

JOSEPH G. JOHNSON and JEROME R. BUSEMEYER

MULTIPLE-STAGE DECISION-MAKING: THE EFFECT OF
PLANNING HORIZON LENGTH ON DYNAMIC
CONSISTENCY

ABSTRACT. Many decisions involve multiple stages of choices and events, and these decisions can be represented graphically as decision trees. Optimal decision strategies for decision trees are commonly determined by a backward induction analysis that demands adherence to three fundamental consistency principles: dynamic, consequential, and strategic. Previous research (Busemeyer et al. 2000, *J. Exp. Psychol. Gen.* 129, 530) found that decision-makers tend to exhibit violations of dynamic and strategic consistency at rates significantly higher than choice inconsistency across various levels of potential reward. The current research extends these findings under new conditions; specifically, it explores the extent to which these principles are violated as a function of the planning horizon length of the decision tree. Results from two experiments suggest that dynamic inconsistency increases as tree length increases; these results are explained within a dynamic approach–avoidance framework.

KEY WORDS: Approach-avoidance conflict, Dynamic consistency, Multi-stage decision-making

1. INTRODUCTION

Multiple-stage decisions refer to decision tasks that consist of a series of interdependent stages leading towards a final resolution. The decision-maker must decide at each stage what action to take next in order to optimize performance (usually utility). One can think of myriad examples of this sort: working towards a degree, troubleshooting, medical treatment, scheduling, budgeting, etc.

Decision trees are a useful means for representing and analyzing multiple-stage decision tasks (Figure 1), where decision nodes [X] indicate decision-maker choices, event nodes (Y) represent elements beyond control of the decision-maker, and terminal nodes ● represent possible final consequences (cf. Gass 1985, Chapter 23).

In this example, the pursuit of a graduate student towards a Ph.D. is represented as a multiple-stage decision tree. The first decision



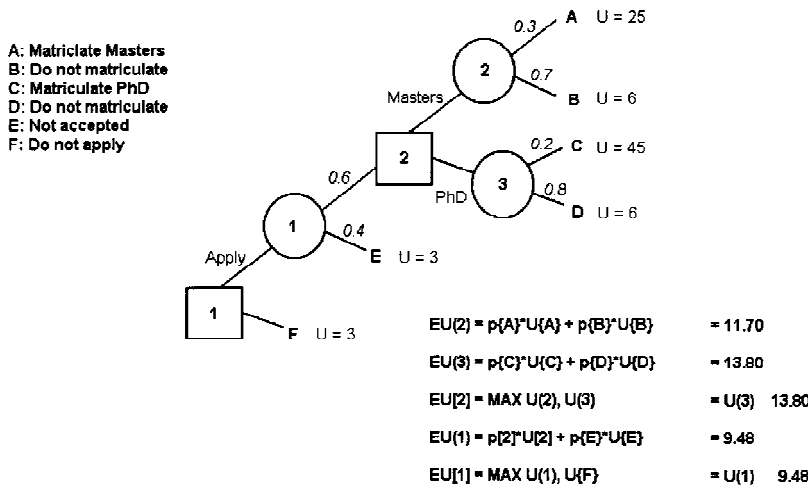


Figure 1. Example of a real-life situation represented as a decision tree and solved using the dynamic programming method.

node concerns whether or not to apply to graduate school, which leads to the event node of being accepted. If accepted, a second decision is required concerning which degree to pursue, leading to probabilistic event nodes dictating the decision-maker's chances of success for each. While optimal navigation of this rather small decision tree may not seem so overwhelming, one can imagine the difficulty in comprehending the different scenarios involved with larger trees, such as a foreign policy decision task. Based on elements of utility theory for single-stage gambles, backward induction (also known as dynamic programming) is an accepted method of selecting the optimal path of decision tree navigation (see, Bertsekas 1976; DeGroot 1970; and Raiffa 1968).

The method of backward induction is applied to the graduate student example at the bottom of Figure 1. First, the decision-maker assigns subjective utility values to all terminal nodes, reflecting his/her satisfaction with the final alternatives. Next, the decision-maker specifies the probabilities at the event nodes, to the best degree possible. For example, by using enrollment and matriculation rates one could assign meaningful values to the event nodes in Figure 1. Using backwards induction, one can then compute the optimal path for any given decision tree. As in SEU theories, the expected utility for event nodes (2) and (3) are determined by weighting the utility of each outcome (A, B for (2)) by its probability of occurrence (0.3,

0.7), resulting in $EU(2) = 11.70$ and $EU(3) = 13.80$. Then, the tree is effectively ‘pruned’ at the preceding decision node, removing the option with the lower expected utility, (2). Thus, whenever a decision-maker reaches node [2], s/he should always choose to proceed to event node (3) – effectively assigning the utility for (3) to [2]. This reasoning is continued down the tree, computing a probabilistic utility for (1) based on the probabilities of the terminal node {E} and the newly defined [2]. Finally, if the value $EU(1)$ exceeds $EU\{F\}$, the student should apply to graduate school.

2. CONSISTENCY PRINCIPLES

The backward induction analysis necessitates three fundamental consistency principles for maintenance of optimization (Hammond 1988; Machina 1989; Sarin and Wakker 1998). As long as these consistency principles hold, backward induction or dynamic programming can be applied and an optimal decision strategy ascertained. The first, called *dynamic consistency*, requires the planned decision strategy to be followed throughout the tree, otherwise defeating the purpose of the backward induction process. In the previous example, if the decision-maker uses dynamic programming to plan a decision strategy as explained above, but chooses to deviate from this plan by going for a Master’s degree when s/he actually reaches [2], this is a violation of dynamic consistency. *Consequential consistency* assumes the decision-maker will not be affected by past events, but that instead only future events and final consequences will be considered at any node. Violation of this principle would undermine the estimation of node probabilities and utilities, since these could change if they were a function of previous outcomes. If the student feels ‘lucky’ to have been accepted, and decides not to risk ‘looking bad’ by attempting a more rigorous course of study, s/he may opt for (2) at [2] as a result of redefining probabilities and/or utilities, violating consequential consistency. Finally, *strategic consistency* assumes that both dynamic and consequential consistencies are fulfilled.

Given the importance of these consistency principles, it is surprising that so little research has been done to empirically test them – especially considering the breadth of literature aimed at disproving




Figure 2. Example (two-stage) experimental decision tree with reward (R) = \$1.20, punishment (P) = 30 math problems, sure thing (S) = \$0.50, early payment (t) = \$0.04, gamble probability (p) = $0.50 = (1-p)$, and cost to advance one node equals one cent.

SEU tenets and assumptions. An initial study by Cubitt et al. (1998) found large violations of dynamic and strategic consistency. This finding was replicated and expanded by Busemeyer et al. (2000), by using the experimental decision tree in Figure 2.

The experimental decision tree (Figure 2) provides empirical tests of the three consistency principles under examination. In this tree, the numbered decision nodes represent a choice of either stopping and taking the monetary payment t , or paying an insignificant amount to try and work up the tree towards the final gamble, [D]. By choosing to continue, the decision-maker is faced with an event node with known probability of success, a , allowing continued navigation; and known probability $(1-a)$ of stopping navigation with no consequence (gain or loss). As long as the event nodes allow continued navigation, the decision-maker must repeatedly choose between continuing up the tree, or stopping early and taking t . If the decision-maker chooses (and is allowed by chance) to proceed to [D], a final decision is made between receiving a 'sure thing' payment of s , or choosing to instead take a final gamble (G). If chosen, this gamble contains a probability of 0.50 of receiving some monetary reward, R , and a probability of 0.50 of facing punishment,

P. Since the only meaningful decision node is [D] (due to the insignificance of *t*) only ‘pruning’ behavior at this point should be considered. That is, maintenance of consistency will be expressed in terms of the decision(s) regarding [D]: gamble vs. sure thing. Furthermore, participants make two different types of choices concerning [D]: a *planned* choice about [D] while in state [1], and a *final* choice made after navigating up to [D]. Also, the final stage [D] is presented in isolation, and participants make an *isolated* choice in this situation.

Using this experimental paradigm, consistency principles can be tested by comparing various pairs of participant choices. Dynamic consistency requires a planned decision to be fully carried out, and thus the planned choice regarding [D] should be equal to the final choice regarding [D]. Planning to take the gamble while at [1], then reversing strategy and deciding to take the sure thing once [D] is reached would be dynamically inconsistent. Consequential consistency requires a decision-maker to consider only successive nodes when making a choice. If a decision-maker has worked up the tree to [D], we need a measure to determine if the final choice made is independent of the previous nodes. By comparing this final choice to the isolated choice (which is the same decision in the absence of navigating the previous nodes) we obtain a test of consequential consistency. Specifically, the final choice should equal the isolated choice to maintain consequential consistency. Strategic consistency is upheld when both dynamic and consequential consistency are satisfied. If the planned choice equals the final choice (dynamic), and the final choice equals the isolated choice (consequential), then the planned choice will equal the isolated choice – which provides the test of strategic consistency. Each of these consistency measures can be compared with *choice inconsistency*, a baseline measure of a participant's tendency to vacillate in decision-making, determined by the proportion of decision reversals on the exact same (planned–planned; final–final; isolated–isolated) choice.

3. DECISION FIELD THEORY

The findings of Busemeyer et al. (2000) supported those of Cubitt et al. (1998), revealing empirical violations of dynamic and stra-




Figure 3. Illustration of the goal gradient hypothesis applied to different lengths. The horizontal axis represents [decision node], with progress to the left; the vertical axis represents valence strength. Note the increasing distance of [1] from [D] as the number of stages, n , increases.

tegic consistency as described above. Furthermore, Busemeyer et al. (2000) made predictions for the violation of dynamic consistency and manipulated the attractiveness of the gamble to test these predictions. These predictions were based on Decision Field Theory (DFT), a dynamic approach to human decision-making (Townsend and Busemeyer, 1989; Busemeyer and Townsend, 1993). The key concept of DFT as it relates to inconsistency in multiple-stage decision tasks is that of the goal-gradient hypothesis, originally developed in the approach-avoidance conflict theory of Lewin (1935) and Miller (1944). Figure 3 illustrates the concept as applied here.

Each decision is based upon an approach tendency, which is determined from potential gains; and an avoidance tendency, which is determined from potential losses. According to the goal gradient hypothesis, the strengths of the approach and avoidance tendencies decrease with increasing distance between the current state and the final decision. Furthermore, the gradient or slope of decrease may differ for approach and avoidance tendencies. Figure 3 illustrates the case where the avoidance gradient is steeper than the approach gradient. The horizontal axis in the figure represents distance (in nodes) from the final decision node [D], and the vertical axis represents the strengths of the approach and avoidance tendencies associated with particular gains (v_R) and losses (v_P). At stage [1], the decision-maker is far from the final consequences and the

approach tendency is greater than the avoidance tendency – when removed from the final consequences of an action, the potential gains are considered more than the potential losses. However, at node [D], the final consequences are impending, potential losses become more salient, and so the avoidance tendency exceeds the approach tendency.¹

Determining a course of action based on DFT requires computing valence differences (δ) between the approach and avoidance tendencies. Specifically, the valence difference is assumed to determine the probability of choosing the gamble (over the sure thing). According to Busemeyer et al. (2000), the valence difference is given by Equation (1):

$$\delta(n) = [(0.50) \cdot g_R(n) \cdot u(R) - (0.50) \cdot g_P(n) \cdot u(P)] \\ - g_R(n)u(S)],$$

In this equation, n is the number of stages separating the planned and final decision, $u(R)$ represents the attraction of the gain, $u(P)$ represents the aversion of the punishment, $u(S)$ represents the attraction of the sure thing, and $g_R(n)$ and $g_P(n)$ are the weights for gains and losses produced by the goal gradient. It follows that the probability of choosing the gamble systematically differs between a planned choice (with $n > 0$) and a final choice (with $n = 0$), depending on the payoffs R , S , and P . In fact, Busemeyer et al. (2000) manipulated these values and indeed found significant differences between planned and final choices. Thus, it seems that DFT provides a plausible explanation, and actually *predicts* dynamic inconsistency under certain conditions.

The manipulation of decision tree length allows us to further test the predictions of Decision Field Theory as an explanation for the violations in dynamic consistency. Recall that DFT posits an approach-avoidance gradient to explain decision reversals – as one gets closer to the final decision, the ‘approach’ and ‘avoidance’ characteristics become more salient and thus more heavily weighted. According to this line of reasoning, longer decision trees increase the distance between the initial and final decision stages, and thus produce greater differences in valences between planned and final choices (see Figure 3).

If the payoffs R , S , and P are held constant, and only the tree length is varied, then $\delta(n)$ will change systematically as a function of tree length, n . For small lengths (e.g., $n = 1$), the difference $\delta(1) - \delta(0)$ corresponding to planned and final choices is predicted to be small, but for large lengths (e.g., $n = 5$), the difference $\delta(5) - \delta(0)$ is predicted to be much larger. Therefore, DFT predicts higher dynamic inconsistency rates for long as compared to short trees.

The importance of examining the effects of this manipulation is also supported by previous research, which has shown that decision-makers may intentionally choose to employ alternate strategies for similar decision types of various lengths (Beach and Mitchell 1978; see Ford et al. 1989, for a review). While much of the existing work manipulating the number of stages has focused on, for example, the changes in subjective conditional probabilities (Savage 1954; Zimmer 1983), DFT instead makes formal *a priori* predictions for the effects of this manipulation on the underlying decision process. The present experiments were designed to test the DFT predictions regarding the effects of increasing the number of stages in a multiple-stage decision task.

4. EXPERIMENTS

Two experiments are reported that test the change in dynamic consistency rates as the planning horizon (length of decision trees) increases. Experiment 1 was designed to hold the expected value of the gamble at the initial decision node constant across lengths so that the gambles were equally attractive from the point of view of the initial decision node. However, this means mathematically that the internode probabilities must increase with length (the internode probability was $0.5^{1/n}$ so that path probability $(0.5^{1/n})^n = 0.5$ was constant).

This design raises a possible concern because much research has shown that participants do not follow normative rules in their use of conjunctive probabilities (e.g., Savage, 1954; Bar-Hillel, 1980; Gneezy, 1996). In particular, Gneezy (1996) showed how participants in multiple-stage tasks might anchor on individual node probabilities when determining the compound probability for the task (tree). Bar-Hillel (1982) has also shown anchoring and adjustment effects

in conjunctive events. Additional research that suggests participants may have systematic biases in dealing with probabilities include study of the immediate gratification effect – the inability to correctly compound choices (e.g. Rabin, 1998); as well as classic phenomena such as base rate neglect (Bar-Hilel, 1980) and the gambler’s fallacy (Jarvik, 1946). The motivation for Experiment 2 was to control for this concern by holding the internode probabilities constant (at 0.80) across lengths. The disadvantage of this design is that the expected values of the gambles at the initial decision node now vary depending on length.

A pilot study was run to determine if the experimental parameters adapted from Busemeyer et al. (2000) were properly chosen to allow for testability in this slightly different domain. For example, the pilot study showed that under Experiment 2 conditions, there were not enough subjects willing to progress towards the final node in a decision tree of length 5 (due to a very low expected value), and thus trees of this length were replaced with trees of length 4 in Experiment 2. Detailed differences between the two experiments will be highlighted as they are introduced.

4.1. *Method*

4.1.1. *Decision trees*

Trees were constructed of varying lengths to test the effect of planning horizon on dynamic consistency. They were similar schematically to the trees used in Busemeyer et al. (2000) and introduced above (Figure 2). In order to determine the desired effects, decision trees of varying lengths $0 \leq n \leq 5$ were constructed, where n represents the number of decision nodes not including [D], the final node. Terminal node values were chosen considering those used in Busemeyer et al. (2000) and the results of the pilot study. For all trees used in data analysis, terminal node values remained constant across *experiments*, *trials (lengths)*, and *participants* at reward $R = \$1.20$ payment; punishment $P = 30$ boring arithmetic problems; sure thing $S = \$0.50$ sure payment; and cost $t = \$0.04$.

Node probabilities (Appendix A) were determined for Experiment 1 such that $\Pr(\text{success at node } x \text{ of } n) = \Pr(x) = (0.50)^{1/n}$. This assignment was employed to ensure equal \Pr_n (reach final gamble) for all n . Trees of all lengths $0 \leq n \leq 5$ were included in Experiment

1. Each session contained four trees for each $n > 0$, two of which required only a planned decision for [D] while at [1], and two of which required only a final decision if node [D] was eventually reached. This produced a total of $(5) \text{ Length} \times (2) \text{ Choice Type} \times (2) \text{ Replication} = 20$ trials per session that entered into the data analysis. There were two trials with the $n = 0$ (isolated) decision; and a total of eight ‘filler’ trees were used that were never to be included in data analysis.²

For Experiment 2, node probabilities were held constant to test whether differences due to length in Experiment 1 (if any) could be contributed to different node probabilities. An intermediate value was chosen and seemed to work well in the pilot study, an internode probability equal to 0.80. Note this holds all experimental parameter values – R , P , S , t , p , and thus $(1-p)$ – constant across lengths. However, this results in a very low expected value for trees where length $n = 5$, and there were not enough subjects in the pilot study willing to progress towards the gamble in these trees. Therefore, trees of length $n = 5$ were replaced with trees of length $n = 4$ to keep the same number of overall trials (30) between experiments and still allow for enough power to test for differences between longer trees and shorter ones. Thus, each session contained four trees for each $0 < n < 4$ and eight trees for $n = 4$, where half of the trees for each n included a planned decision for [D] while at [1], and half required only a final decision if node [D] was reached. There were two trials with the $n = 0$ (isolated) decision; and the same eight ‘filler’ trees that were used in Experiment 1. Appendix A details the exact trees (and their order) for both experiments.

4.1.2. *Participants*

Participants in Experiment 1 were 79 undergraduate students who volunteered for course credit in addition to payment contingent upon their performance on the task (see below). Each participated in one session lasting approximately 1 h, and received payment at the conclusion of the experiment. These participants were randomly split into two groups for counterbalancing of presentation order, resulting in $N_{1A} = 40$ and $N_{1B} = 39$. Experiment 2 employed 76 undergraduate students under the same conditions, counterbalanced across two groups where $N_{2A} = 39$ and $N_{2B} = 37$ (due to subjects who did not

attempt the task, excluded before data analysis). The modal student (Experiment 2) was a female (64%) Caucasian (86%), with an average age of 19.86 years and math experience with at least some calculus (89%).

4.1.3. Procedure

Participants were initially given detailed computer instructions regarding navigation of decision trees and the associated probabilities and payoffs. Also, they were given 11 practice trials of varying R , P , and S , some of which ensured familiarity with the punishment (arithmetic problems) by containing terminal nodes with arithmetic problems only. Probabilities assigned to the practice trees were varied to develop participant experience with the computer spinner that randomly determined 'chance', the outcome of the event nodes. Computer presentation of the trees was sequential with immediate feedback, and completion of the arithmetic problems (punishment) was required before presentation of the next trial, although the monetary reward was delayed until the end of the session. The computer presentation procedure is adapted from Busemeyer et al. (2000). Presentation order of trees for each phase of each condition is available in Appendix A. The presentation order listed is for the first phase of the participant group A: the order is reversed for group B, and the second phase for each group is identical in type and order appended to the first phase. Note that trials used to perform tests of consistency are purposely separated in order to eliminate the dependency that could result from recently seeing decision trees of the same length.

A typical session in these experiments consisted of pre-experimental consent and a survey to record individual participant statistics such as gender, age, and math experience. After a brief overview by the experimenter, participants were seated at individual computer terminals and progressed through the experiment at their own pace. Detailed instructions were on-screen initially and remained in an open window throughout the experiment for reference. The instructions (Appendix C) included detail on concepts such as tree navigation, planned choices, and the determination of chance outcomes. Eleven practice trials followed the instructions to familiarize participants with the probabilistic spinner and the type of arithmetic

problems (addition of two numbers between 1 and 99) used as punishment. For example, some practice trials included *R*, *S*, and *P* that all involved arithmetic problems. Subjects were made aware that practice trials would not be included in calculation of final payment. Also, subjects were told the practice trials were over and subsequent trials would be included in final payment, although this was mentioned two trials prior to the presentation of trials used in data analysis to reduce possible serial position or transition (from practice) effects. In other words, subjects were paid for trials 10 and 11, which resembled the experimental trials but were not included in data analysis.

For each experimental trial, the tree was presented on-screen with the first node (current position) highlighted. For planned decision trials, a message box appeared at the start of the problem which asked the subject to make a choice regarding the final node [D] for that trial beforehand, and the subject was told that his or her choice would be executed automatically if they happened to reach the final node. For final decision trials, no message appeared, and subjects simply navigated down the tree and made a choice if and when they reached the final node. Subjects navigated through the tree by selecting between possible paths at each choice node and witnessing the spinner-determined outcome of chance nodes. If a subject chose to stop and take an early payment, then \$0.04 was added to their running total and the trial ended. If a subject reached the final node [D] and chose the sure thing (*S*), or chose the gamble and won, then \$0.50 or \$1.20, respectively, was added to their running total and the trial ended. If a subject reached the final node [D], chose the gamble, and lost, the subject had to complete 30 arithmetic problems to end the trial. Although subjects were aware of the outcome of each trial upon its conclusion, they were not explicitly provided with information such as their running total, choice history, or outcome history. All trials proceeded in this manner until the conclusion of the final trial, when subjects were informed online of their final payment, then debriefed and paid. Subjects were aware prior to the experiment that payment would immediately follow the experiment, but were not told the number of trials that would be presented (although they were told the experiment would last a little over an hour).

TABLE I
Summary of inconsistency rate data.

Measure	Experiment 1 (ICR)	Experiment 1 (N)	Experiment 2 (ICR)	Experiment 2 (N)
P ₁ – P ₂	0.35	395	0.31	380
F ₁ – F ₂	0.25	56	0.35	82
I ₁ – I ₂	0.25	79	0.46	52
P ₁ – F ₁	0.41	150	0.39	148
P ₁ – F ₂	0.44	135	0.44	154
P ₂ – F ₁	0.41	150	0.46	148
P ₂ – F ₂	0.38	135	0.40	154

Above data are pooled across all lengths.

Experiment 1: C_pICR = 0.33 (N = 530)

D_pICR = 0.41 (N = 570); $z = 3.80, p < 0.001$.

Experiment 2: C_pICR = 0.33 (N = 514)

D_pICR = 0.41 (N = 604); $z = 3.81, p < 0.001$.

4.2. Results

The results from the two experiments are reported in two parts. The first part presents a comparison of dynamic and choice inconsistency rates, pooled across the length manipulation. The second part presents an analysis of the effect of length on the inconsistency rates for each experiment.

4.2.1. Dynamic consistency

The first analysis concerns the comparison of dynamic versus choice inconsistency rates. Table I presents the inconsistency rates (ICRs) determined between all possible combinations of Choice Type and Phase, pooled across Length (see operational definitions of consistency measures; and Appendix B). The first column indicates the type of inconsistency rate, where P = Planned choice, F = Final choice, and I = Isolated choice; the subscripts represent which block was used in the calculation. The first three rows represent measures of choice inconsistency by comparing choices of the same type, and the last four rows contain all possible measures of dynamic inconsistency. For example, the fifth row represents the dynamic inconsistency rate obtained by comparing the Planned choice in the

TABLE II
Length effects on dynamic consistency.

Length (n)	Experiment 1 (D_n ICR)	Experiment 1 (N)	Experiment 2 (D_n ICR)	Experiment 2 (N)
0	0.25	79	0.32	76
1	0.35	116	0.39	214
2	0.40	142	0.39	152
3	0.42	108	0.50	102
4	0.38	88	0.43	136
5	0.50	116	n/a	n/a

Above data are pooled across two blocks.

first trial block with the Final choice in the second trial block. The remaining columns illustrate the observed inconsistency rate and associated N for Experiments 1 and 2, respectively. Thus, the first row in the table conveys that the frequency of change between the first and second Planned choice in Experiment 1 was 0.35, based on 395 observations, and for Experiment 2 was 0.31, based on 380 observations. Each measure was computed for each length, then pooled.

As can be seen in the table, the Pooled Choice ICR (C_p ICR, first three rows) of 0.33 ($N = 530$) was far below the Pooled Dynamic ICR (D_p ICR, last four rows) of 0.41 ($N = 570$) in Experiment 1. Similarly, the C_p ICR of 0.33 ($N = 514$) was far below the D_p ICR of 0.42 ($N = 604$) in Experiment 2. For such large N , a z -test is appropriate for comparing the difference (D_p ICR - C_p ICR), which produced a significant difference for each experiment: $z = 3.80$, $p < 0.001$, for Experiment 1, and $z = 3.81$, $p < 0.001$, for Experiment 2. Note that in Table I, the $n = 0$ case is necessarily treated as an Isolated case to maintain and test proper definitions of consistency. This analysis shows that dynamic inconsistency does reliably exceed choice inconsistency for both experiments; the striking agreement between the two experiments in terms of D_p ICR and C_p ICR will be discussed momentarily.

4.2.2. Length effects

Observing the significant dynamic inconsistency rates above leads to extending our analysis to determine the nature of this phenomenon. Specifically, the next analysis examines length effects on dynamic inconsistency rates. Table 2 presents the Dynamic ICR as a function of Length (D_n ICR), pooled across Phase, for each experiment. A glance at this data suggests that dynamic inconsistency does tend to increase with increasing length in both experiments. The null hypothesis, $H_0: E(D_1\text{ICR}) = E(D_5\text{ICR})$, was tested in Experiment 1 by computing a z -score; the null hypothesis was in turn rejected ($z = 2.31, p = 0.010$). In Table 2, the $n = 0$ is included for comparative purposes. A test comparing Short ($n = 1, 2$) versus Long ($n = 3, 4$) trees resulted in $\chi^2(2) = 5.97, p = 0.0504$, which suggests there may have been a more subtle treatment of Length in Experiment 2, but one that had an effect nonetheless.

A final analysis was performed to determine the differences, if any, of the effect of increasing length between the two experiments. The primary difference between the two experiments was the construction of internode probabilities, the motivation for which was given in the Introduction. A categorical data analysis was performed to determine if the effect of increasing length had different effects depending on whether expected value (Experiment 1) or internode probabilities (Experiment 2) were held constant across lengths. This analysis included tests for the main effects of Experiment, and Length, and the interaction effect of Length \times Experiment. The pooled Length effect in this analysis was significant, $\chi^2(4) = 11.81, p < 0.02$, which suggests the treatment of Length was significant overall. There was no significant difference between experiments ($\chi^2(1) = 2.52, p = 0.1124$), and the Length \times Experiment interaction also was not significant ($\chi^2(4) = 1.47, p = 0.8318$). The results of this analysis suggest that the length effect was robust across the two experiments, regardless of the internode probabilities.

5. DISCUSSION

The current research sought to achieve two distinct goals – to provide further empirical support for Decision Field Theory's predictions regarding dynamic inconsistency; and the second, to extend this ap-

proach by further specifying DFT predictions under a new set of variables. Specifically, planning horizon (as measured by number of stages in a multiple-stage decision task) was manipulated and predicted to be in a direct relationship with dynamic inconsistency rates: as the length of the planning horizon (number of stages) increases, the dynamic inconsistency rate should increase. The results of this study are discussed in three parts – the findings concerning dynamic consistency rates, followed by the moderating effects of tree length on these rates, and finally a note on the subjective use of probability in a multiple-stage decision task.

5.1. *Dynamic consistency*

First, consider the findings regarding dynamic consistency. Applying our empirically defined measures of choice and dynamic consistency provides more convincing support of the inconsistencies found by Busemeyer et al. (2000). Table II provides a summary of inconsistency rate data, pooled across length, for each experiment. These provide a measure (as calculated in Appendix B) of Pooled Choice ICR = 0.33 ($N = 530$), and Pooled Dynamic ICR = 0.41 ($N = 570$) for Experiment 1, and Pooled Choice ICR = 0.33 ($N = 514$), and Pooled Dynamic ICR = 0.42 ($N = 604$) for Experiment 2. The z -test for each difference ($D_p\text{ICR} - C_p\text{ICR}$) was found to be significant, supporting the conclusions concerning violations of dynamic (and thus strategic) consistency proposed by Busemeyer et al. (2000). When considering the strong similarity in these results of the two experiments, there seems to be a reliable ($N > 1000$) violation of dynamic consistency. The choice data provide some support for DFT predictions regarding the direction of change, as well. In both experiments, there was a decrease in preference for the gamble when comparing Planned to Final choices. That is, subjects planned to take the gamble more (0.515) than they decided to actually take the gamble when faced with the decision immediately (0.470). The DFT account makes just such predictions, since the positive characteristics evaluated at a distance from the goal would tend to promote a choice for taking the gamble, whereas the increased influence of the negative characteristics of the gamble when directly faced with the consequences leads subjects to prefer the sure thing. On a side note, the Pooled Choice Inconsistency Rate ($C_p\text{ICR} = 0.33$) computed

from over 1000 observations suggests that human decision-makers, when faced with the same decision (or decision process) at different points in time, may change their minds one-third of the time. Admittedly, this generalization is a precarious one, but is indeed interesting.

While many models of decision-making require adherence to the principle of dynamic consistency (Hammond, 1988; Machina, 1989; Sarin and Wakker, 1999), Decision Field Theory actually predicts the violations that occurred in the present set of experiments. The characteristic of DFT that is responsible for such predictions is the approach-avoidance gradient hypothesis. This property assumes that the decision-maker is more responsive to the positive aspects of an outcome when distant from the realization of the outcome, but becomes increasingly sensitive to the negative aspects as the point of realization approaches. The theoretical approach and avoidance curves are manifest in the decision-maker's value weighting function in determining a valence difference between the gamble and the sure thing (Equation (1)). It is imperative that we clarify the changes that occur with progress towards the 'goal' or realization of the outcome. In the present task, objective probabilities and terminal node values were held constant during each trial as subjects progressed through each decision tree. Therefore, the dynamic inconsistency observed cannot be attributed to a change in the task characteristics such as the receipt of new information, or a change in objective probabilities, that spurs reevaluation of $u(R)$, $u(S)$, or $u(P)$. Instead, our framework explicitly incorporates 'proximity to the goal' into the weighting function via calculation of the valence difference. This component provides a dynamic quality that obviously outperforms any static account of the phenomena (such as context-independent evaluation of Planned and Final choices neglecting dynamic proximity, which would predict identical evaluation of Planned and Final choices and maintenance of dynamic consistency). The fact that this dynamic element allows DFT to accurately predict the observed effects provides support for the use of dynamic models in general.

5.2. *Planning horizon*

In addition to providing an accurate model for general violations

of dynamic consistency, the current research extends the power of DFT predictions in this realm by specifying the effect of lengthening the planning horizon. The main purpose of the current experiments was to vary the length of the planning horizon to extrapolate DFT predictions to the case of n stages. Figure 3 illustrates the rationale of the following argument from a DFT standpoint. Increasing the number of stages effectively moves the initial decision node, [1], further from the final decision node, [D]. The difference between the position of [1] when $n = 1$ and the position of [1] when $n = 5$ leads to larger δ_{plan} , and thus (from Equations (2) and (3)) a larger magnitude difference between $\delta_{\text{final}} - \delta_{\text{plan}}$ since δ_{final} does not vary with n , contributing to greater dynamic inconsistency. The extent of this effect is mediated by the individual approach-avoidance curves associated with a particular decision-maker's navigation of the multiple-stage task. The tendency for this change to occur increases as n increases, because of the increased distance $[n] - [D]$, leading to larger δ_{plan} . In other words, the larger the planned valence difference becomes (due to increasing n), the decision-maker is more prone to reversing the planned decision, resulting in a greater degree of dynamic inconsistency at [D]. Since the planned valence difference for lower n is (relatively) smaller such that it is closer to δ_{final} , the decision-maker is more apt to preserve consistency (Planned Choice = Final Choice). Note in Figure 3 that if [1] is sufficiently close to [D], there may be no preference reversal because the initial starting point may already dictate δ_{plan} has the same sign as δ_{final} such that the associated choices are the same.

The results of the current experiments, Table III, support these theoretical predictions since $D_n\text{ICR}$ does seem to increase as a function of length. The non-monotonicity caused mainly by the discrepant $n = 4$ case in Experiment 1 and $n = 3$ case in Experiment 2 may be a result of the relatively lower N associated with these lengths, although the general trend is still quite clear. Also, the categorical data analysis based on these $D_n\text{ICR}$ pooled across both experiments further supports that Length does indeed have a significant effect on dynamic inconsistency. The primary conclusion is that DICR increases as a function of the length of the decision, defined by number of stages n . Other approaches fail to account for dynamic inconsistency at all, let alone make predictions as to the behavior of DICRs.

The application of Decision Field Theory, however, to the study of inconsistencies yields many fruitful results. DFT explains dynamic inconsistency in terms of competing valence differences δ_{plan} and δ_{final} . The two should be equal to maintain dynamic consistency, and DFT uses an approach-avoidance framework to explain and predict why they are not (Busemeyer et al., 2000). Furthermore, due to increasing δ_{plan} as n increases based on these approach-avoidance gradients, and δ_{final} being independent of n , the difference between the two increases and thus dynamic inconsistency increases with n .

5.3. *Subjective probability*

The final contribution of the present research involves the manner in which decision-makers use probability in multiple-stage decision tasks. The purpose of Experiment 2 was to ensure that increased dynamic inconsistency as a function of length was not due solely to the increased internode probabilities associated with longer decision trees in Experiment 1. Gneezy (1996) suggests that in multiple-stage tasks, participants may use individual node probabilities as an anchor when computing compound probabilities, insufficiently adjusting this anchor in determining $\text{Pr}(\text{reach final gamble})$. Recall that in Experiment 1 the probability of success for each intermediate node was increased as the length increased to keep the compound product of probabilities (i.e., the probability of reaching the final node) constant across n . Participants in Experiment 1 may not have realized that $\text{Pr}(\text{reach node [D]})$ was equal across n , and may have instead evaluated the increasing node probabilities (with length n) as relatively better gambles in planning, per Gneezy (1996). Perhaps, however, after progressing through the tree participants ‘felt’ the effect of compounding the intermediate probabilities, or excluded them from the final choice process altogether. This could result in similar evaluations at the final choice for all lengths (since the expected values were indeed equal), whereas the initial evaluations of longer trees would be inflated, contributing to greater dynamic inconsistency for those trees.

The results of Experiment 2 refute this alternative claim, and lend further support for the DFT account of dynamic inconsistency. Holding node probabilities constant across lengths in Experiment 2 provides a common basis for comparison across lengths. Although

different concerns are raised if node probabilities are constant (resulting in different expected value for trees of different length), the fact that Planned vs. Final choices were compared within a given length should control for this discrepancy. Under these conditions, there was still generally an increase in dynamic inconsistency with length. Also, the statistical test including the pooled data from both experiments suggests a significant length effect. This test furthermore suggests similar behavior in response to length in both experimental conditions, due to the high p -value associated with the Length \times Experiment interaction effect. A final verification comes from the pooled Choice ICRs and Dynamic ICRs, which were virtually the same in both experiments. Had subjects anchored on the individual node probabilities in Experiment 2, then the evaluation of $\Pr(\text{reach node [D]})$ should be similar for all lengths, leading to similar choices for all lengths – this was not the case. In sum, it seems that whether participants experienced different tree lengths with identical expected values (Experiment 1), or identical node probabilities (Experiment 2), they responded by increasing dynamic inconsistency with increasing length.

5.4. *Extensions and implications*

Other explanations may contribute to a refinement of the DFT-based explanations provided herein, such as the effort vs. reward tradeoff in decision tasks (Payne et al., 1993), and specifically in dynamic decision tasks (Kertsholt, 1996). The DFT framework applied to dynamic inconsistency can also be considered in relation to other phenomena such as overconfidence and hindsight bias. DFT predictions about a change from planned to final choices could be responsible for overconfidence effects, where a decision-maker unreliably predicts future states or choices. Also, hindsight bias could arise due to evaluation of ‘reverse-planned’ choices as final choices, which differ from the original planned choice. To clarify, if a decision-maker makes a planned choice, or prediction of future action, but is dynamically inconsistent when the time comes to perform the action, this could resemble overconfidence in the planned decision at the time it was made. Conversely, if the decision-maker makes a (dynamically inconsistent) final choice regarding an action, s/he may exhibit hindsight bias by incorrectly identifying the planned

choice, because it is being 'reverse-evaluated' in the goal state where the crossing of approach and avoidance curves has already occurred.

There are many other issues to consider in this line of research, some of which have been examined elsewhere. The Busemeyer et al. (2000) study varied reward and punishment values in examining dynamic inconsistency, and provides a solid theoretical treatment of the DFT explanations in relation to other accounts. Also, Barkan and Busemeyer (1999) consider the effects of the reference point on dynamic inconsistency by using a two-gamble paradigm. If subjects won the first gamble, they were more inclined to forgo the second gamble because the gain from the first positively shifts the reference point. Losing the first gamble has the opposite effect, negatively shifting the reference point and increasing preference for the second gamble. If the planned choice regarding the second gamble is made prior to the outcome of the first, then the planned choice would be based on a neutral reference point and dynamic inconsistency could be explained in terms of the shift in the reference point for evaluating the second gamble (final choice). Other possibilities for future research include determining the interaction between elements such as decision importance (payoff matrix) and the planning horizon on dynamic inconsistency. Would decision-makers be less prone to changing their minds for decisions of great importance and, if so, how would this effect be modeled? Furthermore, what would the effect be of specifying explicit benefits for maintaining dynamic consistency?

The practical implications regarding dynamic inconsistency in human decision-making are widespread. In one respect, maintaining dynamic consistency seems beneficial in that it allows for derivation of 'rational' preferences across time. On the other hand, decision flexibility over time is not necessarily a harmful or 'irrational' trait – it could instead have adaptive benefits. In consumer-oriented tasks such as budgeting resources (i.e., money, time), there could be dire consequences for violations of dynamic consistency if final expenditures exceed planned expenditures. However, strict maintenance of dynamic consistency may result in suboptimal performance when gauged according to preferences calculated at the final state. Furthermore, in affect-rich or heavily emotional decisions, dynamic inconsistency could occur even after very small changes

in time. In such decisions, dynamic inconsistency as formulated in the present analysis could reduce to typically 'impulsive' behavior, and as such could be modeled explicitly. Even more so than in experimental tasks, in realistic situations the environment itself may dictate the 'appropriate' degree of dynamic inconsistency that should be allowed.

6. CONCLUSION

It is apparent from the current research that Decision Field Theory is a powerful tool in predicting human decision-making behavior in interdependent multiple-stage decision tasks. It is shown that violations of dynamic consistency in such tasks may not be irrational behavior, but rather adherence to an approach postulated by DFT and mediated by the length of the planning horizon. However, since questions still remain concerning participant behavior in multiple-stage decision tasks, further research is necessary to fully justify the DFT framework as applied to these tasks.

The study of human decision-making behavior has come a long way since von Neumann and Morgenstern's (1947) pioneering research. Yet, only recently have researchers begun to stray from various SEU approaches (e.g., Prospect Theory) in order to try and model decision-making behavior, as opposed to simply enumerating the inconsistencies and violations of SEU postulates. To this end, Busemeyer and Townsend (1993) provide a clear argument for the adoption of dynamic-probabilistic models that more closely represent real decisions than other (i.e., deterministic-static) approaches. DFT is one such dynamic-probabilistic model that has been supported by the present study and others. Only further rigorous testing of different DFT predictions across more variables and task environments can ultimately build support for DFT as a useful and parsimonious replacement to SEU theories as the standard in decision-making research.

NOTES

1. Note that 'length of the decision process', in this context, refers to the length in terms of the number of stages in the multiple-stage decision task, as opposed

to the decision-maker's deliberation time. While the two phrases can often be regarded as synonymous, necessary distinction is made here.

2. The 'filler' trees are detailed in Appendix A. The *R*, *P*, and *S* in these trees were manipulated in order to provide variety for the participants, hence their exclusion from data analysis. Another reason for the use of these trees was to temporally space experimental trees to reduce the possible dependence in observations.

ACKNOWLEDGEMENTS

This research was supported by National Institute of Mental Health (NIMH) Perception and Cognition Grant R01 MH55680; National Science Foundation (NSF) Methodology, Measurement, and Statistics Grant SES-0083511; and NIMH Modeling in Cognition Grant MH19879-09.

APPENDIX A: PRESENTATION ORDER AND TYPE OF DECISION TREES

PARAMETER VALUES FOR ALL TREES USED IN DATA ANALYSIS:

(R) Gamble payoff =	120 cents
(P) Punishment =	30 arithmetic problems
(S) Sure thing =	50 cents

DETERMINING NODE PROBABILITIES BASED ON TREE LENGTH:

Experiment 1:

Length (not counting gamble)	Success probability per node
1	0.50
2	0.71
3	0.79
4	0.84
5	0.87

Experiment 2:

Length (not counting gamble)	Success probability per node
1	0.80
2	0.80
3	0.80
4	0.80

TREE ORDER FOR TRIALS USED IN DATA ANALYSIS
(GROUP A):

Experiment 1:

<i>Trial</i>	<i>Type</i>
1 – 11	PRACTICE
12	Planned, 1 node
13	<i>Filler, 2 nodes, ($R = 120 / P = 50$, $S = 20$)</i>
14	Final, 5 nodes
15	Planned, 2 nodes
16	<i>Filler, 3 nodes, 120/50, 30</i>
17	Final, 1 node
18	Planned, 3 nodes
19	Isolated
20	Final, 2 nodes
21	Planned, 4 nodes
22	<i>Filler, 2 nodes, 120/10, 40</i>
23	Final, 3 nodes
24	Planned, 5 nodes
25	<i>Filler, 3 nodes, 120/10, 50</i>
26	Final, 4 nodes

Experiment 2:

<i>Trial</i>	<i>Type</i>
1–11	PRACTICE
12	Planned, 1 node
13	<i>Filler, 2 nodes, 120/50, 20</i>

14	Final, 4 nodes
15	Planned, 2 nodes
16	<i>Filler, 3 nodes, 120/50, 30</i>
17	Final, 1 node
18	Planned, 3 nodes
19	Isolated
20	Final, 2 nodes
21	Planned, 4 nodes
22	<i>Filler, 2 nodes, 120/10, 40</i>
23	Final, 3 nodes
24	Planned, 4 nodes
25	<i>Filler, 3 nodes, 120/10, 50</i>
26	Final, 4 nodes

Trials 27–41 are the exact same type/order as listed above for both Experiments. Trees for Group (B) are the exact reverse of those listed: e.g., Trial 12 (Group A) = Trial 41 (Group B)

APPENDIX B: MATHEMATICAL FORMULATION OF CONSISTENCY MEASURES

Let r_n^{UV} denote the number of pooled decision reversals between two decision types U and V, each representing a single phase $\{a, b\}$ and type $\{P, F\}$, for a given length n . The *choice inconsistency rate* (C_pICR) used to test for violation of dynamic consistency in Table I pooled the choice inconsistency rates C_PICR , C_FICR , and C_0ICR :

$$P_1 - P_2 = \frac{r_1^{PP} + r_2^{PP} + r_3^{PP} + r_4^{PP} + r_5^{PP}}{N_1 + N_2 + N_3 + N_4 + N_5} = C_PICR \quad (1)$$

$$F_1 - F_2 = \frac{r_1^{FF} + r_2^{FF} + r_3^{FF} + r_4^{FF} + r_5^{FF}}{N_1 + N_2 + N_3 + N_4 + N_5} = C_FICR \quad (2)$$

$$I_1 - I_2 = \frac{r_0}{N_0} = C_0ICR \quad (3)$$

In a similar manner, the four measures of type-phase dynamic inconsistency (DICR) were obtained. Using summation notation:

$$DICR = U_a - V_b = \sum_{n=1}^5 \frac{r_n^{U_a V_b}}{N_n} \quad (4)$$

P₁ – F₁
 Averaging across the four distinct pairs: P₁ – F₂
 P₂ – F₁
 P₂ – F₂

yields the *dynamic inconsistency rate* (D_pICR) reported in Table I.

The length-specific dynamic inconsistency rates (D_nICR) provided in Table II were determined by considering only the appropriate n terms in computing DICR, then averaging across pairs (except for $n = 0$) as in computing the DICR – effectively removing the summation from Equation (4) of the above process.

APPENDIX C: INSTRUCTIONS TO PARTICIPANTS

C.1. Introduction

First read over all of the instructions in this window. To read all of the instructions, you will need to scroll down by clicking the down-arrow with the mouse, and you can scroll up by clicking the up-arrow with the mouse. During this first reading, you will not be able to carry out the operations described by the instructions. When you have finished reading the instructions the first time, click the start button with the mouse. You will see these instructions again for reference, but during the second reading, you can carry out the operations indicated by the instructions.

C.2. Decision trees

The basic unit of the experiment is a decision problem. A decision problem is represented by what we call a ‘tree’ or more clearly a ‘decision tree’. Points where you will make a decision on the tree are represented by squares and the choices at each such nodes are

marked by outgoing lines called branches. At each decision node you make a choice among the branches of that node. A branch may have a number or a label at its end or a circle attached to it. In the first case, if you choose the branch with the number attached to it you will be paid the amount represented by that number (negative numbers mean you pay us). Otherwise, that is if a label is attached, you will be asked to perform the number, indicated with the label, of addition problems. On the other branches, that have circles attached to them, the circle marks a 'chance' node. Here, the choice of which branch is taken next is determined by the spinner, according to the frequencies (in percents) written on its outgoing branches.

C.3. Practice and real problems

The first 11 problems are practice decision problems. You do not really win or lose any money (but you may perform the addition problems) on these decision problems. These problems are designed to familiarize you fully with the task. Later you will face the real problems during which you actually win or lose money. At that time you will be fully comfortable with, and knowledgeable of the task involved. The number and type of real problems is set in advance and do not depend on what you do or how quickly you go. Take your time.

C.4. Some details

C.4.1. Gambles

A gamble is an outcome which depends on chance. The possible outcomes of a gamble are marked by the branches coming out of its circle. There are two types of gambles: 'purely monetary' gambles and 'hybrid' gambles. The purely monetary gambles have branches labeled by the number of CENTS you gain (positive number) or lose (negative number) if the chance event falls on them. The branches of hybrid gambles are marked by the number of CENTS you gain or the number of addition (of a two-digit and a one-digit number) problems you perform if the chance event falls on them. These addition problems are easy but tedious. The computer will present you with the problems one at a time. It will wait on you to answer each one correctly. Nothing will proceed any further until all problems are answered correctly. Branches are colored and also labeled by the

frequency (in percent) they are chosen by the spinner. The spinner's face is divided into colored sectors with size proportional to the frequency. The colored sectors correspond to the colored branches.

C.4.2. Early decision

On many (*not all*) decision problems, the first thing you have to do is to tell us what you want to choose if you arrive at the *last stage* at the end of the tree (indicated by the square node with a green circle around it). You may never get there or you may choose not to go there, but if you do, tell us what you want to choose at that point. If you do, in fact, eventually come to this node the computer will carry out your wish exactly as you had indicated and you will not need to make this choice again.

This last stage consists of a choice between: (1) Going up and choosing the gamble marked by a circle, or

(2) Going down and choosing the sure thing. The amount (IN CENTS) that you receive is indicated by the amount shown at the end of the branch. This is guaranteed money you gain.

C.4.3. Making decisions

After you make the early decision (if required) you begin making choices starting from the first stage (the flashing square). You have a choice between: (1) Going up and continuing on to the second stage. You have to pay one cent to continue in the game. If you choose to go up, then the next stage is selected by the outcome of the spinner. The spinner may land on green and allow you to continue to the next decision stage. Or the spinner may land on red and force you to stop with the amount shown at the end of the bottom branch extending out of the circle. (2) Going down, which stops the game, and you receive the amount indicated at the bottom of the first branch.

Choose to go up on this practice tree by clicking the first circle node and paying the one cent for this. At this point, the next stage is determined by the spinner.

For this practice problem, continue making choices going up until either you reach the final decision stage, or you are stopped by the spinner and the decision task ends.

If you have any questions at this point, read over these instructions again, and if you still have trouble, ask the experimenter for help.

REFERENCES

- Bar-Hillel, M. (1980), The base-rate fallacy in probability judgments, *Acta Psychologica* 44, 211–233.
- Barkan, R. and Busemeyer, J. R. (1999), Changing plans: Dynamic inconsistency and the effect of experience on the reference point, *Psychological Bulletin and Review* 6, 547–555.
- Beach, L. R. and Mitchell, T. R. (1978), A contingency model for the selection of decision strategies, *Academy of Management Review* 3, 439–449.
- Becker, G. M. and McClintock, C. M. (1967), Value: Behavioral decision theory. *Annual Review of Psychology* 18, 239–286.
- Bertsekas, D. P. (1976), *Dynamic Programming and Stochastic Control*. New York: Academic Press.
- Busemeyer, J. R. and Townsend, J. T. (1993), Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment, *Psychological Review* 100, 432–459.
- Busemeyer, J. R., Weg, E., Barkan, R., Li, X. and Ma, Z. (2000), Dynamic and consequential consistency of choices between paths of decision trees, *Journal of Experimental Psychology, General* 129, 530–545.
- Camerer, C. F. and Ho, T. H. (1994), Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty* 8(2), 167–196.
- Cubitt, R. P., Starmer, C. and Sugden, R. (1998), Dynamic choice and the common ratio effect: An experimental investigation, *The Economic Journal* 108, 1362–1380.
- DeGroot, M. H. (1970), *Optimal Statistical Decisions*. New York: McGraw-Hill.
- Fishburn, P. C. (1970), *Utility Theory for Decision Making*. New York: John Wiley and Sons.
- Ford, J. K., Schmitt, N., Schechtman, S. L., Hults, B. M. and Doherty, M. L. (1989), Process tracing methods: Contributions, problems, and neglected research questions, *Organizational Behavior and Human Decision Processes* 43, 75–117.
- Gass, S. (1985), *Decision Making, Models and Algorithms: A First Course*. New York: John Wiley and Sons.
- Gneezy, U. (1996), Probability judgments in multi-stage problems: Experimental evidence of systematic biases, *Acta Psychologica* 93, 59–68.
- Hammond, P. J. (1988), Consequentialist foundations for expected utility, *Theory and Decision* 25, 25–78.
- Kahneman, D. and Tversky, A. (1979), Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.
- Kertsholt, J. H. (1996), The effect of information costs on strategy selection in dynamic decision tasks, *Acta Psychologica* 94, 273–290.
- Lewin, K. (1935), *A Dynamic Theory of Personality*. New York: McGraw-Hill.
- Loomes, G., Starmer, C. and Sugden, R. (1991), Observing violations of transitivity by experimental methods, *Econometrica* 59(2): 425–439.
- Luce, R. D. (2000), *Utility of Gains and Losses*. New York: Erlbaum.

- Luce, R. D., Krantz, D. H., Suppes, P. and Tversky, A. (1990). *Foundations of Measurement*, Vol. 3: *Representation, Axiomatization, and Invariance*. San Diego: Academic Press.
- Machina, M. J. (1989), Dynamic consistency and non-expected utility models of choice under uncertainty, *Journal of Economic Literature* 27, 1622–1668.
- Miller, N. E. (1944), Experimental studies of conflict. In: J. McV. Hunt (ed.), *Personality and the Behavior Disorders*, Vol. 1 (pp. 431–465). New York: The Ronald Press.
- Payne, J.W., Bettman, J. R. and Johnson, E. J. (1993), *The Adaptive Decision Maker*. New York: Cambridge University Press.
- Raiffa, H. (1968), *Decision Analysis*. London: Addison-Wesley.
- Rapaport, A. (1975), Research paradigms for studying dynamic decision behavior. In: D. Wendt and C. Vlek (eds.), *Utility, Probability, and Human Decision Making*, Vol. 11. Dordrecht, The Netherlands: Reidel.
- Sarin, R. and Wakker, P. (1998), Dynamic choice and non-expected utility, *Journal of Risk and Uncertainty* 17, 87–119.
- Savage, L. J. (1954), *The Foundations of Statistics*. New York: John Wiley and Sons.
- Taylor, S. E. (1991), Asymmetrical effects of positive and negative events: The mobilization-minimization hypothesis, *Psychological Bulletin* 110, 67–85.
- Thaler, R. H. and Johnson, E. J. (1990), Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice, *Management Science* 36, 643–660.
- Townsend, J. T. and Busemeyer, J. R. (1989), Approach-avoidance: Return to dynamic decision behavior. In: C. Izawa (ed.), *Current Issues in Cognitive Processes: Tulane Flowerree Symposium on Cognition*. Hillsdale, NJ: Erlbaum.
- Tversky, A. and Kahneman, D. (1987), Rational choice and the framing of decisions. In: R. M. Hogarth (ed.), *Rational Choice: The Contrast between Economics and Psychology*. (pp. 67–94). Chicago: University of Chicago Press.
- von Neumann, J. and Morgenstern, O. (1947), *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Weber, E. U. (1994), From subjective probabilities to decision weights: The effect of asymmetric loss function on the evaluation of uncertain outcomes and events, *Psychological Bulletin* 115, 228–242.
- Weber, M. and Camerer, C. (1987), Recent developments in modeling preferences under risk, *OR Spectrum* 9, 129–151.
- Zimmer, A. C. (1983), Verbal vs. numerical processing of subjective probabilities. In R. W. Scholz (ed.), *Decision Making under Uncertainty*. Amsterdam: Elsevier.

Address for correspondence: Joe Johnson, Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany
Phone: +49-30-82406350; E-mail: johnson@mpib-berlin.mpg.de