Judgment and Decision Making, Vol. 7, No. 2, March 2012, pp. 207–223

Decision theory as an aid to private choice

Rex V. Brown*

Abstract

A wise decider D uses the contents of his mind fully, accurately and efficiently. D's *ideal* decisions, i.e., those that best serve his interests, would be embedded in a comprehensive set of totally coherent judgments lodged in his mind. They would conform to the norms of statistical decision theory, which extracts quantitative judgments of fact and value from D's mind contents and checks them for coherence. However, the most practical way for D to approximate his ideal may not be with models that embody those norms, i.e., with applied decision theory (ADT). In practice, ADT can represent only some of D's judgments and those imperfectly. Quite different decision aid, including intuition, pattern recognition and cognitive vigilance (especially combined), typically outperform feasible ADT models—with some notable exceptions. However, decision theory training benefits D's *informal* decisions. ADT, both formal and informal, should become increasingly useful and widespread, as technical, cultural and institutional impediments are overcome.

Keywords: decision-aiding, decision theory, private choice, ideal judgment.

1 Introduction

We would surely all agree that most people (especially other people!) make poor decisions and suffer the consequences. We marry abusive spouses and elect incompetent presidents, where we should have known better. We do not seem to have learned much over the ages.

In the 1950s, a new analytic technology, applied decision theory (ADT)—also known as "decision analysis"—promised to revolutionize decision practice (Raiffa & Schlaifer, 1962). I have spent a long career trying to realize that promise. Along the way, I have significantly modified how I view the distinctive-and-useful-province of ADT.

2 Mission and perspective

In this essay, I offer some very personal thoughts on what ADT can or might do to help private deciders. I have discussed elsewhere ADT's application to managerial decisions, where much of my experience lies (Brown, 2005b, 2009). I have found that managers commonly treat their professional decisions as private decisions with institutional consequences—especially when these consequences fall mainly on others (see Appendix A).

My perspective is that of a technician and technologist, aiding deciders as best he can, however imperfectly. It is not that of a scientist, whose findings must pass the much stiffer test of impersonal verifiability. However, my perspective may possibly stimulate useful scientific work. Decision technology can certainly benefit from the feedback of cognitive psychologists and other scientists.

My comments are based on some 50 years of trying to put ADT to productive use on varied live dilemmas (like whether a woman should divorce her husband or a regulator close down a reactor; see Appendix B). Not having been trained in cognitive psychology, nor immersed in the literature, I have consulted with knowledgeable colleagues to help me distill what is relevant to my work. Well-read JDMers may know of findings that underscore—or undermine—my arguments.

I have drawn illustrations largely from my personal experience and, in particular, from one case (detailed in Brown & Baron, 1991, pp. 119–121):

> In 1989, at age 55, I had to decide whether to replace an arthritic hip immediately ("CUT") or to delay surgery for ten years or so ("WAIT"). My immediate impulse was to CUT, but my surgeon advised WAIT. This disconnect moved me to use the dilemma as a case study to exercise ADT methods in a middle-school decision skills course, drawing on the surgeon's factual judgment and on my value judgments. At the end of the course, most students voted for me to CUT, which by then confirmed my own conclusion. The surgeon acknowledged having under-played certain considerations and agreed to CUT.

I thank the following for helpful comments on drafts of this essay: Andrew Kahr, Jon Baron, Danny Kahneman and most especially Arthur Schleifer. They do not necessarily share my views.

*School of Public Policy, George Mason University, 2018 Lakebreeze Way, Reston, VA, 20191. Email: rexvbrown@gmail.com.

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

After surgery, the surgeon advised me to avoid jolting my hip for a while. I fully concurred, but did not act on my decision. Two months later I took a high-impact aerobics class. It dislodged my new hip and 20 years later I still limp painfully.

3 The decision aiding task

3.1 Raw material: D's mental resources

Improving decider D's decisions involves considering what is in his[1] mind at the time and how well he uses it. His mind-content is untidy, complex, unstructured and unstable. It includes thoughts, anticipations, recollections, misgivings and subliminal feelings, and possibly knowledge of potential information he could add to his mind-content.

My own initial mind-content, when I was first inclined to CUT, included visions of life with and without an early hip replacement, a visceral revulsion at my thigh bone being sawn through and some knowledge of sources of additional medical expertise.

How D processes his mind-content depends on the decision. Civic decisions (such as whom to vote for) are mainly clear-cut choices. Personal decisions (such as whom to marry) are usually unruly, but often effective, processes, which are elusive to describe, let alone to improve on. They are normally fluid, adaptive, parallel, often sub-conscious and prone to change.

Mary first wonders whether to marry Sid or George, and eventually gives up on marriage and decides to enter a nunnery.

D can learn a good deal from research psychologists about how unaided deciders *in general* decide, say, by "satisficing" or succumbing to cognitive biases (Simon, 1979; Gilovich et al., 2002). D can also draw on his own understanding of "human nature" and of his own idiosyncrasies. For example, he may believe that his judgment is clouded by emotion more than most. The worse D thinks his own judgment is, the more he can expect to gain from *any* decision aid—if only he were capable of profiting from it.

[1]For "his", read "his or her", etc.

3.2 Ideal choice: The judgment for D to aim for

My working hypothesis is that, in any situation, D normally has an ideally wise[2] personal[3] decision, if he but knew it. I interpret **D's ideal judgment as that which makes the best use of his entire mind-content, at a given moment, to advance his interests, as he understands them.** (All phases of D's decision process, such as specifying options and acting on choices have their ideal judgments, but choosing among specified options is our focus now.)

The chaotic element in D's mind-content makes it difficult to turn this loose interpretation of ideal choice into an operational definition (though I will I make some attempt later). Nevertheless, most of us talk as if we were convinced that D generally has something close to an ideal choice. We say "D should have known better than to do X." This surely implies that we think that D had some ideal decision, which he failed to act on. The ideal may not be a precise spot, but rather a vicinity.

We use a compass to plan travel, although we know the magnetic north pole is shifting.

Interpreting the *factual* component of an ideal choice, i.e. ideal probability, is relatively straightforward. In principle, facts can be empirically verified and, arguably, their reality could be determined by everything D has ever learned (negligibly influenced by D's subjective judgment). On the other hand, a noted decision theorist has objected that ideal *value* judgments cannot exist, since D does not have any inherent values: these are constructed, not uncovered. I believe, however, that ideal choice does not require that values be stable, only that they exist at the moment of choice (whether constructed or not).

Judging how far D's actual choices fall short of his ideal choice is an evaluation of his room for improvement. A *specific* decision tool can be evaluated by how much improvement it can *actually* achieve.[4]

In deciding between CUT and WAIT, I could conceivably take infinite pains to pursue all possible arguments and use every scrap of material in my mind. This would produce an ideal choice (but not necessarily what I did choose). However, I did not pursue any further

[2]I prefer the term "wise" to "rational" to characterize how effectively an action furthers D's ends (Zeckhauser, 1996). "Rational" is ambiguous: it can also describe decision"means", which may or may not produce wise action, as in "Effective deciders use 'correspondence', not 'coherence' rationality."

[3]Distinct from an *impersonal* ideal, which is based on accessible evidence, independent of any particular assessor (Brown, 1993a).

[4]This is analogous to the familiar ADT use of "perfect information" for evaluating imperfect information.

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

study, since I did not expect it to get me closer enough to that ideal for it to be worth the trouble.

3.3 Bridging the gap between actual and ideal choice

Academic research has made major contributions to prescription. Most of it has been "mono-disciplinary". Psychologists' work on biases, for example, has given rise to "cognitive vigilance", which averts cognitive flaws by being alert to them (Bazerman, 2005). Mathematical statisticians have developed statistical decision theory, as the logic of sound prescription (Raiffa & Schlaifer, 1962), which is the focus of this essay. However research on decision prescription, as such, has been a major gap. Moreover, little *interdisciplinary* research (pure or applied) appears to have been directed at decision prescription.

A number of practically useful decision aids have emerged, including pattern recognition, various heuristics and even "tools" that do not involve conscious reasoning (such as hypnosis). However, most such aid has come from the ad hoc tinkering of professional decision aiders, rather than from academics.

3.3.1 Information gathering

Whether to gather more information—and what information—before making the main decision is a distinctive decision aiding problem.

In particular, such information decisions merit less effort than the main decision. The case has been made that it is worth D about $1000 to improve a decision with stakes of the order of $100,000. (1% rule-of-thumb, attributed to Ron Howard.) Then it clearly would not make sense to spend $1000 to help decide how to spend that $1000, including what new information to seek.

However, information decisions are more complex to think through, and ADT modeling is not normally effective. The best D can do may be to bear informally in mind the cost of a mistake and how much the chance of that cost can be reduced.

4 Statistical decision theory

4.1 Decision theory norms

Statistical decision theory is a logic for coherently structuring judgment, by specifying logical relationships between quantified judgments, according to norms based on persuasive axioms (Savage, 1954). A "personalist" version refers to a specific human's judgments.

The theory identifies judgments as incoherent if the logical relation among them does not hold. For example, if D's utility (suitably defined) for one option is assessed as greater than for another, but its *calculated* probability-weighted "expected" utility is *not,* D is incoherent. However, decision theory cannot directly resolve such incoherence.

Uncertain inference is a part of decision theory. It checks coherence among related probabilities (unconditional, conditional and joint), according to the classical statistics calculus. In particular, Bayes' theorem deals with updating probabilities, in the light of new evidence. (Decision theory has been misleadingly referred to as "Bayesian statistics". However, Bayes' theorem is not limited to a personalist interpretation of probability, and decision theory also addresses value judgments.).

4.2 Applied decision theory

ADT (Applied Decision Theory), introduced in the 1950s by Howard Raiffa (2002) is the practical adaptation of statistical decision theory. Its scope corresponds closely to that of the Decision Analysis Society and the journal *Decision Analysis.*

ADT typically involves constructing quantitative models that equate or approximate a choice or other judgment as a function of other quantified judgments. Any single thread of informal reasoning can generally be expressed as an ADT "mono-model" of D's mind-content and any number of mono-models can evaluate the same judgment.

4.2.1 Consequence projection

The most familiar ADT paradigm is projecting option consequences, for example in the form of a decision tree. D evaluates an option based on the consequence to which it might lead and how much he would like each. Formally, he assesses probabilities and utilities for option consequences and, if he is coherent, prefers the option with the highest mean utility.

> One CUT/WAIT argument hinged on the chance that, if I CUT, I would need a re-operation, which I considered unpleasant but unlikely. If I WAIT I projected continuing pain for now. Informally, I judged that this WAIT prospect was worse overall, and so I favored CUT. Possible quantification: Probability of having to re-operate if I CUT, 40%; utility (on some scale), 50. CUT utility with no re-operation, 90. Implied mean CUT utility, 74[5]. WAIT utility, 60, with re-operation probability 0. 74 > 60 implies CUT is preferred to WAIT.

[5] (.4 x 50) + (.6 x 90)

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

4.2.2 Other prescriptive paradigms

Other familiar decision paradigms tap into different parts of D's mind-content and can also be represented quantitatively and ADT modeled .

- *Multi-criterion estimation.* D assesses how each option affects various criteria and takes account of their relative importance. Quantification: D assesses option impacts on each criterion, evaluates their option utilities and compares options according to an approximate utility function reflecting criteria importance (Keeney & Raiffa, 1993). A common simple variant is importance-weighted criterion evaluation (using a linear additive multi-attribute utility function).

- *Analogy:* D notes comparable past dilemmas and their outcomes, and adjusts for differences. Quantification: option value equals value of analogous option, times ratio of actual to analogous value. Uncertainty about these quantities is quantified as the statistical product of probability distributions. This is a very simple form of ADT modeling, but conforms to the same norms.

- *Consultation.* D anchors his choice to another's judgment and adjusts for differences (formally comparable to analogy).

4.2.3 Uncertainty paradigms

Uncertain inference is a phase of the decision process, with distinctive paradigms.

Conditioned assessment corresponds to the observation "It all depends":

> The chance I assess for re-operation if I CUT depends on whether I avoid jolts. I was confident (misguidedly) that I would avoid jolts, which made my chance of re-operation very low, implying that, overall, re-operation was quite unlikely. Quantification: Conditional on avoiding jolts, probability of re-op 10%; re-op conditional on jolts 50%. Unconditional probability of jolts 20%. Implied unconditional re-op probability 0.18.[6]

Bayesian updating revises uncertainty based on how surprising new evidence is, given that one hypothesis or another is true. Quantification: a "posterior" probability is derived by Bayes' theorem from a "prior" probability and from diagnostic conditional probability distributions ("likelihoods").

[6](0.8 x 0.1) + (0.2 x 0.5).

For any given decision, D can implement any number (potentially thousands) of arguments, which are variants of these paradigms and others. They can be of various structures and sizes, access different regions of D's mind-content and potentially all of it. Quantification can be a calculation on the back of an envelope or a model with hundreds of variables taking months to analyze (justified in some institutional cases).

4.3 Decision theoretic interpretation of ideal choice

In 3.2 I discussed ideal choice, as the imaginary target of a prescriptive effort. Statistical decision theory can interpret ideal choice as super-coherence. Decision theorist Dennis Lindley has observed that "inside every incoherent person there is a coherent person struggling to get out." D's mind-content is transformed into a coherent and comprehensive set of related ideal probability, utility and choice judgments.[7]

However, coherence is a necessary, but not sufficient condition for ideal choice, since there are many candidate sets of coherent judgments that imply a given target choice. The ideal set is the one that "best fits" D's mind-content. "Best fit" is a problematic concept, but it might be interpreted in terms of the "firmness" of D's raw incoherent judgments, expressed as second-order probabilities[8] (Brown, 1990.)

Ideal *factual* judgment is easier to interpret than ideal *value* judgment, because it is anchored to observable reality. It could correspond to progressive Bayesian updating of some primitive prior judgment in the light of everything D ever learns, updated with evolving evidence. (The source of the original prior and exact interpretation may not matter much, since D normally accumulates enough information over time to swamp any plausible prior.)

5 Limitations of ADT modeling

5.1 Disappointed expectations

Since ideal choice conforms to decision theory norms, one might suppose that D gets closest to that ideal using an ADT model that also obeys those norms. That is not necessarily—indeed, not usually—the case. (However, competent ADT should promote some coherence.)

[7]Watson and Brown (1978) propose a theoretical treatment, involving ideal choice, of the practical value of ADT modeling, as a special case of prescriptive method.

[8]The second-order probabilities would themselves need to be ideal, which raises theoretical difficulties. I. J. Good has suggested a resolution involving the limit of an infinite series of progressive best fit operations (personal communication).

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

Our initial expectation that models would revolutionize decision-aiding has certainly not yet come to pass, even in government and business, where most of the earliest applications were attempted.

The experience of psychologist Danny Kahneman is typical. "Around 1973 both Amos Tversky and I thought ... that DA (decision analysis) would conquer the world ... I tried to implement a version of it for ... the Ministry of Foreign Affairs in Israel.... The abject failure of that effort was quite informative, and I have been skeptical ever since ... Relative to what we believed forty years ago, DA has certainly failed. But of course many of the ideas of DA are part of good practice in decision making."[9]

A number of other eminent decision scholars, such as Stephen Watson, Richard Zeckhauser, and even a founding father of ADT, have made similar comments.[10] My own experience resonates with theirs.

ADT still plays a significant role in business and government, but much more limited than we had originally thought (Brown 1992, 2009). Even Harvard Business School, the cradle of ADT teaching in the 1960s, had abandoned it by the 1990s. A number of ADT business "success stories" have been reported (Keefer, Kirkwood, & Corner, 2002). However, I have not seen documented their impact on decider actions, or the real motivation for the exercise. Some seemingly successful ADT applications have also been reported in medical decision-making (Weinstein, 1996). However, very few have been reported in private decision-making.

Why has ADT modeling so far not proven more successful in practice, especially for private decisions? I believe that ADT modeling, as commonly practiced, has enough passing and lasting limitations that, without prohibitive cost and training, it has so far rarely brought D nearer to his ideal choice than his unaided judgment—with some significant exceptions.

5.2 Unfounded criticism

Some criticisms leveled at ADT are not well-founded. For example, some argue that ADT is *inherently* ill-adapted to decision-aiding, because human psychology does not fit its norms. In fact, ADT simply aligns D's choices with related quantified judgments.

> Organizational decision theorist James March argued at an international decision conference[11], that a certain public figure could not be induced by ADT to adopt a cynical self-serving policy, because "she is not that kind of a person." I hold that her commendable values would be reflected in an ADT model that would lead her to reject ignoble options
>
> One of our adolescent students acknowledged he was deciding whether to become a drug dealer (really!). ADT of his values and factual judgments, as he revealed them, indicated that indeed he should deal drugs to be logically consistent. (He wanted to get rich fast and didn't care whose lives he ruined.)

5.3 Passing limitations

However, there are real impediments to ADT success. Some are temporary, for example, due to defects in the state-of-the-art[12] which will doubtless mature, or to decision aiders falling to keep up with the art.

5.3.1 Untapped mind-content

An ADT model (such as a decision tree) can only be a partial, imperfectly measured, coherence check on D's mind-content, since it can only address a few of the countless relevant questions that could draw important material out of D's mind. If he bases his choice on a single model or another approach, D may miss large parts of his mind-content. Much of this pitfall can be avoided by plural evaluation and synthesis, i.e., making a judgment several different ways—including by intuition—and reconciling any inconsistency (Brown & Lindley 1978).

> A *cautionary fable.* Throughout 1995, TV viewers worldwide were regaled by "the trial of the century", which, in the face of damning evidence, acquitted sports star O.J. Simpson of murdering his wife and a friend. Now, imagine that he had defended his murder decision with the following argument. "When I decided on the murders, I took into account the value of a human life, which environmental regulations have put at $20M, making the social cost of my murders $40M. I predicted (accurately) that I would be put on trial, and generate entertainment worth $10 each to 100M TV viewers. I made a public-spirited decision, which promised a social return of $1Billion on an investment of $40 million."

The fatal flaws in this implicitly ADT argument—in particular, the disastrous social precedent of legally condoning murder—are not immediately apparent to many, and might not be noticed if the action implications were

[9] Personal communication.
[10] Personal communication.
[11] SPUDM, Helsinki, 1978.
[12] Influential decision analysts have contested this view (Howard, 1992).

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

not intuitively ridiculous. This may not be the case for more realistic, less obvious, dilemmas.

> In one of my first projects, in 1961, I advised the vice-president for Europe of Stanley Tools to locate a new regional headquarters in Southern Germany, based on a model that evaluated only the economic criteria, which he had specified. He decided instead on Geneva, because it had an International School he could send his children to! If I had considered my assignment to be serving the VP personally, rather than the institution that was paying me (an ethically suspect position), I had originally missed personal convenience as a decisive criterion for him.

5.3.2 Mis-measured judgments

Even when D is aware of a consideration, he may not elicit it accurately from his mind-content. In particular, factual judgments are commonly distorted due to emotions (such as vanity, lust, laziness and delusion).

> Wishful thinking led me to mistakenly project that I would stick to a low-jolt post-surgery regimen.

However, value judgment misperceptions may be more serious than factual distortions. In particular, D often misjudges how happy he will be in future situations (Gilbert, 2006).

> Newly-weds may misguidedly avoid having children due to "temporal myopia", which, for example, may overweight the negatives of diaper-changing relative to later compensations of parenthood.

However, where emotions accurately reflect D's core *values* (like greed), rather than being a source of factual confusion, they are *properly* represented in ADT.

> ADT could *logically* support Hitler's decision to exterminate Jews, if it reflected his evil values.

5.3.3 Misleading simplification

Modeling often requires making simplified assumptions, such as "equivalent substitution". These can mislead if D does not recognize and adjust for any over-simplification. To effectively use an ADT model, D must have a rare understanding of any mismatch with his perception of reality and how to adjust for it, for example with implicit "sensitivity analysis".

D may mistakenly assume that he is indifferent between an uncertain future cash flow and a more conveniently analyzed "expected present value", which disregards risk aversion. This can be handled by assessing and subtracting a risk premium.

> I analyzed my surgery choice as if my only WAIT option were to delay 10 years, although I could, in fact, revisit the decision at any time (say, in response to hip surgery advances). This simplification penalized WAIT. Adjusting for this distortion of reality improved the case for WAIT, but not enough to switch my decision

"Cop-out" simplifications, which avoid the effort or embarrassment of "subjectivity", usually disregard critical complexities (which intuition may mysteriously deal with quite well). Decision analysts often treat unidentified possibilities as impossible. ("What could possibly go wrong?")

> Safety regulations require that the probability of human intrusion at a nuclear dump-site over 10,000 years be assessed. A risk analyst argued for one chance in a million, because that was the probability of disruptive minerals mining, the only intrusion scenario *he could think of* with a validated probability. No doubt, a Native American risk analyst 10,000 years ago, might have made a similarly optimistic assessment of the probability of human intrusion into a ancestral burial site proposed on the then desolate island of Manhattan! As Hamlet cautioned Horatio, "There are more things in heaven and earth, Horatio, than are dreamt of in your philosophy."

Disregarding value and uncertainty dependencies can lead to grave mistakes.

> A World War II story had the German army shipping left boots in one plane and right boots in another. One plane was shot down.
>
> You might think "Mahomet Chen" was the most common name in the world, if you disregarded probabilistic dependence between the two names.

5.3.4 Confusing language

Language that confuses or misleads the lay decider limits the use and usefulness of ADT modeling (Howard, 2004, Brown, 2004). For example, "Decision analysis" implies that ADT is the only way to analyze a decision (Howard 1968, Brown 2008).

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

> An early consulting client, president of a major company, interrupted an ADT presentation I was proud of. "Gobbledygook!" he muttered, stormed out and had nothing further to do with ADT.

In this essay, I am largely sticking with terms familiar to JDM readers.

5.4 Enduring limitations

Some ADT limitations are essentially inherent and not readily corrected.

5.4.1 Misallocated deliberative effort

The inescapably burdensome challenge of modeling a choice diverts attention from critical decision processes *other than* choice, such as gathering information, identifying new options and acting on the decisions made. Within the choice process, constructing an ADT model can divert effort away from developing sound input. Einstein said "imagination is more important than knowledge", implicitly dismissing as relatively unimportant how soundly the imagination and knowledge are *processed*. When ADT first came to scholarly attention, classical statistician Andrew Ehrenberg disparaged it as "sonking"—from "SONK, the Scientification Of Non-Knowledge". He had a point: mathematics can lend misleading authority to a subjective argument and encourage over-spending resources on it.

> I was running late for class and had to decide whether to illegally take a traffic lane reserved for high occupancy vehicles. I did a hurried ADT in my head, in the five minutes before the decision point. I judged that avoiding a 50% risk of being late was worth incurring a 5% chance of a $50 fine, so I took a chance. Too late, I realized I had over-looked another cost of getting caught: it would raise my chance of being late to virtually 100%. Moreover, I had not considered my distaste for illegal action (regardless of consequence). If I hadn't been distracted by the primitive modeling effort, I would have taken these factors into account, informally or intuitively, and wisely avoided the fast route. (As it happened, I made class on time, i.e., my bad choice luckily had a good outcome.)

5.4.2 Distorting wise normal practice

D often naturally uses wise decision strategies that are not readily modeled. For example, he takes a series of small incremental steps toward—or away from—a major commitment, interleaving them with gathering information.

> You put your toe in the water, then possibly your leg, before committing to taking the plunge.
>
> You get progressively more involved with a date and back off if it begins to look unpromising, before winding up at the altar—or not.

In theory, D could prescribe an incremental strategy with a complex "dynamic programming" ADT model, which evaluates possible strategies, step-by-step. This is prohibitively burdensome, but I am not really satisfied with any more tractable approaches that I have explored (Brown, 1993b). These include substituting a simplified "once-and-for-all" option for the real evolving decision process, and heading in the direction of the favored option, until developments indicate otherwise. An alternative is to look only a few steps ahead and compare the situations they take you to. (In chess, you might project the value of pieces left on the board.)

> Careers can be planned *as if* D were about to make an irreversible commitment to one career or another. His real options would be a complex sequence of small, partially reversible, moves spanning many years, covering schooling, job successes and failures, experimental jobs, etc.

Evaluating several independent reasoning approaches, including intuition, can always improve on unaided judgment, *provided that* synthesis of the plural evaluations is handled properly, which is still an undeveloped art.

5.4.3 Ill-fitting estimation models

When an ADT model is approximation, not an identity (as in "population ≡ whites + non-whites"), an error term is needed to make it an identity. If the approximation is poor, allowing for error may be too difficult to be useful. For example, in additive importance-weighted criteria evaluation (noted earlier) criteria may be dangerously incomplete or value-dependent.

6 ADT contributions

Most of ADT's most valuable contributions to D's private decision making are not by basing immediate decisions on ADT models.

6.1 Quantitative models

However, there are significant exceptions.

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

6.1.1 Prescriptive use

An ADT model can be directly useful in making a decision if the decision hinges critically on a question of *degree* (i.e., on quantification).

> My daughter Karen and I used a simple ADT model to decide whether she should deliver twin babies by C-section (Brown 2005b, Prolog). We reduced her choice to comparing her probability of successful natural births with a ratio of consequence utilities (which reflected how much worse a C-section after delivering one baby naturally would be than a C-section for both babies from the start). She was confident that the probability of a successful natural delivery would be high enough to try, even if she researched the issue more. She successfully delivered both babies naturally.

The main deterrent to ADT modeling may not be the quantification itself, but specifying precise numbers. This may not be necessary.

> In the hip dilemma, we assessed and summed the differential impact of options on each criterion simply as 1, 2, 3 pluses or minuses. This was still quantification, but cognitively easier and, in this case, a good enough approximation to actual numbering.

Graphic analogs can avoid numbers entirely (Brown & Pratt, 1996). For example, we found that middle schoolers could compare two options effectively by manipulating a balance beam (Martin and Brown, 1991). They set and adjusted weights and distances to correspond to relative probabilities and utilities, and predicted accurately which side of the beam would tip down, and thus which option they preferred. Implicitly, they were comparing probability-weighted mean utilities.

6.1.2 Uncertainty assessment

ADT modeling has unquestionable usefulness in assessing uncertainty. In particular, conditioned assessment is usually simple enough for laymen to apply unaided.

> Target assessment: Probability of a stock boom next year. Input: probability of Republican electoral win, 20%. If so, probability of a stock rise, 90%; if not, 50%. Implied unconditional stock rise probability, 58%.[13]

Somewhat more complex Bayesian updating is used very effectively, especially on institutional uncertainties (McGrayne, 2011).

[13](.2 x .9) + (.8 x .5)

6.1.3 Use for multiple deciders/occasions

Some decisions interest enough deciders or occasions to justify substantial group effort, say, on career, health, finance and inter-personal relationship choices.

> The career advisory program at the Educational Training Service (ETS) has explored adding an ADT module to their SIGI+ career program. It would integrate information resources, such as employment opportunities and qualifications needed, with clients' career priorities.

"Parametric" models, with resettable inputs, are used on classes of medical decision (Weinstein, 1996).

6.1.4 Post-decision validation

Even when the best way for D to make a decision is not (yet?) with ADT, it may nevertheless help him *validate* his choice (heeding the Red Queen's command in *Alice in Wonderland*: "Verdict first—trial second!"). An ADT model can more readily be retrofitted to a decision than constructed to make the decision in the first place. A plausible model lends credibility to a decision.

> (Hip dilemma). The surgeon who advised me to WAIT accepted my contrary input judgments as reasonable and agreed to CUT. However, it would have been impractically burdensome for me to model his judgments to help him make his *own* mind up originally.
>
> My 13-year-old neighbor Katie had to decide whether to attend a science-oriented "magnet" school. Her intuition was to turn it down, against the strong advice of her parents. We constructed a simple ADT model, based on her judgments (such as lack of interest in science and reluctance to travel two extra hours a day), which confirmed her inclination to turn the school down. We went over the analysis with her parents, who then understood and accepted her position. (A possible alternative outcome to this exercise would have been that Katie's parents persuade her to change her inputs and reverse her choice.)

With some public policy controversies, an ADT model can pin-point sources of disagreement.

> The US Senate Judiciary Committee was deciding whether to support a Community Anti-Crime Bill. A simple ADT model, with importance-weighted estimates of impact on crime, public cost, etc. convincingly argued

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

for the bill's passage. The influential Association of Chiefs of Police (whose political support Committee chairman Biden depended on) strongly opposed the Bill, which would have diverted public funding from the police. We added an "electoral security" criterion to the model and gave it to the chairman to supply his own inputs *privately*. The Bill failed to pass. . . .

6.2 Informal versions of ADT paradigms

Virtually *any* defensible line of decision reasoning can be represented formally or informally by an ADT argument. The informal versions may sharpen reasoning, without the burden and some of the pitfalls of ADT modeling. Although formal modeling is sensibly preceded by informal ADT, the decision process need go no further than informal analysis.

6.2.1 Consequence projection

At the outset of the 2008 housing crisis, many home owners who predicted (accurately) that house prices would continue to drop refused to sell for less than they had originally paid, disregarding the irrelevance of "sunk costs" to projecting consequences. ADT logic could have saved them great losses when they eventually accepted to sell at a still lower price.

Misguided action often results from failure to distinguish facts from values.

Australian authorities cancelled a Beatles tour, on the grounds that Lennon's values were objectionable, after he claimed that more people knew his name than Christ's. He was merely expressing a factual (if possibly mistaken) value-neutral judgment. The Pope might sadly agree with him, without outraging the faithful.

However, recognizing the difference between fact and value can be elusive without ADT logic.

You might interpret the proverb "You catch more flies with honey than with vinegar" as advising you to always use honey when hunting flies and, by extension, to treat antagonists gently. This inference would be sound, if catching flies—or winning over antagonists—were your only criteria. ADT logic alerts you to check whether other criteria (such as deterring adversaries in future) make a difference.

6.2.2 Other decision paradigms

Multi-criteria evaluation.

ADT-based argument for CUT: "Surgery now will be more inconvenient and unpleasant now, but that is really less important to me than later satisfaction, so I prefer to CUT."

Analogy.

In the 2008 US fiscal crisis I argued: "I am not competent to judge the merits of a multi-billion dollar bank bailout. The Washington Post, whose editorial views on such matters I generally trust over my own, strongly favors a bail-out. I suspect that the Post is influenced by bank pressures, but not enough to dissuade me from supporting a bailout."

Consultation.

My surgeon cautiously advises me to WAIT. He may tend to over-prescribe surgery (as many surgeons do), so his untainted WAIT advice should be even stronger. I see no reason why my best judgment should differ from his, which encourages me to WAIT (although, for other reasons, I finally opted to CUT).

Focusing on relevant ADT logic questions helps D to evaluate the persuasiveness of opposing consultants' arguments.

Voter D's dilemma is whether to support gun control, given conflicting pros and cons. Con: "Guns don't kill people, people do!" The relevant consequence projection question is "Will potential murderers kill more people if they have easier access to guns?" Con: "There is no proof that gun control reduces homicide". Pro: "Absence of evidence for is not evidence against". The appropriate multi-criteria question is "Is respect for the right to bear arms so dominant a criterion that saving lives does not count against it, however strong the case, short of proof?"

6.2.3 Uncertain assessment

Bayesian updating.

Until I learned better, I believed Saddam had weapons of mass destruction. My prior judgment was 50:50. Then I learned that he was refusing UN inspection. That news would

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

surprise me more if Saddam *did* than if he *didn't* have the weapons. So, by Bayesian updating logic, I concluded that he probably does. (My judgment was coherent—if factually inaccurate—but I had to use my judgment at the time).

Conditioned assessment.

Re-operation if I CUT would depend on whether I avoid jolts afterwards, which I was confident (again mistakenly) I would. My chance of re-operation would then be very low, but high otherwise. This implied re-operation was quite unlikely.

Analogy.

D is considering buying a house and reasons: "Houses like this have previously gone for about $300k. In the present downturn, I expect I could get it for about 20% less."

6.2.4 Gathering new information first

Consequence projection is logically appealing for evaluating information-gathering decisions. However, it is difficult and cumbersome to implement, even informally. Consultation and adjusting conventional wisdom is usually more practical.

Before I decided to CUT, I could have read up on relevant medical literature. However, I reasoned: "With the little I know now, I may regret my CUT choice. However, I cannot learn much that would change my mind and what I could learn would be too much trouble."

Proverbs are often over-simplified to the point where they mislead, without qualification.

Seemingly contradictory proverbs address whether to get information before acting: "He who hesitates is lost" and "Look before you leap." ADT logic says that this decision depends on additional judgments that make the difference, such as information cost, probability of switching choice and the cost of mistakes.

6.2.5 Hybrid paradigms

Effective informal argument need not correspond to a single ADT paradigm. The following case combined Bayesian updating, consultation and projection.

In mid-2011, I reported to my doctor that I had lost ten pounds over the past year. He re-assured me that such weight swings were not unusual and recommended "watchful waiting", which was normal medical practice in such cases. However, I had heard that cancer was commonly associated with weight loss, so I decided to look at the issue through an ADT lens, although I have no medical background. Should my probability of having cancer be sufficiently higher than medical practice implies to justify costly diagnostic tests?

I reasoned along Bayesian updating lines. I knew cancer is more common at my age (77) than in the population at large, so my prior probability of cancer was higher than normal. My specific weight-loss pattern, including a steady one pound loss a month over ten months, struck me as most surprising if I did not have cancer, but not at all if I did, which was very diagnostic. Without using any numbers, my intuition was that these two observations raised my cancer probability enough to justify diagnostic tests. My doctor approved them and they were done. They revealed aggressive pancreatic cancer, but at an early enough stage to be successfully operated on. (It is very rarely caught early enough to save the patient's life.) I now have no detectable cancer, a success that would have been unlikely without my drawing on ADT logic, however loosely.

(My opting for cancer screening may have been a wise private choice, but socially irresponsible. In the competition for scarce health resources, precedence should surely go to a younger patient, whose life it may extend by 30 or 40 years more than mine. This kind of projection reasoning might have persuaded a nobler soul to pass up testing.)

Formalizing an informal ADT argument in a case like this would make it firmer, more precise, informative, conclusive and more adaptable to circumstances. However, informal ADT is more communicable and "good enough" in most cases (as here) to lead an ADT-trained D (like myself) to much the same conclusions as a competent ADT model, but with much less effort.

6.3 Personalist philosophy

In addition to helping D make specific decisions, ADT provides a valuable philosophy, often at odds with conventional wisdom. In particular, it is *personalist*, i.e., it organizes personal judgment.

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

6.3.1 Research for useful vs. scientific purposes

Although this essay is intended to be practical, not scientific, ADT logic can help descriptive research contribute better to decision aiders' knowledge of D's decision processes. Bear in mind that D needs and uses whatever information he can get, no matter how subjective or tentative

Acceptable, i.e., publishable, scientific research findings need to be of enduring interest, demonstrably impersonal and replicable

> In 1990, I was organizing a number of research projects to help a Democratic Congress gauge if the 1970 Clean Air Act was worth its cost. None of the project researchers would pass me any findings until these cleared scientific publication standards. Before that happened, Republicans took control of Congress and cancelled the project, with nothing to show that policy-makers—or an interested public—could use.

As I understand it, the core of a typical scientific experiment is something like this. A hypothesis about a population is tested against a random sample from that population. The null hypothesis is rejected at, say, a .01 significance level, if the frequency of a divergence as great as observed is less than 1%. This finding may be of only marginal interest to D. In particular, a large enough random sample will reject the null hypothesis, at any significance level, no matter how small the divergence (Bayesian psychologist, Ray Bauer, personal communication).

In a classic psychological experiment (Phillips & Edwards, 1966), student subjects assessed the color mix of balls in an urn based on a random sample. They made much weaker inferences than Bayes' theorem implied. The hypothesis that the parent population of students subjected to this treatment would respond "conservatively" was confirmed—i.e., the null hypothesis was rejected—at, say, the .01 level.

An interesting question here to deciders and scientists is: what can one infer from this experiment about whether people *generally* process evidence too conservatively. The researchers surmised that they do (Edwards,1968). My ADT perspective prompts me to critically assess "sampling error" (Brown, 1969). I suspect that subjects make unfamiliar inferences about balls in an urn, from random samples quite differently (and more conservatively) than people making everyday inferences from the happenstance evidence they normally observe,

This was a cheaply administered sample, large enough to promise "significant" results. If the decision-aiding objective of the experiment research were information about decider conservatism, I would favor using the research budget to give fewer subjects a less convenient inference task, but more representative of their everyday experience. (However, any findings would still need scientific validation.)

Decision aiding practice need not be subjected to scientific documentation constraints. Environmental policy specialist Grainger Morgan (1978) has convincingly argued that bad science can be good policy. Policy pundit George Kennan observed "Tentative answers to important questions are more useful than tentative answers to less important questions."

A useful decision tool-kit can include lo-tech devices, such as "sleeping on it", exercise, alertness pills and even drinking coffee, that have not been validated by reviewable research, although an eminent psychologist has faulted such a list as "too eclectic". As it happens, "sleeping on it" is supported by research (Killgore et al., 2006). While not necessary for practical purposes, such credentials do enhance a decision tools claim to be effective.

6.3.2 Interpreting pseudo-objective communications

A widespread "judgment-free" ideology has trickled down from scientific research (where it belongs) to practical decision-making (where it does not). ADT logic can counter spurious insistence on "objectivity", which is often used to advance special interests.

Environmental regulation. As commonly practiced, "Cost-benefit Analysis"(CBA) omits benefits that cannot be authoritatively documented, thereby ignoring impacts requiring personal judgment. This allows industry costs to dominate environmental benefits and discourage regulation.

> A government agency had us evaluate whether a BP oil-drilling development in the Arctic should be permitted (Brown et al, 1997). Our study was initiated by an oil and gas trade association that, no doubt, expected a conventional CBA. When it became clear that we would take into account informed judgment about oil-drilling impacts (say, on fish and wildlife habitat), our contract was terminated.

Optimistic Risk assessment. Some environmental regulations require reports on the safety of hazardous facilities (such as nuclear plants) to adopt "Probabilistic Risk Assessment" (PRA) procedures (US NRC, 1984). PRA, as required and practiced, normally takes account only of well-documented sources of risk (such as experiments) and disregards other evidence, however relevant (such as unscheduled inspections). The public is misled into believing facilities are safer than realistic judgment would indicate.

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

> We reviewed one such PRA (paid for by the regulated company), which treated more than half the risk at a reactor (e.g., from earthquakes) as zero, because "we don't have the data" (Brown & Ulvila, 1988).

Sample survey reporting. Estimates based on sample surveys routinely report "a margin of error" of, say ± 5%. Reassured clients may not be alerted that this addresses only random sampling error (the sole source of unambiguously calculated error). Total error (including mismeasured samples and unrepresentative sampling frames) is usually much greater. ADT can readily account for it (Brown, 1968).

> In 1936, the Chicago Tribune miscalled a presidential election result ("Landon defeats Roosevelt"), apparently because it had sampled from a frame tilted toward well-to-do voters who mostly supported Landon.
>
> Social surveys, based on personal interviews, have reported that men *on average* have more heterosexual relationships than women. Arithmetically (i.e., logically) they should be virtually the same, so there must be error in one or both averages. I suspect measurement error: men overstate and/or women understate how many partners they have had.

6.4 Educating intuition

It has been said: "Education is what you're left with when you've forgotten everything you've learned." Training in decision theory and its logic educates both intuition and informal decision making guidelines. Physicist Richard Feynman could intuitively approximate the results of complex mathematical operations, without calculation, possibly as a result of having performed many comparable calculations in the past.

All the early ADT pioneers that I know who have left academia to become influential deciders[14] tell me that they never use ADT explicitly, but benefit informally from the training. Experience applying ADT to one instance of a common dilemma can enhance informal ADT or intuition on others, such as career planning or investment.

Economist Herb Simon[15] has observed that, although experts (such as chess grandmasters) commonly decide by pattern recognition, novices develop expertise in ways akin to ADT, such as consequence projection.

However, poorly understood ADT can misdirect judgment. For example, ADT beginners learn that "Bad outcomes don't mean bad decisions", but may overlook the informative clues outcomes can provide.

> Losing a game of bridge doesn't *necessarily* imply that you played poorly, but it begins to cast doubt, which becomes stronger as your losses mount up.

7 ADT prospects

It will take the development of a major "decision prescription" enterprise for ADT—formal or informal—to become part of a decision-aiding tool-kit that will enhance people's lives.

That will require initiatives from several overlapping communities, which cannot be taken for granted.

7.1 Prescriptive scientists

A stable cadre of researchers dedicated to prescriptive research, and empowered to conduct it, requires that researchers can count on attractive long-term career opportunities and resources. Currently, basic prescriptive research is conducted almost entirely within university units with a separate primary emphasis, such as business, psychology, economics and statistics. A researcher's career advancement depends on giving these units primary allegiance and meeting their professional criteria. It is, moreover, rarely advanced by an interdisciplinary orientation.

An infrastructure of institutions dedicated to decision prescription is needed. These would includes a professional association (narrower than SJDM and broader than the Decision Analysis Society), a journal,[16] dedicated university departments and a research institute. Activities pursued should balance normative, descriptive and interdisciplinary interests, and draw on live case material to generate research ideas.

There have been unsuccessful precedents. In the 1980s, I served on an NSF panel that supported a failed $10M proposal from Duke University to fund a Decision Science Institute.

In the 1970s, my efforts to introduce a prescriptive doctoral program at the University of Michigan, were blocked by established departments (whose favorite students showed interest in switching). Resistance to new institutional orientations may not give way until a new generation of scholars replaces the old.

[14] Including an Under-secretary of the Treasury, a congressman, a prominent financier, a business owner and a head of the Federal Price Control Board.

[15] Personal communication.

[16] I appreciate JDM's hospitality for this ugly duckling essay.

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

7.1.1 Research topics

Many topically important (but not timeless) prescriptive research issues are likely to remain neglected in the current institutional environment[17] (Brown, 2006). Journals normally publish only authoritatively supported findings, which encourages research along well-established discipline-specific lines. It discourages productive exploratory research and promising hypotheses—perhaps suited to a mythical "Journal of Partly-Baked Ideas" proposed by Philosopher I. J. Good (1962).

However, some prescriptively relevant, but so far largely neglected, research topics have enough academic appeal to assure they will eventually be pursued, e.g.

- How to measure the quality of decisions and decision processes (Brown, 1994).
- Longitudinal empirical research on whether those trained in ADT decide better than than others.

7.2 Decision-aid technologists

Pure research and applied decision aid design have different imperatives and the latter need not wait on the former. As Tversky has remarked, "You don't need to finish the basement before starting on the roof".[18] Moreover, pure and applied research need not be co-located (Lorsch & Lawrence, 1968). ADT methodology can be focused on a single domain (as it already is by the Society for Medical Decision Making). Private decisions (this essay's focus) could be the province of a personal counseling organization (which perhaps exists).

As ADT technology and acceptance evolves, enterprising technologists/entrepreneurs may develop new marketable products. Public interest groups could create parametric models to help citizens make controversial decisions, such as whether to support legislation on health policy, abortion and economic stimuli, entering their own input. Emerging socio-technological trends may exploit ADT.

> Computerized on-line dating services might integrate currently competing matchmaker judgment and statistical data in predicting romantic success, using Bayesian updating algorithms.

7.3 Decision aiders

Natural selection should winnow out lags between ADT practice and state-of-the-art ADT technology. However, aiders and deciders have some entrenched priority mismatches (Brown, 2005a). For example, aiders may resist plural evaluation of options—i.e., trying different approaches to a dilemma—to avoid professionally embarrassing discrepancies. Deciders may have to take the initiative to overcome such conflicts of interest.

Decision aid practitioners have tended to specialize by technique (such as ADT), but to generalize by application domain. Practitioners should gain competitive advantage by broadening their decision-aiding technique to include other types of aid (such as cognitive vigilance), but focusing on limited problem areas (such as private choice).

A reformed decision glossary, legitimized by an authoritative professional body, could cure much current decider confusion with practitioners' decision language.

7.4 Method disseminators: Teachers and writers

Recent best-sellers have catered to public interest in how people *do* make decisions (Ariely, 2008). Before long, comparable interest may develop in how people *should* make decisions, what ADT pioneer Robert Schlaifer used to call "mental hygiene".

As the field matures, educational institutions and students may press for the teaching of practical decision and other reasoning. Critical Thinking is a broad emerging field of logical enquiry and reasoning, which appears to include decisions and judgments (Hunter, 2009) and may generate useful vehicles and precedents for ADT-focused education. I expect training in ADT (along with other decision skills) will become an integral part of all education, as trained teachers and instructional materials become available.

> In the 1970s, Venezuela's first (and only) Minister of Intelligence, Luis Machado, mandated decision skills as a required part of the country's educational system, but the requirement was retired with him.

In primary school or middle school, pupils may learn to ask questions along ADT lines. ("What does Mommy think? Why? Do I believe Mommy? So what?"). In high school students may learn basic decision skills, including the psychology and logic. (Such skills could arguably even take priority over hallowed academic requirements, such as regional geography, calculus or Shakespeare). At college level, more ambitious decision skills courses would be universally offered. ADT material is already being integrated into some existing "judgment and decision making" psychology courses (Baron, 2003).

Societal resistance may impede educational progress. Disinterested rational deliberation can threaten all kinds of vested interests who fear "undesirable" behavior.

[17] Some years ago, NSF hired a director from industry (Eric Bloch) charged with emphasizing *useful* research, but for lack of internal NSF support, the appointment was short-lived.

[18] Personal communication.

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

Fundamentalist parents have complained that teaching decision skills to grade-schoolers will weaken their religious faith. Environmental Science faculty at a British university resisted teaching ADT, fearing it would lead students to the "wrong conclusions". A respected surgeon admitted to me that it would "cost me money", if patients made ADT-informed decisions about costly surgery.

8 Conclusions

ADT may never enjoy the widespread—indeed revolutionary—success that we once forecast. However, as ADT becomes integrated into the culture, I am cautiously optimistic that people will make wiser choices—and act on them to enhance the quality of their lives. In everyday discourse, we should hear more comments like "on the other hand", as people weigh pros and cons with less thoughtless certainty. People who already structure their thinking along ADT lines naturally may still have an edge—but less so.

Citizens who have absorbed some ADT logic will vote, based on projected consequences and their own basic priorities, and less on impulses engineered by others. Public servants who expect voters to serve their own interests wisely may be moved to further the public good and reduce the risk of bequeathing a destitute nation to their grand-children.

The judicial system may benefit from jurors deciding lawsuits more soundly.

In a recent trial, Casey Anthony was acquitted of murdering her daughter, in the face of damning evidence. If jurors had some ADT priming (or were guided by a judge who had), I expect they would have appropriately found her guilty "beyond reasonable doubt".

It should come as no surprise that a man-made technology like ADT should be slow in out-performing a human reasoning capacity that has been forged over millennia of evolution. However, other "unnatural" technologies have competed successfully with nature. It took many decades before the Wright Brothers' efforts at applying aero-dynamic theory to man-made flight first emulated and ultimately dominated the flying performance of birds. Perhaps we can expect Raiffa and Schlaifer's decision theory innovations to eventually revolutionizing human decision-making, after all.

References

Ariely, D. (2008). *Predictably irrational: the hidden forces that shape our decisions*. New York: Harper Collins.

Baron, J. (2003). *Thinking and deciding, 3d ed.*. New York: Cambridge University Press..

Baron, J., & Brown R. V.(1991). *Toward improved instruction in decision making to adolescents: conceptual framework and pilot program.* In Baron, J., & Brown, R. V. (Eds.), *Teaching decision making to adolescents*. New York: Lawrence Erlbaum Associates,.

Bazerman, M. E. (2005). *Judgment in managerial decision making*. New York: Wiley.

Brown, R.V. (1968). Evaluation of total survey error . (Oswald George Prize Paper). *Statistician, 17*, 335-355.

Brown, R. V. (1969). *Research and the credibility of estimates*. Boston: Harvard University, Graduate School of Business Administration, Division of Research.

Brown, R. V. (1978). Heresy in decision analysis: Modeling subsequent acts without rollback. *Decision Sciences, 9*, 543–554.

Brown, R. V. (1990). Assessment uncertainty technology for making and defending risky decisions. *Journal of Behavioral Decision Making, 3*, 213–228,

Brown, R. V. (1992). The state of the art of decision analysis: A personal perspective. *Interfaces, 22 (6, Nov-Dec.)*, 5–14.

Brown, R. V. (1993a). Impersonal probability as an ideal assessment based on accessible evidence: A viable construct? *Journal of Risk and Uncertainty, 7,* 215–235.

Brown, R. V. (1993b). *Managing incremental commitment.* Working paper. School of Public Policy, George Mason University, Mason. http://mason.gmu.edu/~rbrown/docs/Managing.pdf

Brown, R. V. (1994). The role of statistical decision theory in decision aiding: Measuring decision effectiveness in the light of outcomes. In Freeman, P. R. & Smith, A. F. (Eds.). *Aspects of uncertainty: A tribute to Dennis Lindley.* New York: Wiley.

Brown, R. V. (2000). *Fitting decision aids to institutions: Organizational design issues*. http://mason.gmu.edu/~rbrown.

Brown, R. V. (2004). Naming concepts worth naming. *Decision Analysis, 1*, 86–88.

Brown, R. V. (2005a). The operation was a success but the patient died: Aider priorities influence decision aid usefulness. *Interfaces, 36,* 22–31.

Brown, R. V. (2005b) *Rational choice and judgment: Decision analysis for the decider.* New York: Wiley.

Brown, R. V. (2006). Making decision research useful—not just rewarding. *Judgment and Decision Making, 1*,

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

162–173.

Brown, R. V. (2008). Decision analysis: Applied Decision Theory? *Decision Analysis Today, 6 (May)*, 10–12.

Brown, R. V. (2009). Working with policy makers on their choices: A decision analyst reminisces. *Decision Analysis, 6*, 14–24.

Brown, R. V., Flanders, N. E., & Larichev, O. I. (1997). Decision science for regulating the arctic environment. *Arctic Research of the U.S., 10 (fall/winter)*, 24–33.

Brown, R. V.,, Kahr, A. S., & Peterson, C. R. (1974). *Decision analysis for the manager*. New York: Holt, Rinehart, and Winston,.

Brown, R. V., & Lindley, D. V. (1978). *Reconciling incoherent judgments—Toward principles of personal rationality.* McLean, VA: Decisions and Designs, Inc. (NTIS No. AD A059639), July.

Brown, R. V., & Pratt, J. W. (1996). Normative validity of graphical aids for designing and using estimation studies. In R. Zeckhauser, R. L. Keeney, & J. Sebenius (Eds.). *Wise choices: Symposium in honor of Howard Raiffa.* New York: Wiley.

Brown, R. V., & Ulvila, J. W. (1988). Does a reactor need a safety backfit? Case study on communicating decision and risk analysis information to managers. *Risk Analysis, 8*, 271–282.

Clemen, R., & Reilly T. (2000). *Making hard decisions with decision tools*. Belmont, CA: Wadsworth.

Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmutz (Ed.), *Formal representation of human judgment.* New York: Wiley.

Gilbert, D. (2006). *Stumbling onto happiness.* New York: Knopf,

Gilovich, T., Griffin, D. W., & Kahneman, D. (2002). *Heuristics and biases: The psychology of intuitive judgment.* New York: Cambridge University Press.

Good, I. J. (1962). *The scientist speculates: an anthology of partly-baked ideas.* London: Heinemann.

Howard, R. A. (1968). Decision analysis: applied decision theory. In *Proceedings of the Fourth Conference on Operational Research*, New York: Wiley.

Howard, R. A. (1992). Heathens, heretics and cults: the religious spectrum of decision analysis. *Interfaces, Nov-Dec, 22*, (6), 15–17.

Howard, R. A. (2004). Toward precision in decision language. *Decision Analysis, 1*, 71–78.

Hunter, D. (2009). *A practical guide to critical thinking: Deciding what to do and what to believe*. New York, Wiley.'

Kahneman, D. (2012). *Thinking fast and slow.* New York: Farrar, Strauss.

Keefer, D. L., Kirkwood, C. W., & Corner J. L. (2002). *Summary of decision analysis applications in the operations research literature, 1990–2000.* Dept. of Supply Chain Management Arizona State University.

Keeney, R. L., & Raiffa, H. (1993). *Decisions with multiple objectives: Preference and value tradeoffs.* New York: Cambridge University Press. (Originally published, 1976.)

Killgore, W. D., Balkin, R. J., & Eisenstein, N. J. (2011). Impaired decision making following 49 hours of sleep deprivation *Journal of Sleep Research, 15,* 7–13.

Lindley, D. V., Tversky, A., & Brown, R. V. (1979). On the reconciliation of probability assessments. *Journal of the Royal Statistical Society* A. *142*, 146–180 (with commentary).

Lorsch, J. & Lawrence, P. (1968). *Differentiation and integration in industrial organizations.* Boston: Harvard Business School.

McGrayne, S. B. (2011). *The theory that would not die: how Bayes' theorem cracked the enigma code, hunted down Russian submarines and emerged triumphant from two centuries of controversy.* New Haven: Yale University Press.

Martin, A. W., & Brown R. V. (1991). Analog devices for teaching decision skills to adolescents. In J. Baron and R. V. Brown (Eds.), *Teaching decision making to adolescents*, pp. 207–235. Hillsdale, NJ: Lawrence Erlbaum Associates.

Morgan, M. G. (1978). Bad science and good policy analysis. *Science. 201*, 971.

Phillips, L. D. & Edwards, W. (1966). Conservatism in a simple probability inference task. *Journal of Experimental Psychology, 72*, 346–355.

Porter, H. (1987). Why the Super-powers need this Englishman. Cover article *Illustrated London News.* October.

Raiffa, H. (2002). Decision Analysis: A personal account of how it got started and evolved. *Operations Research. 50*, 179–185.

Raiffa, H., & Schlaifer, R. (1962). *Applied statistical decision theory.* Boston, MA: Harvard University Graduate School of Business Administration, Division of Research.

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

Schlaifer, R. O. (1969). *Analysis of decisions under uncertainty.* New York: McGraw-Hill.

Simon, H. A. (1979). Rational decision making in business organizations. *American Economic Review*, *64*, 493–513.

Tolcott, M. A., & Holt, V. (Eds.) (1988). *Impact and potential of decision research on decision aiding.* American Psychological Association.

U. S. Nuclear Regulatory Commission. (1984). *Probabilistic Risk Assessment (PRA): Status report and*

Guidance for Regulatory application., NUREG-1050. Washington, DC: USNRC.

Watson, S. R., & Brown, R. V. (1978). The valuation of decision analysis. *Journal of the Royal Statistical Society, Series A*, *141*, 69–78.

Weinstein, M. C. (1996). Decision analysis in health and medicine: Two decades of progress and challenges. In R. J. Zeckhauser, R. L. Keeney, & J. Sebenius (Eds.). *Wise Choices.* Boston, MA: Harvard Business School Press.

Zeckhauser, R. J., Keeney R. L., & Sebenius J. (Eds.) (1996). *Wise Choices.* Boston, MA: Harvard Business School Press.

Appendix

A: Private vs. institutional choice

This essay has referred primarily to individual decisions, although most explicit uses of ADT—including my own—have been for institutions (whose mistakes are more costly). Indeed, schools of business and public policy are the main market for ADT courses.

Business ADT pioneer Art Schleifer illustrates how special-purpose ADT adaptations often apply to personal decision-making: "When should I leave for the airport for my flight? Time to airport is uncertain. If you arrive too late, you miss the flight (very bad), but if you arrive too early, you waste time in the airport (bad, but not as bad). This is a critical-fractile (inventory control) problem which one can analyze informally—generally, leave a safety margin that balances the two bad outcomes appropriately.... I am being sued for $x. Should I settle for $y (less than $x)? It depends on my assessment of the chances of winning. An informal example of option pricing." The stakes in private choice are rarely high enough to justify comparable modeling.

Private citizens evaluate the same public policy issues as government officials, but quite differently.

> Government may devote many man-years deciding what national health plan to propose to Congress (supported by much factual research). Few voters will spend more than an hour or so on such an issue.

Institutional and private decisions have distinctive features. For example, successful decision aid to institutions depends on organizational more than on logical or cognitive issues (Brown, 2000). However, the distinctions sometimes blur. Institutional managers can act as private individuals, whose criteria include institutional interests.

> A senior nuclear regulator acknowledged that he was influenced by the "hullaballoo factor" (professional embarrassment), when deciding if a reactor was safe enough to operate. (His predecessor lost his job after the 1979 "Three Mile Island" accident, which produced negligible health impact, but much unwelcome media attention).

To say that an institution itself "has" an ideal choice, we must treat it as a "unitary actor" and be prepared to say things like "It was not rational of AOL to buy Time-Warner Inc." This requires attributing values to an organization, which is certainly a stretch. (Economist Amartya Sen has argued against the existence of institutional values, but acknowledges that some kind of "revealed preference" can be inferred from institutional action.[19]) Factual judgments would have to be attributed to a human assessor (possibly technical staff), not to the institution.

B: How experience has shaped my views

My cautious view of what ADT can and cannot usefully do has evolved over some 50 years of applying and adapting it for use on varied decisions, including: a teenager on what high school to attend; a woman on whether to leave her husband; an auto-executive on whether to get into the tire business; a regulator on whether to close down a reactor.

As a youth, my interest in decision aiding was driven by acknowledged deficiencies in my decision-making. ADT has greatly improved my "slow", but not my "fast", decisions (Kahneman, 2012). My daughter Karen will not allow me to drive her 12-year old twins, knowing how badly I react to traffic incidents. This is the daughter who, some 12 years earlier, had trusted me to guide her decision, when I had time, on whether to have these same twins delivered by C-section (see 6.1.1).

I went to university intending to learn how to make my mind up effectively. I soon abandoned a logic and philosophy program as unusably abstract. I switched to economics, then social anthropology. My first job, with a consulting firm, was to help managers make marketing decisions. Since I had no business experience, I felt I could only contribute some kind of logical decision process. Not finding any relevant activity in England, I arranged a short stint with the Management Studies Group at Cambridge. I first learned some statistics by teaching it there and qualified as an Incorporated Statistician. Then I tried to develop an, inevitably primitive, statistics-based method for quantifying judgment.

Then I learned that Robert Schlaifer and Howard Raiffa at Harvard Business School were already adapting, for managerial use, an advanced "statistical decision theory", with basically the same objectives. I sent them

[19] Personal communication.

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press

a paper (Brown, 1968), which received an applied statistics award, on the strength of which they invited me to join their group. I spent a few years at HBS learning and teaching ADT modeling. Around this time, a few ADT enthusiasts and I founded the Decision Analysis Society. To balance the normative with the descriptive, I spent a few more years at the University of Michigan, to learn how to account for human considerations in ADT modeling from behavioral decision theorists Ward Edwards and Cam Peterson.

I spent most of the next 20 years in Washington trying to put what I had learned at the service of government and business executives, eventually at the highest levels.

I got to use ADT models on major policy issues (for example, for the National Security Council, on whether the US should "tilt" toward Arab interests at the time of the 1970s oil embargo). My work attracted media attention, including public radio interviews and a magazine cover (Porter, 1987). However, my ADT analyses—at least quantitative versions—do not seem to have changed actual decisions much, but served more to provide reassurance and justification. This has also been the case with my private decision aiding.

In 1978, two behavioral decision theorists and I founded Decision Science Consortium, Inc. We used our own staff time on decision consulting and developing methodology, funded by defense, energy, environmental and other "mission" agencies. In parallel, we developed working relationships with academics in supporting disciplines, such as psychology and statistics, arranging joint projects funded by government research agencies.

Our decision technology strategy was basically the familiar engineering design practice, build-test-build-test. That is, we tried our methods out on live problems, identified deficiencies, cured them as best we could, applied any improvements to the next problem, and so on. We organized a short conference of leading decision researchers (including Simon and Tversky) and practicing decision aiders on the practical impact of existing decision research and found very little impact (Tolcott & Holt, 1988).

I interleaved my consulting with faculty appointments in management (Cambridge), statistics (University College London), organizational psychology (LSE), social and decision sciences (Carnegie-Mellon), information technology and public policy (both at George Mason). I mainly conferred and collaborated with colleagues with complementary backgrounds.

A standing ONR grant funded me to turn methodological consulting issues into researchable hypotheses. For example, I had observed that effective deciders (unlike most analysts) look at a difficult choice several different ways before making up their minds. I arranged for a psychologist and statistician to work with me on diagnosing and treating the plural evaluation problem (Lindley et al. 19769.) Over the next ten years, we published between us more than a dozen archival papers on derivative issues (Brown, 1992). Since then I use plural evaluation to make any non-trivial judgment.

My practical experience has led me to change radically the decision analysis procedures (but not the underlying logic), that I and others had originally advocated (Schlaifer, 1969; Brown et al., 1974; Brown, 2005b). However, old traditions die hard,[20] and the original ADT practices are still widely taught (Clemen & Reilly, 2000).

I have also changed the application focus of my work, which this essay reflects. Working with elected politicians convinced me that we have a helpless citizenry and since the late 1980s I have switched my main effort from helping institutional managers to private individuals.[21]

[20]My first published heresy took the intervention of a world-class decision theorist to overrule the strenuous objections of the journal's decision analysis area editor (Brown, 1978).

[21]I would be glad to hear from colleagues interested in pursuing issues raised in this essay, and especially in teaching ADT and other decision aid. Adapting decision theory to the needs of private deciders is a lonely perch and I would welcome supportive company before I fall off. My URL is http://mason.gmu.edu/~rbrown

https://doi.org/10.1017/S1930297500003041 Published online by Cambridge University Press