American Psychological Association

Decision


2325-9965/20/$12.00

2020, Vol. 7, No. 1, 20–42
http://dx.doi.org/10.1037/dec0000113

# Decision Field Theory-Planning: A Cognitive Model of Planning on the Fly in Multistage Decision Making

Jared M. Hotaling
University of New South Wales

The world is full of complex environments in which individuals must plan a series of choices to obtain some desired outcome. In these situations, entire sequences of events, including one's future decisions, should be considered before taking an action. Backward induction provides a normative strategy for planning, in which one works backward, deterministically, from the end of a scenario. However, this model often fails to account for human behavior. This article proposes an alternative account, decision field theory–planning (DFT-P), in which individuals plan future choices on the fly through repeated forward-looking mental simulations. As they imagine the possible outcomes of their actions, decision makers simulate their future choices moment to moment. A key prediction of DFT-P is that payoff variability produces noisy simulations and reduces sensitivity to value differences. In two experiments, a robust multistage payoff variability effect was found, with preferences becoming weaker as variability increased. A formal comparison showed that DFT-P provided a good account of people's behavior, while a heuristic model and a flexible version of the backward induction model did not. These results confirm a fundamental prediction of DFT-P, and demonstrate its utility as a tool for understanding how people plan future choices and allocate cognitive resources in multistage decision making.

*Keywords:* dynamic decision making, planning, cognitive models, decision trees, DFT



A risky decision problem present an individual with a choice between uncertain alternatives. Although most studies of risky decision making have used single-stage gambles, real world decisions often involve planning across sequences of interdependent events and actions. For instance, imagine that an office worker, Emma, is faced with the multistage decision problem shown in Figure 1. Emma is unhappy with her current level of pay, and is trying to decide whether or not to ask for a raise. However, before she makes a decision she must consider the consequences of each action. This is simple if she takes the safe option and chooses not to ask for a raise; she will experience no change and will continue earning her current salary for certain. If she instead chooses the risky option, things are more complicated. Now Emma's boss has an impact on the eventual outcome. On the one hand, if he says "yes" and accepts her request, he will offer her two additional options to choose from: she can receive stock in the company, or she can receive increased wages. Her choice will determine her final outcome, though some uncertainty remains. For instance, if Emma chooses company stock, the end result will depend on factors influencing stock prices. On the other hand, if Emma's boss says "no" and denies her a raise, Emma will be faced with another choice between two uncertain alternatives: she can quit her job and begin searching for a new one, or she can continue with her current job and deal with the fallout from her failed request. Again, her choice here constrains, but does

This article was published Online First June 20, 2019.

Thank you to Katherine Harlan and Daniel DeCaro, who provided valuable feedback on earlier drafts.

Correspondence concerning this article should be addressed to Jared M. Hotaling, School of Psychology, University of New South Wales, Sydney, New South Wales 2052, Australia. E-mail: jaredhotaling@gmail.com

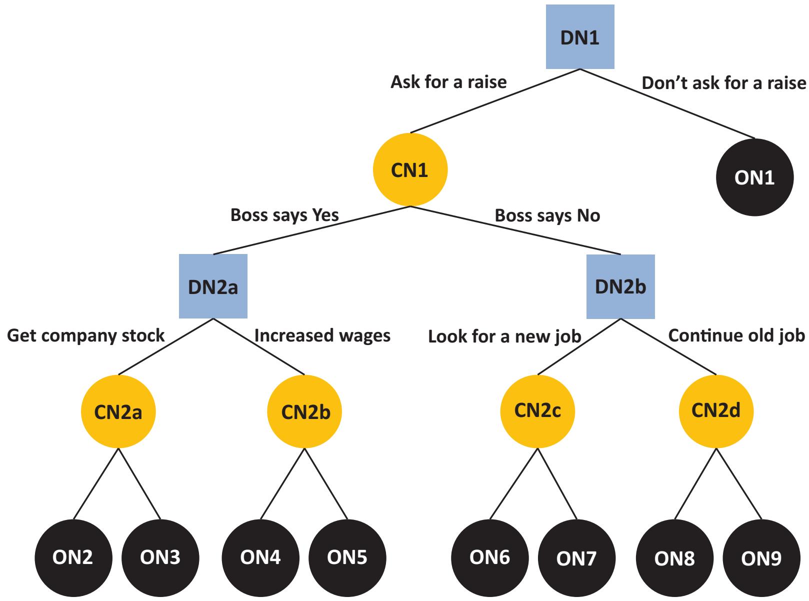


*Figure 1.* A multistage decision scenario, represented as a decision tree. Decision nodes (DNs) represent points where decision maker chooses an action. Chance nodes (CNs) represent points were an uncertain external event occurs. Outcome nodes (ONs) represent final outcomes or payoffs. See the online article for the color version of this figure.

not determine her final outcome. For example, if Emma chooses to look for a new job, the consequences of this action will also depend of the uncertain state of the job market she enters.

Figure 1 uses a decision tree to more formally represent the scenario described above. The tree consists of three kinds of nodes: decision nodes (DNs), chance nodes (CNs), and outcome nodes (ONs). DNs represent points where Emma chooses an action. CNs denote random events that occur in the world, which Emma cannot control. Finally, each branch of the tree terminates with an ON representing a final consequence. Given this decision tree formalization, Emma's goal is now to make choices such that she reaches an attractive ON and avoids the unattractive ONs. To solve her multistage decision problem, Emma must first decide at DN1 whether it is better to keep her current salary, the sure outcome ON1, or take a gamble by asking her boss for a raise and moving to CN1. If she chooses to gamble (represented as moving left in the tree from DN1 to CN1), she is unsure how her boss will react. This uncertainty can be represented as the probability of moving left or right at CN1. The value of this probability would therefore correspond to Emma's belief regarding how likely her boss is to say "yes" or "no." Each possibility branches to an additional DN, either DN2a or DN2b. The former represents Emma's subsequent decision between company stock and increased wages if her boss says "yes." The latter represents her choice between staying at her current job or quitting to look for a new job if her boss says "no." Because Emma cannot be sure of the exact outcome of each of these final actions, CN2a-d represent her uncertainty about how external factors will ultimately influence which ON she reaches, and therefore which outcome she experiences. For example, the probability of moving left or right at CN2a might correspond to Emma's belief about the likelihood that company stock will increase or decrease in value. This decision tree formalism allows for the quantitative study of multistage decision making, and this article proposes a new theory of how people solve these sorts of complex multistage decision problems.

## Thinking Backward

Before introducing the new theory, let us consider a more traditional model for how Emma should approach her decision problem, backward induction. According to the typical version of backward induction, choices are deterministic, in that the decision maker is assumed to always choose the option with the highest utility. Importantly, the theory construes multistage decision making as occurring across two distinct phases. The first involves collapsing the features of each node down to a single utility value. These values are compared at each future DN, allowing the decision maker to explicitly commit to a series of future choices. In the second phase the decision maker implements her planned choices at each DN. In later sections we will contrast this two-phase account with a new model of multistage decision making that posits a single, online, forward-looking process.

To better understand the predictions of the backward induction model, let us return to the problem depicted in Figure 1. Imagine that Emma first inspects the ONs and identifies a few particularly attractive outcomes, say, ON1, ON3, and ON8. Before she makes her first move, Emma must be careful. If she is to reach one of the attractive nodes, she should engage in some planning before making a choice at DN1. According to backward induction, Emma should do this by first considering the likelihoods of reaching ON3 and ON8 (indicated by CN2a and CN2d, respectively). Moving up another level in the tree, Emma next imagines what choices she will make in the future if she reaches DN2a or DN2b (collectively denoted DN2a/b below). Because she hopes to reach either ON3 or ON8, she forms a plan to move left (i.e., choose company stock) at DN2a and move right (i.e., choose to continue her old job) at DN2b. Having fixed these plans, Emma should now move up another level in the tree to consider the likelihood of her boss saying "yes" versus "no" at CN1. With this information Emma can form an expectation of what she will receive if she chooses to ask for a raise (i.e., move left at DN1), based on her planned choices (at DN2a/b) and her beliefs about the external events that might impact the scenario (at CNs). All that remains is for Emma to compare this expectation to the value of ON1 and to choose the larger. If this choice eventually leads to a subsequent DN, Emma must execute her planned choice.

This method of working backward from the bottom of a decision tree forms the basis of most traditional approaches to planning and multistage decision making (see Bertsekas, 2017; DeGroot, 1970; Keeney & Raiffa, 1976; Raiffa, 1968; Von Winterfeldt & Edwards, 1986). More formally, the process begins by assigning a utility value to each ON, representing its worth to the decision maker. In our example decision tree, ON1–9 would first be assigned the values $u_1$–$u_9$, respectively. Next, an expected utility is computed as the weighted average utility for each of CN2a–d. For instance, the expected utility (EU) for CN2c, EU(CN2c), is calculated by weighting $u_6$ and $u_7$ by their probabilities, $p_6$ and $p_7$, respectively, and summing the results. Third, final DNs (DN2a/b in Figure 1) are each assigned an expected utility by taking the maximum value of the CNs branching out of them. All other branches are pruned and subsequently ignored. This pruning represents the notion that the decision maker will choose the alternative with the higher EU at all future DNs, and therefore need not consider any information from the nodes she will choose to avoid. This process of calculating EU values for CNs and pruning foregone branches at DNs can be repeated indefinitely for multistage decision trees of any length, until utilities are propagated to the first DN. At this point, the decision maker need only choose the alternative with the highest expected utility (after pruning), and execute her planned decisions upon reaching later DNs.

The theory of backward induction has been used successfully to examine choices in contexts such as sequential games (Fudenberg & Tirole, 1991), investment (Dixit & Pindyck, 1994), business valuation (Anderson, 2009), and economic growth and taxation (Ljunggvist & Sargent, 2000). As a normative model, it can be extremely useful because following its prescriptions yields a strategy that maximizes one's long-term expected utility. However, as a descriptive model of human decision making, it faces several challenges. For one, the dynamic programming methods typically used to compute optimal solutions do not scale well to large or complex decision tree scenarios (see van Rooij, Wright, Kwisthout, & Wareham, 2018).

How such extensive calculations could be carried out by humans when they must be approximated with modern computers remains a challenge to anyone proposing backward induction as a descriptive theory of human behavior.

## Dynamic Consistency

A central tenet of planning via backward induction is known as dynamic consistency, which essentially prescribes that a decision maker follow through on her plans (Machina, 1989; Sarin & Wakker, 1998). The normativity of backward induction is critically dependent on this assumption because commitments made during the planning phase lead the decision maker to prune the decision tree and ignore irrelevant information related to forgone prospects. Failing to follow through on one's planned choices during the execution phase jeopardizes the entire strategy by changing the relevant information. For example, if upon reaching DN2a Emma chooses to move right instead of left, as she had planned, the values she induced earlier for DN2a and CN1 will be inaccurate. Because she already used these inaccurate values to make her choice at DN1, she undermines the optimality of her policy.

Many such violations of dynamic consistency have been demonstrated in multistage decision tasks. Cubitt, Starmer, and Sugden (1998) tested for dynamic inconsistency with a between-subjects design in which each individual made a single choice between the same two gambles presented in one of several forms. In the most relevant conditions, participants were told that the choice trial would begin with a lottery. With probability $p$ they would receive nothing; with probability $1 - p$, they would be allowed to choose between the two gambles. In the planned choice condition, participants were asked which gamble they would choose if they won the lottery. In the final choice condition, participants were told that they had already won the lottery, and were asked to make a choice. Cubitt et al. found a significant difference between planned and final choices, suggesting that participants were not dynamically consistent. Busemeyer and colleagues extended this work with a series of experiments employing within-subject designs to demonstrate various violations of consistency principles (Barkan & Busemeyer, 2003; Busemeyer, Weg, Barkan, Li, & Ma, 2000). Johnson and Busemeyer (2001) showed that dynamic choice inconsistencies increase with the length of the decision tree, suggesting that planning becomes more difficult as the planning horizon moves further into the future.

Some more recent research has attempted to directly investigate people's planning in multistage decision making (Bone, Hey, & Suckling, 2003, 2009; Carbone & Hey, 2001; Dimperio, 2009; Hey & Knoll, 2007, 2011). Carbone and Hey (2001) conducted a study using large three-stage decision trees comprised of interleaved binary DNs and CNs, with 64 terminal ONs. They used process tracing techniques to identify planning strategies. Participants were given unlimited time to explore the tree and examine node values. Carbone and Hey found that most people did not prune the tree according to their planned choices. That is, although participants did look ahead to future ONs to inform their strategies, they seemed unaware that they could make optimal choices in the future. Similarly, Bone et al. (2003) found that decision makers would often pass up a sure payoff at an early DN, only to settle for a smaller sure payoff at a later DN. This too suggests that decision makers did not form explicit plans for future choices.

Multistage decision trees have also been used to uncover individual differences in planning. Hey and Knoll (2011) used process tracing data to identified several distinct strategies.[1] For example, some participants approximated an optimal strategy, while others minimized effort by ignoring information. Hey and Knoll (2007) ran a three-stage decision study in which some participants appeared to plan all three stages completely, while others did not plan at all. Dimperio (2009) found similarly mixed results across several experiments, with some people carefully planning future choices and others treating future DNs as if they were random CNs, outside of their control. Taken together, these findings suggest that, even without time pressure or memory limitations, participants rarely formulate explicit plans in accordance with the backward induction model.

[1] Note that interpretation of these results is limited because Hey and Knoll (2011) did not provide formal details regarding how people's strategies were classified using process tracing data.

## The Payoff Variability Effect

The so-called payoff variability effect (Busemeyer, 1985; Erev & Barron, 2005; Myers & Sadler, 1960; Myers, Suydam, & Gambino, 1965) is an especially important challenge to the backward induction model. To understand why, imagine a task where participants choose between pairs of the following options: a low risk gamble (LR) to gain or lose \$0.05, with equal probability, a high risk gamble (HR) to gain or lose \$0.33, with equal probability, a certain loss (CL) of \$0.01, and a certain gain (CG) of \$0.08. The key finding is that participants typically choose LR from the set {LR, CL} more than they choose HR from the set {HR, CL}, suggesting that LR has greater utility than HR (i.e., risk aversion). In contrast, participants choose LR from the set {LR, CG} less than they choose HR from the set {HR, CG}, implying that HR has greater utility than LR (i.e., risk seeking). This reversal cannot be easily explained by strong utility models, such as prospect theory (Kahneman & Tversky, 1979) and ratio of strength models because they represent subjective utilities as static and context independent.

Below, we demonstrate that the payoff variability effect also occurs in multistage decision making, where outcomes related to the risky option are only experienced after an additional future choice. According to backward induction, people should be insensitive to this future outcome variability because they work backward through the tree. As they plan future choices and prune unnecessary branches, they collapse all outcome information down to an expected utility. Thus, utility, but not variability, is propagated back through the branches of the tree. Our findings point to a fundamentally different planning mechanism. This article introduces an alternative theory in which individuals plan future choices dynamically through repeated, forward-looking mental simulations. Unlike the backward induction model, this theory addresses the role of attention in directing one's limited cognitive resources to particular information. By focusing on the underlying cognitive processes involved, it accounts for the multistage payoff variability effect and offers a more psychologically plausible account of planning and dynamic decision making.

## Decision Field Theory–Planning

Judgment and decision making researchers have traditionally developed high-level, abstract theories, opting to understand and predict behavior using simple, intuitive concepts like subjective utility, preference consistency, and maximization. This approach is typically agnostic to the underlying mechanisms that give rise to behavioral tendencies. This article takes a different approach, and introduces a new model of planning and dynamic decision making, decision field theory–planning (DFT-P), which focuses on how preferences are constructed from internal cognitive processes. As its name implies, it is based on decision field theory (DFT; Busemeyer & Townsend, 1993; Diederich, 2003; Hotaling, Busemeyer, & Li, 2010; Roe, Busemeyer, & Townsend, 2001). DFT-P adopts the attention-driven information sampling and deliberation mechanism of DFT, and extends it to situations where an individual must plan decisions in a multistage decision tree. Although the model can, in principle, be applied to situations where there are many alternatives, the focus here is on a version of DFT-P for decision trees with binary DNs.

### Planning as Mental Simulation

The central claim of DFT-P is that people sample information about complex decision scenarios through repeated, rapid mental simulations. As they deliberate over which action to take at present, they imagine possible sequences of events, and simulate their own future decisions. Returning to the earlier example, when Emma imagines the consequences of asking for a raise (DN1), she does so by simulating a likely sequence of events. She might first imagine that her boss will say "yes," and will offer a choice between the stocks and wages. She then envisions choosing the stock option, which she imagines will yield a large payoff (e.g., ON2) because of the strong stock market. She compares this simulated outcome to the certain outcome of not asking for a raise (ON1) to produce momentary evidence in favor of the riskier option (assuming ON2 $>$ ON1). On the basis of this evidence, Emma's preference evolves toward asking for a raise. However, at the next moment Emma's attention may focus on less rosy possibilities. Now she imagines that her boss rejects her request, leaving her to decide whether to quit her job. She imagines quitting and then being unable to find work (ON7) due to a bad job market. Comparing this outcome to ON1 produces evidence pushing Emma's preference in the opposite direction, toward the safe alternative (assuming ON7 < ON1). In this fashion, she repeatedly simulates sequences of events, based on what she thinks will happen in the future. Simulations differ from moment to moment as she attends to different possibilities. Crucially, Emma's attentional biases may lead her to over- or undersample certain outcomes relative to their objective probabilities, for example, due to optimism. Over time, she accumulates evidence until her preference is sufficiently strong to make a choice.

The reader will note that—although the example above portrays mental simulations as proceeding via conscious thought—DFT-P does not propose that decision makers are explicitly aware of the simulations that drive their choices. This is not a unique feature of the model, as the presence of automatic, subconscious processes is a property shared with most cognitive theories of behavior. In fact, mental simulation in DFT-P is a natural extension of the sampling mechanisms at the heart of most accumulator or sequential sampling models (SSMs) of decision making. According to the SSM framework, during deliberation individuals sample information related to choice alternatives. In many models—for example, DFT and the leaky competing accumulator (Usher & McClelland, 2004)—these samples are drawn from the possible outcomes associated with each alternative, and can be interpreted as mental simulations of each choice. Thus, when considering a choice between two risky prospects, these models posit that decision makers repeatedly imagine and compare the outcomes of each option. With DFT-P, the same logic applies to deliberation and information accumulation in multistage decisions. In this sense, mental simulation is simply the multistage extension of sequential sampling.

## Formal Model Specification

DFT-P belongs to the broad class of SSMs, which posit that choices are made by accumulating information in real-time, until preference for one option reaches some critical threshold (see Ratcliff & Smith, 2004 for a review). It is built on DFT (Busemeyer & Townsend, 1993), and likewise proposes that individuals collect information by thinking about the possible outcomes of each available action. From moment to moment, different outcomes come to mind and are compared, producing fluctuations in preference. As attention switches stochastically, the strengths and weaknesses of each alternative cause the preference state to evolve in a noisy manner. Eventually, sufficient evidence accumulates in favor of one option that the corresponding decision threshold is reached and a choice is made.

More formally, DFT-P posits that at each moment in time a mental simulation is run for each alternative, $i$ and $j$ (again assuming binary DNs), and that the resulting outcomes, $v_i$ and $v_j$, are compared to produce a momentary valence in favor of $i$:

$$V = v_i - v_j. \tag{1}$$

The preference state at time $t$ is defined as the sum of the previous preference state and the new valence:

$$P(t) = P(t-1) + V(t-1). \tag{2}$$

$P(0)$ is the initial preference state, and can be used to explain initial bias and carry-over effects from past experience or previous decisions.

The primary focus of this article is the simulated version of DFT-P because it provides a more intuitive algorithmic representation of the proposed decision process. To compute the predictions of this model, one additional component is needed. A simulated value function, S, defines the probability of sampling each outcome when deliberating at each DN (see the Section 2.2.1 for more details). S determines the momentary values of $v_i$ and $v_j$, while Equations 1 and 2 describe how evidence accumulates across time. A decision is made as soon as the absolute value of the preference state exceeds some evidence threshold, $\theta$. According to DFT-P, Emma selects a $\theta$ value before deliberation, based on her level of caution.

As an alternative to the simulation version of DFT-P, quantitative predictions for choice proportions can be calculated analytically without the need for computer simulations. For this, a few additional calculations are needed. Across time, the mean input to the sampling process is the expectation of the valences. For a specific decision tree (e.g., Figure 2), this value can be calculated by taking the difference of mean inputs across alternatives:

$$\mu = \bar{v}_i - \bar{v}_j. \tag{3}$$

For each choice alternative (at any level of the tree) the mean input is simply the sum of all possible outcomes of choosing that alternative, weighted by the likelihood of mentally simulating each outcome according to S. That is for each alternative, $i$, leading to $k$ possible outcomes, $o_{ik}$, the mean input to the deliberation process is

$$\bar{v}_i = \sum_k S(o_{ik}). \tag{4}$$

As random fluctuations in attention produce different simulations, momentary valences also differ. The variance in valences is defined as

$$\sigma^2 = E[V - \mu]^2. \tag{5}$$

Finally, the probability of choosing alternative $i$ is[2]

$$\frac{1}{1 + e^{-\frac{\theta\mu}{\sigma}}}. \tag{6}$$

This equation highlights a key prediction of DFT-P that sensitivity to value differences ($\mu$) decreases as payoffs become more variable ($\sigma$ increases) or when a lower decision threshold is set.

Indeed, choices will generally favor the higher value option, but certain factors may reduce one's sensitivity. The most important of these factors for the present work is payoff variability. When ON values are extreme (i.e., highly variable), simulations will sometimes produce valences that strongly favor one option and other times strongly favor the alternative. This produces large fluctuations in preference and more random decisions when these fluctuations exceed $\theta$. In contrast, with low ON variability, each simulation will reach a similar payoff and produce a similar valence. Preference will evolve more smoothly toward the higher value option. The second factor influencing sensitivity is the decision threshold. DFT-P posits that individuals may set their decision thresholds so as to conserve cognitive resources while simulating future choices. Using a low threshold, $\theta_1$, when planning ahead at DN1 produces fast but more random decisions based on fewer simulations. After making an initial choice, individuals may set a higher threshold, $\theta_2$, for their final decision at DN2a/b because this allows for more careful deliberation based on many simulations.

Below, these fundamental predictions are tested in two multistage decision experiments. To preview the results, participants in both experiments displayed payoff variability effects, and made choices consistent with the predictions of DFT-P. A formal model comparison confirmed that DFT-P provides a better account of people's choice behavior compared to a flexible version of backward induction or a model built on simple heuristics.

**Simulated value function.** At each moment, DFT-P simulates one outcome for each choice alternative and compares these outcomes to compute a valence. A simulated value function, S, guides this process by determining which outcomes are simulated. Building on recent work showing that attention—typically defined in terms of visual fixation—can influence both preferential choice (Franco-Watkins & Johnson, 2011b; Krajbich & Rangel, 2011) and risky decision making (Franco-Watkins & Johnson, 2011a), DFT-P posits that selective attention shapes choices by focusing simulations on some outcomes more than others. Rather than faithfully simulating each outcome proportionate to its objective probability, an individual's imagination might, for example, optimistically oversimulate gains relative to losses or exaggerate the likelihood of extreme, rather than moderate payoffs.

In naturalistic contexts any number of environmental, personality, or cognitive factors (e.g., salience, memory strength, or emotional significance) may bias mental simulations away from an accurate reflection of expected likelihoods. However, at present minimal assumptions are made regarding this mechanism. The simulated version of DFT-P used to analyze behavioral data below

[2] Derived mathematically according to Busemeyer and Diederich (2010, chapter 4).

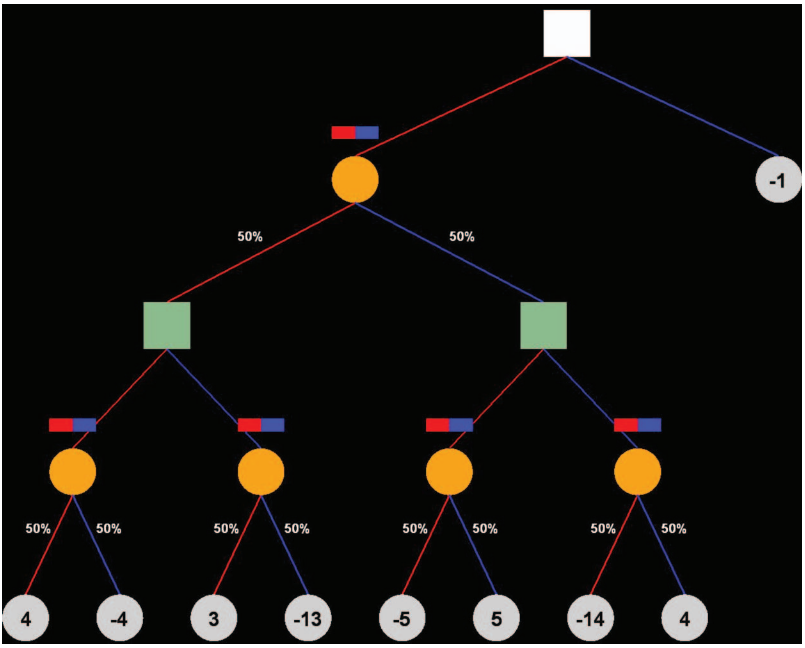


*Figure 2.* An example decision tree used in Experiment 1. Here, Chance Node (CN) 2 and CN4 are dominant CNs and CN3 and CN5 are dominated CNs. This tree has a certain payoff of −1, low payoff variability, and low negativity. The white square indicates that the current position is Decision Node 1. See the online article for the color version of this figure.



posits that daydreaming and outcome extremity are the only factors biasing attention. According to the model, people may sometimes lose focus and let their minds wander to other matters. This daydreaming is not limited to the outcomes in the decision tree, as people may imagine impossible outcomes. We represent this as random noise entering the decision process. With probability $\delta$ the model samples uniformly from the range of values present in the decision tree. When this happens, the momentary valence is simply

$$V \sim U(n, m), \tag{7}$$

where $n$ and $m$ are the minimum and maximum from ON1–9. Limiting these daydreams to the range of values found in the tree represents the intuition that people do not consider unrealistically extreme outcomes. Note that this mechanism is conceptually similar to the random sampling component of the distracted version of DFT (Bhatia, 2014).

When not daydreaming, outcomes are sampled as the decision maker runs a series of mental simulations tracing paths from the current DN to different ONs. For choices at DN2a/b this is simple. When a simulation arrives at one of CN2a–d the more extreme outcome (defined as having the higher absolute value) is sampled with probability $\varphi$, and the more moderate value is sampled with probability $1 - \varphi$. For choices at DN1, simulating an outcome for the gamble (choosing CN1) requires that the decision maker imagine what they will choose in the future if they reach DN2a/b. Returning to the example in Figure 1, when simulating an outcome for gambling at DN1, Emma first imagines what will happen at CN1.[3] Depending on which event she simulates for this node she will then imagine herself descending to either DN2a or DN2b. She must now simulate what she will do in the future if she reaches that DN. Simulated choices are modeled using the same process described above for actual choices (i.e., the deci-

[3] For the purposes of this article, decision makers are assumed to use the objective probabilities for CN1 when simulating events at this node.

sion process described by DFT-P). After simulating this future choice, Emma imagines an event for the last CN, leading finally to an ON. The value of this node represents what Emma momentarily imagines she will receive if she moves to CN1. She compares this value to the value of ON1 according to Equation 1 and the resultant valence is integrated into her preference state via Equation 2. Importantly, according to DFT-P, as Emma deliberates over what to do at DN1, she does not explicitly plan to make a decision at DN2a/b. Instead, moment to moment she simulates paths through the decision tree and imagines making future choices, without committing to a particular action.[4]

## Experiment 1

One unique aspect of the present task is that the gambles at DN1 involved a second decision between two CNs at DN2a/b. By independently manipulating the value of the certain payoff (ON1) and the riskiness of the gamble (i.e., the variability across ON2–9) one can assess the impact of variability across different multistage choice contexts. The primary research questions is this: Will preferences in a decision tree paradigm show a payoff variability effect, as DFT-P predicts, or will people be insensitive to variability and only attend to expected utilities, as prescribed by backward induction? More specifically, DFT-P predicts that greater payoff variability will produce noisier sampling during mental simulations, which will move choice proportions closer to 0.50. Thus, for trees with a certain loss (ON1 $<$ 0)—where one would expect participants to prefer the gamble—increased variability should decrease preference for the gamble (i.e., implying risk aversion[5]). For trees with a certain gain (ON1 $>$ 1)—where one would expect participants to prefer the safe alternative—increased variability should increase preference for the gamble (i.e., implying risk seeking). In contrast, backward induction predicts no relationship between variability and preference.

### Method

Ethical approval for all experiments was obtained through the institutional review board of Indiana University.

**Participants.** Forty-three students from Indiana University were recruited via posted flyers in the psychology building to participate in a paid decision making study for an average of \$8, with actual payment to be determined by their performance in the experiment (range \$5–\$11). We had intended for 40 participants, but a scheduling mistake led to three more being included.

**Procedure.** After obtaining informed consent, the experimenter seated participants at a computer where they began by watching a tutorial familiarizing them with the task. They were told that they would earn points based on the outcomes they received from each trial, and that those points would be converted to money for their final payment.

On each trial in the experiment a decision tree like the one shown in Figure 2 was presented. Each tree had the same structure, beginning at DN1 with a choice between a certain payoff and a complex multistage gamble. If participants chose the gamble, they made an additional choice (DN2a or DN2b), depending on the event at CN1. At all times, a white square indicated the participant's position in the tree (e.g., the marker is at DN1 in Figure 2). Participants selected the node they would like to move to next by clicking it with the mouse. The marker was then shown moving along the tree branch to the indicated node. When the maker reached a CN an urn appeared onscreen from which a marble was drawn. Drawing a red marble would send the marker to the left, while a blue marble would send it to the right. If the marker reached DN2a/b, the participant made a second choice by clicking on one of the remaining CNs. Each trial finished when the marker reached an ON, either by choosing the certain payoff at DN1 (ON1) or by completing the gamble (ON2–9). The points associated with this node were added to the participant's total

[4] Under the present formulation of DFT-P, Emma's imagined future choices do not constrain the decision she will make when she actually reaches DN2a/b. However, memory could be incorporated into models assumptions to represent partial dependencies between planned and actual choices.

[5] Note that we define risk aversion to be a negative relationship between the probability of choosing to gamble and the variance of the gamble, rather than a preference for a gamble with utility equal to that of a certain payoff.

score displayed on screen before the next trial began.

**Materials and design.** Each decision tree had the structure shows in Figure 2. If a participant chose to gamble at DN1, the marker would move to the second decision stage (DN2a or DN2b), where they would be faced with a second choice between two risky options: a dominated CN, which had a negative expected value (EV), and a dominant CN, which offered a gain or loss of some number of points, X, with equal probability (i.e., EV = 0; see Figure 2 for an example decision tree). Three critical factors were manipulated across trials in a 3 (Certain Payoff) × 3 (Variability) × 3 (Negativity) factorial design. The certain payoff at ON1 was 8, 1, or −1 points. These values were selected to assess the variability effect in circumstances where participants would typically prefer to gamble (i.e., ON1 = −1) and in those where participants would typically prefer the safe option (i.e., ON1 = 8). Variability was only applied to the dominant CNs, and was high (e.g., gain or lose 33), medium (e.g., gain or lose 19), or low (e.g., gain or lose 5). Negativity was only applied to the dominated CNs, and was high (EV = −15), medium (EV = −10), or low (EV = −5).[6] A factorial combination of these three variables produced 27 unique tree types. Each appeared once per block, in a random order.[7] Participants completed four blocks, for a total 108 trials. No additional measures were collected, nor were any additional conditions removed from analyses.

## Results

Because our research questions primarily dealt with how individuals planned future choices, the analyses were focused on behavior at DN1. Overall, participants chose the risky multistage gamble (CN1) on 48% ($SD$ = 0.31) of trials. However, Figure 3 shows that choice proportions at DN1 were heavily influenced by variability and certain payoff. As expected, as variance increased, choice proportions moved toward 0.50. For example, when the certain payoff was a loss of −1, participants appeared risk-averse, and the proportion of risky choices decreasing from 0.76 ($SD$ = 0.30) to 0.63 ($SD$ = 0.36). However, when the certain payoff was a gain of 8, participants appeared risk-seeking, and the proportion of risky choices increasing from 0.17 ($SD$ = 0.33) to 0.27 ($SD$ = 0.36). These seemingly contradictory trends constitute a multistage payoff variability effect. To test the robustness of this effect, a contrast score was calculated for each participant by comparing the difference between choice proportions when the certain payoff was −1 versus 8, across low versus high variability, or $(\psi_{-1,Low} - \psi_{8,Low}) - (\psi_{-1,High} - \psi_{8,High})$. A positive contrast value indicates that an individual's choices showed a payoff variability effect in that the value of certain payoff had a greater effect when variability was low, compared to when it was high. A two-tailed $t$ test showed that contrast values were significantly greater than 0 ($M$ = 0.23, $SD$ = 0.31), $t(42) = 4.96$, $p < .001$. This finding was consistent across individuals, with 30 participants (70%) having a positive contrast value.

A 3 (Variability) × 3 (Negativity) × 3 (Certain Payoff) repeated-measures analysis of variance (ANOVA) also supported this observation. A significant interaction between variability and certain payoff, $F(4, 84) = 10.60$, $MSE = 1.88$, $p < .001$, confirmed that the effect of variability depended on the value of ON1. A main effect of negativity, $F(2, 42) = 12.39$, $MSE = 2.28$, $p < .001$, was also found, with risky choices decreasing (from 52% to 46%) as the negativity of the dominated CNs increased.

An analysis of response times revealed that participants made much slower decisions at DN1 ($M$ = 3.14 s, $SD$ = 1.86 s) than DN2a/b ($M$ = 0.63 s, $SD$ = 1.36 s), $t(42) = 6.45$, $p < .001$. There was a trend indicating that higher variability led to slower responses at DN1 ($M_{low}$ = 3.08 s vs. $M_{high}$ = 3.28 s), but a one-way repeated-measures ANOVA showed that this was not significant, $F(2, 42) = 1.26$, $MSE = 21.33$, $p = .288$.

[6] Dominated CNs were included in the design of Experiment 1 to reduce the difficulty of choices at DN2a/b because participants could easy identify that one option was superior to the other. This simplification was removed in Experiment 2.

[7] To add visual variety to the stimulus set, and to make the experimental manipulations less obvious, ON values beneath DN2a and DN2b were made slightly different (see Figure 2). Also, for each block of trials, a superficially unique version of each tree was created by rearranging ONs, while preserving the underlying structure of the decision problem.

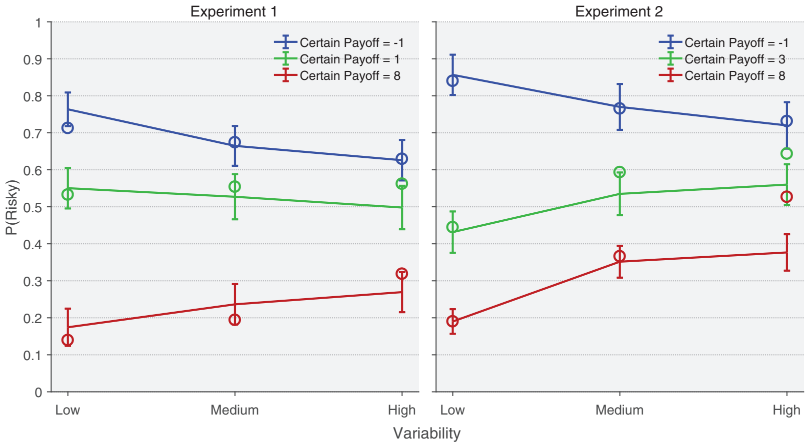


*Figure 3.* Mean proportion of choices in favor of the gamble (Chance Node 1) at Decision Node 1, as a function of dominant chance node variability and certain payoff. Lines are observed behavioral means and circles are the predictions of decision field theory–planning. Error bars are between-subjects standard errors. See the online article for the color version of this figure.

## Discussion

Results from Experiment 1 confirm the fundamental predictions of DFT-P. Not surprisingly, choices at DN1 showed sensitivity to outcome values. When ON1 was a gain, participants preferred it to the gamble; and when ON1 was a loss, participants preferred the gamble over the certain payoff. However, when the riskiness of CN1 was increased by increasing payoff variability, choice proportions moved toward indifference. This seemingly contradictory behavior—simultaneously displaying characteristics of both risk-seeking and risk-aversion—is easily explained by DFT-P's forward-looking simulation process. According to this account, increased payoff variability leads to noisier sampling of outcome values. These more variable samples increase the likelihood that the preference state exceeds the "incorrect" decision threshold (i.e., the threshold for the option with lower EV) due to noise. Thus, the model provides an intuitive and psychologically plausible account of people's behavior, which runs contrary to the prescriptions of backward induction.

The slower responses observed for DN1 compared to DN2a/b suggest that participants engaged in more elaborate deliberation when considering their future choices. This fits with DFT-P's proposal that choices at DN1 involve additional processing, in the form of multistage simulations involving imagined future choices. The model also predicts that higher levels of payoff variability will lead to earlier terminations (due to more variability in the preference state). The observed effect, though nonsignificant, suggests that this may have occurred for some individuals.

The frequency with which participants chose to gamble at DN1, despite CN1 typically having a lower EV than ON1, suggests that people had a bias in favor of gambling. However, an alternative interpretation is possible if one considers the visual salience of the two option on screen. Figure 2 shows that the gamble (i.e., the branches descending from CN1) was much larger than the safe option (ON1). This greater visual prominence may have led participants to choose CN1 more than they would have otherwise. Fortunately, this potential demand characteristic does not impact the key findings described above.

## Experiment 2

Experiment 2 aims to replicate the multistage payoff variability effect observed in Experiment 1.

## Method

**Participants.** Forty students from Indiana University were recruited and paid using the same procedure as in Experiment 1.

**Procedure.** Experiment 2 used a procedure identical to that of Experiment 1.

**Materials and design.** The negativity factor was removed from Experiment 2, resulting in a 3 (Variability) × 3 (Certain Payoff) design. Each of CN2-5 had similar EV (range −.5 to .5). Variability was manipulated as in Experiment 1, except that—due to the absence of the negativity manipulation—each of CN2-5 had the same level of variability and approximately the same EV. To equate the stimulus set sizes across experiments, and to add variety to the stimulus set in an effort to obscure the experimental design, three versions of each tree were created. The versions were structurally equivalent, but used slightly different ON values. Each of the 27 tree types appeared once per block, with participants completing five blocks for a total of 135 trials. Across blocks, trees were again made superficially distinct by rearranging ONs while preserving tree structure.

## Results

Analyses were again focused on decisions at DN1. Overall, participants chose the risky option (CN1) on 53% ($SD = 0.28$) of trials. Figure 3 shows that choice proportions followed a similar pattern to that observed in Experiment 1. Choice proportions again moved toward indifference (0.50) as payoff variability increased. With a certain payoff of −1, participants appeared risk-averse, and the proportion of risky choices decreasing from 0.86 ($SD = 0.21$) to 0.72 ($SD = 0.31$) as variance increased. In contrast, with a certain payoff of 8, they were risk-seeking, and the proportion of risky choices increasing from 0.19 ($SD = 0.34$) to 0.38 ($SD = 0.40$). Contrast scores were calculated as in Experiment 1, and showed a reliable multistage payoff variability effect ($M = 0.32$, $SD = 0.34$), $t(39) = 5.94$, $p < .001$. Twenty-eight participants (70%) showed the effect at the individual level, as indicated by positive contrast scores. A 3 (Variability) × 3 (Certain Payoff) repeated-measures ANOVA showed an interaction confirming that payoff variability moved preferences in opposite directions, depending on the level of certain payoff, $F(4, 78) = 18.68$, $MSE = 4.84$, $p < .001$.

Response times showed a similar pattern to Experiment 1, though the effect was smaller. Participants made slower decisions at DN1 ($M = 1.99$ s, $SD = 1.62$ s) than DN2a/b ($M = 1.26$ s, $SD = 1.60$ s), $t(39) = 1.86$, $p = .07$. Unlike in Experiment 1, no consistent trend was found in DN1 response times as a function of variability. Medium variability produced the slowest responses ($M_{medium} = 2.15$ s), followed by low variability ($M_{low} = 1.99$ s) and high variability ($M_{high} = 1.82$ s).

## Discussion

Experiment 2 replicates the multistage payoff variability effect. Unlike the backward induction model, participants showed sensitivity to the riskiness of CN1. When variability across ON1-9 was low, participants showed strong preferences; choosing the certain payoff when it was worth 8 points, and choosing CN1 when the certain payoff was −1. However, high outcome variability attenuated these preferences, resulting in more similar behavior across levels of certain payoff. Again, these simultaneous patterns of risk-seeking (certain payoff = 8) and risk-aversion (certain payoff = −1) naturally fall out of the noisy simulation process described by DFT-P.

# Comparing Models of Multistage Decision Making

Having established the multistage payoff variability effect in Experiments 1 and 2, we now examine three competing models of multistage decision making. Each represents a fundamentally different approach to solving a multistage decision problem. In comparing each model's ability to account for behavior, we hope to gain additional insight into the cognitive processes that participants used when planning and making choices. To preview our results, we find overwhelming support for DFT-P over models designed around static utility maximization and simple heuristics. Moreover, DFT-P is the only model that accounts for the multistage payoff variability effect observed in both experiments. All necessary modeling files are available at the following link: https://osf.io/fj725.

## Static Model

The static model tests whether it is possible to account for the observed payoff variability effects without the dynamic simulation process proposed by DFT-P. The model is a generalization of backward induction that assumes individuals choose probabilistically according to

$$\frac{1}{1 + e^{-sd}} \tag{8}$$

where $d$ is the difference in expected utility across alternatives and $s$ is a sensitivity parameter. ON values, $x$, are transformed into subjective utilities using standard power functions, like those used in prospect theory (Kahneman & Tversky, 1979), where $u(x) = x^{\alpha}$ for $x > 0$ and $u(x) = -\lambda|x|^{\beta}$ for $x < 0$. The parameters $\alpha$ and $\beta$ represent risk attitude for gains and losses, respectively, with values below 1 indicating risk-aversion and values above 1 indicating risk-seeking; $\lambda$ is a loss-aversion parameter representing the greater impact of losses relative to gains.

For choices at DN2a/b, the static model simply multiplies utilities by their corresponding CN probabilities to compute the expected utility for each option. For DN1, the model uses a flexible version of backward induction to compute the expected utility of CN1. Unlike the deterministic version of backward induction introduced earlier, the static model incorporates uncertainty about future choices. Rather than assuming that all future actions will maximize expected utility, it weights CNs according to their likelihoods of being chosen (i.e., according to the previously computed choice probabilities for DN2a/b). Thus when $s$ is large the static model reduces to the optimal backward induction model. Critically, the static model is backward-looking, in that it works backward from the end of the decision tree. It is also insensitive to payoff variability because it necessarily collapses the subtree descending from CN1 to a single expected utility value. The model uses four free parameters: $\alpha$, $\beta$, $\lambda$, and $s$.

## Min–Max Model

The min–max model investigates a simple heuristic strategy for multistage choice. It formalizes the intuition that while deliberating at DN1 decision makers attend to two features of the decision tree: the maximum gain, $g$, and minimum loss, $l$. When the former is large and the latter is small, the model predicts high rates of risk taking. When gains are small and losses are large, the model predicts fewer risky decisions. For example, returning to the scenario in Figure 1, Emma will consider the most extreme positive and negative outcomes that could result from asking for a raise. If the positive extreme is sufficiently better than the status quo (ON1) and the negative extreme is not significantly worse, she will choose to take the gamble. This model bears some similarity to the minimax and maximin strategies developed in game theory (see Savage, 1954). Unlike the static model, the min–max model incorporates sensitivity to local context because it compares outcome values to the certain payoff ($u_1$) on each trial. More formally, at DN1 the model evaluates the statement:

$$(g > u_1 + x) \;\&\; (l > u_1 - y),$$

where $x$ and $y$ are free parameters that determine whether $g$ and $l$, respectively, are considered sufficiently high. If the above statement is true, the model predicts that the individual will choose the risky option with probability $1 - \varepsilon$, where $\varepsilon$ is a parameter indicating the probability of making a "trembling hand" error. If the above statement is false, probability of making the risky choice is $\varepsilon$. At DN2a/b, the model assumes that individuals choose the higher expected value option with probability $1 - \varepsilon$.

The min–max model uses three free parameters: $x$, $y$, and $\varepsilon$. Unlike DFT-P and the static model, it does not consider future choices nor does it use the dependencies represented in the structure of the tree to inform its behavior. It does however provide a mechanism for divergent behavior across levels of certain payoff—a hallmark of the payoff variability effect—because its predictions depend on the value of $u_1$.

## Decision Field Theory–Planning

DFT-P operated as described earlier, with one exception. The model allows for the possibility that an individual has some a priori preference for taking risks, regardless of the outcomes on offer. In the context of Experiments 1 and 2, people may simply enjoy the activity of gambling—taking pleasure in the suspense and excitement—or they might prefer the safer option. To formally represent this, DFT-P uses a response bias at DN1. For choices at DN1, $P(0) = \rho$, with $\rho > 0$ indicating a bias in favor of risk, and $\rho < 0$ indicating a bias in favor of the certain payoff. This brings the total number of free parameters in DFT-P to five. Response bias $\rho$ and decision thresholds $\theta_1$ and $\theta_2$ determine the initial preference state and the choice criteria for planned and final decisions, respectively. Daydreaming probability $\delta$ and sampling bias $\varphi$ determine the simulated value function.

## Quantitative Model Comparison

To assess the goodness of fit of a model for a given response, we calculated the probability of making that response according to the model. We measured the ability of a model to account for people's behavior across an entire experiment by summing the logarithms of these probabilities across all responses and individuals.[8] We estimated model parameters that maximized this summed log-likelihood using a bounded Nelder–Mead optimization routine (Nelder & Mead, 1965). To avoid the problem of local minima, a grid of equally spaced parameter start points was defined. The optimization routine was run from each point in the grid in search of the globally optimal parameters.

To correct for the difference in number of parameters across models we computed the Akaike information criterion (AIC; Akaike, 1998) and Bayesian information criterion (Schwarz, 1978). These measure goodness of fit, while penalizing models based on the number of free parameters. Because both measures support that same conclusions, we focus on AIC for the remainder of this article. Using AIC values, we computed AIC weights corresponding to the probability of each model according to the observed data (Wagenmakers & Farrell, 2004)[9]. Table 1 summaries the results from this quantitative comparison, and shows overwhelming evidence for DFT-P. In fact, AIC weights for DFT-P are nearly 1 for both experiments, with essentially no support for the static and min–max models. Thus, DFT-P clearly provides the best fit to people's behavior.

Figure 4 compares the predictions of each model to participants' behavior for each DN in Experiments 1 and 2. Each point represents a DN for one decision tree, with each tree therefore contributing three points to the plot. In both experiments DFT-P's predictions most closely match observed mean choice proportions, and its predictions fall closest to the diagonal that represents perfect accuracy. The min–max model clearly performs the worst, showing insensitivity to most experimental factors, and making the same prediction for many DNs. The static model fares better, with predictions closer to the observed means, but still notably worse than DFT-P. This advantage is particularly evident in Experiment 1, where the static model significantly underestimates several values in the middle of the scale. Most importantly, Figure 4 show that DFT-P can qualitatively account for participants' choice behavior across a range of unique choice problems.

The failure of the static model presents a strong challenge to backward induction as a descriptive theory. The model represents a flexible, stochastic version of backward induction involving a psychologically inspired nonlinear utility function with several free parameters. Its inability to compete with DFT-P suggests that its fundamental assumptions cannot be supported. That is, decision makers do not work backward from the end of a decision tree, as prescribed by the optimal model. Additionally, the poor performance of the min–max model helps to rule out the hypothesis that participant relied on simple heuristics based on salient outcomes. These findings lend support to DFT-P's central claim that people plan future choices via a process of rapid forward-looking simulations of potential outcomes.

## Qualitative Model Comparison

Having established that DFT-P provides the best quantitative account of participants' overall choice behavior, we now turn to the multistage payoff variability effect. Because

[8] Although this method carries a risk of distorting results, estimating parameters separately for each individual was not feasible due to the sparsity of data at the individual level.

[9] This assumes that the true model is included in the test set.

Table 1
*Model Fitting Results*

| Model | Parameters | Log-likelihood | | AIC | | AIC weight | |
|---|---|---|---|---|---|---|---|
| | | Exp. 1 | Exp. 2 | Exp. 1 | Exp. 2 | Exp. 1 | Exp. 2 |
| DFT-P | $\rho$, $\theta_1$, $\theta_2$, $\delta$, $\varphi$ | 3,188 | 5,603 | 6,387 | 11,217 | 1 | 1 |
| Static | $\alpha$, $\beta$, $\lambda$, $s$ | 3,308 | 5,845 | 6,625 | 11,698 | 0 | 0 |
| Min–max | $x$, $y$, $\epsilon$ | 3,938 | 5,786 | 7,882 | 11,578 | 0 | 0 |

*Note.* Exp. = Experiment; AIC = Akaike information criterion; DFT-P = decision field theory–planning.



this pattern naturally emerges from the dynamic forward-looking simulation process upon which DFT-P is built, we expected the model to reproduce the pattern accurately. Figure 3 shows that this was indeed the case. Although the parameters of DFT-P were optimized with regard to the entire dataset, mean choice proportions demonstrate that this theoretically important effect emerges nonetheless. In both Experiments, the model yields a payoff variability effect—characterized as risk-aversion when ON1 was low and risk-seeking when ON1 was high—similar to that displayed by participants.

Despite the poor performance of the static and min–max models in the quantitative analysis, we were also interested in how well they could approximate the payoff variability effect. Figure 5 shows that neither model captures people's preferences. The static model appears incapable of producing anything like the convergence of choice proportions seen in the data. In fact, it is constrained to produce parallel lines in Figure 5. This is not surprising because, according to the model, expected values are propagated backward through the tree in a way that does not preserve variability information. This means that, as the model prepares to make a choice at DN1 it collapses the subtree descending from CN1 to a single utility value. As a result, the model's nonlinear utility function produces only limited sensitivity to payoff variability.

It was possible that the min–max model—which involves some degree of context sensitivity—might fare better; however, it too fails to capture the payoff variability effect. The model displays a different pattern of mis-

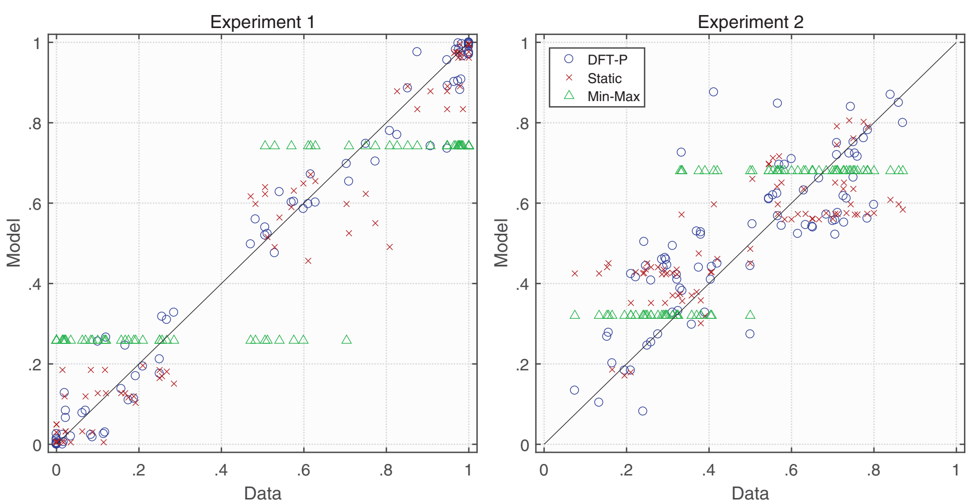


*Figure 4.* Scatter plots of model predictions and observed behavioral mean choice proportions for each decision node. DFT-P = decision field theory–planning. See the online article for the color version of this figure.

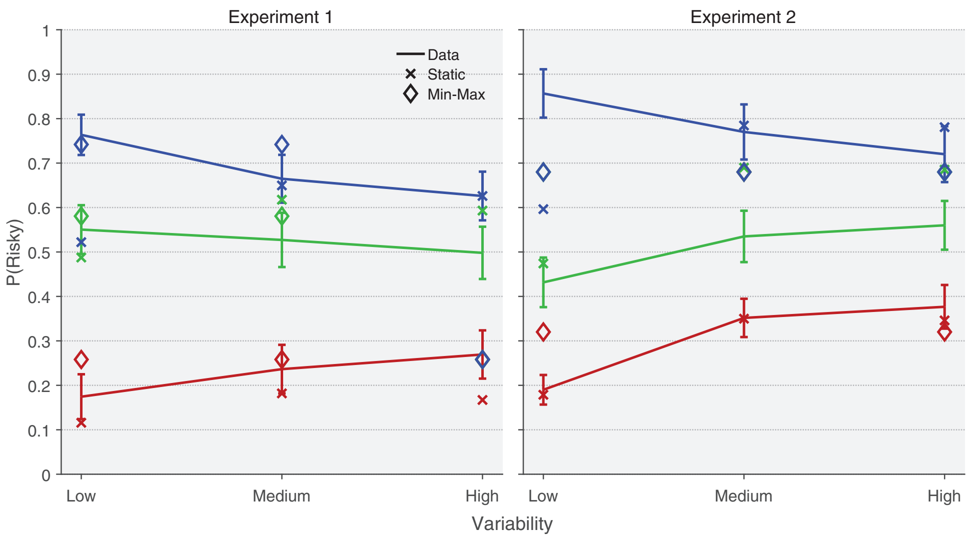


*Figure 5.* Mean proportion of choices in favor of the gamble (Chance Node 1) at Decision Node 1, as a function of dominant chance node variability and certain payoff. In Experiment 1, trees had a certain payoff of −1 (top/blue lines), 1 (middle/green lines), or 8 (bottom/red lines). In Experiment 2, trees had a certain payoff of −1 (top/blue lines), 3 (middle/green lines), or 8 (bottom/red lines). Error bars are between-subjects standard errors. See the online article for the color version of this figure.

fits, compared to the static model. As demonstrated in the left panel of Figure 5, it is capable of producing convergence in choice proportions as variability increases. However, the pattern bears little resemblance to that displayed by participants. For example, in Experiment 1, the model predicts dramatic risk aversion for high and medium levels of certain payoff, while in Experiment 2, the model predicts complete insensitivity to variability for these conditions.

It is possible that the static and min–max models could be modified to incorporate sensitivity to outcome variability. For example, a variance term could be added to the static models to improve its ability to produce payoff variability effects. This would involve changing Equation 8 to be more similar to the analytical version of DFT-P (Equation 6). However, the theoretical motivation for this addition would be unclear because the static model provides no mechanism for calculating variance. In contrast, for DFT-P sensitivity to outcome variability is explained as an emergent and natural consequence of a psychological process. In this sense, DFT-P is the only quantitative modeling framework for understanding the cognition underlying people's multistage decisions.

**Psychological interpretation of DFT-P parameters.** Another virtue of DFT-P is that its parameters can be interpreted in terms of psychological processes. Table 2 gives the optimal parameter values estimated during model fitting (see Tables A1 and A2 in the Appendix for parameters of the static and min–max models). As expected, in both experiments $\theta_1 < \theta_2$, indicating that participants collected relatively little evidence when planning decisions at DN1, but later raised their thresholds to carefully choose between options at DN2a/b. The substantially higher value of $\theta_2$ estimated in Experiment 2 is also sensible considering that choice alternatives at DN2a/b had nearly equal EV in that experiment. That is, participants in Experiment 2 would have needed to set higher thresh-

Table 2
*Decision Field Theory–Planning Parameter Values*

| Experiment | $\rho$ | $\theta_1$ | $\theta_2$ | $\varphi$ | $\delta$ |
|---|---|---|---|---|---|
| 1 | .33 | 7.37 | 15.82 | .14 | .50 |
| 2 | .50 | 8.04 | 39.87 | .50 | .36 |

olds than their counterparts in Experiment 1 to maximize EV at DN2a/b; ρ was positive in both experiments, showing that participants were biased in favor of gambling at DN1, and required more evidence to choose ON1 than CN1. The model's simulated value function also sheds light on the people's mental simulations. Sampling bias (φ) values indicate that simulations tended to favor extreme outcomes. This was particularly true in Experiment 2, where EVs were similar across CNs. Daydreaming (δ) also played a key role in people's mental simulations, with a substantial proportion of simulations resulting in random samples. This suggests that sampling of random information is a crucial mechanism for incorporating error into DFT-P's behavior (see also Bhatia, 2014).

**Response times.** DFT-P's mental simulation process makes predictions regarding the time it takes to make a decision, and some of the result from Experiments 1 and 2 support these predictions. For example, observing slower responses at DN1 compared to DN2a/b fits with the idea that simulations at DN1 involve long paths and require the simulation of future decisions. However, we are limited in our ability to test DFT-P's response time predictions at present. Additional work is needed to understand the proper interpretation of time across decision stages in DFT-P. This is illustrated by the fact that each simulation at DN1 requires an entire simulated decision at DN2a/b. It is unlikely that each sample within these simulated decisions takes the same amount of time as a sample within a real decision, so some scaling factor will be needed. Luckily, DFT-P provides a pathway for future investigations. Studies directly manipulating factors affecting the duration of simulated decisions—such as EV difference or stochastic dominance— can aid further development of DFT-P's response time component. This would also benefit from an experimental design focused on collecting many responses from each individual, allowing parameters to be simultaneously fit to individual choices and response time distributions.

## General Discussion

How people form plans and make decisions in dynamic environments has received relatively little attention from decision scientists despite the ubiquity and importance of such decisions in everyday life. This article introduces a new model, DFT-P, which posits a set of cognitive mechanisms to explain how people solve multistage decision problems. In contrast to the optimal backward induction model—which proscribes that one work backward from the end of a decision tree, making explicit plans for future choices—DFT-P proposes that individuals plan future choices on the fly. According to the model, individuals do not make explicit plans for future DNs, but rather use repeated forward-looking mental simulations to imagine possible sequences of events that might result from their actions. These simulations produce evidence, in form of momentary valences, which accumulate until a decision threshold is reached and choice is made.

DFT-P can be viewed as a theory of bounded rationality (Simon, 1982). Whereas backward induction specifies the optimal method for maximizing payoffs given unlimited time and cognitive resources, DFT-P provides an alternative account. It stipulates that planning is limited in several ways. First, the decision maker has limited attention and must therefore focus on the most important or salient outcomes, as represented in the model's simulated value function.[10] Second, the decision maker is unable to follow through on preplanned choices, and instead uses mental simulations to imagine what they will choose in the future. Third, the decision maker's deliberation is limited by their decision threshold representing the amount of evidence deemed sufficient for making a choice. Thus, despite the failure of the heuristic min–max model, our findings support the view that planning and multistage decision making is boundedly rational.

In two multistage decision experiments we tested a key prediction of the model—that noisy mental simulations should produce sensitivity to payoff variability—and found strong support for DFT-P. In both experiments, increased payoff variability attenuated preferences, producing both risk-seeking and risk-aversion effects. A formal comparison of DFT-P with two competitor models showed that neither a flexible ver-

[10] See Pachur, Schulte-Mecklenbeck, Murphy, and Hertwig (2018) for recent work investigating the links between biased attentional processes and static utility model parameters.

sion of backward inductions (the static model) nor a heuristic model focused on extreme ON values (the min–max model) provided an adequate account of these behavioral patterns. Using DFT-P as an analytical tool, we showed that its parameters can shed light on how people solved decision trees. We found a tendency to simulate extreme outcomes over more moderate ones, a finding in line with previous demonstrations of "overweighting of extremes" (Lieder, Griffiths, & Hsu, 2018; Ludvig, Madan, McMillan, Xu, & Spetch, 2018; Madan, Ludvig, & Spetch, 2014; Vanunu, Hotaling, & Newell, 2019). We also observed that participants simplified planning by setting a lower evidence threshold for planned choices relative to final choices. These findings provide a springboard for future investigations into the role of guided attention in dynamic decision making.

## Additional Evidence for Mental Simulation

This article is not the first to propose mental simulation as a means of making decisions. For instance, Kahneman and Tversky (1982) argue that people use the simulation heuristic to assess the likelihood of various outcomes when making predictions. Gigerenzer and colleagues have proscribed fast-and-frugal decision trees for effective medical decision making (Gigerenzer & Kurzenhaeuser, 2005), and as a means of achieving high performance in signal detection tasks (Luan, Schooler, & Gigerenzer, 2011). In machine learning, Monte Carlo tree search is a heuristic method for deriving optimal choices by randomly simulating paths through multistage decision trees (see Browne et al., 2012). This approach has been successfully applied to several problems within artificial intelligence, most notably game playing (e.g., Lee, Müller, & Teytaud, 2010; Pepels, Winands, & Lanctot, 2014). The ability of Monte Carlo tree search programs to compete with and beat champion Go players (see Silver et al., 2016) suggests that mental simulation can be an effective method for collecting information about a decision scenario. Indeed, it has been shown that making decisions on a basis of a limited number of samples, as DFT-P proposes, can approximate the behavior of an Bayesian ideal observer (Vul, Goodman, Griffiths, & Tenenbaum, 2014), lending further credence to the notion of mental simulation as a potent tool for complex choice problems.

There is also good evidence that individuals spontaneously construct the kinds of tree-like subjective representations needed to support planning via mental simulation. Much of this work involves the use of mental models to facilitate inferences in causal scenarios. For example, Sloman (2005) proposes that individuals naturally and automatically view the world within a causal framework. They construct tree-like causal models representing the relationships between relevant factors, events, and actions. As with DFT-P, these knowledge structures allow decision makers to think and act in complex choice scenarios. In fact, there is evidence that people use tree-like subjective representations to support various tasks, including judgment (Krynski & Tenenbaum, 2007), inference (Rottman & Hastie, 2014; Sobel, Tenenbaum, & Gopnik, 2004; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003), prediction (Fernbach, Darlow, & Sloman, 2010, 2011), counterfactual reasoning (Meder, Hagmayer, & Waldmann, 2009; Rips & Edwards, 2013), learning (Gopnik et al., 2004; Tenenbaum, Griffiths, & Kemp, 2006), and even the interpretation of causal verbs (Sloman, Barbey, & Hotaling, 2009; Walsh & Sloman, 2011).

Growing evidence also suggests that mental simulation plays a significant role in dynamic decision making. Klein and colleagues (Klein, 1993, 1998) proposed the recognition-primed decision model to describe how people use mental simulations to make quick noncompensatory decisions in complex real-world environments. In the field of neuroscience, Suzuki et al. (2012) conducted an fMRI study in which individuals were shown to use "direct recruitment" of their own mental processes to predict the choices of others. Physiological evidence for forward-looking mental simulations comes from research into so-called hippocampal place cells. Using experiments where rats learned to navigate branching mazes in search of food, A. Johnson and Redish (2007) and Dragoi and Tonegawa (2011) found neural evidence that rats simulate potential future outcomes while paused at the junction of two paths. Building on this work, Pezzulo and colleagues (Chersi & Pezzulo, 2012; Pezzulo, Rigoli, & Chersi, 2013) developed a computational model of maze navigation in which simulations are used to virtually "walk in the hippocampus."

In sum, previous research has shown that people construct tree-like representations of interconnected events, decisions, and outcomes to support a wide range of behaviors, and there is evidence that individuals use structured mental representations to imagine how the various elements of a scenario or problem will interact in the future. These findings broadly support the claims made by DFT-P.

## Accounting for Dynamic Inconsistency

DFT-P provides a new tool for modeling dynamic inconsistency. Because it does not explicitly commit to future choices it does not obey the principle of dynamic consistency, and can produce inconsistencies in multiple ways. The first source of inconsistency stems from the stochastic nature of the deliberation process. That is, simulations differ from moment to moment (and stage to stage) as different outcomes are sampled. Thus, imagined future choices will differ from actual choices when $\theta$ is not very large.

The second source of dynamic inconsistency relates to the changes in information processing across decision stages. As individuals move from planning to executing their decisions, the model posits that they may shift their attentional focus. Returning to the earlier example, Emma may initially give equal attention to the potential risks and rewards of asking for a raise (DN1), but later focus primarily on potential gains when choosing between company stock and increased wages (DN2a). With simulations at DN2a/b now based on a different set of information from that used at DN1, her preferences may change in a dynamically inconsistent way. Individuals may also set different decision thresholds across stages. By collecting different amounts of information for planned and immediate choices, preference reversals can occur across decision stages.

## Using DFT-P to Model Group and Individual Differences

Much of the scientific utility of cognitive modeling lies in its ability to shed light on the unique strategies that individuals employ. The present article does not model individual differences, but future work could exploit hierarchical modeling techniques to address individual differences despite limited individual-level data. DFT-P provides a new framework within which individual differences in multistage decision making can be identified. For instance, one might posit that some individuals simplify planning by ignoring outcome magnitudes when thinking ahead. This could be implemented with DFT-P by using a simulated value function that ignores ON value magnitudes during planning; attending only to the sign of each ON. Comparing the fit of this model to one that does not simplify planning would provide a convenient test of the simplification hypothesis. Hotaling and Busemeyer (2012)—using a model similar to DFT-P—demonstrated that this method can be used to identify groups of optimal planners, myopic planners, and nonplanners (see also Chapter 6 in Hotaling, 2013 for a demonstration of fitting DFT-P to individuals). Their success suggests that the DFT-P framework could prove a valuable tool for understanding how various naturalistic groups (e.g., clinical populations, age groups, drug users, etc.) differ in their capabilities and strategies for planning and multistage decision making.

## Conclusion

DFT-P is a new, psychologically grounded approach to understanding how individuals cope with the complexities of dynamic decision making. The experimental results and modeling presented here sheds new light on these topics, and provide strong support for the model's key predictions. Future work will explore ways of representing different decision making strategies, with a focus on identifying and characterizing the unique ways in which groups and individuals solve multistage decision problems.

## References

Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov & F. Caski (Eds.), *Selected papers of Hirotugu Akaike* (pp. 199–213). New York, NY: Springer.

Anderson, P. L. (2009). The value of private businesses in the United States. *Business Economics, 44,* 87–108. http://dx.doi.org/10.1057/be.2009.4

Barkan, R., & Busemeyer, J. R. (2003). Modeling dynamic inconsistency with a changing reference point. *Journal of Behavioral Decision Making, 16,* 235–255. http://dx.doi.org/10.1002/bdm.444

Bertsekas, D. P. (2017). *Dynamic programming and optimal control* (4 ed., Vol. 1). Belmont, MA: Athena Scientific.

Bhatia, S. (2014). Sequential sampling and paradoxes of risky choice. *Psychonomic Bulletin & Review, 21,* 1095–1111. http://dx.doi.org/10.3758/s13423-014-0650-1

Bone, J. D., Hey, J. D., & Suckling, J. R. (2003). Do people plan ahead? *Applied Economics Letters, 10,* 277–280. http://dx.doi.org/10.1080/1350485032000056882

Bone, J. D., Hey, J. D., & Suckling, J. R. (2009). Do people plan? *Experimental Economics, 12,* 12–25. http://dx.doi.org/10.1007/s10683-007-9187-8

Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., . . . Colton, S. (2012). A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games, 4,* 1–43. http://dx.doi.org/10.1109/TCIAIG.2012.2186810

Busemeyer, J. R. (1985). Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 538–564. http://dx.doi.org/10.1037/0278-7393.11.3.538

Busemeyer, J. R., & Diederich, A. (2010). *Cognitive modeling*. New York, NY: Sage.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100,* 432–459. http://dx.doi.org/10.1037/0033-295X.100.3.432

Busemeyer, J. R., Weg, E., Barkan, R., Li, X., & Ma, Z. (2000). Dynamic and consequential consistency of choices between paths of decision trees. *Journal of Experimental Psychology: General, 129,* 530–545. http://dx.doi.org/10.1037/0096-3445.129.4.530

Carbone, E., & Hey, J. D. (2001). A test of the principle of optimality. *Theory and Decision, 50,* 263–281. http://dx.doi.org/10.1023/A:1010342908638

Chersi, F., & Pezzulo, G. (2012). Using hippocampal-striatal loops for spatial navigation and goal-directed decision-making. *Cognitive Processing, 13,* S125–S129. http://dx.doi.org/10.1007/s10339-012-0475-7

Cubitt, R. P., Starmer, C., & Sugden, R. (1998). Dynamic choice and the common ratio effect: An experimental investigation. *Economic Journal, 108,* 1362–1380. http://dx.doi.org/10.1111/1468-0297.00346

DeGroot, M. H. (1970). *Optimal statistical decisions*. New York, NY: McGraw-Hill.

Diederich, A. (2003). MDFT account of decision making under time pressure. *Psychonomic Bulletin & Review, 10,* 157–166. http://dx.doi.org/10.3758/BF03196480

Dimperio, E. (2009). *A dynamic model of planning behavior in multi-stage risky decision tasks* (Doctoral dissertation). Retrieved from ProQuest Dissertations and Theses Database. (UMI 3380074)

Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton, NJ: Princeton University Press.

Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature, 469,* 397–401. http://dx.doi.org/10.1038/nature09633

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112,* 912–931. http://dx.doi.org/10.1037/0033-295X.112.4.912

Fernbach, P. M., Darlow, A., & Sloman, S. A. (2010). Neglect of alternative causes in predictive but not diagnostic reasoning. *Psychological Science, 21,* 329–336. http://dx.doi.org/10.1177/0956797610361430

Fernbach, P. M., Darlow, A., & Sloman, S. A. (2011). Asymmetries in predictive and diagnostic reasoning. *Journal of Experimental Psychology: General, 140,* 168–185. http://dx.doi.org/10.1037/a0022100

Franco-Watkins, A. M., & Johnson, J. G. (2011a). Applying the decision moving window to risky choice: Comparison of eye-tracking and mouse-tracing methods. *Judgment and Decision Making, 6,* 740–749.

Franco-Watkins, A. M., & Johnson, J. G. (2011b). Decision moving window: Using interactive eye tracking to examine decision processes. *Behavior Research Methods, 43,* 853–863. http://dx.doi.org/10.3758/s13428-011-0083-y

Fudenberg, D., & Tirole, J. (1991). *Game theory*. Cambridge, MA: MIT Press.

Gigerenzer, G., & Kurzenhaeuser, S. (2005). Fast and frugal heuristics in medical decision making. In R. Bibace, J. D. Laird, K. L. Noller, & J. Valsiner (Eds.), *Science and medicine in dialogue: Thinking through particulars and universals* (pp. 3–15). Westport, CT: Praeger.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111,* 3–32. http://dx.doi.org/10.1037/0033-295X.111.1.3

Hey, J. D., & Knoll, J. A. (2007). How far ahead do people plan? *Economics Letters, 96,* 8–13. http://dx.doi.org/10.1016/j.econlet.2006.12.015

Hey, J. D., & Knoll, J. A. (2011). Strategies in dynamic decision making: An experimental investigation of the rationality of decision behavior. *Journal of Economic Psychology, 32,* 399–409. http://dx.doi.org/10.1016/j.joep.2011.02.011

Hotaling, J. M. (2013). *Decision field theory-planning: A cognitive model of planning and dynamic decision making* (Doctoral dissertation). Indiana University, Bloomington, IN.

Hotaling, J. M., & Busemeyer, J. R. (2012). DFT-D: A cognitive-dynamical model of dynamic decision making. *Synthese, 189,* 67–80.

Hotaling, J. M., Busemeyer, J. R., & Li, J. (2010). Theoretical developments in decision field theory: Comment on Tsetsos, Usher, and Chater (2010). *Psychological Review, 117,* 1294–1298. http://dx.doi.org/10.1037/a0020401

Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *The Journal of Neuroscience, 27,* 12176–12189. http://dx.doi.org/10.1523/JNEUROSCI.3761-07.2007

Johnson, J. G., & Busemeyer, J. R. (2001). Multiple-stage decision-making: The effect of planning horizon length on dynamic consistency. *Theory and Decision, 51,* 217–246. http://dx.doi.org/10.1023/A:1015507110940

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision making under risk. *Econometrica, 47,* 263–292. http://dx.doi.org/10.2307/1914185

Kahneman, D., & Tversky, A. (1982). The simulation heuristic. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 201–208). Cambridge, United Kingdom: Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511809477.015

Keeney, R. L., & Raiffa, H. (1976). *Decision analysis with multiple conflicting objectives*. New York, NY: Wiley& Sons.

Klein, G. A. (1993). A recognition-primed decision (RPD) model of rapid decision making. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsambok (Eds.), *Decision making in action: Models and methods* (pp. 138–147). Westport, CT: Ablex Publishing.

Klein, G. A. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences of the United States of America, 108,* 13852–13857. http://dx.doi.org/10.1073/pnas.1101328108

Krynski, T. R., & Tenenbaum, J. B. (2007). The role of causality in judgment under uncertainty. *Journal of Experimental Psychology: General, 136,* 430–450. http://dx.doi.org/10.1037/0096-3445.136.3.430

Lee, C.-S., Müller, M., & Teytaud, O. (2010). Special issue on Monte Carlo techniques and Computer Go. *IEEE Transactions on Computational Intelligence and AI in Games, 2,* 225–228. http://dx.doi.org/10.1109/TCIAIG.2010.2099154

Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review, 125,* 1–32. http://dx.doi.org/10.1037/rev0000074

Ljungqvist, L., & Sargent, T. J. (2000). *Recursive macroeconomic theory*. Cambridge, MA: MIT Press.

Luan, S., Schooler, L. J., & Gigerenzer, G. (2011). A signal-detection analysis of fast-and-frugal trees. *Psychological Review, 118,* 316–338. http://dx.doi.org/10.1037/a0022684

Ludvig, E. A., Madan, C. R., McMillan, N., Xu, Y., & Spetch, M. L. (2018). Living near the edge: How extreme outcomes and their neighbors drive risky choice. *Journal of Experimental Psychology: General, 147,* 1905–1918. http://dx.doi.org/10.1037/xge0000414

Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature, 27,* 1622–1668.

Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin & Review, 21,* 629–636. http://dx.doi.org/10.3758/s13423-013-0542-9

Meder, B., Hagmayer, Y., & Waldmann, M. R. (2009). The role of learning data in causal reasoning about observations and interventions. *Memory & Cognition, 37,* 249–264. http://dx.doi.org/10.3758/MC.37.3.249

Myers, J. L., & Sadler, E. (1960). Effects of range of payoffs as a variable in risk taking. *Journal of Experimental Psychology, 60,* 306–309. http://dx.doi.org/10.1037/h0042499

Myers, J. L., Suydam, M. M., & Gambino, B. (1965). Contingent gains and losses in a risk-taking situation. *Journal of Mathematical Psychology, 2,* 263–370. http://dx.doi.org/10.1016/0022-2496(65)90010-6

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *The Computer Journal, 7,* 308–313. http://dx.doi.org/10.1093/comjnl/7.4.308

Pachur, T., Schulte-Mecklenbeck, M., Murphy, R. O., & Hertwig, R. (2018). Prospect theory reflects selective allocation of attention. *Journal of Experimental Psychology: General, 147,* 147–169. http://dx.doi.org/10.1037/xge0000406

Pepels, T., Winands, M. H., & Lanctot, M. (2014). Real-time Monte Carlo tree search in Ms. Pac-Man. *IEEE Transactions on Computational Intel-*

*ligence and AI in Games, 6,* 245–257. http://dx.doi.org/10.1109/TCIAIG.2013.2291577

Pezzulo, G., Rigoli, F., & Chersi, F. (2013). The mixed instrumental controller: Using value of information to combine habitual choice and mental simulation. *Frontiers in Psychology, 4,* 92. http://dx.doi.org/10.3389/fpsyg.2013.00092

Raiffa, H. (1968). *Decision analysis: Introductory lectures on choices under uncertainty*. Reading, MA: Addison Wesley.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review, 111,* 333–367. http://dx.doi.org/10.1037/0033-295X.111.2.333

Rips, L. J., & Edwards, B. J. (2013). Inference and explanation in counterfactual reasoning. *Cognitive Science, 37,* 1107–1135. http://dx.doi.org/10.1111/cogs.12024

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review, 108,* 370–392. http://dx.doi.org/10.1037/0033-295X.108.2.370

Rottman, B. M., & Hastie, R. (2014). Reasoning about causal relationships: Inferences on causal networks. *Psychological Bulletin, 140,* 109–139. http://dx.doi.org/10.1037/a0031903

Sarin, R., & Wakker, P. P. (1998). Dynamic choice and nonexpected utility. *Journal of Risk and Uncertainty, 17,* 87–120. http://dx.doi.org/10.1023/A:1007769628257

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

Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics, 6,* 461–464. http://dx.doi.org/10.1214/aos/1176344136

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., . . . Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature, 529,* 484–489. http://dx.doi.org/10.1038/nature16961

Simon, H. A. (1982). *Models of bounded rationality*. Cambridge, MA: MIT Press.

Sloman, S. (2005). *Causal models: How people think about the world and its alternatives*. New York, NY: Oxford University Press. http://dx.doi.org/10.1093/acprof:oso/9780195183115.001.0001

Sloman, S., Barbey, A. K., & Hotaling, J. M. (2009). A causal model theory of the meaning of cause, enable, and prevent. *Cognitive Science, 33,* 21–50. http://dx.doi.org/10.1111/j.1551-6709.2008.01002.x

Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science, 28,* 303–333. http://dx.doi.org/10.1207/s15516709cog2803_1

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science, 27,* 453–462. http://dx.doi.org/10.1207/s15516709cog2703_6

Suzuki, S., Harasawa, N., Ueno, K., Gardner, J. L., Ichinohe, N., Haruno, M., . . . Nakahara, H. (2012). Learning to simulate others' decisions. *Neuron, 74,* 1125–1137. http://dx.doi.org/10.1016/j.neuron.2012.04.030

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences, 10,* 309–318. http://dx.doi.org/10.1016/j.tics.2006.05.009

Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review, 111,* 757–769. http://dx.doi.org/10.1037/0033-295X.111.3.757

van Rooij, I., Wright, C. D., Kwisthout, J., & Wareham, T. (2018). Rational analysis, intractability, and the prospects of 'as if'-explanations. *Synthese, 195,* 491–510. http://dx.doi.org/10.1007/s11229-014-0532-0

Vanunu, Y., Hotaling, J. M., & Newell, B. R. (2019). *Elucidating the differential impact of extreme-outcomes in perceptual and preferential choice*. Manuscript under review.

Von Winterfeldt, D., & Edwards, W. (1986). *Decision analysis and behavioral research* (Vol. 604). New York, NY: Cambridge University Press.

Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science, 38,* 599–637. http://dx.doi.org/10.1111/cogs.12101

Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review, 11,* 192–196. http://dx.doi.org/10.3758/BF03206482

Walsh, C. R., & Sloman, S. A. (2011). The meaning of cause and prevent: The role of causal mechanism. *Mind & Language, 26,* 21–52. http://dx.doi.org/10.1111/j.1468-0017.2010.01409.x

(*Appendix follows*)

# Appendix

## Model Parameters

Table A1
*Static Model Parameter Values*

| Experiment | α | β | λ | *s* |
|---|---|---|---|---|
| 1 | 1.46 | 1.71 | .39 | .10 |
| 2 | 1.41 | 1.80 | .21 | .10 |

Table A2
*Min–Max Model Parameter Values*

| Experiment | *x* | *y* | ε |
|---|---|---|---|
| 1 | 2.01 | 24.09 | .26 |
| 2 | 1.45 | 46.53 | .32 |

Received January 21, 2015
Revision received May 13, 2019
Accepted May 18, 2019 ■