Decision Theory Made Relevant:

Between the Software and the Shrink[1]

by

Itzhak Gilboa[2], Maria Rouziou[3], and Olivier Sibony[4]

January 2018

Decision theory offers a formal approach to decision making, which is often viewed and taught as the rational way to approach managerial decisions. Half a century ago it generated high hopes of capturing and perhaps replacing intuition, and providing the "right" answer in practically all managerial situations. Today it seems fair to say that decision theory has not lived up to these expectations. Behavioral science provides ample evidence that managers fail to follow the dicta of decision theory, even when these are explained to them. As a result, executives often find decision theory frustrating and useless and prefer to rely on their intuition. This paper suggests that this extreme conclusion is unwarranted and calls for a re-appraisal of decision theory. We propose that it should not always be regarded as a mathematical tool that produces the *answer; rather, it can be viewed as a framework for a dialog between the decision maker and the decision theorist. In one extreme, the decision theorist studies the problem and provides the "correct'' answer. But in another, the decision theorist only challenges the decision maker's intuition and logic. In between, a whole gamut of possible dialogs exists, in which decision theory doesn't replace intuition, but supports and refines it.*

[1] We are grateful to the editors for comments and suggestions. Financial support from ISF Grant 1077/17 (Gilboa) is gratefully acknowledged.
[2] HEC, Paris, and Tel-Aviv University. tzachigilboa@gmail.com
[3] Owen School of Management, Vanderbilt University. maria.rouziou@owen.vanderbilt.edu
[4] HEC, Paris. oliviersibony@gmail.com

People and organizations have been conscious of their decision making at least since early history. Whether under certainty or uncertainty, as individuals or as groups, for a short or a long horizon, decisions needed to be made and some of them required conscious deliberation. Indeed, some of the basic ideas that decision theory has to offer can be found in ancient texts. For example, the idea that one's preferences might change and that, being aware of that, one may opt for pre-commitment, namely for less freedom, goes back to Homer's Odyssey, where Odysseus ties himself to a mast in order to enjoy the sirens' singing without being tempted to a lethal attempt of reaching them. Similarly, the notion that one should diversify one's portfolio can be found in the Old Testament, where Jacob explains why he splits his camp before meeting his brother Esau (who is powerful and has good reasons to bear a grudge).

The Age of Enlightenment has seen formal, analytical discussions of decisions. In the mid-17th century, Pascal described his famous "wager", which argues that even though the existence of God may have a low probability, the payoff of eternal life is infinite, which makes believing a rational choice. This line of argument introduced such ideas as the decision matrix, dominant strategies, subjective probabilities, and expected utility maximization.[5] A century later the Marquis de Condorcet discussed his well-known "paradox" for social choice, showing that majority voting may generate cyclical decisions. Yet, it is probably fair to say that decision making has not been the topic of scientific inquiry until the mid-20th century.[6] The mathematization of economics, the rise of operations research, the invention of game theory, and the development of mathematical tools of convex analysis and optimization merged to give birth to a new field, attempting to capture human decision making mathematically, be it for individuals or organizations. The general conceptual tools included the distinction between objective and constraints; the notion of constrained optimization; axiomatically-based paradigms for decision making (such as expected utility for dealing with uncertainty and discounted utility for decisions over time); and related ideas that also border on operations research, game theory, and microeconomics.

[5] See Hacking (1975, pp. 63-72). Expected utility maximization was explicitly suggested first by Bernoulli (1738), whereas the notion of subjective probabilities and their updating – by Bayes (1763).

[6] The founding fathers of the field in modern times are Ramsey (1926), de Finetti (1931,37), von Neumann and Morgenstern (1944), Savage (1954), and Anscombe and Aumann (1963). See also Arrow (1951).

There appear to have been good reasons for optimism in the 1950s. It certainly seemed that the field had come up with a general-purpose theory, which can be an aid both for descriptive purposes, describing how people *do* make decisions, and for normative ones, that is, recommending how they *should* make decisions.[7] Before proceeding with the story, it may be useful to digress and remind the reader what the main building blocks of a decision model are.

What is decision theory? …

A textbook definition of decision theory highlights some fundamental conceptual distinctions. First, one has to distinguish between feasibility and desirability; that is, between what one *can have* and what one *wants*. Confounding feasibility and desirability is typically viewed as a rationality "sin", as in the case of wishful thinking (believing that an outcome is possible only because it is desirable), or of "sour grapes" (believing that an outcome is not desirable because it is not possible). Next, most decision situations involve uncertainty, be it regarding other decision makers' choices, about external, non-strategic sources of uncertainty that one may refer to as "Nature", or both. It is then important to distinguish between choices that are under the decision maker's control, and those that aren't, and that are therefore random as far as she is concerned. The decision maker should attempt to generate beliefs over the latter, but not over the former.

Decision theory offers a few basic building blocks for the formal representation of a decision and how to think about it. First, there are *outcomes*, which are supposed to capture all that matters to the decision maker as a final result of the process. Outcomes may be more or less desirable, and the degree of desirability is often measured by a function, referred to as *the utility function*. For example, consider an insurance problem where the decision maker should decide whether or not to buy insurance for her house against damages. The possible outcomes should specify whether or not damages occurred, and if so their extent, as well as whether the decision maker paid an insurance premium, received compensation etc. Presumably, the decision maker's first best outcome would be not to have damages and not to spend any money on insurance; and then to pay the premium but still have no damages, and so forth.

[7] See Howard (1966, 1988) for a roadmap for decision analysis in practice.

Second, there are *states of the world*, referring to possible scenarios that may unfold in the problem. In line with the distinction mentioned above, a state of the world does not specify any choices of the decision maker herself. Rather, it describes the way any *other* uncertainty is resolved. For example, in the insurance problem a state of the world should specify whether a hurricane occurs, but not the decision maker's choice whether to buy insurance against damages. States of the world may be more or less likely, and their degree of likelihood may be measured by probabilities. However, there is no room to discuss their desirability, as it's not up to the decision maker to choose among them. In the example above, the decision maker may certainly prefer that a hurricane does not occur. But this type of preference should not interfere with the decision making process, as it would be a categorical mistake to attempt to select among the states of the world. Or, put differently, if the decision maker has some control over the state that obtains, the problem isn't properly defined. For example, if one considers insurance against fire, and it turns out that one may take precautionary measures and thereby affect the probability of fire, the states of the world shouldn't simply be "fire" and "no fire", but more sophisticated reaction functions, describing how the decision maker's environment would respond to his measures. For example, it is possible that there be no fire no matter what the decision maker does, or that fire should erupt only if precautionary measures were not taken, and so forth.

Finally, the decision maker's feasible choices are referred to as *acts*. Given the decision maker's choice of an act, and the environment's "choice" of a state, a unique outcome is determined. Often one models all three entities in a *decision matrix*, whose rows designate the available acts, whose columns list the possible states, and whose entries are outcomes. For each act (row) and each state (column) there is a unique outcome.

For example, in a very simple version of the insurance problem above, we might have a matrix such as

	Hurricane	No Hurricane
Insure	House destroyed, compensation received	No damages, premium paid
Not Insure	House destroyed, no compensation	No damages, no premium paid

The convention in decision theory is that by the description above one means to convey that

- the choice problem is between the acts “Insure” and “Not Insure”;
- the two possible states of the world are “Hurricane” and “No Hurricane”;
- the decision maker does not have control over the two states above;
- the outcomes that matter to the decision maker are given in the matrix.

A decision maker should then specify her desirability (of outcomes) by a utility function (which is not necessarily the direct monetary value of each outcome), and her likelihood judgments (over states) by a probability distribution. She may then follow a variety of decision criteria based on these concepts. The most popular one, justified by remarkably compelling axiomatic foundations, is the *expected utility paradigm*. It suggests that each act be evaluated according to a single number, which is the mathematical expectation of the utility of the outcome it may generate, where the expectation is taken relative to the probabilities of the states of the world. For example, if the probability of a Hurricane in the example above is 0.01, the expected utility of “Insure” is

$$0.01 * u(\text{House destroyed, compensation received}) + 0.99 * u(\text{No damages, premium paid})$$

and that of “Not Insure” is

$$0.01 * u(\text{House destroyed, no compensation}) + 0.99 * u(\text{No damages, no premium paid})$$

Expected utility theory (EUT) then suggests that an act with the highest expected utility be chosen. In this example, a rational decision maker will decide to insure if the utility loss incurred by paying the premium (in case of no hurricane) is less than 0.01/0.99≈1% of the utility loss incurred by not receiving compensation in the event of a hurricane. [8]

… and why is it not used?

This "textbook model" is elegant and enticing. However, in the form described above, decision theory appears to be largely absent from business decision making today – a far cry from its original aspirations. With notable exceptions such as research departments in financial institutions and in some capital-intensive industries, practitioners do not make much explicit use of decision theory. The proximate cause is probably that it raises very practical challenges. But there are also more fundamental reasons to challenge this textbook model of decision theory.

Practical issues

There are several reasons for which decision makers may not be able to follow the algorithm suggested by the expected utility paradigm.

- Often, *probabilities* cannot be assigned to the states of the world. When insuring against a hurricane or a fire, one may rely on statistical data to find out the probabilities of the states. But when thinking about financial crises or wars, global warming or a cure for a type of cancer, there aren't sufficient data about similar and causally-independent past problems that can be used to compute probabilities by empirical frequencies. Every financial crisis is unique, and so is every political conflict, to the extent that statisticians would shy away from providing probabilistic assessment to these one-shot, unique events. In these situations the classical paradigm suggests the use of *subjective* probabilities, but it does not offer the tools to find such probabilities, and many decision makers might feel that their beliefs are too vague and too ill-structured to be put into the straightjacket of probability theory. In addition, in business situations involving more than one decision maker, agreeing on subjective probabilities may soon feel like an

[8] For textbooks, see Fishburn (1970), Kreps (1988), Gilboa (2009), and Wakker (2010). Non-technical introductions are to be found in Gilboa (2010a,b).

exercise in futility or a source of endless disagreements. Indeed, these probabilistic assessments are *subjective*. We expect them to differ across individuals, and it is then difficult to implement the normative recommendation of using them in the expected utility formula. A decision maker may feel some tension between using probabilities that are inherently subjective, and at the same time claiming perfect rationality in using them.[9]

- Along similar lines, decision makers might have a problem in coming up with a *utility function* as a way to quantify their judgments on the desirability of outcomes. In some cases, decision theorists can offer helpful questionnaires that can "calibrate" one's utility function. For example, if the outcomes are only monetary (as in the insurance example above), one can ask simple questions about decision under risk from which the decision maker's utility function can be "measured". However, there are decision problems in which the outcomes are to be judged according to a multiplicity of criteria, and these do not always lend themselves to simple mathematical tradeoffs. For example, in making a career choice one has to weigh future income vs. personal satisfaction, work conditions vs. job prestige. How can a decision maker summarize the desirability of the outcomes by a single number?
- Sometimes, the very structure of the decision matrix is difficult to imagine. There are many problems – including, arguably, most problems of business strategy – in which the set of *states of the world* is large and unwieldy. The decision maker may be at a loss trying to come up with a reasonably exhaustive list of scenarios.
- Similarly, the set of *acts* available to the decision maker isn't always naturally and simply given in the description of the problem. It is relatively simple in the insurance problem above, but consider the assignment of workers to jobs, or a financial investment strategy. In such cases, listing all possible acts is theoretically possible, but the set of choices can be intractably large. Even more problematic are problems of strategy in which the set of possible choices is theoretically infinite: in such cases, creativity in generating "acts" is in itself a large part of the question.

[9] On different notions of rationality, see Arrow (1986), Eztioni (1986), Simon (1986), Gilboa and Schmeidler (2001), and Gilboa, Maccheroni, Marinacci, and Schmeidler (2010).

Behavioral challenges to decision theory

The difficulties that real-life decision makers face in using decision theory are supported by considerable evidence from behavioral psychology. In fact, if one wanted to come up with a list of the areas in which behavioral sciences have shown that real humans suffer from bounded rationality, the three pillars of decision theory would feature prominently.[10]

- Humans find it impossible to define the *utility* of various outcomes in a way that conforms with the axiomatic principles of rational decision making[11]. Even when evaluating outcomes that can be fully quantified in monetary terms, we are consistently inconsistent in assigning values to outcomes that entail perceived losses, which makes us prone to the sunk cost error and to the endowment effect[12]. Mental accounting can make us value equivalent outcomes very differently depending on the way they are framed[13]. An even more fundamental problem is that, as a result, our preferences are not fixed: depending on how the problem is framed, our preferences can change and we can rank them in different ways[14]. Last but not least, we are also aware that outcomes will be evaluated in hindsight. Especially in the context of an organization, in which we know that credit and blame will be attributed after the fact, this makes us loath to make risky decisions, even when they have a positive expected utility[15].
- We are hopeless at quantifying *probabilities*: we find salient but infrequent events vastly more probable than they really are. We believe plausible conjunctions of events to be more probable than components of the conjunction, which violates a basic principle of probabilities. We tend to neglect base rates, resulting in widely distorted estimates when attempting Bayesian inferences[16]. In addition, we consistently overestimate our chances of success in our endeavors, and even the probability of lucky events outside our control[17]. These problems in our assessments of probabilities are not circumscribed to

[10] For some of the pioneering work in psychology, see Preston and Baratta (1948), Edwards (1954), Simon (1957), Tversky and Kahneman (1973, 74, 81).
[11] Kahneman and Tversky (1979)
[12] Kahneman, Knetsch, and Thaler (1991)
[13] Thaler (1985)
[14] Tversky (1969), Kahneman and Tversky (1979)
[15] Kahneman and Lovallo (1993)
[16] Tversky and Kahneman (1974)
[17] For a review, see for instance Moore and Healy (2008).

uneducated or inattentive subjects: medical decision making, for instance, has been fertile ground for documenting them among intelligent, thoughtful, and well-intended physicians[18].

- Finally, we are not very effective at defining the *acts* we should consider. Inertia – and sheer lack of imagination – can lead us to be overly satisfied with the status quo and to fail to generate alternative acts altogether[19]. When we do generate options, we tend not to generate enough of them, and to satisfy ourselves prematurely with the first one that seems acceptable: confirmation bias makes us prone to disregard negative data about the option we are considering. And the options we do generate are highly susceptible to framing effects, which can make us reject acts we would consider attractive if viewed in a different light.

In short, when Kahneman and Tversky wrote that invariance of preferences is "normatively essential, intuitively compelling, and psychologically unfeasible"[20], they might have been writing more broadly about all the components of decision theory. It seems that managers who have tried decision theory and found it unusable are not lazy: they are human. The very idea of "Behavioral Strategy" as a field of research implies a critique of the idea of a rational strategist; and since a strategist practicing decision theory is the epitome of rationality, it is not surprising that it is a rare breed indeed[21].

The idealized model of rational decision making seems so thoroughly discredited that one could wonder how humans survive at all, and how they managed to obtain such remarkable scientific, artistic, and technological achievements. As many psychologists, including Kahneman and Tversky themselves, have observed, the prevalence of modes of thinking that sometimes lead to irrational behavior can be explained by their superiority in evolutionary terms[22]. Taking this line of thought seriously, an alternative school has emerged. Psychologists led by Gerd Gigerenzer claim that, in naturally occurring setups, people's intuition is much better than psychological experiments would lead us to believe.[23] This argument is often taken to say that in reality people

---

[18] Croskerry (2003)
[19] Nutt (1993)
[20] Kahneman and Tversky (1984)
[21] Powell & al. (2011)
[22] For a review see Evans (2008).
[23] See, for example, Gigerenzer and Hoffrage (1995) and Gigerenzer & Goldstein (1996).

do not commit errors and therefore do not need burdensome tools such as those of decision theory. An abundant and successful subgenre of popular management literature exalts the virtues of intuition as a guide to managerial decision making[24], and misses no opportunity to ridicule the Dr. Spocks who attempt to reduce complex problems to equations.

There are, therefore, two extreme views holding that decision theory is useless for managers:

- The hopeless view: as shown by behavioral studies, people do not and will never be able to think along the lines of the theory in any problem of interest. In this view, only "engineering" problems that lend themselves to algorithmic solutions can be addressed by decision theoretic tools. But these problems do not tend to be strategic ones, and can be solved by programming the right software.
- The cheerful view: as argued by others, intuition works just fine in natural, ecologically valid circumstances, and the few failures that might occur in lab experiments are no reason to worry.

The bottom line common to these views is that (luckily?) managers and executives are exempt from the unpleasant task of studying the ideas of decision theory. It is this bottom line we take issue with.

Whither decision theory?

Despite this bleak picture, we believe there is reason, more than ever, to use decision theory.

We start with the observation that, at times, even the most virulent critiques of decision theory admit that it is of some help. Operations research offers algorithms to find shortest paths in graphs, statistics provides fast and useful techniques to analyze data, and securities analysts develop increasingly complex models to optimize their portfolios based on historical data. This seems to suggest that, when it is possible and practical, decision makers do attempt to maximize expected utility in a purely rational manner.

Indeed, virtually all executives would probably agree, *in principle*, with the axioms of rational decision making and with the basic logic of the axioms underlying expected utility

[24] For instance, Gladwell (2007) and Klein (2007).

maximization. It is hard to imagine a serious businessperson who does not agree that if choice A is better that choice B and B better than C, then A is better than C – the transitivity axiom of decision theory. It is equally unlikely that a decision maker would not want to pick a dominant bet (i.e., one that results in higher payoffs in all situations) when such an option exists. Even if managers find DT impractical, and even if behavioral economists have shown that it is rather poor as a descriptive tool, it remains a serious contender for the ideal of how we *should* decide.

This is especially important in the context of strategic decisions made by managers and executives, who might need to present and justify those decisions to supervising third parties. For example, a manager might need to convince her superiors that a decision she plans to make is the right one; or, after a decision she took resulted in a poor outcome, she might need to justify the decision making process post-hoc, and defend the decision as a reasonable one at the time. Such post-hoc justifications might also be needed when reporting to governance bodies, outside regulators, or even court of law. In all of these situations, the decision maker is an "agent" reporting to a "principal"; and following basic DT principles seems to be a sound starting point for the agent to justify proposed future and actual past decisions to the principal. For instance, having to justify a decision that ended up with a poor outcome, the agent would rather not have to explain that her decision making process had intransitivities and wandered around in cycles. Not all of DT's principles are as compelling as transitivity, and, as we comment below, there is also a question about the appropriate decision making procedure for various problems. Yet, flagrant violation of the most obvious principles puts one in a position of risk in case of bad outcomes.

More broadly, the systematic shortcomings of a descriptive model do not automatically make it irrelevant as an ideal for normative purposes. In fact, the opposite is the case: a normative theory that only describes the way people behave is useless. A normative theory that is useful to practitioners is one that leads them to behave differently from the way they would have behaved before hearing about it. True, a normative theory should be practicable. If one cannot follow the theory's prescriptions, it becomes useless even as a normative one. But a descriptive failure is only the beginning of the normative story.

The challenge thus becomes to bridge the gap between the descriptive reality and the normative ideal. We believe the tools of decision theory are well suited to do so, provided they are used in the right way: as a paradigm, or a conceptual framework, rather than a "theory" in the standard sense. We argue that, in order to be useful, decision models need not provide a concrete, final answer; rather, in many cases, they are a platform on which a discussion can take place, and thus help analyze options and make decisions. Decision theory – and decision models – can be useful in more than one way, and one should not be disappointed if one knows what to expect.

Decision theory as a dialog

As a general framework, one may think of a dialog between the decision maker and the decision theorist. But the nature of this dialog can and should be very different depending on the type of decision at hand. One can broadly distinguish among three types of decisions, leading to different dialogs.

- In the first type, the problem is well specified, and lends itself to an *algorithmic solution*. Thus the authority lies with the decision theorist. She knows what's best: the decision maker only needs to provide her with the required data about his problem and then listen to her astute analysis, or, more simply, just obediently follow her bottom-line recommendations.
- The second type of decisions are those in which theory cannot supply a unique "best" answer to the decision maker. However, it can *test the consistency* of his reasoning and help him reach a decision that's better in his own eyes. In this type of dialog, the decision maker's intuition plays a prominent role, and the theorist's job is to make sense of this intuition, by formulating the decision maker's goals, beliefs, and constraints. Importantly, the authority in these cases is split: the decision maker is responsible for suggesting the solutions, but it's the theorist who decides whether a line of reasoning makes sense or not.
- Finally, there are also cases where a decision maker might not be able to describe his problem in the language imposed by the theory, not just because data is missing, but because the problem is so ill-structured that the decision maker finds it hard to articulate the logic of the proposed decision. In these cases the dialog will serve to clarify that

logic. To do so it needs not only to identify goals, constraints, and beliefs, but also to select the decision model in which they are to be analyzed. Thus, more authority is transferred from the theorist to the decision maker.

In the following sub-sections we provide examples that illustrate these three cases.

1. Type I: Algorithmic solutions

The case that best conforms to the vision of the founding fathers in the 1950s is the one that corresponds to operations research or statistical software today (or, in the realm of mass market applications, to apps such as Google Maps or Waze). In such cases, the outcome that the decision maker is trying to optimize for is clear and unambiguous, and all relevant information and inputs are known or knowable. Reaching the best decision is only a matter of using mathematical analysis.

For instance, for a job assignment problem the dialog between decision maker and theorist might go as follows:

Decision maker: I need to decide on the best way to assign new territories to salespeople. I assigned them sort of intuitively, but I'm not sure that's the best way to do it.

Theorist: What's your goal?

Decision maker: Since all these territories are new for us and we designed them to have the same potential, every salesperson's sales productivity is likely to be equivalent. So my primary objective is to minimize their travel costs, including the costs of initially relocating.

Theorist: And what do you know about these costs?

Decision maker: It varies a lot depending on the territory and the salesperson's current location. But I can pretty much figure it all out.

Theorist: That's good. How many salespeople and how many territories are there?

Decision maker: 12. Of each.

Theorist: Well, there are just over 479 million possible combinations. You should be pretty lucky to have found an optimal one.

Decision maker: So what do I do?

Theorist: Luckily there are fast algorithms that can do that. You don't have to just sit there and try out all these 479 million possible solutions. Just give me the data.

Decision maker: The data about the cost of relocating each salesperson to each territory?

Theorist: Yes. We're talking about 144 numbers. You said you had pretty good estimates of these.

Decision maker: I sure do. It's a one-page table and I'm sending it to you as an Excel attachment as we speak.

Theorist: Thanks. You can have a coffee. I'll get you the answer in no time.

[Coffee break]

Theorist: I analyzed your problem and here's the best match.

Decision maker: Are you sure it's the best decision?

Theorist: Yep. It's a well-known problem. Back in the '50s it was considered difficult, but since then people have developed fast algorithms for it. And there are actual mathematical proofs that the algorithm implemented by my software finds an optimal solution. Trust me. Given your goals and means, that's the thing to do.

This is an example in which decision theory is extremely successful at solving a problem that is analytically much more complex than it appears. In fact, decision theory is so successful that it becomes almost irrelevant to decision makers: the decision maker in the dialog above need not study the algorithm that the software uses. Similarly, a user of a navigation smartphone application need not study the shortest-path algorithms that the software uses. It suffices that an algorithm is found and analyzed by a mathematician, and that it is implemented by a programmer, and from that point on the rest of the world can use their product without having to delve into its nuts and bolts.

In these cases the decision maker's intuition is irrelevant: should the decision maker feel that she has a better idea to follow, the theorist should politely explain to her that she is simply wrong – and should have no difficulty showing that the optimal solution the model provides is at least as good as (and typically superior to) hers.[25] *Given her definition of her problem*, the software provided a best solution, full stop. Thus, in such problems theory replaces intuition so well that there is no room for intuition. At the same time, there is no need for the decision maker to study the theory in depth: a basic awareness of the existence of DT tools and an understanding of the types of problems it can solve algorithmically should be sufficient.

2. Type II: Objective consistency tests

A second family of cases concerns decisions that are not sufficiently well-defined to identify a best solution. This can be the case because, say, the objective function is explicit, but not all the necessary data is available. Often, it is hard to imagine all possible scenarios, and even harder to figure out what are their respective probabilities. Sometimes, even spelling out all possible decisions is a challenge, especially in a multi-stage decision problem.

Consider, for instance, optimizing a portfolio of investments. In this case acts and objectives are easy to state: the decision maker knows that she wants to maximize the value of her portfolio. Moreover, she can easily describe the set of portfolios she can purchase, as well as the set of scenarios, consisting of all possible prices of equities. (These are large sets, of course, but they can readily be described by mathematical models.) However, there is no objective way to assign probabilities to scenarios. In this case the theorist will not be able to algorithmically find the "correct" answer. Yet, the theorist will be able to test whether a certain decision can be supported by an "objective" line of reasoning, as in classical decision making models. Consider the following dialog.

> Decision maker: I feel like buying equity X. Is this a good idea?
> Theorist: Hard to tell. Do you think it would go up?
> Decision maker: I'm asking you. You're the expert.
> Theorist: I'm an expert on decision making, not on investments.
> Decision maker: What do you mean by that?

[25] A decision problem can have more than one solution, where these solutions are equivalent in terms of the specified objective function. It is therefore possible that the decision maker found another solution, which is also optimal; but it is impossible that she found a strictly better one.

Theorist: I'm good at fitting a problem to a model, but I don't know what's going to happen.

Decision maker: You don't?

Theorist: Nope, sorry. You probably have a much better intuition than I do. Fact is, you have money to invest and I don't.

Decision maker: OK, then. It was nice talking to you.

Theorist: Wait a minute, though. I am sometimes of help.

Decision maker: How so?

Theorist: I can help you see what one needs to assume to justify the decision you're about to make. So when people are just about to push the button, they often like to see how I would phrase their decisions.

Decision maker: And what good does it do me?

Theorist: First, you can avoid certain things that are clearly mistakes. Psychologists have found, for example, that people can be influenced by the way options are represented. They call it *framing effects*. For instance, in your business, some people have a tendency to ignore sunk costs; or to hold on too long to losing positions. When faced with these decisions, people sometimes consider their own choices a mistake. And my models help people avoid that.

Decision maker: That sounds silly. I'm an investment professional, and I don't think I'd be falling prey to these framing effects.

Theorist: Good for you. But there are other problems. People are prone to status quo biases; they are weak at manipulating probabilities; they are exposed to availability biases, which, even when we're aware of them, are hard to fix. And then people have been shown to be over-optimistic and too self-confident in their portfolio selections.

Decision maker: Are we really that stupid?

Theorist: Not stupid at all. Only we have tons of decision to make in very little time, based on very little data, and our minds do the best they can. Only the things that make sense in general, or on average, can lead us to wrong decisions in concrete situations.

Decision maker: OK, go on.

Theorist: So now let's see whether your proposed decision can be justified by the standard models.

Decision maker: Fine. I want to buy this stock.

Theorist: So we're going to take your current portfolio, add this transaction and see what you need to assume about the market to make this decision a smart one.

Decision maker: Do you mean, if this equity will go up?

Theorist: Yes, among other things. But not only "if". What is the probability that it would go up by this much and by that much, etc.

Decision maker: I don't think I have all these probabilities.

Theorist: And I'm sure *I* don't. But we can just see if there are probabilities that justify this choice and that you feel comfortable with.

Decision Maker: So you cannot tell me the probabilities, but you will tell me what I would need to believe the probabilities are for my decision to be the best one?

Theorist: Exactly!

Decision maker: And how does this help me?

Theorist: First, we'll be able to rule out a bunch of bad decisions that people tend to make because of the biases I mentioned. Not that it applies to you, of course, as you pointed out.

Decision maker: And second?

Theorist: You're running a fund, aren't you?

Decision maker: Yes, one of the funds in a large investment firm.

Theorist: Well, in case the portfolio loses in value your boss might wish to know why you did what you did. Of course it's the results that count, but showing an explicit model with reasonable probabilities would look better than saying it was your gut feelings.

Decision maker: Yes, I guess it would.

Theorist: OK, now let's put some numbers on this intuition of yours.

Clearly, the dialog may proceed to find a justification for the decision maker's proposed decision, but then again it may not. It is possible that the decision maker will be confronted with the assumptions needed for such a justification and decide that they are too extreme, or indefensible, and therefore decide not to implement the proposed decision. For example, if the decision maker has to specify the past cases on which she bases her probability estimates, and proceeds to analyze the similarity of the current problem to these past cases, she may find out that some cases influenced her intended decision more than they should, while others might have been ignored. Similarly, she may find that she ignored certain patterns of change, or was too quick to identify patterns that are not there.[26]

To sum, there are many cases in which the decision model cannot replace intuition. But it can *supplement* it by testing the implicit assumptions underlying it, and thereby limit the impact of cognitive biases that intuition alone would not detect. In other words, the tools of decision theory can help overcome many of the biases that sometimes make decision theory a poor description of people's behavior.[27]

3. Type III: Subjective consistency tests

We believe that sometimes decision makers who are perfectly willing to use decision theory models find it hard to do so because their problems are too ill-structured and too remote from the decision theory jargon for models to be useful. Moreover, decision makers might also question the presumably-objective logic of classical decision models. In such cases, dialogs as in the previous sub-section would break down, leaving the decision maker frustrated by the attempt to formalize the vague and quantify the unquantifiable. However, the developments in decision theory over the past few decades allow the theorist to use a variety of possible models, depending on the problem at hand. The failure of the classical model need not render all of decision theory useless.

Possible failures of the classical model may be the inability to capture the decision maker's beliefs by a single probability, or her goals by a single utility function. For instance, in choosing

[26] See Gilboa and Schmeidler (2012) and Gilboa, Samuelson, and Schmeidler (2015).
[27] Howard (1966, 1988) also emphasizes the role of the dialog between the decision maker and the decision theorist. Howard (1988) begins with a quote of Laplace (1812), highlighting the role of theory as a safeguard against "the illusions that often mislead us" "even when dealing with things that cannot be subjected to this calculus".

between candidates for an executive position, or in evaluating possible targets for a strategic acquisition, decision makers will be hard pressed to define unambiguously (let alone quantitatively) the result they are trying to optimize. Furthermore, the list of possible choices cannot be well defined, as the set of potential moves is theoretically infinite. In other cases, such as developing a marketing strategy, coming up with sufficiently creative options might well constitute the harder part of the problem. In these cases there are too many ways in which the model can be specified, and too many vastly different options compete for the title of "the correct decision".

In such situations, decision makers who have a vague recollection of their studies might rightly feel that the model they learned at school cannot be applied. It is therefore very tempting to rely solely on intuition as a guide. However, here again, this would expose the decision maker to a variety of biases and mistakes. Instead, we hold that an open-minded dialog with a theorist who is familiar with classical theory as well as with its generalizations and alternatives might still be of help.

Consider for instance the case of Facebook deciding to acquire Instagram – a strategic move that intrigued many observers. The dialog between executives and decision theorists might unfold along the following lines:

> Theorist: So what's on the agenda today?
>
> Decision maker: Instagram. Should we buy it?
>
> Theorist: What are you trying to do? What's your goal?
>
> Decision maker: You see, I'm not sure.
>
> Theorist: Do you want to maximize profit?
>
> Decision maker: I guess so. It would be silly to minimize profit, wouldn't it?
>
> Theorist: So how much more profit will you have with Instagram bought?
>
> Decision maker: I don't know. They're not making any money at present.
>
> Theorist: So why buy them?
>
> Decision maker: Well, for one, I don't want to have them out there pulling my users. Before I know it they'll add some text capacities and everyone will be spending time on their phones with Instagram rather than Facebook.

Theorist: It seems to me that you're worrying about your market share, which is probably a proxy for long-run profit.

Decision maker: Yes, sure, that's it.

Theorist: OK, so let's try to write down a decision model and see what you need to assume to make this purchase worthwhile.

Decision maker: What do you mean in this "decision model"? I don't like models. They never work. If I were to use models I wouldn't be where I am.

Theorist: I'm not trying to replace your intuition with the model. I'm only trying to make you see what is needed to justify the decision.

Decision maker: What do you mean?

Theorist: For instance, you could tell me that you don't care so much about more profits, and what you really want is to be the most famous and influential businessperson on Earth.

Decision maker: Well, I wouldn't mind that [smiling].

Theorist: Good. And then I'd ask you how much future profit you're willing to give up for this goal.

Decision maker: I see. But I'm not an ego-maniac. I do care about profit, only I don't know what the long run market will look like.

Theorist: OK. So we're making progress. We'll try to justify the decision by assuming that you only want to maximize profit, but that you do it in an uncertainty environment, not knowing what the market will look like 5, 10 years from now.

Decision maker: Exactly.

Theorist: So let's try to jot down some scenarios that you consider plausible, and put probabilities on them.

Decision maker: Probabilities?

Theorist: Why not?

Decision maker: I barely know what my goals are, and even less what the scenarios are, so how would I come up with probabilities? I think I'll just stick to my gut feeling.

Theorist: OK, OK, I didn't mean to push you. We'll use the model that you feel comfortable with. If you don't want probabilities, we won't have probabilities.

Decision maker: So what will we do?

Theorist: We can just jot down the scenarios and see what can happen.

Decision maker: And how do you make a decision with such a model?

Theorist: Well, we could think of the maxmin criterion, that suggests that you think of the worst case and try to make it as good as possible.

Decision maker: The worst case? Isn't it a bit paranoid? If I were always thinking of the worst cases I wouldn't be where I am.

Theorist: OK, there are other criteria as well. We can think of the best case, or of some combination of these. And we can have ranges of probabilities, so that, without committing to a number, you can still say what's more likely than what.

Decision maker: Sorry for being stubborn. But sometimes I can't even specify the possible scenarios. Too often I've seen people constructing such models, having 10-20 scenarios, and then being surprised by something that they didn't even take into account, let alone put probabilities on.

Theorist: Yes, I know. I've also seen that happening.

Decision maker: So? We part amicably?

Theorist: Not quite. I can also be stubborn. You see, there are models that don't require scenarios. Some models only look at the past, at cases you already have seen, and try to judge what's the best thing to do based on them.

Decision maker: Do I need you to help me with that?

Theorists: Sometimes. It's useful to spell out the past cases that you think of, how similar you think they are to your current problem, etc. Often people notice that they forgot cases or put too much emphasis on some prominent ones.

Decision maker: A little bit like what you said the other day about estimating probabilities?

Theorist: Precisely. Only this time it's for models with no probabilities, where cases are directly given you recommendation for or against certain choices.

Decision maker: Do you have such models?

Theorist: Don't worry, I have more models than you have decisions.

Decision maker: So how will you decide which model to use?

Theorist: I'll do what you feel comfortable with. You see, it's your decision, and you should like it, as well as the reasoning process that supports it. Again, I'm not trying to replace your intuition. I'm trying to put it into some framework that would allow you to test its coherence. Sometimes people try to put their intuition into words and find that they learn something from the process.

Decision maker: Like going to a shrink?

Theorist: Sort of. A shrink who uses formal models.

Notice that in the above, in order to support the management of Facebook as it considers this decision, the theorist first pulled up one of the more flexible models that DT has to offer, namely a model in which an event may have a range of probabilities, rather than a single probability number.[28] Similarly, there are decision models that allow for multiple utility functions, suggesting a way to deal with the problem of multi-criteria decision making. And there even are models that do not assume a complete, well-defined decision matrix, such as models of decisions that are based on past cases rather than on future scenarios.[29] But these more modern decision theories are naturally even less familiar to business people, and hence require an active participation of decision maker and decision theorists.

Observe that the dialog above took place at two levels. There is – as in sub-section 2 above – the level at which the decision theorist challenges the decision maker to formulate beliefs about utilities and probabilities that make his recommendation internally consistent. But there is also a meta-level, where the decision maker might feel that the choice of a classical decision theory model is too restrictive or not best suited for his needs.

[28] See Gilboa and Schmeidler (1989).

[29] See Gilboa and Schmeidler (2001).

In sum, the three types of dialogs span a spectrum: at one extreme the decision maker gives the decision theorist a problem as input, and gets back a decision as output. At the other extreme, the decision maker provides both the problem and a suggested solution, and expects to receive back a justification for the suggested choice (and even the notion of "justification" is negotiable).

Observe that these dialogs, at least for the more complex problems which encompass the majority of "strategic" issues, require decision makers to be able to explain their problems and to understand the analysis conducted in the terms of the formal model. For that reason, we hold that decision makers should be aware of the fundamentals of decision theory in order to be able to converse with decision theorists, as well as with computer-aided decision software, and reach a decision that would be good in their own eyes. Moreover, they should be aware of alternative decision theories, as the very language in which the dialog takes place also needs to be chosen jointly by the decision maker and the theorist.

Conclusion

There is a natural tendency to expect science to solve our problems. We want physicists to provide the theory that allows engineers to build bridges, and we expect biochemists to develop medications that physicians can prescribe.

There was once a similar expectation of decision theory; and that expectation is still sometimes met, as in the case of our first example. Indeed, a growing number of "solvable" problems are now encapsulated in black-box software: no knowledge of decision theory is required to solve the shortest-path problem or to recognize one song among millions.

Precisely because problems that have a single best answer have been automated in this way (or soon will be), executives will need to deal almost exclusively with ambiguous problems for which decision theory cannot compute "the" answer. It will therefore be increasingly tempting to dismiss theory as useless and to rely on intuition alone.

But this would be dangerous. Even when theory cannot provide a single, well-defined answer, it can be useful to challenge intuition, or at a minimum to guide and support it. If executives become familiar with the range of ways in which decision theory can help them deal with difficult decisions, they will find that the process of doing so helps them avoid many cognitive traps – indeed, the same cognitive traps that invalidate decision theory as a descriptive model. In other words, if used in the right way, the tools of decision theory can be useful not in spite of the limitations that behavioral sciences have shown, but precisely because of them.

References

Anscombe, F. J. and R. J. Aumann (1963), "A Definition of Subjective Probability", *The Annals of Mathematics and Statistics*, **34**: 199-205.

Arrow, K. J. (1951), "Alternative Approaches to the Theory of Choice in Risk-Taking Situations", *Econometrica*, **19**: 404-437.

——— (1986), "Rationality of Self and Others in an Economic System", *Journal of Business*, **59**: S385-S399.

Bayes, T. (1763), "An Essay towards solving a Problem in the Doctrine of Chances." Communicated by Mr. Price, *Philosophical Transactions of the Royal Society of London* **53:** 370-418.

Bernoulli, D. (1738), "Exposition of a New Theory on the Measurement of Risk", *Econometrica* **22** (1954): 23-36.

Croskerry, P. (2003), "The Importance of Cognitive Errors in Diagnosis and Strategies to Minimize Them", *Academic Medicine*, **78**: 775-780.

de Finetti, B. (1930), Funzione caratteristica di un fenomeno aleatorio. *Atti Accad. Naz. Lincei Rend. Cl. Sci. Fis. Mat. Nat.* **4** : 86-133.

——— (1937), "La Prevision: Ses Lois Logiques, Ses Sources Subjectives", *Annales de l'Institut Henri Poincare*, **7** : 1-68.

Edwards W. (1954), "The Theory of Decision Making", *Psychological Bulletin*, **51**: 380-417.

Etzioni, A. (1986), "Rationality is Anti-Entropic", *Journal of Economic Psychology*, **7**: 17-36.

Evans, J. S. B. (2008), "Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition", *Annual Review of Psychology*, **59**: 255-278.

Fishburn, P.C. (1970), *Utility Theory for Decision Making*. John Wiley and Sons.

Gigerenzer, G. and U. Hoffrage (1995), "How to Improve Bayesian Reasoning without Instructions: Frequency Formats", *Psychological Review*, **102**: 684-704.

Gigerenzer, G. and D. G. Goldstein (1996), "Reasoning the Fast and Frugal Way: Models of Bounded Rationality", *Psychological Review*, **103**: 650-669.

Gilboa, I. (2009), *Theory of Decision under Uncertainty*. Cambridge University Press (Econometric Society Monograph Series)

——— (2010a), *Rational Choice*. MIT Press.

——— (2010b), *Making Better Decisions*. Wiley-Blackwell.

Gilboa, I., F. Maccheroni, M. Marinacci, and D. Schmeidler (2010), "Objective and Subjective Rationality in a Multiple Prior Model", *Econometrica*, **78**: 755-770.

Gilboa, I., L. Samuelson, and D. Schmeidler (2015), *Analogies and Theories: Formal Models of Reasoning*. Oxford University Press.

Gilboa, I. and D. Schmeidler (1989), "Maxmin Expected Utility with a Non-Unique Prior", *Journal of Mathematical Economics*, **18**: 141-153.

——— (2001), *A Theory of Case-Based Decisions*. Cambridge University Press.

——— (2012), *Case-Based Predictions*. World Scientific Publishers, Economic Theory Series (Eric Maskin, Ed.).

Gladwell, M. (2007), *Blink: The Power of Thinking without Thinking*, Back Bay Books.

Hacking, I. (1975), *The Emergence of Probability*. Cambridge University Press.

Howard, R. A. (1966), "Decision Analysis: Applied Decision Theory", *Proceedings of the Fourth International Conference on Operational Research*, 55-71, Wiley-Interscience.

Howard, R. A. (1988), "Decision Analysis: Practice and Promise", *Management Science*, **34**: 679-695.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1991), "The Endowment Effect, Loss Aversion, and Status Quo Bias: Anomalies," *Journal of Economic Perspectives*, **5**: 193-206.

Kahneman, D. and D. Lovallo (1993), "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking", *Management Science*, **39**: 17-31.

Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk", *Econometrica*, **47**: 263-291.

Klein, G. (2007), *The Power of Intuition: How to Use Your Gut Feelings to Make Better Decisions at Work*. Crown Business.

Kreps, D. (1988), *Notes on the Theory of Choice*. Westview Press: Underground Classics in Economics.

Laplace, P. S. (1812), *Philosophical Essay on Probabilities.* Paris, Chapter 18.

Moore, D. A. and P. J. Healy (2008), "The Trouble with Overconfidence", *Psychological Review*, **115**: 502-517.

Nutt, P. C. (1993), "The Identification of Solution Ideas during Organizational Decision Making", *Management Science*, **39**: 1071-1085.

Powell, T., Lovallo, D., & Fox, C. (2011). Behavioral strategy. *Strategic Management Journal*, 32(13), 1369-1386.

Preston, M. G. and P. Baratta (1948), "An Experimental Study of the Auction Value of an Uncertain Outcome", *American Journal of Psychology*, **61**: 183-193.

Ramsey, F. P. (1931), "Truth and Probability", in *The Foundation of Mathematics and Other Logical Essays*. New York, Harcourt, Brace and Co.

Savage, L. J. (1954), *The Foundations of Statistics*. New York: John Wiley and Sons. (Second addition in 1972, Dover)

Simon, H. A. (1957), *Models of Man*. New York: John Wiley and Sons.

——— (1986), "Rationality in Psychology and Economics", *Journal of Business*, **59**: S209-S224.

Thaler, R. (1985), "Mental Accounting and Consumer Choice", *Marketing Science*, **4**: 199-214.

Tversky, A. (1969), "Intransitivity of Preferences." *Psychological Review*, **76**: 31-48.

Tversky, A., and D. Kahneman, (1973), "Availability: A Heuristic for Judging Frequency and Probability", *Cognitive Psychology*, **5**: 207-232.

——— (1974), "Judgment under Uncertainty: Heuristics and Biases", *Science*, **185**: 1124-1131.

——— (1981), "The Framing of Decisions and the Psychology of Choice", *Science*, **211**: 453-458.

von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*. Princeton, N.J.: Princeton University Press.

Wakker, P. P. (2010), *Prospect Theory*. Cambridge University Press.