Applying Decision Analysis to Human Factors in Decision Making at Stanford University Medical Center

John Celona

Abstract Decision-making is a fundamental life skill, yet people reliably and predictably make decisions which, upon a more careful consideration, they no longer agree with. Developed over the last 50 years, decision analysis drew from statistical decision theory, cognitive psychology, economics, and system engineering to comprise a philosophy, theory, and methodology for making better decisions in complex and uncertain situations. We briefly discuss decision analysis, then show how it was applied to making better decisions in medical malpractice cases at Stanford University Medical Center.

Keywords Decision analysis · Human factors · Medical malpractice · Insurance reserve decisions

1 Introduction

As mankind greatly altered its living environment over the last few 1000 years, some human traits well honed by evolution in the preceding millennia appear now to be maladaptive—and evolution has not yet had time to catch up.

For example, metabolisms and calorie storage capacities which worked well when brief spurts of plenty were inevitably followed privation result in excess fat accumulation in the modern situation of perpetual plenty. Health suffers as a result.

Other potentially helpful traits have yet to develop. The ability to know which way is up when the horizon is not visible is basic to birds but absent in humans, who never had an issue knowing which way the ground was. This deficit challenges pilots who must consciously ignore intuition and rely on instruments to avoid crashing.

J. Celona (✉)
Decision Analysis Associates, LLC, San Carlos, CA, USA
e-mail: JCelona@DecisionAA.com

J.I. Kantola et al. (eds.), *Advances in Human Factors, Business Management, Training and Education*, Advances in Intelligent Systems and Computing 498,
DOI 10.1007/978-3-319-42070-7_39

So it is with decision-making. Our thought processes and decision-making heuristics are well-tuned for avoiding being eaten by predators and evaluating possibly hostile intentions when encountering another human.

However, as humanity has developed its environment and economy for greater comfort, safety, convenience and wealth, it has created complex social, economic, and technological systems which defy purely intuitive evaluation.

This development has also created new and greater sources of uncertainty for which we have not yet developed reliable intuition. Greater development magnifies the impacts of natural sources of uncertainty (storms, floods, earthquakes, etc.). Economies and financial markets turn rapidly and unpredictably on the collective sentiments of large numbers of people about whom we have some or no information, and no direct contact to evaluate their judgments or intentions.

Decision analysis was developed to address these human factors in decision-making through providing a philosophy, theory, and methods of making better decisions in even the most complex and uncertain situations. Decision analysis provides a path to clearly understanding which alternative is best and why which stands up to careful scrutiny and maximizes your odds of enjoying a favorable outcome.

2 How Decision Analysis Works

Thankfully, evolution has equipped us with two very powerful and complimentary thinking and decision-making modes which can interact to develop and reprogram how we make decisions and to make better decisions. Decision analysis structures the action and interaction of both these modes.

This first mode is the fast, intuitive mode which I call the “Intuitive Supercomputer” and which Daniel Kahneman refers to as “Type 1” thinking [1]. This is the way we make most decisions: we just sort of think about it a bit (or not!) and then decide.

The Intuitive Supercomputer is fast, easy, intuitive, and works well most of the time. It employs subtle and complex reasoning processes which are difficult to trace or describe. It is highly specialized for understanding other people and governs emotions, trust, empathy and—most importantly—action! People are never convinced or act unless their intuition is on board.

However, this mode of thinking and decision making reliably and predictably goes astray with even a little uncertainty or complexity [1]. Hence the need to supplement it.

Fortunately, we are also equipped with another mode which I call the “Logical Co-processor” and which Kahneman refers to as “Type 2” thinking. It is slower and more effort to invoke, but less prone to error. It requires concentration and focus—at the risk of possibly solving the wrong problem. It uses transparent, logical reasoning, but does not lead to action unless trusted and persuasive. Hence the need to convince the Intuitive Supercomputer.

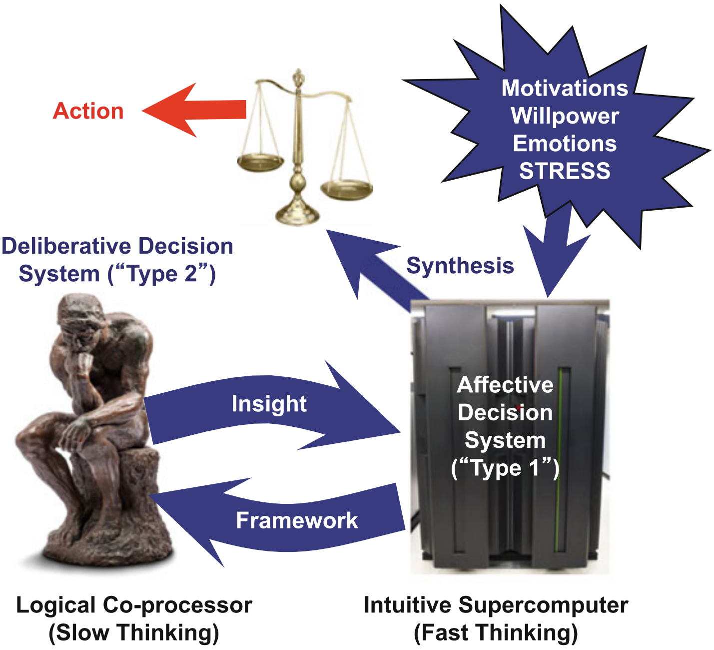

Fig. 1 Decision analysis frames the interchange between your Intuitive Supercomputer and logical co-processor to create a clear and persuasive rationale for which course of action is best and why

Decision analysis works by using intuition to construct the frame for the decision: the alternatives, preferences, and information. A great alternative is like a great idea: there is no set process for coming up with one; you just do. Intuition works great for this. There are also brainstorming techniques which can help.

The logical co-processor can then consider this frame more carefully. What could happen? What is that worth? What are the chances of that or something else? In decision analysis we construct a quantitative model for analysis, including evaluating alternative scenarios. We refer the reader elsewhere for details of how to undertake and evaluate formal analysis [2–4].

The results of the analysis generate insights which feed back to the Intuitive Supercomputer for synthesis: do the results make sense? Are we missing something? Do we have a cogent and clear understanding of which alternative is best and why which can be briefly, persuasively communicated?

This powerful, structured interaction of our two innate thinking modes is the strength of decision analysis. These relationships are shown in Fig. 1.

When Einstein created his Theory of Relativity, he was asked “Where is your lab work?” He didn’t get there via physical experimentation. Instead, he conducted his famous “thought experiments” (the trains moving past each other, etc.) to create the theory [5].

In the same way, decision analysis allows you to create thought experiments as to which course of action could be best and why *before* acting. This kind of forethought is absolutely critical for most important decisions on life in which you get only one chance to act and cannot “rewind the tape” to try a different plan.

3 The Elements of Decision Analysis

One of the elements of a decision analysis (and decisions in general) is the alternatives: what are the different possible things you could do? One of the key distinguishing features of decision analysis is we always examine multiple alternatives, including the “do nothing” alternative against which we can compare the possible risks and benefits of doing something. This is in stark contrast to the more common advocacy mode of decision making in which there one proposed course of action. The complete set of elements is shown in Fig. 2.

A second element is information, which includes all the data one can muster, plus quantification of the possible things that could happen. In cases where there is (1) a good historical data set; and (2) a stable system, statistical analysis of historical data can furnish decent prospects for future uncertainty. One example is pricing life insurance. There is a lot of data and people aren’t suddenly living longer or dying sooner, so analysis of historical data is a good guide.

However, these two conditions are often lacking and this shortfall is typically ignored: analyzing data is the hammer people know how to pound quite well. The results can be catastrophic, as in the use of historical data to price mortgage-based securities leading to the 2008 financial crisis (as related to the author by people at the ratings agencies). Hence the need to supplement data analysis with subjective, Bayesian probabilities.

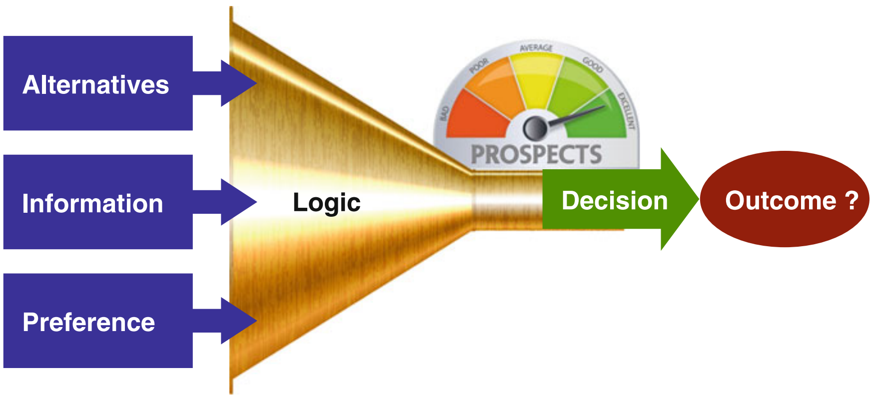

Fig. 2 The elements of a good decision

Bayesian probabilities date back over three centuries to the Reverend Thomas Bayes and most people are unfamiliar with them though they see them every day. Check your weather forecast: the chance of rain a week from tomorrow is a Bayesian probability expressed as a percentage (e.g., 40 %). Check it in 2 days and the number will likely be different—though the available historical data has not changed a bit. Bayesian probabilities are a person or persons' view as to the likelihood of a particular event based on what they know at the time. When the information or your evaluation of it changes, we would expect the probability to likewise change.

Use of Bayesian probability is one of the key distinguishing features of decision analysis and why it works better. It is also one of the most controversial, but it is a battle well worth fighting. The fact that it works so well is the reason decision analysis has become standard in risky, high-stakes industries like drug development and upstream oil and gas. Before a drug is developed at Pfizer, Merck or Genentech or a project is undertaken at Exxon or Chevron, there is a decision analysis of its prospects [6].

A third element is preferences, of which there are two parts. A time value of money is used to convert funds flows occurring over time to net present values for ready comparisons to other flow patterns. Second, a risk preference is used to translate uncertain prospects to what the equivalent certain amount would be for the decision maker(s).

This use of a quantified attitude towards risk as another distinguishing feature of decision analysis. It contrasts with the "risk tolerance" approach commonly applied where it is simply specified "we won't tolerate X"—with no inquiry as to what the costs of preventing X or allowing a little X would be.

These elements are put into a quantitative model (the "thought experiment") to analyze the potential outcomes. The results typically include values for individual scenarios and a probability distribution on total value to show the complete set of prospects for a particular alternative. Distributions are often summarized with a mean value.

These results drive the ultimate decision, following which is an outcome: what actually happens once the decision has been made and acted upon.

Decision analysis also furnishes a definition of a good decision: one that is logically consistent with everything known *at the time the decision is made*. We distinguish a good outcome as what you hope will happen.

This approach is in stark contrast to how people typically evaluate decisions: it was a good decision if things turned out well, and a bad decision if they don't. Promotions and compensation plans as well as political fortunes typically work this way. In many cases, they reward luck rather than a well-considered decision.

Decision analysts take this different perspective and, in fact, document the basis for the decision at the time it is made for later review in light of how things turn out. This saves jobs and careers as well as improving the quality of future decisions.

4 Applying Decision Analysis at Stanford University Medical Center

Stanford Hospital self-insures medical malpractice liability through a Bermuda-based captive insurance company, the Stanford University Medical Indemnity Trust (SUMIT). Accordingly, for each case SUMIT must manage the decisions facing client, counsel, and the insurance company, namely:

- Whether to settle or continue litigating;
- How much to offer or pay in settlement negotiations;
- Where to direct legal and factual investigations;
- How much to reserve for potential damages (indemnity);
- How much to reserve for costs (legal fees and experts).

Amounts required to fund SUMIT are charged as insurance premiums to Stanford Hospital.

Our involvement had a number of objectives:

- Insure that amounts reserved were enough to cover potential exposure (avoiding unpleasant surprises and the need to put in more money) without being over-reserved (which means Stanford Hospital was charged more than necessary for insurance premiums).
- Ensure that total reserves met regulatory and prudential standards for captive insurance companies.
- Set early and stable reserves to eliminate "stair step" increases in reserves.
- Minimize the total cost of cases by resolving them for less than exposure at trial.

Applying decision analysis furnished the means of achieving the objectives. The question was how to apply it routinely and economically for each and every case. To do so, we thought like a decision analyst and considered a number of alternatives.

The first alternative was continuing the status quo of making educated guesses based on experience. However, typical prior practice was starting reserves at a token amount and continually increasing them as a case progressed. This situation was flagged as a problem because reserves at any one time were neither a good picture of the outstanding exposure nor sufficient provision to cover it. In addition, there were a number of bad surprises when cases required far more to resolve than what was reserved.

The second alternative was to develop and fill out templates for each case, as is commonly done for drug and oil development projects. The concern here was that the templates would freeze the level of analysis at a fairly basic level not adequate for the variation in the cases and which didn't provide needed guidance for case management.

The third possibility was developing an expert system ("analyst in a box") for case analysis, as I have worked on for portfolios of R&D projects. However, the cost and time required to develop such a system are considerable.

The last alternative was to develop the “wetware”: the expertise of staff assisted by an expert analyst (myself) to quickly conduct custom analysis for each case. The issue here was the steep learning curve for staff attorneys, doctors, and claims managers with no prior experience in decision analysis.

The first alternative was unattractive because it merely continued the issues leading to this effort. The second and third required a considerable up-front investment along with their other issues. We decided to try the last alternative both to get started right away and because, if it proved unsatisfactory, the work completed would provide a necessary basis for switching to the second or third alternatives.

5 Getting Started

We began with creating a business process to implement the last alternative of conducting a rapid, custom analysis of each case. This process is called the Decision Analysis Reserve and Trial Strategy (DARTS) process. Following an initial 2-day training session, it involved:

- Convening regular meetings to discuss, analyze, and update cases.
- Developing and analyzing a custom model for each case in real time (2–3 h for most cases).
- Immediately feeding back the results: Do they make sense? Are we missing something? What if…?
- Continually developing staff understanding and expertise.

For each required decision, we specified the decision analysis result which would guide it, as described in Table 1.

Initial efforts were focused on introducing staff to thinking in terms of possibilities, probabilities, and alternative scenarios. Figure 3 shows the decision tree from one of the very early analyses.

Note that the Special Damages appearing after General Damages not connected implies node replication: no matter how General Damages turns out, there is still an uncertainty on Special Damages with possible outcomes ranging from $5–15 K.

Table 1 Required decisions and guidance from decision analysis results from DARTS

Decision	DARTS output guidance
Indemnity reserve	Mean value of exposure at trial
Cost reserve	Mean value of total cost exposure
Settle or litigate	Go to trial if no settlement for less than trial exposure
Settlement offer	Less than mean value of trial exposure
Legal and factual investigations	Key issues identified in discussion and sensitivity analysis of trial exposure

Fig. 3 Decision tree from an early DARTS case analysis

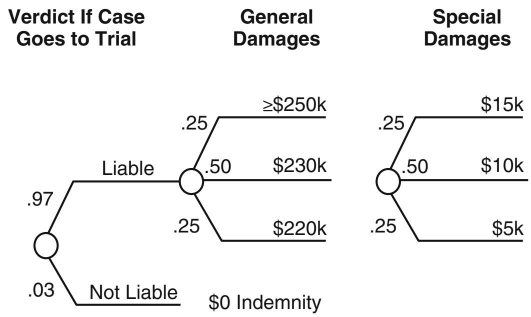

This early tree illustrates the usual difficulties in beginning to think about ranges of uncertainty. The ranges on both General Damages and Special Damages are very likely much more narrow than what the actual uncertainty was in this case. Thankfully, assessing Bayesian probabilities and thinking in terms of alternate possibilities is a learned skill and people get better over time with practice. Figure 4 shows a tree from a more recent case and the vast difference in how staff's thinking about cases changed over time.

Here we see a much more sophisticated approach to modeling the issues and uncertainties in a case. In particular, there are many conditional probability assessments, such whether Dr. X's testimony for the defendant is interpreted by the jury as an admission of fault affecting the "Balance of the Experts": which expert witness are persuasive to the jury. The conditional assessments in this tree are so

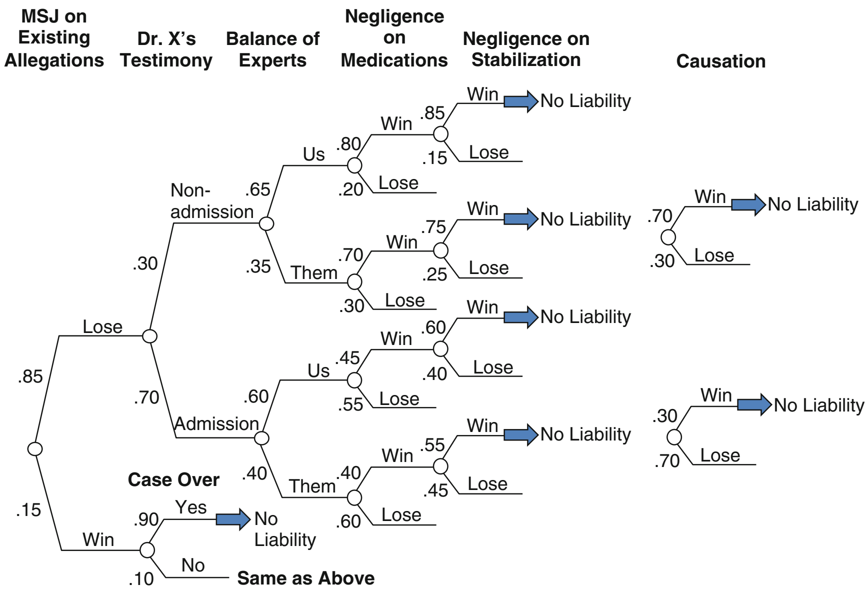

Fig. 4 Decision tree from a recent DARTS case

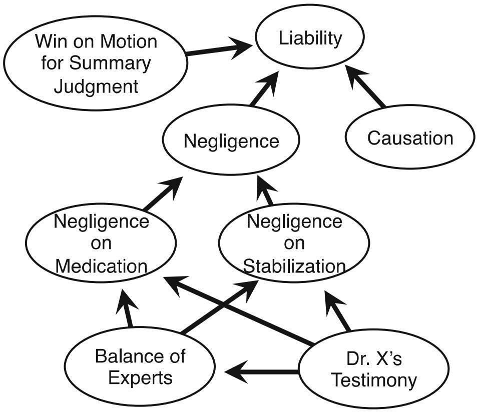

Fig. 5 Relationships from decision tree shown in relevance diagram form

complex it is perhaps easier to view those dependencies in the form of a relevance diagram (Fig. 5).

In my experience, dependencies between decisions and uncertainties are often present, usually critical when they are, and often missed when using standard analysis templates. Capturing the dependencies in a particular case is one of the great strengths of doing a custom but "as simple as possible but no more" analysis.

This greater sophistication in thinking about possibilities and probabilities in uncertain and complex cases is not only reflected in better analysis results, but also in staff's perception of the value of applying decision analysis. They uniformly agree that much of the value comes from the improved understanding of the case that the process yields. This better informs them in managing the case and in dealing with opposing parties.

There is no magic in the decision analysis process. It just supplies a structured and theoretically sound process for making better use of your innate cognitive abilities (intuitive and logical) to reach a better understanding of what is best to do and why.

6 Results

For a single case or decision, it can be difficult to evaluate how good the decision analysis process was. Quantifying and analyzing the uncertainty does not make it go away, and a good decision can still result in an unfortunate outcome.

However, over multiple decisions (claims in this case), it is possible to compare predicted against actual results and see how well your decision analysis process is guiding your decisions. Pharmaceutical and oil and gas companies do this routinely, and we were able to do this for Stanford.

Over the first 2 years, we analyzed fifty cases. Twenty-five have been resolved and the rest were still pending as of when we looked at the results. Seven cases were

Table 2 Results and predicted probability of liability at trial for cases resolved with no indemnity payment

How case was resolved	Predicted probability of liability at trial (%)	Indemnity reserve
Defense verdict at trial	25	$106 K
Case dismissed	5	$40 K
Summary judgment for defense	5	$76 K
Statute of limitations expired Before case filed	34	$873 K
Case dismissed	24	$75 K
Case dismissed	20	$59 K
Summary judgment for defense	5	$83 K

resolved with no indemnity payment at all. The first question, then, was how those case were resolved and if we correctly predicted a low probability of liability if the case went to trial. The results are shown in Table 2.

Note that only the first case actually went to trial, and that the amounts reserved for indemnity accounted for the predicted probability of liability if a case went to trial. The mean indemnity given a finding of liability at trial is the indemnity reserve divided by the probability of liability (e.g., 106 ÷ 0.25 = $424,000 for the first case).

Our working conclusion (always pending more data) is that the decision analysis was able to correctly predict low prospects for liability if a case went to trial in cases resolved with no indemnity payment.

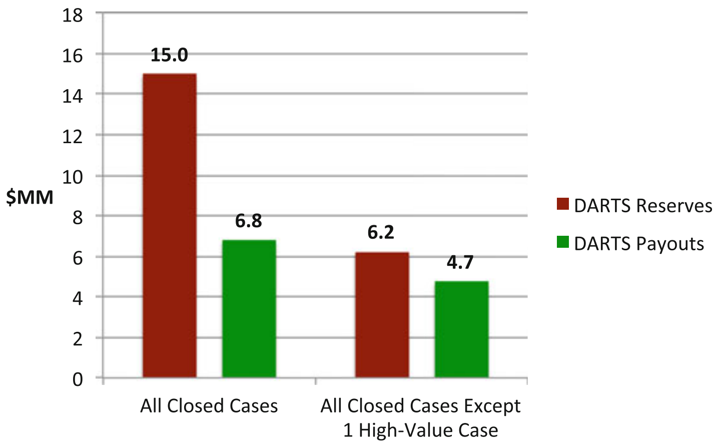

Fig. 6 Amounts reserved versus actual payouts for all closed cases

Fig. 7 Payouts versus reserves for cases with an indemnity payment excluding one high-value case

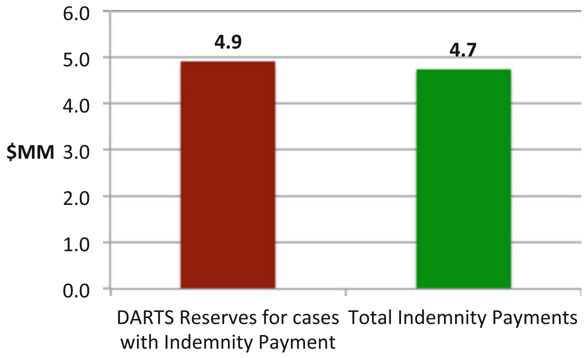

Next we looked at total payouts versus reserve amounts for all closed cases (Fig. 6).

These results include the cases resolved for no indemnity payment. In total, we were successful in resolving cases for less than the predicted exposure at trial. The next view was looking only at cases for which there was an indemnity payment except the one high-value case in which an exceptionally favorable settlement was reached (Fig. 7).

This chart shows close matching between predicted exposure at trial and amounts paid to resolve cases, and success in resolving case for less than exposure at trial. We next looked at variance between the amounts to resolve cases and the mean exposures from the decision analysis to see if there were systematic biases in the estimates (Fig. 8).

The distribution looks roughly uniform, indicating no systematic biases in over- or under-estimating the exposure. We would expect variation from the calculated mean because analyzing the uncertainty does not make it go away.

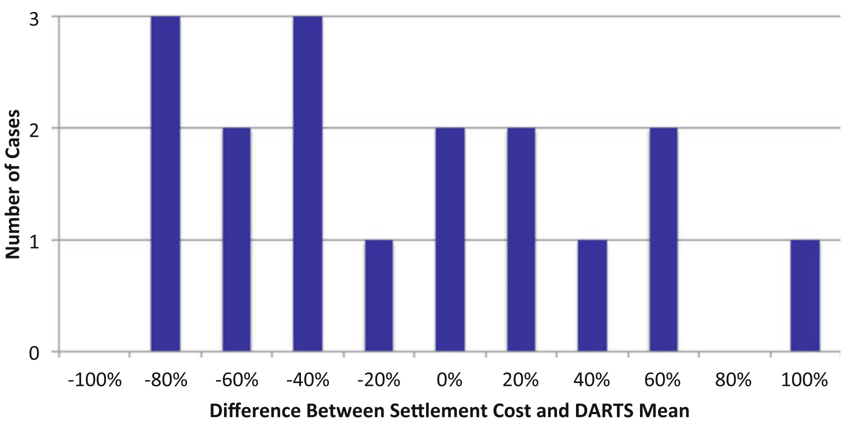

Fig. 8 Differences between amounts to resolve cases and analyzed mean exposures at trial as a percentage of the mean exposure at trial

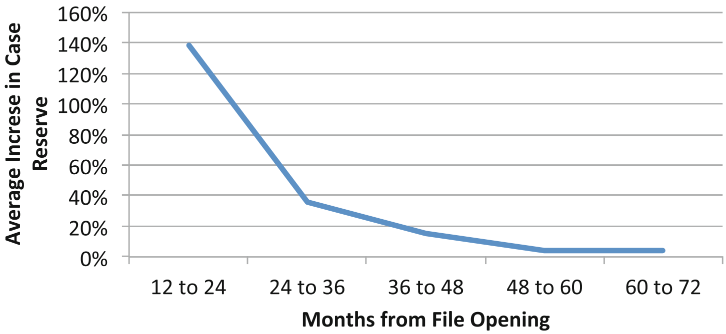

Fig. 9 Prior average change in case reserves over time

Fig. 10 Changes in case reserves with revised decision analyses

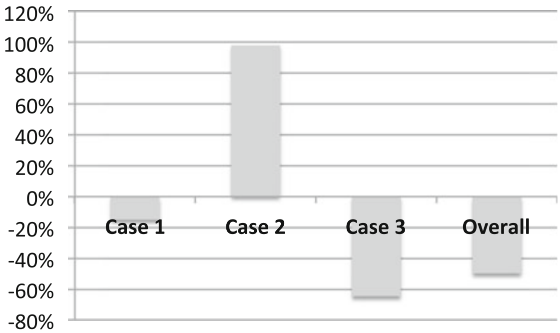

The last question was how well we did in making an early and stable evaluation of case exposure to avoid the prior "stair step" increases in case reserves. The changes in amounts reserved are known as "loss development." Following is Stanford's prior and typical loss development from actuarial methods (analysis of historical data) (Fig. 9).

Our experience with cases we analyzed was quite different. For almost all the cases, there was no change at all in case reserves—despite repeated reviews as a case developed and new information came in. In only three cases was the reserve changed.

Overall, the change in total reserves was a decrease. In actuarial terms, this would be known as negative loss development and is entirely contrary to usual experience (Fig. 10).

7 Conclusions

We were successful in achieving our objectives of resolving cases for less than the potential exposure at trial, and in achieving early and stable evaluation of case exposure. This was despite the great uncertainty, complexity, variation, and paucity of information for many of the cases. Staff learned that even when "there was no information," they actually had quite a lot of information for evaluating case prospects.

In addition, we were successful at identifying which issues and uncertainties were critical in determining exposure, and in directing subsequent legal and factual inquiries into those issues. This resulted in considering savings in legal fees, discovery, and expert witness costs.

Lastly, we would note that, although this work was for cases in which Stanford was a defendant, the methods work equally well for plaintiffs and for insurance companies on both sides of a case.

Acknowledgments The author would like to acknowledge the expert and dedicated staff at Stanford Hospital, the Stanford University Medical Indemnity Trust, and The Risk Authority, Stanford for their efforts in making this work possible. Also, many thanks to Joseph Celona for reviewing the manuscript.

References

1. Kahneman, D.: Thinking, Fast and Slow. Farrar, Straus and Giroux, New York (2011)
2. McNamee, P., Celona, J.: Decision Analysis for the Professional, 4th edn. SmartOrg., Inc., Menlo Park, California (2007)
3. Howard, R.A., Abbas, A.E.: Foundations of Decision Analysis. Pearson, New York (2016)
4. Celona, J.: Winning at Litigation through Decision Analysis. Springer, New York (2016)
5. Einstein, E.: Relativity. Crown Publishers, Inc., New York (1961)
6. https://www.youtube.com/watch?v=JRCxZA6ay3M