On Coconuts in Foggy Mine-Fields: An approach to studying future-choice decisions

Robin M. Hogarth
Universitat Pompeu Fabra, Barcelona, Spain

Paper prepared for
"The Irrational Economist" Conference
and Book Writing in honor of Howard Kunreuther

December 4-5, 2008

Chapter to appear in
The Irrational Economist: Future Directions in Behavioral Economics and Risk Management
E. Michel-Kerjan and P. Slovic (eds.) *Public Affairs Press.* (forthcoming 2009)

On coconuts in foggy mine-fields: An approach to studying future-choice decisions

Robin M. Hogarth

Universitat Pompeu Fabra, Barcelona, Spain

robin.hogarth@upf.edu

November 7, 2008

One of the outstanding features of Howard Kunreuther's work has been a simultaneous willingness to both use theory and to reject the orthodox. I would like to pay tribute to Howard by building on this tradition to consider what I believe is a critically important conceptual and practical problem in the field of decision making. This is the issue of what I call "future-choice" decisions – decisions that have three important sources of complexity. First, actions taken today can have unknown consequences at future horizons that are difficult to specify. Second, decisions imply difficult inter-temporal tradeoffs. And third, it is problematic to specify relevant states of the world let alone assess their probabilities.

Consider, for example, taking decisions today that will affect the layout of a city. How far into the future should such a decision look? What volumes of traffic are likely to develop on alternative highways? What currently unknown technological advances could change costs and benefits? Will public preferences remain stable? What are the possibilities of different natural or social disasters? And so on. The list of possibilities and complications is endless. Imagine further a 23-year old planning a career and ways of acquiring capital across her life. How can she evaluate different career paths? How can she predict changes in her personal situation as well as health? How can she assess tradeoffs between "capital" and "consumption" at different times?

Today, the "standard" economic model for dealing with these situations is the discounted utility (DU) model. However, from both descriptive and prescriptive perspectives, extensive empirical evidence shows that this model has little to recommend itself (see, e.g., the comprehensive review by Frederick, Loewenstein, & O'Donoghue, 2002).

The plan of this paper is as follows. I first want to salute the considerable advances made in our understanding of these issues that have been gained from decision theory. In so doing, however, one needs to keep clearly in mind just what a good decision analysis can and cannot take into account. Second, I will argue that – despite many valiant attempts – much of the phenomena that occur in the socio-economic world cannot be predicted with any accuracy. Third, I will go on to distinguish two types of uncertainty that characterize decisions labeled *subway* and *coconut* uncertainty, respectively. Fourth, the presence of coconut uncertainty essentially implies the breakdown of decision-theoretic and forecasting models and demands a new approach for future-choice decisions. Whereas I cannot offer a new approach – or an intellectual breakthrough – I can suggest a heuristic principle and metaphor that may help us deal with some of the issues. Perhaps by discussing their advantages and disadvantages, we might illuminate some of the problems of future-choice decisions and, at the least, set an agenda for future research.

The beauty of decision theory

Like many graduate students of my generation – and some of our predecessors – I was totally seduced by statistical decision theory. Based on the rigorous analytical framework provided by such eminent scholars as John von Neumann, Oskar Morgenstern, Bruno de Finetti, and L. J. Savage – and interpreted so brilliantly by Robert Shlaifer, Howard Raiffa, Ralph Keeney and others – we had in our hands a tool for making decisions that was not only theoretically sound but capable of synthesizing both the hard and soft data that characterize most naturally occurring decisions. True, it was often hard to obtain all necessary inputs to make the mechanism work but the mere exercise of going through the analysis could yield many insights.

But statistical decision theory never really became the universal tool that many of us imagined it would. This is not to say, however, that it did not develop. Indeed, scholars are still finding ways of making it better or more practical and I applaud these developments.

Nor do I believe that decision theory has been held up by the fact that many of its axioms are violated by people's choices in the so-called choice paradox literature. At one level, these so-called paradoxes illustrate precisely why a good theory is necessary.

My belief is that statistical decision theory fails in many important problems – and particularly future-choice problems – because humans are incapable of characterizing the uncertainties of the world in which they operate. Decision theory strictly speaking only applies to the "small worlds" discussed by Savage (1954), the limitations of which he was so clearly aware.

The predictability of socio-economic phenomena

There is no doubt that the ability to predict events in the environment has been an important element of human development. Indeed, there is much that can be predicted with stunning accuracy. Think of the tides, solar eclipses, and many related physical phenomena. At the same time, there is much that cannot be predicted, the timing and location of large earthquakes, for example, or certain kinds of weather patterns. When it comes to predicting many socio-economic phenomena, the human predictive track record is frankly dismal. As examples, I won't comment here on the extensive failures to predict human behavior in psychology which are well-known (see, e.g., Dawes, 1994) but instead briefly comment on predictions made in economics and business.

A critical operational concern in economics and business (private and public) is the forecasting of many different time series. Consider, for example, data concerning imports and exports across time, the supply and demand for specific products and classes of goods, inventories, and various economic indicators. Forecasting these variables with a reasonable level of accuracy is essential because, without good forecasts, individuals and firms cannot plan and economic activity suffers.

In 1979, Spyros Makridakis and Michèle Hibon published a paper in which they compared the out-of-sample forecasting performance of 22 forecasting methods on 111 time series obtained from various sources in business and economics.[1] Their methodology was conceptually simple: separate each time series into a fitting phase and predictive phase; fit all models on the fitting data; use the fitted models to make predictions for the predictive phase; and compare predictions with actual outcomes.

[1] For further elaboration on the description of the work by Makridakis, see Hogarth (in press).

The results surprised even the authors: "…if a single user had to forecast for all 111 series, he would have achieved the best results by using exponential smoothing methods after adjusting the data for seasonality" (Makridakis & Hibon, 1979, p. 101). In other words, a remarkably simple model (that essentially combines only the last few observations) outpredicted many complex and statistically sophisticated models that used many variables and provided closer fits to the data in the fitting phase of the analyses. The essential point made by Makridakis and Hibon was also conceptually simple: The real processes underlying time series in business and economics do not conform fully with the assumptions of complex statistical models and thus extreme caution should be taken when predicting out-of-sample. Thus, even though the complex models can fit past data well, their predictive ability in future samples falls short of the performance of their simpler counterparts.

Comments on the Makridakis and Hibon paper were published by the *Journal of the Royal Statistical Society* and make interesting reading today. The theoretical statisticians were not complimentary. However, Makridakis's reactions since 1979 have been exemplary.

In 1982, he published results of the so-called M-competition (Makridakis et al., 1982) in which experts in different forecasting methods were invited to predict 1001 series. In 1993, results of the M-2 competition appeared (Makridakis et al., 1993). This competition was similar to the M-competition in that experts were invited to use their own methods. It differed, however, in that there were fewer forecasts but these were conducted in real time (e.g., participants were asked to provide a forecast for next year). Moreover, forecasters could obtain background and qualitative data on the series they were asked to forecast. Finally, in the M-3 competition (Makridakis & Hibon, 2000), forecasts were prepared for several models using 3003 time series drawn from various areas of economic activity and for different forecast horizons. All of these M-competitions (along with similar studies by other scholars) essentially replicated the earlier findings of Makridakis and Hibon, namely:

> (a) Statistically sophisticated or complex methods do not necessarily provide more accurate forecasts than simpler ones. (b) The relative ranking of the performance of the various methods varies according to the accuracy measure being used. (c) The accuracy when various methods are being combined outperforms, on average, the individual methods being combined and does very well in comparison to the other methods. (d) The accuracy of the various methods depends on the length of the forecasting horizon involved. (Makridakis & Hibon, 2000, p. 452.)

The empirical evidence is clear. Predictions in the socio-economic domain are fraught with all kinds of obstacles involving non-stationarity and the implications of events that could not even be imagined. The big need is to understand just what is predictable and to what extent.

<u>Two types of uncertainty.</u>

I think that what really hampers the use of formal analytical techniques for future-choice decisions is that uncertainty can take two forms. In a forthcoming book

(Makridakis, Hogarth, & Gaba, in press), we refer to these as *subway* and *coconut* uncertainty, respectively.

In explaining these concepts, we cite the example of a mythical and obsessive character called Klaus who everyday charts his commuting times to and from work (using the subway) so that, over a certain period, he accumulates a statistical representation of his travel times in the form of what is essentially a normal distribution – see Figure 1. Indeed, assuming that nothing systematic disturbs the statistical pattern observed, Klaus is able to use the characteristics of the normal distribution to calculate the probability that his arrival time on any given day will fall within specified limits. Moreover, he can actually validate his model's predictions.

By subway uncertainty, then, we mean a source of uncertainty where the statistical properties are well known and where future "surprises" fall within well-specified limits. This kind of uncertainty – or approximations thereof – can be well handled within our decision-theoretic and forecasting models.

Insert Figures 1 and 2 about here

By coconut uncertainty we mean something quite different. Imagine you are sitting under a palm-tree on a South Seas island and a coconut happens to fall on your head causing you considerable distress. Now there are many disasters that you can imagine in life, but probably being hit on the head by a coconut is not one of them. In other words, by realizations of coconut uncertainty we mean events that you probably never even imagined could occur – there are so many different ones and you don't know which particular one will happen. Indeed, you might not even have a good handle on the class of events that could be described as coconuts. And, of course, coconuts could be positive as well as negative in their impact.[2]

Interestingly, in some domains there is a history of coconuts and, even despite this, people are still surprised by their occurrence. Daily returns on the stock market provide a case in point. Figure 2 shows the distribution of daily returns of the Dow Jones Industrial Index (DJIA) for the period from January 1, 1900 to December 31, 2007. As you can see, many observations lie outside the plus/minus three standard deviations limits and this graph does not even contain the observations for the wild market movements that occurred over the last six months. Parenthetically, if you look at the same data as a time series, there are periods – mainly at "crisis" times (e.g., October 1987 or 2001) – where there seems to be serial dependence in the size of fluctuations, rather like the pre- and after-shocks that accompany earthquakes. I deliberately use the stock market data example because, even with a publicly available history of coconuts, people are still surprised when they occur. In principle, in considering the stock market and other financial data one could always model some coconuts; but for some reason (perhaps analytical tractability) many practitioners fail to do this and inevitably suffer the consequences. (Recall the Long Term Capital Management fiasco.)

My point, however, is not limited to the financial markets. My point is that for future-choice decisions many variables are inherently unknowable. We cannot characterize

[2] Nassim Nicholas Taleb (2007) refers to coconuts with large negative consequences as "black swans."

the uncertainty because, if we are honest, we cannot even specify the events that might or might not occur. For example, at the beginning of the 1980s how many foresaw the widespread use of personal computers and the development of the internet? Indeed, if you had been given correct forecasts of these developments at the time, would you have believed them? My contention is that you probably would not have known how to evaluate the forecasts.

What developments will occur over the next twenty-five years? I believe we are all quite blind. Moreover, if we simply extrapolate past trends, my guess is huge errors. In many ways, the path of social and economic development follows an evolutionary trail and, as is well known, although evolution provides a good story for explaining the past, it makes no predictions.

Only commit as far as you can predict

As a social scientist I would love to come up with a "good" or "amended" discounted utility model that could be used for future-choice decisions. However, the complexities are such that I do not think this is feasible. Instead, I would like to suggest a different strategy. This is to formulate a number of feasible heuristic rules that people can use to guide their actions when facing future-choice decisions. Whereas we must accept that no rule can be a guarantee, we could at least test such rules through simulation techniques and get some sense of their possibilities and limitations. Thus the rules would depend on more than just common sense.

The essence of the rule I would like to suggest here is to "only commit as far as you can predict." I would like to give it a fancy name but since I cannot think of one, I will call it *Occam* (in honor of the famous law – there is a weak resemblance). The rationale for Occam was suggested to me by the old joke about a Japanese Airline pilot who ditched his aircraft in the bay at San Francisco airport after a flight from Tokyo. "Why did you miss the airport by 200 yards?" a journalist asked. "Well," replied the pilot, "considering I came all the way from Tokyo, 200 yards was not much of an error."

The insight of this story is that, normally, when making commitments (in this case of where to land), a person (here a pilot) should match actions with the level of uncertainty that he or she is able to manage. Thus, when leaving Tokyo it is ridiculous to commit to the precise parameters for landing in San Francisco and all that the pilot should have done was to select a path that could be adapted as conditions changed. When close to San Francisco, however, a commitment to a landing spot could be selected because there would be considerably less uncertainty about the specific parameters. In other words, the level of commitment should be matched to the level of uncertainty.

Using the metaphor of the pilot, what are the implications of Occam for future-choice decisions?

First, many future-choice decisions have a temporal structure – similar to the Tokyo-San Francisco flight path. Consider, for example, problems involved in urban planning or young people thinking about their careers.[3] However, unlike the Tokyo-San Francisco flight that has a precise end-goal (i.e., reach a specific spot in San Francisco), these future-choice decisions don't necessarily have well-defined end-states. On the other hand, they are undoubtedly driven by a "direction" or values (e.g., to have a viable city or to achieve personal and professional success). The important point is that the time line of all of these decisions can be broken down into periods for which it is reasonable to make commitments that can be evaluated.

To simplify the language, I will use "targets" to refer to end-goals (e.g., touching down in San Francisco) and "sub-targets" as the intermediate goals (e.g., a position to be reached en route after, say, a couple of hours). Thus, in the case of urban development, an administrator might imagine the target to be the state of the city in 100 hundred years as, say, vibrant, prosperous, and easy-to-navigate. However, without being able to specify the uncertainties a century in advance, it would be foolish to make a large commitment. Instead, it would be more appropriate to set, say, a ten-year (or perhaps somewhat longer) horizon for sub-targets and commit to these.

From airplanes to foggy mine-fields

Although it makes a point, the Tokyo-San Francisco metaphor is too simple for many realistic future-choice decisions. So instead, I would like to introduce the metaphor of traveling across a mine-field in a fog. The general goal – as in life – is to get across the field in good shape, and let's assume that in crossing the field there are various positive awards that you can collect. But at the same time, your ability to see where you are going is restricted – randomly – by the density of the fog so that, whereas you might be able to see quite far in some cases, this is not going to be the general rule. The mines too are distributed randomly around the field and vary in how much damage they can inflict. That is, whereas some are "coconuts" others might be quite containable. You also have some diagnostic mine-detecting equipment but this is not 100% reliable and can be biased for some types of mines.

In the foggy mine-field, the target cannot be precise – it's just to get to the other side in the best possible condition. As to sub-targets, they are going to vary considerably depending on the state of the fog when you make any particular commitment. How then should one act?

It should be clear that decision making in the foggy mine-field cannot be modeled easily by, say, some form of dynamic programming. The reason is that the characteristics of the environment are not known in advance but are only revealed as you advance. (Sure, you can have some priors on what the general structure of the field is like, but there will still be events you cannot even imagine.) My contention is that future-choice decisions are often like foggy mine-fields and thus a really important research agenda will involve determining how to handle such environments.

[3] Urban planning provides a great example. Few would have seen the implications of the enormous growth of urban communities over the last one hundred years across the world in both economically rich and poor countries.

What research can be done on this? I argue that a good starting point would be to take rules like Occam described above and see how it performs in simulations that model different variations of the foggy mine-field paradigm. My suspicion – based on the work on heuristics in conventional environments (see, e.g., Hogarth & Karelaia, 2007), is that heuristics that "work well" will be those that match the environment in critical ways.

As an example of such a simulation, imagine the following video-type game which I shall simply call "life." You start with an endowment that has to be allocated across a portfolio of assets and the goal is to cross a foggy mine-field and reach the other side with a portfolio that has increased substantially in value. The assets in your portfolio differ in their characteristics. One (that simulates your health) decreases over time but you can invest resources to reduce the speed of decrease. If this asset hits a specified value, you are out of the game (i.e., death). The other assets are income producing but differ in their rates of return and how much you can invest in ways that are hard to predict – because of the fog. Some, for example, are more predictable than others (less affected by the fog). In the game, time is conceptually continuous but is represented by discrete periods…….if you don't take a specific decision at a given period to manage a specific asset, it will continue to be under the influence of the last time you made a decision concerning the asset. And, of course, there are potential dangers in the form of mines that can explode and destroy your assets (i.e., coconuts). You can however spend to get some (imperfect) information.

There are, of course, clearly many issues involved in setting up the parameters of foggy mine-fields such as this as well as technical issues of creating the actual simulation. However, in my opinion, the most interesting issue to investigate is to understand which decision rules or "heuristics" do well in these kinds of situations, as well as in variations of the game. How would different operational versions of Occam perform and how sensitive would these be to changes in the game's parameters? What other rules might experts in decision making suggest?

I would dearly like to get Howard's ideas and inputs on these ideas.

References

Dawes, R.M. (1994). *House of Cards: Psychology and Psychotherapy Built on Myth*. New York: The Free Press.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, *40*(2), 351-401.

Hogarth, R. M. (in press). When simple is hard to accept. In P. M. Todd, G. Gigerenzer, & The ABC Research Group (Eds.), *Ecological rationality: Intelligence in the world*. Oxford: Oxford University Press.

Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: Matching rules and environments. *Psychological Review*, *114* (3), 733-758.

Makridakis, S., & Hibon, M. (1979). Accuracy of forecasting: An empirical investigation (with discussion). *Journal of the Royal Statistical Society*. *Series A*, *142*, 97-145.

Makridakis, S., & Hibon, M. (2000). The M3-competition: Results, conclusions and implications. *International Journal of Forecasting*, *16*, 451-476.

Makridakis, S., Hogarth, R. M., & Gaba, A. (in press). *Dance with Chance: Harnessing the Power of Luck.* Oxford, UK: Oneworld.

Makridakis, S., Andersen, A., Carbone, R., Fildes, R., Hibon, M., Lewandowski, R., Newton, J., Parzen, E., & Winkler, R. (1982). The accuracy of extrapolation (time series) methods: Results of a forecasting competition. *Journal of Forecasting*, *1*, 111-153.

Makridakis, S., Chatfield, C., Hibon, M., Lawrence, M., Mills, T., Ord, K., & Simmons, L. F. (1993). The M-2 competition: A real-time judgmentally based forecasting study. *International Journal of Forecasting*, *9*, 5-23.

Savage, L. J. (1954). *The Foundations of Statistics*. New York, NY: Wiley.

Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable*. New York, NY: Random House.

Figure 1:   Klaus's travel time to work (from Makridakis et al., in press)

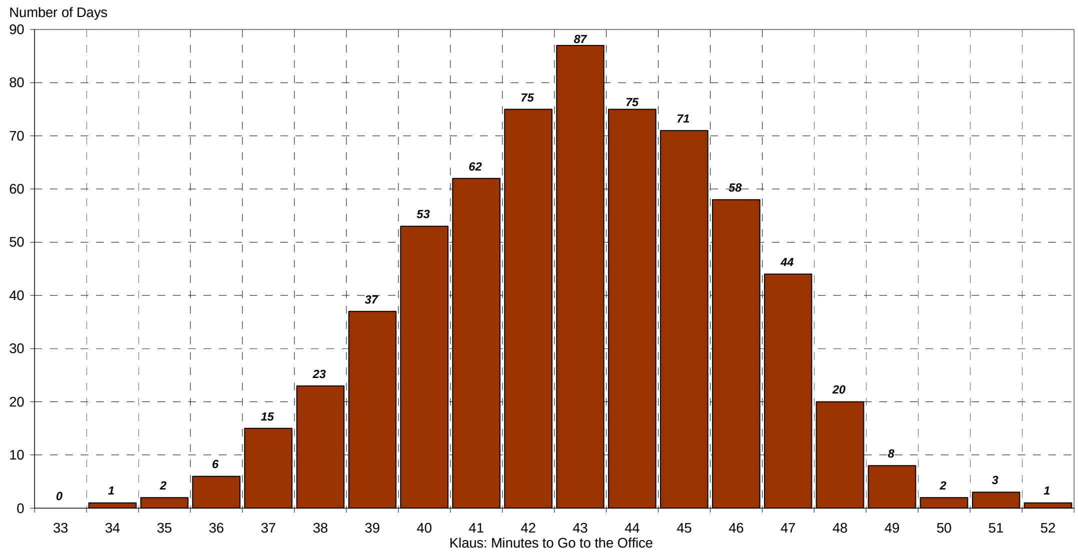

Figure 2: DJIA from January 1, 1900 to December 31, 2007
(from Makridakis et al., in press)