

Guided decisions processes

Manel Baucells · Rakesh K. Sarin

Received: 28 January 2013 / Accepted: 22 March 2013 / Published online: 3 May 2013
© Springer-Verlag Berlin Heidelberg and EURO - The Association of European Operational Research Societies 2013

Abstract The heuristics and bias research program has convincingly demonstrated that our judgments and choices are prone to systematic errors. Decision analysis requires coherent judgments about beliefs (probabilities) and tastes (utilities), and a rational procedure to combine them so that choices maximize subjective expected utility. A guided decision process is a middle-of-the-road between decision analysis and intuitive judgments in which the emphasis is on improving decisions through simple decision rules. These rules reduce cost of thinking, or decision effort, for the myriad decisions that one faces in daily life; but at the same time, they are personalized to the individual and produce near optimal choices. We discuss the principles behind the guided decision processes research program, and illustrate the approach using several examples.

Keywords Decision analysis · Behavioral decision making · Narrow bracket · Insurance · Multi-attribute decisions

Manel Baucells benefited from the Spanish Ministerio de Economía y Competitividad [Grant ECO2010-17145].

M. Baucells
Rand Corporation, Santa Monica, CA, USA
e-mail: manel.baucells@upf.edu

M. Baucells
Department of Economics and Business and BGSE, Universtat Pompeu Fabra, Barcelona, Spain

R. K. Sarin (✉)
Decisions, Operations and Technology Management Area, UCLA Anderson
School of Management, Los Angeles, CA, USA
e-mail: rakesh.sarin@anderson.ucla.edu

Introduction

“Everything should be as simple as possible, but not simpler.” – Albert Einstein

Behavioral decision research has convincingly demonstrated that our decisions are often guided by emotions (feelings of liking and disliking) and that we use mental shortcuts or heuristics to come to an intuitive solution. Sometimes these intuitive solutions are reasonable and sometimes they are quite wrong. For certain decisions, an advisable approach might be to design a good choice architecture, that is, set the environment in a way that the intuitive choice (often the default option) happens to be a good one (Johnson et al. 2012). But quite often decision makers are required to make a conscious choice in less amicable environments, and the intuitive decision may not be a good one. It is in these cases that our decision processes need guidance.

Our research program for guided decision processes is based on providing individuals with *personalized decision rules* satisfying the following two principles:

- First, the decision rule should require only a few meaningful inputs from the decision maker.
- Second, the decision rule used to produce a recommendation should have strong prescriptive properties.

The first principle ensures that the process is personalized with respect to the decision maker’s goals, objectives, and circumstances. To satisfy this principle, the decision process should be closer to how people actually think about and solve the decision problem. The decision rule will be based on a requisite decision model (Phillips 1984) in which a few of the model parameters may be elicited, but some may be set to some prescriptive value. The second principle should ensure that the recommendation avoids biases and common errors, and that it is appropriate in a broad range of circumstances. To implement these personalized decision rules, we have in mind a decision maker with access to a tool, such as smart phone app that allows him to enter a few inputs, and produces a sensible personalized recommendation.

In the next section, we discuss what we mean by strong prescriptive properties. In section “[Universal vs. personalized decision rules](#)”, we propose some personalized decision rules, which we compare with universal decision rules. In section “[Application](#)”, we illustrate the approach with some examples. In section “[Conclusions](#)”, we conclude by offering directions for future research.

Prescriptive properties of guided decision processes

Avoid clear mistakes

A simple question that students often get wrong is the following:

A bat and ball cost \$1.10.

The bat costs one dollar more than the ball.

How much does the ball cost?

Students typically answer 10 cents. Some, of course, get the correct answer: 5 cents. It seems people do not check their answers, so the guidance provided by a rational process is not invoked. The bat and ball problem is harmless with little cost of error, but there are many other significant problems for which the cost of an error is high and for which our intuitive judgments and choices are simply wrong. Some examples are: the conjunction fallacy in which the probability of A and B is judged to be higher than probability of A alone, the hot hand fallacy (belief in lucky streaks), and base rate neglect in updating probabilities. A striking example of base rate neglect was the 1988 Illinois law mandating HIV testing for all couples who applied for a marriage license. Since the base rate of HIV among engaged heterosexual couples at that time was so low, false positives dominated true positives, $P(\text{No HIV}|\text{positive test}) \gg P(\text{HIV}|\text{positive test})$, causing a public outcry. The law was repealed a year later after huge economic and psychological costs on the residents.

Kahneman (2011) beautifully illustrates that these mistakes are made because System 2, the controlled and deliberative process, is not invoked or fails to check the operations of System 1, the automatic and associative process. A recent article by Rubinstein (2012) discusses several examples of choices involving clear mistakes, and shows how taking more time to respond improves accuracy in these circumstances. A minimal requirement for a guided decision process is to ensure that no clear mistakes are made.

Robust to narrow framing

It is well known in both behavioral decision making and decision analysis that a sequence of two decisions considered separately (in a narrow frame) can produce an inferior outcome when compared to a single comprehensive decision (in a broad frame). In a narrow frame, we may inadvertently accept a profile that is dominated in each state. Even when people are explicitly told to examine both decisions before they make a choice, they adopt a narrow frame and consider each decision separately. The classic example is Kahneman (2011):

Decision 1: choose between

- A. Sure gain of \$240, and
- B. 25 % chance to gain \$1,000 and 75 % chance to gain nothing.

Decision 2: choose between

- C. Sure loss of \$750, and
- D. 75 % chance to lose \$1,000 and 25 % chance to lose nothing.

A large majority of people choose A and D. The outcomes for the combined choice are:

- A and D: 25 % chance to win \$240 and 75 % chance to lose \$760.
- B and C: 25 % chance to win \$250 and 75 % to lose \$750.

Clearly, option B and C dominates option A and D as you are leaving money on the table by choosing A and D. The twin preference for being risk averse for gains

and risk seeking for losses produces a costly error. In the above example, choosing B and C over A and D can be universally recommended. If we can somehow persuade a person to adopt a broad frame, then we are done. There is no resistance, as the dominant decision is so transparent.

But in many real life situations, choices arrive sequentially. Even a person of great mental agility cannot aggregate all possible risks, such as when there are six simple binary decisions to be considered simultaneously. At home, should you get insurance or an extended warranty for your: (i) inside telephone wiring, (ii) refrigerator, (iii) television, (iv) washer, (v) dryer, and (vi) water heater. In the broad frame, this is a single decision with 64 options. Do you really know anyone who thinks like that (in the broad frame)? Almost everyone will look at each of these risks on their own merit and decide separately whether the insurance is justified for each case. As Kahneman (2011, p. 336) said, “Humans are by nature narrow framers”.

So we need decision rules that are myopically near optimal. That is, when the rule is repetitively applied to narrowly framed situations, the answer is consistent with what the rule would recommend if applied to one, broadly framed, decision. In this way, we combine the behavioral inclination of thinking about decisions one at a time with the rational approach of decision analysis (broad framing).

Dynamically consistent and other subtleties of normative theory

There are some classic examples, such as the Allais Paradox, the Ellsberg Paradox, and hyperbolic discounting for which even experts violate normative theory. Even Savage and Samuelson ended up making inconsistent choices when presented with these clever problems. Of course to Savage, his errors reinforced his faith in the value of the normative theory. But Samuelson (1952) was more skeptical. He worried his preferences would fall prey to the paradox: “I sometimes feel that Savage and I are the only ones in the world who will give a consistent Bernoulli answer to questionnaires of the type that Professor Allais has been circulating—but I am often not very sure about my own consistency”. Raiffa has given persuasive arguments to get people to change their minds and follow the axioms of subjective expected utility theory. But he admits that his efforts have yielded limited success.

Frankly, some belief or faith in the theory is needed to correct these violations. People do not see what is wrong in choosing \$3,000 up-front over a lottery that yields \$4,000 with probability 0.8 and \$0 otherwise, while simultaneously choosing a lottery that yields \$4,000 with probability 0.04 and \$0 otherwise over a lottery that yields \$3,000 with probability 0.05 and \$0 otherwise. These choices may be inconsistent with Expected Utility Theory, but a subject does not as easily recognize this as a “bat and ball” type mistake. The deliberative System 2 often corrects the errors produced by the intuitive System 1, but for some problems System 2 remains puzzled or clueless, as you cannot correct a mistake that you do not recognize as a mistake.

It is the obligation of academics to develop theories and models that are consistent with observed preference patterns, and several modifications of subjective expected utility (SEU) theory that have been offered. These models have had some success in explaining observed behavior, but have had little impact on improving decision making. We believe that SEU is still the model of choice for

prescriptive decision making. Even Amos Tversky, a founder of behavioral decision theory, regarded SEU as a rational model of choice. His dispute was with economists who regarded SEU as a good model for how people actually behave. To Tversky, SEU is not a good descriptive model of human behavior, but it is a perfectly reasonable model for rational choice. According to Kahneman (2011, p. 314), “Amos called the theorists who tried to rationalize violations of utility theory ‘lawyers for the misguided’”. Therefore, guided decision processes do not need to accommodate non-normative departures from SEU.

In some cases, choices of a person are broadly sensible if one considers all the concerns (e.g., fairness, anticipated disappointment) of the decision maker. Apparent violations of the normative theory occur because the relationship between the theory-based model and the real-world problem is contestable (Phillips 1984). This problem is particularly acute when the decision maker has multiple goals/objectives. An example is the ultimatum game where people reject positive monetary offers because they have legitimate non-monetary objectives (punishment to those who abuse power). Guided decision processes, therefore, need to be appropriate in a broader context of the theory-based model.

Universal vs. personalized decision rules

In MBA classes, most decision theorists teach decision tree analysis for choosing the optimal decision strategy. The criterion of choice is to maximize expected monetary value (EMV) using a bit of hand waving (e.g., looking at risk profiles) to incorporate risk attitudes in the analysis. This approach makes sense in most situations for which consequences are small relative to wealth. A simple analysis using EMV produces good enough solutions. If monetary consequences and probabilities are objective, then EMV is a universal rule because it gives the same recommendation to all users.

But EMV may not be adequate when stakes are not small, when the problem involves probability learning (McCardle and Winkler 1992), or when the decision variable is to choose a diversification strategy among multiple investment possibilities. The advice then is to pick the option that maximizes subjective expected utility. But there is an infinitude of utility functions one can choose from. In other words, when a friend asks you for advice on retirement investments, you may not endear your friend by suggesting that he or she should choose the one that maximizes her expected utility. Our middle-of-the-road and yet personalized approach, which is based on insights from both behavioral and rational decision theory approaches, assume that people engage in narrow bracketing and consider each risk in isolation. When deliberating about whether or not to buy coverage for cell phones people pay attention to the merits of the decision at hand and don't consider potential repair costs of internal telephone wiring or the possibility of a lost gift package to their mother. To be practical, we want a decision rule that is myopic, that is, even when applied to each risk considered in isolation it produces a near optimal global solution.

One and only one utility function, log utility, also called Bernoullian utility, possesses the property of being myopic under a broad set of environments.¹ Bernoullian utility has many desirable properties and yields a near optimal rule (MacLean et al. 2011).

Our middle-of-the-road approach is to adopt Bernoullian utility, $\ln(W + x)$, which requires a simple yet meaningful input, W . This input can be interpreted as the wealth or cash position of the individual and x is the current gain or loss. For some people, W could be their annual income, but for others it may just be their monthly salary. A large percentage of Americans, for example, live month-to-month (monthly inflow approximately equals monthly outflow) without much savings. To operationalize the notion, we define wealth as *the highest financial loss that a person can absorb without causing a major disruption in their life*. In summary, our personalized decision rules for decisions involving monetary consequences will be based on log utility.

For multi-attribute decisions, linear decision rules with equal weighting on the attributes have performed quite well in predictive tasks, such as prediction of academic success and psychological disorders. The concluding sentence in Dawes and Corrigan (1974) illustrates the simplicity of the rule, “The whole trick is to decide what variables to look at and then to know how to add”. In a colorful study, Howard and Dawes (1976) showed that self-ratings of marital happiness is predicted by the formula:

$$\text{Frequency of sexual intercourse} - \text{Frequency of arguments}$$

Happy couples had a positive number and unhappy ones had a negative number for the equation. If the consequence table of a multi-attribute decision is objective, then equal weighting is a universal rule. The universal rule for marital happiness may be valid on average across all couples. But potential failures of such a simple rule are easy to detect. Hypothetically, different partners within the relationship may place different weights on the two components. So the relationship could be satisfactory for one person, but unsatisfactory for the other partner.

Multi-attribute utility analysis in which tradeoffs and utilities are specified and rules of combining these (additive, multiplicative, or a more general rule) are axiomatically supported, is theoretically the best solution in a multi-attribute decision problem. But the cost of effort may not justify going all the way and the approximate solutions obtained by simple heuristics may be good enough. Universal decision rules make sense if the cost of effort is taken into account.

For a multi-attribute situation, a universal rule such as equal weighting may be very simple and lacks personalization. Our middle-of-the-road approach would suggest rules, such as lexicographic elimination by aspect to identify preferred alternatives. A productive debate here is about conditions under which such heuristics provide a reasonable or near optimal solution (Hogarth and Karelaia 2006; Baucells et al. 2008; Katsikopoulos 2011).

¹ Mossin (1968) shows that power and log utility are the only forms that are myopic in multi-period investment situations. Hakansson (1971) shows that log utility is the only utility that is myopic if future returns are serially correlated. Browne and Whitt (1996) show that log utility is also Bayesian myopic, and that exponential or power utility are not.

The personalized approach we propose would still rely on general purpose, near optimal, decision rules. However, we believe guided decision processes should incorporate meaningful individualized information.

Applications

In this section, we will analyze three classes of decision problems: insurance decisions, managerial decisions using decision trees, and multi-attribute decisions. We show how some simple yet personalized decision rules can be used to obtain preferred choices for these problems. These personalized decision rules require only a few meaningful inputs from the decision maker and produce near optimal solutions. Clearly, the guided decision process program is broader than these three classes of illustrative problems and we hope that our results will stimulate further research on developing, analyzing, and evaluating personalized decision rules that assist our decision process for a wider class of situations.

Insurance decision

Most homeowners living in disaster prone areas do not buy insurance. In California, for example, 90 % of homeowners do not carry earthquake insurance, yet about 80 % of its residents live near a fault. Similarly, hurricane Sandy in the Northeast United States in Fall 2012 made vivid the havoc a natural disaster can cause and yet most homeowners in flood-prone areas in the United States do not voluntarily purchase flood insurance. Worse yet, those who purchase flood insurance drop the coverage after only 3–4 years if a disaster did not strike.

On the flipside, people over-insure against risks with relatively small losses. Examples include cell phones, postal service for \$50 gift packages, telephone wires, household appliances, such as refrigerators, and all sorts of extended warranties for products with modest values.

Behavioral decision research has provided valuable insights into why people fail to protect themselves from risks that could lead to financial ruin. One reason is the “availability” bias as earthquakes or floods occur with irregular patterns over unpredictable time intervals. So after a significant event, such as the Northridge earthquake in Los Angeles in 1994, people do purchase earthquake insurance and upgrade their homes to current seismic standards (e.g., install an automatic shut off for gas, strap down water heaters, and stock emergency supplies of food, water, and first aid);² however, the memories of catastrophe fade over time. Just after the 1994 earthquake, words like bedrock and liquefaction became common in conversations. But over time, people became complacent. Disaster preparation is often equated with a nuisance drill in the workplace when you are forced to vacate your office. Emergency supplies dwindle and insurance is not renewed, that is, until the next big one arrives. Kunreuther et al. (2013) have presented data, information, and reasons

² *The Wall Street Journal*, December 13, 2012, headlined an article on the demand for flood insurance after hurricane Sandy “No. 1 Sales Agent for Flood Insurance: Sandy”.

for why people fail to insure against natural disasters and how public policy can encourage prudent behavior.

Back to the flipside (and to the chagrin of both decision analysts and behavioral economists), people over-insure against small risks. Thaler and Sunstein (2008, p. 79) write “extended warranties are plentiful in the real world and many people buy them. Hint: Don’t”. Kahneman (2011, p. 340) recommends “never buy extended warranties”. Of course, decision theorists have argued for a long time that for small risks people should be risk neutral.

So here is the dilemma. In the solemn conclave of assembled scholars of every ilk, the advice is: buy insurance against big risks, and do not buy insurance against small risks. People seem to be doing just the opposite. The key idea of this paper is that the advice on insurance, investments, and other risk taking decisions, as well as on life decisions about health, career, and family, needs to be *personalized*.

Let us illustrate the personalized advice in the context of insurance. The general advice “you don’t want insurance to cover what you could pay for yourself out of pocket” is sensible, but not quite correct. In an actuarial sense, paying \$150 to insure an iPhone that will otherwise cost \$860 to replace if damaged (theft or damage in a fire or earthquake is not covered) may not be a good idea. But what if you have a propensity to leave your iPhone in the pocket of your pants when you toss them in the washing machine? What if your cash flow is such that you live month-to-month with little cushion? Of course, you do not have to buy an iPhone, but if you saved enough to buy one and cannot afford to replace it in the near future, then maybe insuring it is not such a bad idea.

Consider the simplest case in which the consumer faces a binary risk of accident/damage vs. no accident. A myriad of consumer products, such as smartphones, electronic products, appliances, as well as other services (internal telephone wiring, postal packages, and extended warranties for repairs) fall into this simple case. In some other cases, the insurance contract offers only partial coverage. In case of an accident, the customer has to pay a deductible amount, D , and receives $L - D$, where L is the replacement cost of the item. Protection coverage for iPhones in the United States is around $C = \$150$, but the customer pays another $D = \$150$ to replace the damaged phone, that has a replacement cost of $L = \$860$. In general, we will compare more insurance with less insurance. More insurance is a protection plan that costs $C_1 > 0$ and has a small deductible of $D_1 > \geq 0$ in case of damage. Less insurance has cost $C_2 < C_1$ but a larger potential loss of $D_2 > D_1$ in case of damage. We know the actuarial probability of loss has to be less than

$$p = \frac{C_1 - C_2}{D_2 - D_1}$$

because the insurer needs to cover transaction costs and possibly make a profit (in this iPhone case, $p = 150/(860 - 150) = 0.211$). In addition to this factual information, the rule using log utility requires *two* personalized inputs: W , the individual’s wealth, and π , the individual’s perceived probability of an accident.

Having these two inputs, the expected utility of purchasing insurance is $\pi \ln(W - C_1 - D_1) + (1 - \pi) \ln(W - C_1)$. The expected utility of not purchasing the insurance is $\pi \ln(W - C_2 - D_2) + (1 - \pi) \ln(W - C_2)$. We assume that the

decision maker can afford insurance, $W > C_1$. What if he cannot afford the potential loss? If $D_2 \geq W > C_1$, then the consumer should always buy the insurance. This is consistent with the advice: always insure for large losses. Finally, if the consumer can afford the loss, $D_2 < W$, what should we recommend?

According to our approach, the customer should purchase the protection plan if $\pi \ln(W - C_1 - D_1) + (1 - \pi) \ln(W - C_1) \geq \pi \ln(W - C_2 - D_2) + (1 - \pi) \ln(W - C_2)$, or

$$\pi \geq \frac{\ln\left(\frac{W-C_2}{W-C_1}\right)}{\ln\left(\frac{W-C_1-D_1}{W-C_2-D_2}\right) + \ln\left(\frac{W-C_2}{W-C_1}\right)}. \tag{1}$$

It is easy to show that the right-hand side of (1) is less than $(C_1 - C_2)/(D_2 - D_1)$. If the decision maker thinks that his probability of accident is higher than average, $\pi \geq (C_1 - C_2)/(D_2 - D_1)$, then the rule recommends buying the insurance for all values of W .

Table 1 gives the threshold, π , for alternative wealth levels. We make two observations. First, at a low level of wealth, a person can insure even when his subjective probability of loss is significantly below the actuarially fair probability. At high levels of wealth, the recommendation is close to risk neutral (a risk neutral person would not insure if $\pi < p = 0.211$). Second, even a rich person could choose to insure if he judges himself to be especially prone to accident, or the accident supposes a large time cost due to hassle. Therefore, the general advice “don’t buy insurance for small risks” should be more nuanced.

What happens if the person chose not to purchase insurance using our guideline and an adverse outcome occurred. Their wealth is reduced and, therefore, the threshold probability for the next insurance decision will be lower. A similar risk for which the insurance was refused may now be insured. There is no inconsistency here as the decision to insure or not depends on current wealth level. So our myopic rule can be sequentially applied to opportunities to insure as risks arise. Clearly, the rule does not account for background risks formally; however, a person who has assumed several risks without insuring them could adjust their wealth level downward so that the threshold for insuring the next risk is lower. This approach is not entirely satisfactory, but will safeguard against bankruptcy.

We have considered a discrete choice between more insurance and less insurance of a single product. When one faces a simultaneous choice of more insurance vs. less insurance for n products, one can in principle compute the joint distribution of final consequences, and choose the set of products for which to buy more insurance

Table 1 Threshold probability as a function of wealth

W	π (with \$150 deductible)
1,000	0.092
1,500	0.144
2,500	0.174
10,000	0.203
100,000	0.210

by maximizing expected utility over all 2^n combinations. When one pools many of these independent risks together, however, one can take a more holistic approach and insure up to a total dollar value and self-insure the rest of risks. This would be the case when one starts a household and buys a refrigerator, washer, dryer, water heater, and so on. A personalized risk policy may prescribe insuring none of these risks for a person with high wealth, but only a few, but not all, for a person with modest wealth.

In some cases, one can choose the amount of coverage. Suppose the cost of insurance per dollar of coverage is Z , for which $0 < Z < 1$. The maximum possible loss is L and one can take full insurance to cover L by paying LZ . Let I be the amount of coverage and $I Z$ its cost. The optimal amount of coverage, I , is personalized based on wealth level. As wealth level increases, I decreases from L to zero. Suppose a person's subjective probability of loss is π . Assuming the actuarially fair probability of loss, p , is equal to π , then the expected utility of buying coverage I is:

$$\pi \ln(W - ZI - L + I) + (1 - \pi) \ln(W - ZI).$$

By setting the first derivative of EU with respect to I as zero, we obtain

$$\frac{\pi(I - Z)}{W - ZI - L + I} = \frac{(1 - \pi)Z}{W - ZI}.$$

Solving for I gives:

$$I^* = \frac{(1 - \pi)ZL + (\pi - Z)W}{Z(1 - Z)}.$$

Clearly, the optimal coverage depends on the initial wealth, W . By taking the derivative of I^* with respect to W , we can see the direction of the impact of wealth on the optimal coverage.

$$\frac{\partial I^*}{\partial W} = -\frac{Z - \pi}{Z(1 - Z)}.$$

This expression is strictly negative because the insurance company will charge more than the actuarially fair value to cover transaction costs, implying that $Z > \pi$. Thus as expected, wealthier people will desire less coverage and self-insure the rest, taking a deductible of $(L - I)$.

Clearly if $Z = \pi$, then everyone will seek full coverage ($I^* = L$) leading to a universal recommendation. Personalized advice requires the actuarially fair probability of loss (e.g., 1 in 100 for fire risk), the maximum possible loss (e.g., $L = \$300,000$), and the insurance company's profit margin (say 20 % or $Z = 1.2\pi$). These parameters π , Z , and L can be estimated from available data. So the optimal coverage, I^* , for an individual can be recommended by knowing his wealth, W . The rule we propose does not need to elicit any risk aversion parameter, which is hard to measure and seems to be quite fickle. Rabin and Thaler (2001, p. 255) state that "people do not display a consistent coefficient of relative risk aversion, so it is a waste of time to try to measure it". In contrast, people do know their cash positions

and have a reasonably good idea of how much loss they can sustain without a major disturbance to their lives.

Managerial decisions

While EMV is a good decision rule in many contexts, it is important to acknowledge that the rule is not appropriate for mid- to high-stakes decisions, such as those often faced by small businesses or family-owned operations. As with the insurance decision, our personalized decision rule is to use Bernoullian utility.

The approach is personalized, as the first step is to elicit the wealth, W , of the company. Owners should ask themselves the maximum loss their company could bear. Then, replace the gains and losses of the risky decision, x , by $g = \ln\left(\frac{W+x}{W}\right)$. This utility function is equivalent to $\ln(W+x)$, and has an advantage in that the consequences and expected values can be interpreted as growth rates. Once we calculate the growth rate associated with each terminal node, we simply roll back the tree and choose the alternative with the highest expected growth rate. Expected growth rates can be interpreted as the long-run average rate of growth of our initial wealth. For example, if wealth increases by 10% in one year, then $g = \ln 1.1 = 0.0953$, which is the continuous time growth rate of the wealth and is a meaningful measure of return.

The business case *Carter Racing* is an example of a high-stakes managerial decision (Brittain and Sitkin 1986). It describes the situation of a racing team owned by two brothers for whom their next race can either lead to success, failure, or an intermediate outcome. Success means securing the current oil sponsorship of \$500,000 and gaining a new tire sponsorship worth \$1,000,000. In the intermediate scenario, only the current oil sponsorship is maintained. Finally, an engine failure represents the loss of all sponsorships, plus an additional loss of \$20,000 to replace the engine.

The team has the option of not racing, which would secure the oil sponsorship, but have a penalty of \$17,500. The case brings additional intangibles of reputation together with the crucial judgment on the probability of an engine failure. In the base case, the probabilities of the three scenarios can be calculated from historical performance data with results of $\pi_{\text{success}} = 12/24$, $\pi_{\text{intermediate}} = 5/24$, and $\pi_{\text{failure}} = 7/24$.

Using EMV, we conclude that the EMV of racing is \$848,333, where as the EMV of not racing is \$482,500. The EMV of racing is higher, but payoffs range from \$1.5 million to a loss of \$20,000. The case involves a choice of framing financial consequences, as one could include or exclude the oil sponsorship as part of the status quo. In case of excluding the oil sponsorship, the gains and losses in case of racing are $\mathbf{x} = (\$1.5\text{M}, \$0.5, -\$0.02\text{M})$. We then compare the EMV, $E[\mathbf{x}] = \$848,333$, to not racing, \$482,500. The EMV of course favors the risky alternative, but EMV is not adequate here because the continuation of the business after failure is at stake. An intuitive assessment of the risk profile may be subject to biases. A change of frame by including the oil sponsorship (and shifting all payoffs by $-\$0.5\text{M}$) may produce a preference reversal.

The decision rule we propose is to use Bernoullian utility. This requires that we elicit a single personalized parameter: the financial capacity of the company. An engine failure leaves the company in a bad financial position. If that loss were to occur, how much more could the company afford to lose? Say the answer is $W = \$100,000$. Now we can calculate the consequences in terms of growth rates. If choosing to race, the additive consequences of $\mathbf{x} = (\$1.5\text{M}, \$0.5, -\$0.02\text{M})$ are replaced by growth rates, $\mathbf{g} = \ln((W + \mathbf{x})/W) = (2.772, 1.792, -0.223)$, and we conclude that the expected growth rate if racing is $E[\mathbf{g}] = 1.695$. This can be compared to the growth rate of not racing, which is equal to $\ln((W + \$482,500)/W) = 1.762$. Because $1.762 > 1.695$, the personalized advice is not to race.

As is customary, we can perform a sensitivity analysis (Clemen 1996) of the decision with respect to the maximum admissible loss, W . We observe that racing becomes advisable for wealth above $\$150,000$. In fact, using goal seek in Excel, one can find the critical point of wealth for which both options would be deemed indifferent, $\hat{W} = \$133,400$. Thus, above a wealth of $\$133,400$, it is advisable to choose the risky option of racing; and, below this wealth level, it is better to choose the safe option of not racing.

Growth rates of racing vs. not racing as a function of wealth

W	Racing	Not racing
50,000	2.06	2.37
100,000	1.69	1.76
150,000	1.46	1.44
200,000	1.30	1.23
300,000	1.08	0.96

This decision to race or not may be part of a grand tree, in which other risky decisions are involved (e.g., insurance decisions, portfolio choices, or future investment possibilities). Such a grand tree would have a large number of chance nodes, decision nodes, and end points. For a general concave utility, it is necessary to evaluate the grand tree to determine a grand contingent plan. Such plan will then advice us on whether to race or not. For log utility, however, the answer of this grand tree will always agree with the answer to the small tree with two decisions and four endpoints.

One purpose of the grand tree is to account for a string of good or bad outcomes in future decisions. The *raison d'être* of decision trees is to reflect the impact of future decisions and uncertainties on today's choice. Log utility exhibits a myopic property: we can consider each "small" tree in isolation, and still make the "grand" tree optimal decision. The increase or decrease in the wealth after the current decision will automatically reflect which way (to race or not race) our future choices will be swayed. In other words, what is optimal for the "small" tree given current wealth will be consistent with the optimal strategy identified by solving the grand tree. This advantage of log utility is not computational alone, but it also has a psychological appeal. People find it more natural to their way of thinking to solve

the current year's decision problem, to observe the results, and then to move on to sequentially solve successive years decision problems one problem at a time.

Multi-attribute decisions

In many decisions, the consequences are not necessarily financial or monetary. In the selection of a job, one may consider location, reputation of the company, work environment, and salary. Other examples include selection of an apartment, an automobile, a vacation, or a restaurant. Multi-attribute utility is an idealized approach that requires one to provide weights (tradeoffs) or attributes, utility functions (desirability) for each attribute, and a combination rule that is axiomatically supported. Indeed, for very important decisions, the effort required to do a thorough multi-attribute utility analysis is justified.

For many decisions, some simple rules may provide near optimal results. These rules are personalized in the sense that a decision maker may specify some pertinent information and the suggested alternative takes into account this individual information. Universal decision rules such as linear utility and equal weighting work well in some classes of problems (Dawes and Corrigan 1974); however, these rules are inadequate for situations where idiosyncratic preferences are common.

One decision rule that is personalized and yet produces near optimum solutions is *elimination by aspects* (Tversky 1972; Hogarth and Karelaia 2005b; Baucells et al. 2008). The first step to apply this decision rule is to "binarize" the decision (Hogarth and Karelaia 2005a). This means that each entry of the table of consequences is encoded as either a zero (low, absent, or unacceptable) or a one (high, present, or acceptable). Binarization is consistent with Simon's Satisficing Principle.

A decision maker is then required to provide an ordinal ranking of the attributes. Elimination by aspects is a lexicographic rule that considers attributes in order of importance and, in each step, eliminates all the alternatives with zero values on that attribute (we implicitly assume that some alternatives have a one). The procedure stops when either one alternative remains, or the attributes have been exhausted. If two or more alternatives remain, then judgment can be used to select the preferred alternative.

Let us consider an example of choosing the location of a beach vacation. If the consequence table is binary, then Dawes' rule reduces to counting the number of pros of each alternative. In our example, each location has two pros. Therefore, there is a tie between the three locations. The recommendation for anybody going on this vacation would be to choose at random!

Consequence table of choosing a location of a beach vacation

	Maui	San Diego	Miami
Windsurfing	1	1	0
Snorkeling	1	0	1
City life	0	1	1

Let us now apply elimination by aspects. Consider three hypothetical individuals that, based on their goals and objectives, rank attributes in the following way:

Individual 1 : *Windsurfing* \succ *Snorkeling* \succ *City Life*,

Individual 2 : *Snorkeling* \succ *City Life* \succ *Windsurfing*, and

Individual 3 : *City Life* \succ *Windsurfing* \succ *Snorkeling*.

Applying this personalized decision rule, Individual 1 would select Maui and San Diego in the first round, and eliminate San Diego in the second round. The personalized recommendation for Individual 1 is Maui. Similarly, it is easy to verify that the recommendation for Individual 2 is Miami, whereas the decision rule recommends individual 3 to go San Diego.

An additional advantage of elimination by aspects is that it can also produce a ranking of alternatives. To select the second best alternative, the idea is to eliminate the best alternative from the set, and re-run elimination by aspects again. For example, Individual 1 would choose San Diego as his second best alternative. The third choice for individual 1 is Miami. Thus, the personalized output of this decision rule would be the following ranking of alternatives.

Individual 1 : *Maui* \succ *San Diego* \succ *Miami*,

Individual 2 : *Miami* \succ *Maui* \succ *San Diego*, and

Individual 3 : *San Diego* \succ *Miami* \succ *Maui*.

Conclusions

The heuristics and bias research program has been successful in demonstrating that judgements and choices are prone to systematic errors. People use narrow frames and are loss averse. These two biases in combination produce results that no reasonable mind can accept (e.g., preference for a stochastically dominated alternative). Similarly, people have inertia and stay with the status quo even when a change could improve economic or physical well-being. Thaler and Sunstein (2008) have used a strategy of choice architecture to set defaults that produce desirable results. These universal defaults are appropriate in some situations, but may not be advisable in situations in which people genuinely have different goals and objectives (Carroll et al. 2009).

We advocate for *personalized decision rules*. A personalized decision rule possesses two characteristics: (i) it requires only a few meaningful inputs from the decision maker, and (ii) it produces near optimal solutions.

Our approach is particularly useful when the cost of decision effort is substantial. Research is needed to develop personal decision rules for classes of decision problems where full analysis is costly and time consuming. Examples include multi-attribute decisions, group decisions, risk analysis, search problems, and strategic situations. Here are some ideas for future research.

Consider multi-attribute decision problems where the consequence table is available. What would be the minimum set of meaningful inputs to carry out a multi-attribute utility analysis? We would suggest the following: an ordinal ranking of attributes, aspiration levels for each attribute, and one additional parameter, the

compensatory degree. When the compensatory degree is set to zero, we would be employing rapidly decreasing weights, such as the rank ordered centroid (ROC) decision weights (Katsikopoulos and Fasolo 2006).³ When the parameter is one, we would be using equal weights (EW). For intermediate parameter values, the weight would be the average of ROC and EW. Some creative solution needs to be devised to find a simple way to normalize attribute values in a 0–1 scale given the aspiration values. These inputs are sufficient to produce a score or ranking of each alternative.

A second example is a search problem in which the underlying offer distribution is unknown. The personalized decision rule should elicit few inputs, such as the range of possible offers. As offers arrive, one can employ some simple updating rule to update the distribution (e.g., the empirical distribution constrained by the initial range) and determine if the current offer is acceptable.

Future research should explore accuracy of personalized decision rules with respect to various modeling assumptions. For example, bounds on errors could be developed and those in turn can be used to evaluate the effectiveness of alternative decision rules. It is also worthwhile to investigate ideal balance/tradeoff between cognitive investment (simplicity) and the quality of the solution (optimality).

We should emphasize that the gold standard for rational decision making is SEU theory, also known as Bayesian rationality. Application of SEU, however, requires full information on objectives, beliefs, and preferences of the decision maker *and* the assumption that the decision maker follows Bayesian rationality. Personalized decision rules require limited information, and are therefore easier to implement. Nevertheless, these rules are judged by how close the solutions are to those obtained by SEU. Future efforts should improve both the applicability and the quality of these decision rules.

The guided decision process we propose meshes well with new technologies. Smartphone apps could obtain most of the objective information from data sources (consequence table, actuarial probabilities of loss, average claim rate among all subscribers), and request entry by the user for a few individual inputs (financial capacity and subjective probability of accident). Then, it could apply the decision rule and provide a near optimal recommendation or a sorting of the alternatives. Applications, such as Yelp for choices of restaurant or Kayak for choices of airplane tickets and hotels could benefit from this approach.

We have demonstrated how to personalize decision rules for insurance decisions, business investment decisions, and multi-attribute choices. Our approach with appropriate modifications can be adapted to decisions about investments, health, career, and well-being. Our hope is that personalized decision rules will be investigated in real settings to help people make better decisions.

³ Given n attributes with ranks from 1 (most important) to n (least important), the ROC weights, $w_1 > w_2 > \dots > w_n$, are given by:

$$w_k = \frac{\sum_{i=k}^n \frac{1}{i}}{n}, \quad k = 1, \dots, n.$$

Reference

- Baucells M, Carrasco JA, Hogarth RM (2008) Cumulative dominance and heuristic performance in binary multiattribute choice. *Math Soc Sci* 56(5):1289–1304
- Brittain J, Sitkin S (1986) Carter racing: parts A & B. Kellogg School of Management, New York
- Browne S, Whitt W (1996) Portfolio choice and the Bayesian Kelly criterion. *Adv Appl Probab* 28(4):1145–1176
- Carroll GD, Choi J, Laibson D, Madrian BC, Metrick A (2009) Optimal defaults and active decisions. *Q J Econ* 124(4):1639–1674
- Clemen R (1996) Making hard decisions: an introduction to decision analysis. Duxbury Press, Boston
- Dawes RM, Corrigan B (1974) Linear models in decision making. *Psychol Bull* 81(2):95–106
- Hakansson NH (1971) On optimal myopic portfolio policies, with and without serial correlation of yields. *J Bus* 44(3):324–334
- Hogarth RM, Karelaia N (2005a) Ignoring information in binary choice with continuous variables: when is less more? *J Math Psychol* 49(2):115–124
- Hogarth RM, Karelaia N (2005b) Simple models for multi-attribute choice with many alternatives: when it does and does not pay to face tradeoffs with binary attributes. *Manag Sci* 51(12):1860–1872
- Hogarth RM, Karelaia N (2006). Regions of rationality: maps for bounded agents. *Decision Anal* 3(3):124–144
- Howard JW, Dawes RM (1976) Linear prediction of marital happiness. *Personal Soc Psychol Bull* 2(4):478–480
- Johnson EJ, Shu SB, Dellaert BG, Fox C, Goldstein D, Häubl G, Larrick RP, Payne JW, Peters E, Schkade D, Wansink B, Weber EU (2012) Beyond nudges: tools of a choice architecture. *Market Lett* 23:487–504
- Kahneman D (2011) Thinking, fast and slow. Farrar Straus & Giroux, New York
- Katsikopoulos KV (2011) Psychological heuristics for making inferences: definition, performance, and the emerging theory and practice. *Decision Anal* 8(1):124–144
- Katsikopoulos KV, Fasolo B (2006) New tools for decision analysts. *IEEE Transact. Syst. Man Cybern. Part A: Syst. Hum.* 36(5):960–967
- Kunreuther H, Pauly MV, McMorro S (2013) Insurance and behavioral economics: improving decisions in the most misunderstood industry. Cambridge University Press, Cambridge
- MacLean LC, Thorp EO, Ziemba WT (2011) The Kelly Capital growth investment criterion. Cambridge University Press, Singapore
- McCardle KF, Winkler RL (1992) Repeated gambles, learning, and risk aversion. *Manag Sci* 38(6):807–818
- Mossin J (1968) Optimal multiperiod portfolio policies. *J Bus* 41:215–229
- Phillips L (1984) A theory of requisite decision models. *Acta Psychol* 56:29–48
- Rabin M, Thaler R (2001). Anomalies: risk aversion. *J Econ Perspect* 15(1):219–232
- Rubinstein A (2012). Response time and decision making: a “free” experimental study. Working Paper
- Samuelson P (1952) Probability, utility, and the independence axiom. *Econometrica* 20(4):670–678
- Thaler RH, Sunstein CR (2008). *Nudge: improving decisions about health, wealth, and happiness*. Penguin, New York
- Tversky A. (1972) Elimination by aspects: a theory of choice. *Psychol Rev* 79(4):281–299