be captured as value ranges imposed on input and output variables of the simulation. A cognitive map can capture the desirable results and the reasoning behind the decisions made. This can be produced collaboratively.

LOAD RESULTS FROM SIMULATIONS AS A NEW DATASET

From an original Monte Carlo simulation, all results (up to 100,000 results) should be stored as a simple table in CSV or Excel. Each row is an Alternative result from the simulation, and the columns represent the individual Variables.

After data is imported into DACORD, add a date column if one does not exist. The Analysis package requires a timeline to display the graphs.

CREATE ANALYSIS USING THE NEW DATASET

Create a New Analysis and enter the name for the optimisation study. In the list of datasets, select the imported Dataset. The system will load the new dataset, check the settings, and launch the calculation.

If the Dataset is large (more than 10,000 rows) the calculation of the Analysis may take several minutes. The Analysis will load when the calculation is finished. Optionally, it is safe to close the Analysis window and reopen when the My Analysis page shows the Analysis is *Ready*, or at any time before it shows *Ready* and the Analysis results will load after the calculation is finished.

EXPLORE LANDSCAPE OF DATA BY VISUALISING IN EXPLORE TAB

Exploring the Original Data by using the Analysis concentrates on two tabs: Explore and Compare. The Explore Tab shows the multivariate view of the entire dataset using Parallel Coordinates. The Compare Tab uses a scatter graph to plot Variables, one variable on the X axis, and any other variables simultaneously on the Y axis.

Using Explore: Parallel Coordinates Graph

At this stage, it is useful to explore the data to understand the relationships between Variables and the distributions of Variables across the entire dataset and within certain ranges. The user can gain general knowledge about the data, test their assumptions about the results, and validate that the results of the simulation are within sensible ranges.

Brief Guide to Parallel Coordinates main interactions

- Move variable axes (left-right) along the parallel coordinates to position adjacent variables and view different adjacent patterns. Observe which variables are mostly parallel. It can be possible to discount one of these variables from the analysis (because they contain the same or similar information).
- Add Filters to one variable. Draw the filter to view different parts of this variables distribution. See how many Alternatives are visible and how many are removed:
    - Using the upper or lower range of the Variable where the density of the Variable is low.
    - Using the filter as a range to understand the distribution. Create a filter with upper and lower limits within say 20% of the overall range.

Look for patterns between adjacent variables

The relationship between two Variables can be viewed by making the axes adjacent and observing the pattern or direct influence when changing filters.

- Parallel lines between two axes shows a direct correlation; therefore both variables carry the same information and meaning.
- When lines cross in a superposition of X-shape this shows a negative relationship and conflicting constraints.
- When lines cross randomly, there is no particular relationship between the variables.
- When viewing Density of the distribution of the variables, switch between blue and green colours to view a different contrast between high and low density.

Interact with the data

Test your assumptions about the data. If you assume there is a pattern relating two Variables, then drag the variables together to look at the adjacent pattern, i.e. the lines between the two axes for the Variables. Apply filters to one of the Variables: look at the upper and lower ranges; look at the densest range of the Variable.

- Is the relationship what you expected?

If there are outliers, then dig deeper into the data. Apply filters for the axes to capture the outliers and view how this applies to other Variables.

- Do other Variables show outliers, or are the outliers on one Variable within normal ranges on all other Variables?

If outliers on one Variable relate to outliers on some other Variables but are normal in most Variables, then we can assume there is a relationship between the outlying Variables that is not connected to the others. Further questions of the data could be to use the Compare Variables graph to look at distributions of these related variables.

Using Compare: Scatter graph

View the bivariate distribution by comparing two variables in the Compare graph.

- After observing the effects of adjacent distributions, it is useful to compare the variable distributions on the Compare Variables graph. The distribution shapes can be useful when confirming that the Simulation model is performing as expected, and also when the influence of one variable against another can be seen. This is a different view of the same distribution through adjacent axes on the parallel coordinates plot.
- Select the X-axis from the Variables drop-down list. Choose a Variable that is a pivotal input or output, one with most influence on the results.
- Unselect the X-axis Variable from the Variables list to remove the less interesting X=Y plot from the Compare graph.
- Select Variables in the Variables list that are also influential. The bivariate distributions can be seen between the X Variable and the Variables list.

APPLY OBJECTIVE CONSTRAINTS

Viewing the data on the Explore Tab, apply filters to Variables to constrain the inputs and outputs within desired and objective ranges. Start with one objective, say to maximise a Variable, and then choose an upper Range that yields valid results. Apply the objectives to multiple variables to ensure valid ranges.

The upper right shows the number of Alternatives that remains, and the Data Table shows all the Alternatives as a row of values.

The Alternatives can be reduced until the number of remaining data rows is very small, even one row, to find a single desired solution to the original problem.

FIGURE 7: MAP WITH ALL THE FILTER SETTINGS APPLIED TO THE DATASET

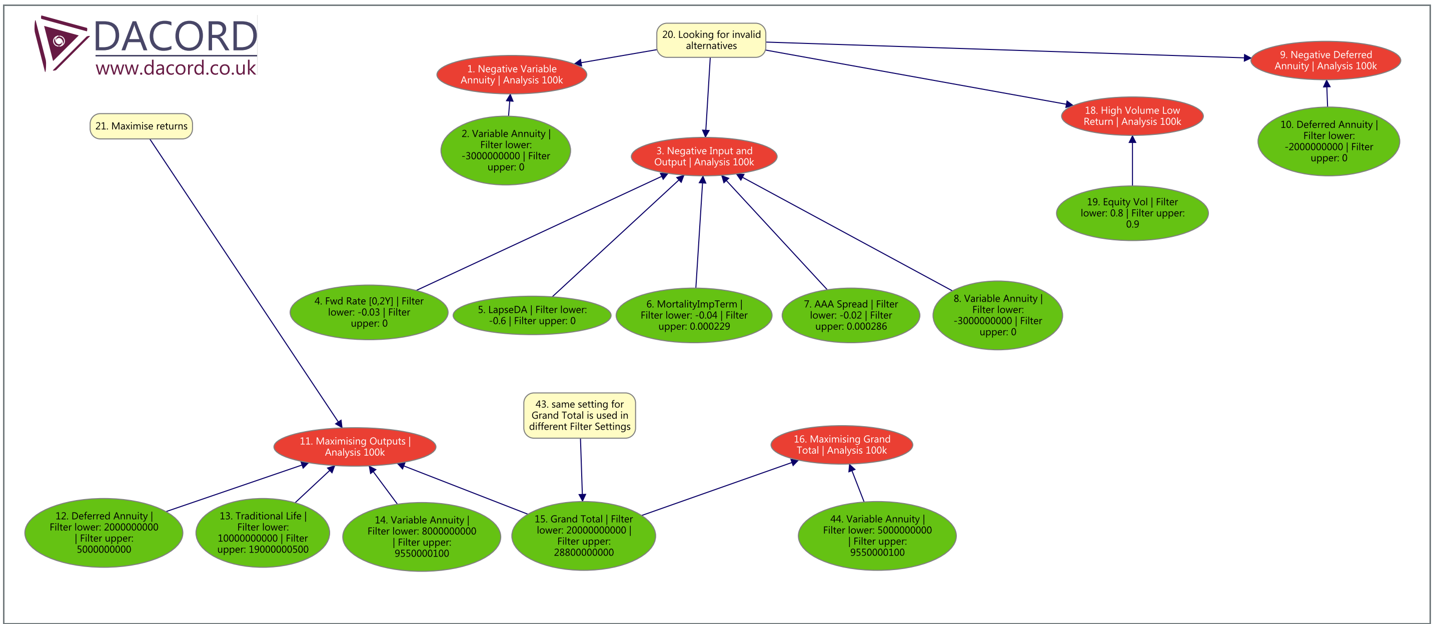

EXPORT AND COLLABORATE

The Analysis package provides several mechanisms to collaborate on the decision-making process for Optimisation. The Analysis can be shared between DACORD users, so more than one person can explore the alternatives. The analysis provides access to the Dataset, Variable Groups, and any Filter Settings saved in the Explore tab.

Additionally, the Explore tab provides three methods to export the Filter Settings and data. *Export->Settings* will provide a CSV of the filter settings that can be used in a report and imported back into the Analysis.

Export->Data will provide a CSV of the data table, as shown below the Parallel Coordinates graph. This is useful to export the selected alternatives after filtering the data.

Export->MAP is the most useful for collaboration and importing. The Filter Settings are exported into a DACORD Map, which is opened in a new window. The Map can be added to with new nodes and annotations to provide a rich commentary on why filter ranges were chosen. For example, annotations to state filter settings are used to accept on valid ranges of a Variable, or add nodes as a Cognitive Map to describe the root cause and reasons behind a group of filter settings.

Using Maps, the Merge function (from My Maps) can be used to combine all the exported Filter Settings into a single Map. The Map can be shared with other DACORD users to collaborate on the commentary of the Filter Settings. An example merged Map is shown in Figure 6.

The nodes in the Map contain the upper and lower filter values for each Variable. The node text can be edited to modify the filter values so long as the structure of the text and the variable name is not changed.

The export Filter Settings in CSV or in Maps can be imported in the Analysis, or any Analysis that uses the same dataset, or even another dataset with the same Variables. This can be a very powerful collaboration tool to exchange filter settings between different Analyses, along with documented reasoning for each filter value.

4. COMMENTARY ON FINDINGS

As noted previously, our case study focuses on result data generated from the ECSight system. The example dataset in this study contains 27 input parameters and four output values for 100,000 alternatives. The data is provided in CSV format, with a header row of Variable names and data rows each containing an alternative.

The aim is to identify key drivers of outcomes from within the complete range of solutions. The exploration of the results using parallel coordinates illustrates one stage in the decision process, where a range of outcomes is selected from thousands or millions of alternatives. The exploration of the results can be run post simulation on very large datasets from Monte Carlo execution to cover a wide range of results. However, the graph can also be populated interactively whilst running simulations and guide execution to the desired results, allowing specific trade-offs between variables to be better understood. The results data concentrates on useful alternatives by rejecting unwanted results during the execution.

In practice, different companies will focus on different decision criteria. For the purposes of illustration, we make the following assumptions after initial data exploration to provide objectives:

Assumption 1: Decisions are primarily informed by outcomes for the level of the overall balance sheet valuation ('Grand Total') as well as the three constituent balance sheet valuations ('Deferred Annuity,' 'Traditional Life,' and 'Variable Annuity'). Filtering on these values provides information about certain regions of the scenario set, for example the tail scenarios that are particularly important to economic capital.

Assumption 2: The input variables Equity Level and Equity Vol have a strong influence on outcomes. Using Equity Level and Equity Vol as filters helps illustrate the relationship between these and other risk drivers. Equity Level and Equity Vol in particular have an inverse relationship.

FILTERING GRAND TOTAL

Filtering on Grand Total above 20 billion yields 191 solutions. The outputs Deferred Annuity, Traditional Life, and Variable Annuity are also in the upper limits of their distribution for these solutions.

The inputs for each Spread and Haircut show common patterns of minimal spread and maximal haircut. These inputs also show that to achieve maximal Grand Total, the Equity Level should be high and Equity Vol should be low.

This filter could be the starting point for formulating an optimal strategy. Other Variables should be constrained to valid ranges, the objective applied.

FIGURE 8: MAXIMISING GRAND TOTAL

VARIABLE NAME	FILTER LOWER	FILTER UPPER
GRAND TOTAL	20000000000	28800000000

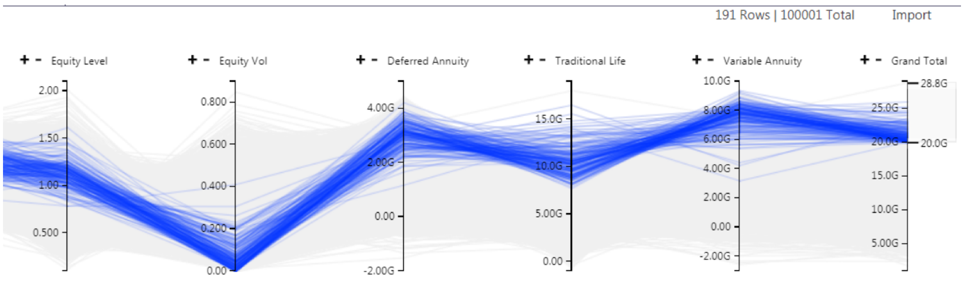

Assuming the four outputs should be maximised, the filters have been set to discover the solutions that correspond to maximising all four outputs.

There are 19 solutions, however the Equity Vol is very low for all solutions.

FIGURE 9: MAXIMISING OUTPUTS

VARIABLE NAME	FILTER LOWER	FILTER UPPER
DEFERRED ANNUITY	2000000000	5000000000
TRADITIONAL LIFE	10000000000	19000000000
VARIABLE ANNUITY	8000000000	9550000000
GRAND TOTAL	20000000000	28800000000

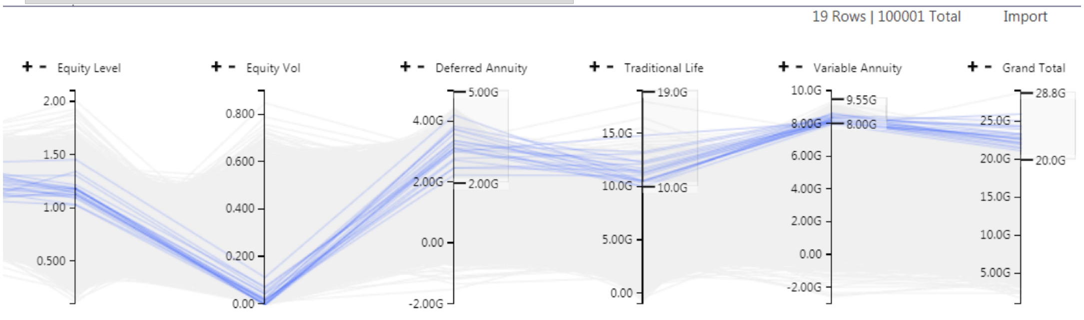

Setting the Equity Vol to choose the maximum appears to be an outlier result. The overall picture shows that most other variables are within their distributions when Equity Vol is at the maximum. However, the exceptions are: Equity Level is very low; and the outputs Traditional Life, Variable Annuity, and Grand Total are low.

The extreme may be an indication that the model is very sensitive to Equity Vol.

FIGURE 10: HIGH EQUITY VOLATILITY

VARIABLE NAME	FILTER LOWER	FILTER UPPER
EQUITY VOL	0.8	0.9

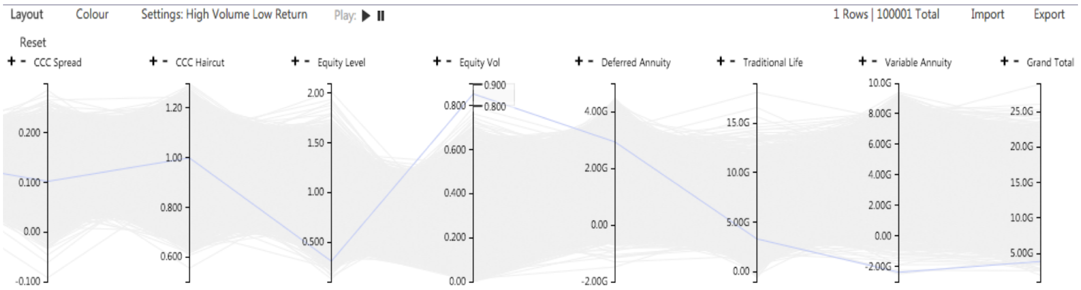

Now we apply the range to Deferred Annuity output to find when this value becomes negative. This reveals that most variables are within the middle of their distribution. Markedly, Traditional Life is also low (indicating correlation between outcomes for these blocks) whereas Variable Annuity is high and contains only positive (non-negative) values, indicating some diversification benefit with the Variable Annuity.

FIGURE 11: NEGATIVE DEFERRED ANNUITY

VARIABLE NAME	FILTER LOWER	FILTER UPPER
DEFERRED ANNUITY	-2000000000	0

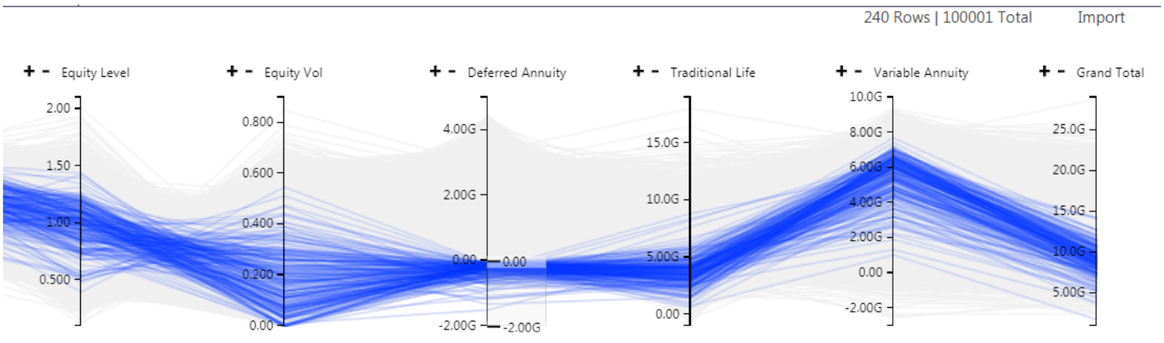

Filtering on the negative values of Variable Annuity reveals most variables are within the middle of their distribution. Markedly Equity Level is low and Equity Vol is high.

FIGURE 12: NEGATIVE VARIABLE ANNUITY

VARIABLE NAME	VARIABLE ID	FILTER LOWER	FILTER UPPER
VARIABLE ANNUITY	31	-3000000000	0

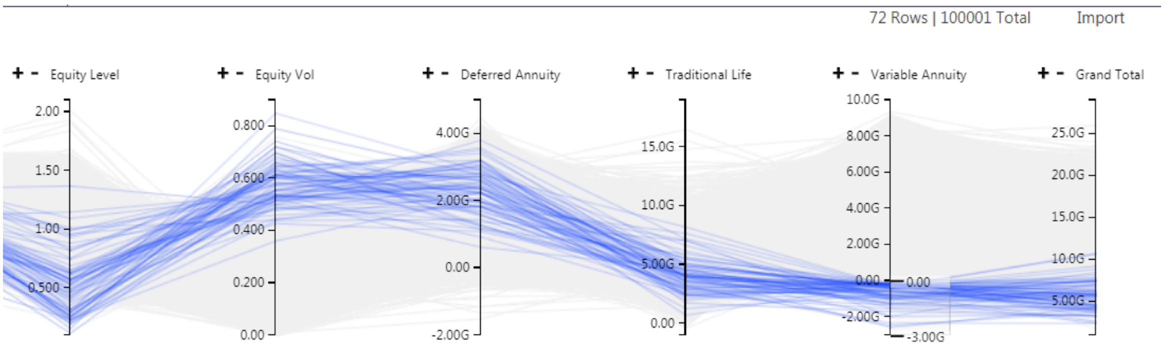

Investigating the negative limits of the variables for forward interest rates (0 to 2yr), deferred annuity lapses, Term mortality improvements, AAA spreads and variable annuity NPV (respectively Fwd Rate [0,2Y], LapseDA, MortalityImpTerm, AAA Spread, and Variable Annuity reveals three solutions at this limit. These scenarios are candidates on which to focus during model validation exercises due to their rarity.

FIGURE 13: NEGATIVE INPUT AND OUTPUT

VARIABLE NAME	VARIABLE ID	FILTER LOWER	FILTER UPPER
FWD RATE [0,2Y]	4	-0.03	-0.000286
LAPSEDA	8	-0.6	0.00343
MORTALITYIMPTERM	12	-0.04	0.000229
AAA SPREAD	15	-0.02	0.000286
VARIABLE ANNUITY	31	-3000000000	120000000

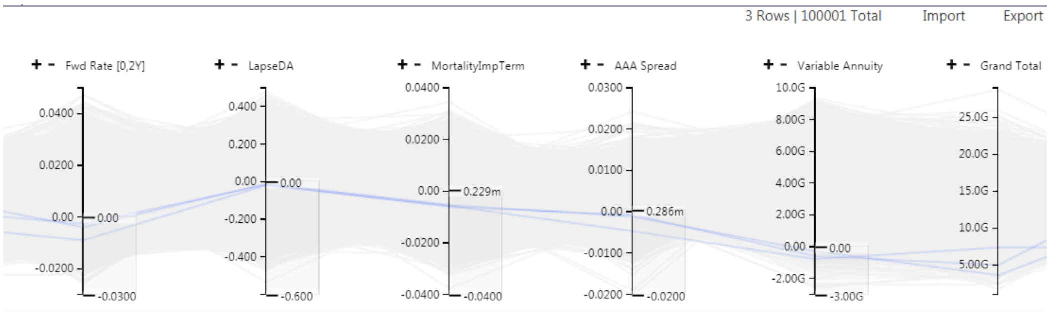

The method has demonstrated it is possible to view and explore a large number of simulation results. In this case, 100,000 alternative results can be explored at the same time with visibility of all input and output variables. Filtering is applied to individual variables rapidly across the whole range of results.

The following interactions with the data proved successful in helping to make sense of it:

- Exploring the distribution and range of a variable's value
- Exploring the relationships between variables
- Rapidly filtering a variable by constraints to see the effect on the range of alternatives and the distribution of other variables
- Rapidly apply and modify multi-criteria constraints across all the variables to select optimal solutions
- Exploring the alternatives produced by modelling and validating the range of results

Using Maps with the Filter Settings produced better results than expected. Maps provided the link between analysis of the data and the context of the problem. Merging the Filter Settings into a cognitive map made it very easy to document and explain the limits used in filtering, and to explain the reasoning behind various sets of results. The Maps also allowed collaboration through the use of sharing the Map, and offered surprising power when reusing (via Import) the Filter Settings in Analyses of different datasets with the same variable names by importing settings.

In the Optimisation Case Study, the optimal result was obtained using the assumptions of maximising the outputs variables. In the Maximising Outputs results (FIGURE 9), there are 19 results. Using these results, there is high risk in the value of Equity Vol. Therefore applying a filter to choose the optimal value of Equity Vol from the 19 filtered results yields the optimal result, as shown in figure 14.

FIGURE 14: OPTIMAL RESULT

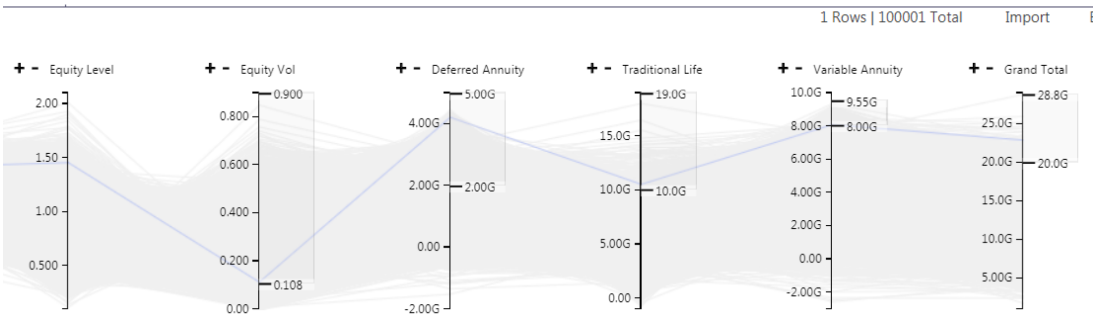

5. FUTURE DIRECTIONS

As presented, the methodology illustrated in this paper could be used as part of a human-and-machine iterative process. Results produced from the model would be analysed as illustrated in this paper, and domain experts reviewing the results of the analysis provide insights into the features that subsequent iterations of modelling should contain. The output of this study demonstrates an investigation to search 100,000 simulation results for optimal solutions. Exploring the results benefits from collaborative access to the data and cognitive maps to capture the selected solutions and the reasons for the selection. During this investigation, the visualisation techniques have been created rapidly, interacting with the large dataset. Whilst machine algorithms are capable of seeking an 'optimal' solution from a given model, the approach we have presented here allows for the possibility that the results of the first iteration of modelling demonstrate that model changes are required and/or a different view of the problem should be explored. This is not yet feasible purely by using machine algorithms, and is unlikely to be for some time. In FIGURE 6: MULTI-CRITERIA DECISION WORKFLOW FOR SELECTING OPTIMAL RESULTS we outlined the process proposed by this paper, whereby results are produced, analysed, and models revised and re-run iteratively until the optimal solution is obtained.

In addition to the approach shown in this study, further enhancements have been researched to extend the capabilities in five key areas:

- Exploring big data (millions of results)
- Interacting with simulation to improve the results sets and steer simulations to optimal results
- Using enhanced visualisation techniques, such as automating variable ordering and visualising distributions and clustering using colour
- Keeping problem in the loop
- Practical applications for insurance ERM

These are described in more detail below.

Exploring Big Data: Exploring Millions of rows instantaneously is the next technical challenge.

Summarization: refers to the computation of aggregated data and usually involves loss of information. DACORD currently does not lose any information when plotting parallel coordinates. This is an important information-seeking mantra—to keep all data. For much larger datasets, there are many aggregation approaches to show high volume information in parallel coordinates, either as additional visual items, or by representing sets of items using alternative visual encodings such as envelopes of lines or density. The next work will be to discover the best techniques in trade-off between speed of pre-processing and retaining important information about trends and distributions.

Interacting with Simulation: The techniques for searching large data can also be used in near real-time to steer simulations and provide dynamic input to simulation control. Automated optimisation algorithms have a large search space and may not find the optimal solution within time. Using an algorithm to execute simulation runs that accepts multi-objective inputs means the user can change preferences for constraints and objectives when viewing the results so far, and the change of preference is used to guide further simulations.

Complex Constraints and Objectives Language: Using complex constraints and objectives is important when Monte Carlo is steering simulations to produce more results in the range of the desired solutions.

The constraints limits allow for upper and lower ranges to be specified for each variable. As explained earlier, constraints can be dependent across variables; therefore constraint rules can be more expressive to define relative dependency. For example, $0.5 < A \leq 1.5$ and $B < (2A)$.

Objectives are target values for a simulation and can be easily added to the exploration plot. When implementing steering input for an algorithm, objectives provide the guide for valid or most desired output values. These are used by the steering algorithm to reject certain individual alternatives and produce a wider range of desired results.

Using steering of simulation modelling will:

- Reduce the time to find an optimal answer
- Reject irrelevant or impossible results early in the simulation process
- Improve resolution, thereby providing more results in the relevant space
- Provide an interactive entry of constraints and objectives, using expert knowledge to guide the algorithm
- Reduce the rules in the simulation algorithm, thus allowing different types of searches without changing code
- Validate the model by exploring results and confirming the model produces realistic results

Using Enhanced Visualisation: In DACORD, the parallel coordinates view is supported by the other plots for exploring the data, including the scatter plot for viewing the bivariate distributions. The following visualisation techniques could be used to show properties on enhanced parallel coordinate plots:

Distribution: Each axis can be used to explore the distribution. The z-value or standard score of a variable can be used to colour the lines on the plot to aide viewing the impact of one variable on all other variables displayed.

Dependency modelling: Dependency modelling is the process of establishing qualitative or quantitative dependencies between variables. Linear correlation between two variables is the most common dependency that can be visualized in parallel coordinates as a result of the point–line duality. The quantification of dependencies is an important measure for determining the relative importance of dimensions that can be used to order axes in parallel coordinates. The adjacency order of variables is important when viewing dependencies. Automatic ordering can be used to show two main dependencies:

- Conflict: A relationship in which performance in one objective is seen to deteriorate as performance in another is improved.
- Harmony: A relationship in which enhancement of performance in an objective is witnessed as another objective is improved.

Classification: Classification is the task of mapping data samples to a set of predefined classes. A typical technique in interactive visualization environments that supports the classification of samples is brushing. Brushing is typically used to select data points that are then subject to further processing, such as learning a classifier. In DACORD, the brushing or filtering of a variable in the parallel coordinates plot can be projected onto the Compare scatter plot to directly project the range of one variable onto the range of another variable.

Clustering: Clustering is the identification of sets of data items exhibiting similar characteristics. There is a wide range of automatic clustering techniques that typically depend on the similarity measure being used. Parallel coordinates can be used for 'visual clustering,' i.e. to find groups of similar points based on visual features such as the proximity of lines or line density. Similar to the Classification enhancement, the clustering would project the brushing onto time series plots to highlight ranges in the results that apply in time. This enhancement can identify when the whole system conforms to similar behaviour.

Change and Deviation Detection: Change detection and deviation detection is the visualization of outliers or other anomalies of the data with respect to some previously known measure. For example, data samples can be classified as outliers using a density estimation based on parallel coordinates of the raw data. The detection of abnormal behaviour using parallel coordinates is also an important task in process control applications. In DACORD, the filters applied to axes are inclusive limits. An enhancement would be to reverse the filters and visualise the excluded values and their effect across all the result to specifically investigate deviation from normal behaviour.

The future directions presented here require development of visualisation and algorithmic data processing, along with some research and testing on specific human interface design to ensure the correct balance in speed, colours, and system behaviour.

Keeping Problem in the Loop: Ultimate improvement in business problems where optimisation issues need to be addressed, will be enabled when the process of problem refinement and problem reconsideration are an integral part of all tools. The suggested use of cognitive mapping techniques for discussion of multi-criteria is a promising start. Further work is required to make this simpler and these ideas more widely accessible. The shift in perspective required is that optimisation becomes one tool, albeit improved and powerful, within the integrated process and art of problem solving.

Practical Applications for Insurance ERM: Visualization and filtering of the output of a multidimensional internal model immediately suggests a number of useful avenues for analysis:

- **Finding patterns in key risks driving capital losses**. By filtering for adverse scenarios in a parallel coordinate visualization one can look for risk drivers where the scenarios pass through a narrow band of values. For example, a distribution of losses on variable annuities might show a majority of adverse scenarios passing through the bottom of the equity level axis. This is an effective way to identify opportunities for hedging, or other mitigation strategies for unhedgeable risks, such as reducing exposure.
- **Understanding diversification**. A well-diversified risk profile will show a variety of scenario paths when filtering for adverse scenarios. Creating parallel coordinate visualizations of the risk profile for subcomponents of a business (e.g. insurance product type, asset type) and comparing them side-by-side can enable quick judgements about the similarities and differences in the risk profiles of different portfolio elements.

6. APPENDIX – CASE STUDY MODEL

Model Lines of Business

DEFERRED ANNUITIES

- Based on a pricing model developed for a reinsurer
- Includes products with minimum crediting guarantees
- $20 billion of Account Value
- Market-consistent BEL of $20.75 billion
- Backed by a portfolio of fixed-rate bonds with an NPV of $23.24 billion
  - The bond portfolio is designed to have slightly more interest rate sensitivity than the liability, reflecting our experience that asset managers tend to go longer than the liability in the current environment.

DEFERRED ANNUITY BOND PORTFOLIO CHARACTERISTICS			
RATING CLASS BREAKDOWN		MATURITY BREAKDOWN	
AAA	23.7%	1	19.1%
AA	29.1%	2	17.2%
A	29.1%	3	16.4%
BBB	10.0%	5	14.4%
BB	5.0%	7	10.9%
B	2.0%	10	10.0%
CCC	1.0%	12	9.6%
		15	2.2%
		20	0.1%
		30	0.2%

TRADITIONAL LIFE INSURANCE

- Based on an appraisal model of a large reinsurer
- Includes term, whole life, and traditional mortality insurance products
- $522B of face amount
- $4.15B of statutory reserves
- Market-consistent BEL of $0.27 billion
  - Note that this small liability in comparison to reserves is due to the fact that US XXX statutory reserve requirements assume lapse after the initial term period, but post-level term profits are included in the liability best estimate.

- Backed by a portfolio of fixed-rate bonds with a book value of $4.15 billion, to match the statutory reserve requirement
  - The bond portfolio is designed to have a duration of 5.0 years, reflecting the anticipated runoff of reserves.

TRADITIONAL LIFE BOND PORTFOLIO CHARACTERISTICS			
RATING CLASS BREAKDOWN		MATURITY BREAKDOWN	
AAA	16.4%	1	19.8%
AA	33.8%	2	17.7%
A	31.6%	3	15.8%
BBB	10.0%	5	11.9%
BB	5.0%	7	8.5%
B	2.0%	10	6.4%
CCC	1.0%	12	3.8%
		15	0.2%
		20	0.0%
		30	16.0%

VARIABLE ANNUITIES

- Based on a solvency II internal model maintained by a large VA writer.
- Includes products with GLWBs, GMDBs, and GMIBs.
- $30 billion of account value.
- Market-consistent BEL (net of Account Value) of -$1.27 billion.
- Hedged with an initial hedge position of futures covering 50% of the NetAV delta and swaps covering 80% of the NetAV rho measured at the 2, 7 and 20 year key rates. Dynamic hedging is not modeled.
- VA reserve capital consists of a $5 billion bond portfolio with the same characteristics as the traditional life bond portfolio.

REAL-WORLD SCENARIOS

Changes in the balance sheet are projected over 100k real-world scenarios representing an instantaneous shock to market conditions consistent with a 1y stress. Baseline market conditions roughly equal to those on 12/31/2013. The real-world scenario generator utilizes stochastic volatility models with jump processes to project the evolution of the risk factors over one year. Parameters for each risk factor's model and the correlations between different risk factors' parameters are calibrated from historical market data.

VALUATION TECHNIQUES

All liability models are evaluated using proxy models developed by sampling existing liability actuarial 'heavy' models. Assets are modeled seriatim using ECSight's asset models.

MODEL RISK FACTORS

The following risk factors are present in the scenarios:

- Market
  - Equity level
  - Equity volatility level
  - AAA-CCC spreads and default/transition haircuts
  - Three interest rate factors, interpolated to form a full curve
- Insurance
  - Lapse by LOB, with baseline/post-level-term factors for traditional
  - Mortality by LOB, excluding deferred annuities, with baseline/improvement factors for traditional

The following table identifies which risk factors are applicable to the underlying balance sheet items in the model:

RISK FACTOR / BALANCE SHEET ITEM	EQUITY LEVEL	EQUITY VOL	CREDIT SPREADS	CREDIT MIGRATION &DEFAULT	INTEREST RATES	LAPSE	MORTALITY
DEFERRED ANNUITIES					X	X	
TRADITIONAL LIFE					X	X	X
VARIABLE ANNUITIES	X	X			X	X	X
FIXED INCOME BONDS			X	X	X		
FUTURES	X				X		
SWAPS					X		

MODEL RESULTS

The following table summarizes the opening balance sheet and the absolute 99.5 percentile undiversified losses by line of business and the diversified aggregate loss based on the real-world distribution of balance sheet changes.

LOB	NPV	99.5 PERCENTILE LOSS
DEFERRED ANNUITY	2.49B	-2.29B
TRADITIONAL LIFE	3.88B	-2.63B
VARIABLE ANNUITY	6.27B	-4.88B
GRAND TOTAL	12.64B	-5.94B

The following chart plots the simulated real-world distribution of losses:

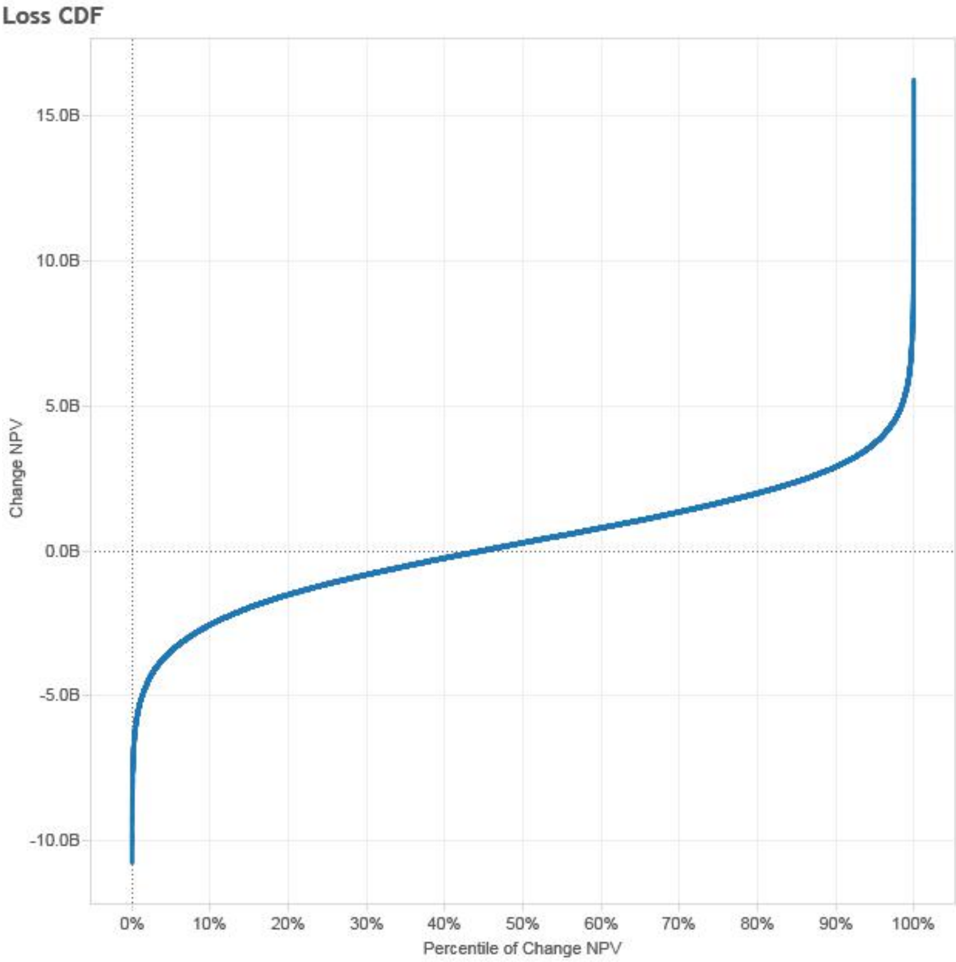

Milliman is among the world's largest providers of actuarial and related products and services. The firm has consulting practices in life insurance and financial services, property & casualty insurance, healthcare, and employee benefits. Founded in 1947, Milliman is an independent firm with offices in major cities around the globe.

milliman.com

CONTACT

Neil Cantle
neil.cantle@milliman.com

Corey Grigg
corey.grigg@milliman.com