CHAPTER FIVE

Dynamic decision making: Empirical and theoretical directions

Jared M. Hotaling[a,*] and David Kellen[b]
[a]Department of Psychology, University of Illinois, Urbana-Champaign, Champaign, IL, United States
[b]Department of Psychology, Syracuse University, Syracuse, NY, United States
*Corresponding author: e-mail address: jaredhotaling@gmail.com

Contents

Abstract

Although most research into risky decision making has focused on simple scenarios—where isolated choices are made independent of one another—many important decisions in life play out across sequences of interdependent events and actions. For example, a student planning for a future career must consider which university to attend, which classes to take, which internships to pursue, and which job opportunities these might eventually lead to. Despite the ubiquity and importance of such decision problems, we know relatively little about how people manage the complexities of dynamic, multistage decisions. The goal of this article is to provide an accessible overview of some of the empirical and theoretical developments taking place in the study of dynamic decision making. We begin with a summary of some general questions being raised, then highlight two important lines of research. The first focuses on testing whether people's choices deviate from the prescriptions of a rational received view. The second instead focuses on individual differences in dynamic decision making, and seeks to classifying people into different types. Finally, we discuss candidate explanations for behavioral findings, alternative avenues of research, and opportunities for their integration.

Psychology of Learning and Motivation, Volume 76
ISSN 0079-7421
https://doi.org/10.1016/bs.plm.2022.03.004

1. Introduction

Many important decisions in life play out across sequences of interdependent events and actions. For example, a student planning for a future career must consider which university to attend, which classes to take, which internships to pursue, and which job opportunities these might eventually lead to. At each step in the sequence, there is risk and uncertainty, some of which depends on the future actions of the student. In these scenarios, decision makers must consider the ramifications of several events *outside* of their control (e.g., uncertain random events), as well as events *within* their control (e.g., their own future choices), and the relation between the two.

Despite the ubiquity and importance of such complex, multistage decision problems, most research into risky decision making has focused on much simpler scenarios. Typically, individuals will choose between two uncertain alternatives, with each choice made in isolation and independent of the next. Although such studies have greatly advanced our understanding of human decision processes, their focus on static, isolated choices constitute a challenge to those interested in how people manage the complexities of more dynamic, multistage decisions.

The goal of this article is to provide an accessible overview of some of the empirical and theoretical developments taking place in the study of dynamic decision making. In the first part, after some of the basic concepts are introduced, we will provide an overview of some general questions being raised, e.g., whether people's choices deviate from a rational received view. We will also discuss how these deviations can be accommodated by the dynamic cognitive processes postulated by *Decision Field Theory* (DFT; Busemeyer & Townsend, 1993). In the second part, we will turn to a parallel line of research that, instead of focusing on developing fine-grained accounts such as DFT, has focused on the individual differences found in dynamic decision making scenarios, classifying people into different types. In the third and last part, we will discuss potential explanations for observed patterns of behavior, as well as different avenues for future research, and opportunities for their integration.

1.1 Dynamic decision making: Some basic concepts

Fig. 1 presents an example of a multistage decision problem of the kind considered in dynamic decision making. This problem is represented in the form of a decision tree comprised of three types of nodes:

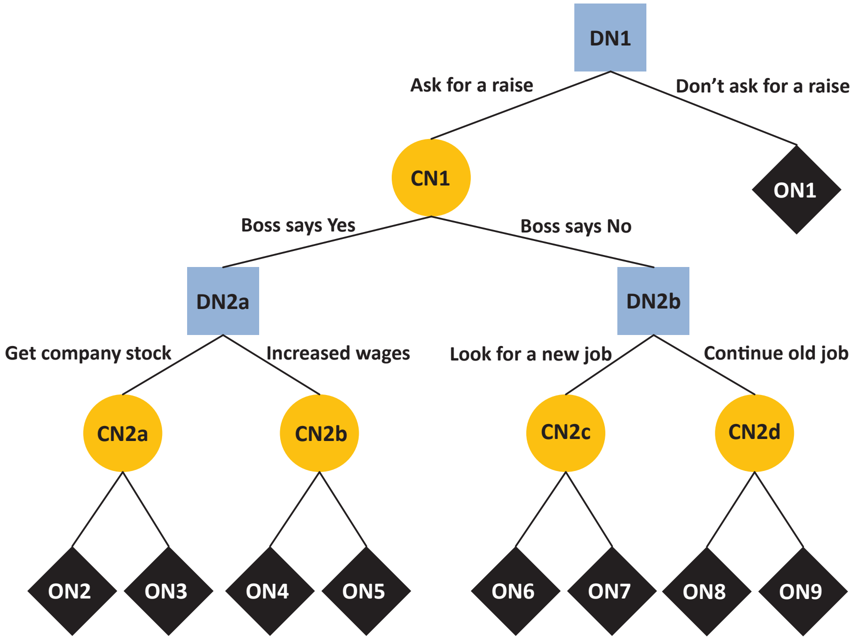

Fig. 1 A multistage decision scenario, represented as a decision tree. Decision nodes (DNs) represent points where decision maker chooses an action. Chance nodes (CNs) represent points where an uncertain external event occurs. Outcome nodes (ONs) represent final outcomes or payoffs.

1. *Decision Nodes* (DNs): Points where the decision maker chooses one of the available alternative event-actions.
2. *Chance Nodes* (CNs): Points where some risky/uncertain event (outside of the decision maker's control) takes place. By risky/uncertain, we mean that the probability with which a given event might occur can be known or unknown.
3. *Outcome Nodes* (ONs): The final consequences that result from specific sequences of events.

As an example, let us consider the case of the hypothetical decision maker, Emma, who wishes to improve her employment conditions. The problem Emma faces—deciding whether to ask her boss for a raise, while considering the implications of her boss's reaction on her outcome—can be represented by the decision tree in Fig. 1. Note that a variety of alternative real-life examples could have been used instead (e.g., deciding whether or to undergo surgery, weighing asking someone out on a date, choosing specific classes or opting for a given career path).

In this scenario, Emma's goal is to make choices such that she reaches an attractive ON and avoids the unattractive ONs. To begin, she must first decide whether it is better to (a) keep her current salary, or (b) take a risk by asking her boss for a raise. This decision point corresponds to DN1. Option (a) leads to a sure outcome ON1, whereas option (b) leads her to CN1, which represents the (uncertain) way in which her boss will react. If she chooses (b), she must next consider how her boss will react to her request. Since this is outside of her control, the best she can do is estimate the probability that they will say "yes" or "no." As a result, the probability of moving left or right at CN1 corresponds to Emma's belief regarding how likely her boss is to say "yes" or "no," respectively. Each of these possibilities leads to different DNs. If her boss says Yes, she moves left to DN2a, where she can choose between company stock (CN2a) and increased wages (CN2b). If her boss says No, she moves to DN2b, and is faced with the unenviable decision between staying at her current job (CN2c) or quitting to look for a new one (CN2d). Since Emma cannot be certain of the outcome of each of these final actions, CN2a–d represent her belief that external factors will ultimately influence which outcome she will experience. For instance, the probability of moving left or right at CN2a might depend on Emma's beliefs about the likelihood that company stock will increase or decrease in value over a given time period.

Emma's example scenario above and the decision tree structure that can be used to represent it provide a useful *mise en scene* when describing some of the different lines of research on dynamic decision making. As discussed below, some of the research being conducted focuses on the empirical violation of formal axioms, and how we can develop theoretical accounts that can accommodate these violations. In this context, the emphasis is placed on the development of "single" theoretical accounts that make a number of claims at the level of the cognitive processes that underlie choices. In scenarios such as Emma's, one is interested in a detailed understanding of the deliberation processes and how these processes can be affected by different factors (e.g., attention, complexity of decision problem). This type of research can be contrasted with contemporary work that places its focus on better understanding individual differences and the identification of different types of decision makers while being agnostic to the causes behind them (e.g., the possibility that each type is better described by different theories). The goal is not to make claims about how deliberative processes take place, or what kinds of processes are necessary and/or sufficient to explain the observed choices, but rather to identify the different ways in which people engage with these decision problems, not only in terms of the choices that they make, but also in terms of how information is pursued and how inconsistencies are handled. Here the aim is to figure out what kind of decision maker Emma appears to be, and how representative she is (i.e., is she of a different kind than Mary or Peter?).

2. Empirical review

2.1 Micro-process theories and the testing of rational axioms

In the study of human judgment and decision making, it is common for researchers to rely on a normative "received view" as a reference point. Empirical violations of this received view are used to motivate new theoretical accounts (for an overview, see Kellen, in press). For example, in the study of *decisions under risk and uncertainty*, researchers have long used the *expected utility model* as a reference point.[a] Empirical demonstrations of its axioms (e.g., independence) being violated has led to the development of many different theoretical accounts such as Prospect Theory (Kahneman & Tversky, 1979) or Regret Theory (Loomes & Sugden, 1982). In the context of *probability judgments*, Tversky and Kahneman (1983) demonstrated the occurrence of conjunction fallacies or "Linda Effects" that violate the axioms of probability theory and have led to the development of numerous alternative accounts (Costello & Watts, 2017; Crupi, Fitelson, & Tentori, 2008; Tversky & Kahneman, 1983).

A similar situation is found in the case of dynamic decision making, where *backward induction* is taken as the received view. One prominent branch of research into dynamic decision making has focused on testing the degree to which human behavior conforms to rational norms laid out by backward induction. More specifically, researchers have looked into two of its axioms, namely *dynamic consistency* and *consequential consistency*. Systematic violations of these axioms, and the resulting divergence between

[a] Let u(x) be the *subjective utility* of a given monetary gain x. The utility function u() is assumed to be monotonically increasing, such that $u(x) \geq u(y)$ if and only if $x \geq y$. According to the expected utility model, the expected utility of an option that yields mutually-exclusive gains x1, x2, …., xn with probabilities p1, p2, …, pn is given by the weighted sum u(x1)p1 + u(2)p2 + … + u(*xn*)pn. The model postulates that individuals prefer the option with the greatest expected utility. For a comprehensive yet accessible treatment, see Wakker (2010).

normative predictions and observed behavior, are often used to inform general-purpose models that seek to describe and explain decision making across a wide array of multistage scenarios.

2.1.1 Backward induction

It is generally accepted that the best solution in dynamic decision problems like the one represented by the decision tree in Fig. 1 can be achieved using *backward induction* (Bertsekas, 2017; DeGroot, 1970; Keeney & Raiffa, 1976; Raiffa, 1968; von Winterfeldt & Edwards, 1986). Intuitively, it proposes that individuals should first consider the outcomes represented at the end of the tree, and should work backwards, planning future actions in reverse chronological order until they reach the current DN. According to the standard version of backward induction, choices are deterministic, in that the decision maker always chooses the option with the highest expected utility. The model's decision strategy involves two distinct phases. First, the decision maker works backward through the tree, collapsing the features of each node down to a single utility value. At each DN, utility values are compared, and the decision maker explicitly commits to choosing the higher utility option. After working backward to the beginning of the decision tree, the second phase begins, wherein the decision maker executes their planned choices until reaching a terminal ON. Under standard assumptions, it can be shown that backward induction is the optimal strategy, such that following its prescriptions yields a policy that maximizes one's long-term expected utility.

Returning to our earlier example, the application of backward induction would require Emma to begin by inspecting the different outcome and identifying the few that strike her as particularly attractive, say ON1, ON3, and ON8. But before making her first move at DN1, Emma must formulate a plan for reaching one of the attractive ONs. First, she should consider the likelihoods of reaching ON3 and ON8, taking into consideration CN2a and CN2d. Next, moving up another level in the tree, Emma needs to imagine what choices she will make in the future if she reaches DN2a or DN2b. Since 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 planned these decisions, Emma now moves up another level in the tree to consider the likelihood of her boss saying "yes" versus "no" at CN1. Based on her determination (which might be an educated guess), Emma can use these probabilities (along with her beliefs about the other external factors) to form an expectation of what she will get if she chooses to ask for a raise (i.e., move left at DN1) and sticks to her planned choices (at DN2a and DN2b). All that remains for Emma is to compare this expectation (i.e., a subjective expected value) with the value of ON1 (i.e., not asking for anything) and to choose the option with the greater value. She then proceeds down the tree, and upon reaching a second DN, simply executes her planned choice.

Backward induction is a powerful tool for analyzing and improving dynamic decision making and has been successfully applied to a broad range of problems, including sequential games (Fudenberg & Tirole, 1991), investment (Dixit & Pindyck, 1994), business valuation (Anderson, 2009), and economic growth and taxation (Ljungqvist & Sargent, 2012). Despite its utility as a normative model, a growing body of evidence challenges its ability to describe how people make dynamic decisions. For instance, modelers typically use dynamic programming techniques to compute the predictions of backward induction. However, modern computers often struggle to calculate optimal solutions for complex decision trees, instead relying on methods that approximate backward induction (see van Rooij, Wright, Kwisthout, & Wareham, 2018). Such findings suggest that, while people may be capable of using backward induction for very simple problems, it is implausible that they would for many real-world problems.

Another set of challenges to backward induction come from violations of the consistency principles that form its foundation. First, the model assumes that people are *dynamically consistent*. Intuitively, this stipulates that a decision maker must follow through on planned decisions (Machina, 1989; van Rooij et al., 2018). That is, when working backward through a decision tree, the calculations of backward induction assume that future decisions can be accurately predicted. Branches of the tree are pruned according to the paths the decision maker will not take. If a decision maker fails to execute a planned decision, these pruned branches become relevant, meaning that the calculations of backward induction are no longer valid, and the optimality of the decision policy collapses. For instance, in the earlier example, Emma plans to move right at DN2b, and therefore prunes the branch linking this node to CN2c. If Emma later arrives at DN2b and changes her mind by deciding to look for a new job, she will undermine the optimality of her plan.

The second key assumption of backward induction is *consequential consistency*, which prescribes that a decision maker should only consider the future consequences of their decisions. This means that, all else being equal, it should not matter what events (including decisions) preceded a DN. That is, a person should take the same action when making an isolated decision as they would if this DN was encountered after a sequence of earlier events. If behavior differs as a function of past events, people are said to be consequentially inconsistent because their choices depend not only on the consequences of each option on offer, but also on their history of actions and experiences. To better understand the importance of consequential consistency in backward induction, consider the example in Fig. 1. When working backward from the end of the tree, Emma plans decisions for DN2a and DN2b without knowing what will happen at CN1 (or what she will choose at DN1). Working backward from the end the decision tree prevents her from accessing information about the past when planning for the future. Thus, the only way to satisfy dynamic consistency is to base her decisions purely on future consequences. To do otherwise would make her dynamically inconstant because planned choices would depend on future consequences alone, but actual choices would also be affected by our history.

2.1.2 Violations of dynamic and consequential consistency

Several studies have focused on directly testing the consistency principles of backward induction. Rather than compare human performance to optimal performance at a global level, as several earlier studies have (Brehmer, 1992), this research aims to avoid confounds that might introduce alternative explanations for sub-optimality (e.g., computational errors, memory limitation, insufficient learning, etc.). They instead use simple experimental paradigms with only enough complexity to test the assumptions of backward induction. Much like earlier research that used simple, binary, single-stage risky decisions to test the axiom of expected utility theory, this work employs simple decision trees, typically with binary DNs.

To better understand the insights that this research provides, consider the approach by Busemeyer, Weg, Barkan, Li, and Ma (2000). Importantly, this study used a within-subjects design, which allowed the authors to compare each person's rate of choice inconsistency—the rate at which actions change across repetitions of the same decision—to their rates of dynamic and consequential inconsistency. Only if the latter two exceed that of the former, can behavior truly be said to violate the prescriptions of backward induction. Busemeyer et al. employed simple asymmetric trees structures with a sequence of four DNs. Each of first three DNs offered a choice between a sure payoff and a CN, which had a 50/50 chance of continuing to the next DN or exiting the tree and ending the trial with no payoff. The final DN, D, offered a choice between a sure payoff, S, and a gamble, G, involving a 50/50 chance of receiving reward, R, or punishment, P.

To test the assumption of dynamic consistency, on some trials, at the beginning of the tree participants were asked to plan what they would choose if they found themselves at D. The computer would then automatically carry out their plan upon reaching D. On other trials, participants did not explicitly plan a decision, but simply made a final choice upon reaching D. Comparing planned and actual choices, Busemeyer et al. found that people were dynamically inconsistent. When thinking ahead about the choice that they would make at D, participants typically planned to gamble, but when they actually reached D, they often changed their minds, and opted to take the sure payoff, S.

To test the assumption of consequential consistency, Busemeyer et al. (2000) compared actual choices at D to choices made in another condition, where D appear in isolation (i.e., as a simple, single-stage choice between S and G). They found that rates of consequential inconsistency did not significantly exceed rates of choice inconsistency, meaning that people were just as likely to change their decision across repetitions of the same choice as they were to change their decision across conditions. Since pure choice variability was sufficient to explain apparent inconsistencies across isolated and multistage scenarios, the authors concluded that participants were consequentially consistent.

Barkan and Busemeyer (1999, 2003) found a similar pattern of inconsistencies in scenarios where individuals received the outcome of an initial gamble before deciding whether to take a second. In their experiments, each trial consisted of an obligatory gamble, followed by a second identical gamble that was optional. Before playing the first gamble, participants had to plan whether or not to take the second gamble. Two plans were made for each trial: one contingent on winning the first gamble, and another contingent on loosing. After making these plans, participants observed the outcome of the first gamble, and decided whether to take the second. Barkan and Busemeyer found that in approximately 20% of trials, participants failed to carry out their planned choices. Moreover, these inconsistencies systematically depended on the experienced outcomes. If the first gamble was won, participants showed risk-averse inconsistencies, with many who had planned to gamble now choosing not to take the second gamble. On the other hand, if the first gamble was unsuccessful, participants showed risk-seeking inconsistencies, with many of those who planned to forego the second gamble now choosing to take it.

Johnson and Busemeyer (2001)—building on earlier work by Busemeyer and colleagues—found that dynamic consistency was systematically affected by *planning horizon*. Using a design similar to Busemeyer et al. (2000), they asked people what they planned to choose in the event that they reached a final DN and compared these plans to the actual choices people made upon reaching the final node. However, Johnson and Busemeyer (2001) also manipulated the length of decision trees, such that the number of intermediate choices between the beginning and end differed across trials. They found that as the length of trees increased—and participants were asked to plan further ahead—they were less likely to execute their planned choice. For the shortest trees—where only a single CN separated participants' current position from the final DN—participants typically followed through on their plans, and mean inconsistency rates were around 25–30%. However, with longer trees—where four or five DNs separated participants current position from the final DN—inconsistency rates were higher (40–50%).

The empirical studies reviewed here provide us with some insight into the *overall* occurrence of dynamic and consequential inconsistencies and some of the factors (e.g., planning horizon) that affect its relative frequency. In the section that follows, we will take a somewhat different perspective and focus on individual characterizations. For example, can we identify a subset of "inconsistent" individuals—i.e., those that do not engage in backward induction—and if so, how large is this subset?

2.2 Characterizing individual differences in dynamic choices

In the field of judgment and decision making, behavioral "paradoxes" or "stylized effects," such as the violations of dynamic and consequential consistency discussed earlier, provide researchers with evidence necessary to indict a received view such as expected utility. By demonstrating that some people are at odds with this received view, the researcher can then motivate the development of new theoretical accounts that (hopefully) can accommodate them. One of the primary roles of any new model being proposed is to organize the existing corpus of empirical results in terms of its theoretical components. These components provide an explanation for the observed phenomena, but also relate them in a principled way (e.g., Loomes, 2010). For instance, Decision Field Theory (DFT) can explain the observed violations of dynamic and consequential consistency by postulating a dynamic deliberation process in which attention guides how information about different alternatives is gathered (Johnson & Busemeyer, 2001).

The complexity of the theoretical components being proposed often compromises the ability to study them in detail, at least using a single experimental paradigm or domain. To put it in simple terms, the data available are often too coarse to support any kind of fine-grained characterization (for relevant discussions, see Kellen, 2019; Navarro, 2019; Shiffrin & Nobel, 1997). Fortunately, this limitation can be offset by an appeal to generality: Because these theoretical components are assumed to play a major role in the judgments and decisions that are made across a wide range of contexts, one can continuously expand their domain of application. In other words, these theoretical components issue a number of explanatory "promissory notes" to be cashed out in future research endeavors (Soyland, 1994). In the case of DFT, the same mental simulation process assumed to underlie people's deliberations when choosing among a wide variety of risky options (e.g., Kellen, Steiner, Davis-Stober, & Pappas, 2020) is also called upon when accounting for the aforementioned violations of consistency in dynamic choices (Hotaling, 2020; Hotaling & Busemeyer, 2012), the effects of time pressure in people's inclinations toward rewards and punishments (Diederich, 2003), or the influence that the characteristics of choice alternatives play (e.g., Berkowitsch, Scheibehenne, & Rieskamp, 2014; Busemeyer & Goldstein, 1992), among many other phenomena.

One of the consequences of developing increasingly general and fine-grained accounts is that it shifts researchers' attention away from any specific set of phenomena and toward the theoretical components assumed to underlie them. But although these shifts are not only desirable but a necessary step in the development of causal explanations in psychology (see Cummins, 1983, 2000), they are not without costs. One of these costs is that the resulting theoretical models become increasingly impractical when dealing with a number of relevant scenarios (e.g., applied research settings; see Gonzalez, Fakhari, & Busemeyer, 2017). In these cases, researchers are better off relying on simplified models that gloss over many—if not all—of the underlying processes (Batchelder & Riefer, 1999; van Ravenzwaaij, Donkin, & Vandekerckhove, 2017), to the point that this kind of approach has led to the emergence of a new research subfield known as cognitive psychometrics (e.g., Oravecz, Muth, & Vandekerckhove, 2016; Riefer, Knapp, Batchelder, Bamber, & Manifold, 2002).

In the context of dynamic decision making, many modeling efforts of this kind have been spearheaded by John D. Hey and colleagues (e.g., Hey, 2005; Hey & Knoll, 2011; Hey & Panaccione, 2011), who have developed a number of accounts that focus on the classification of individuals into different types, rather than on the processes that underlie said differences (for similar modeling approaches, see Birnbaum & Schmidt, 2010;

Bott, Kellen, & Klauer, 2021; Hilbig & Moshagen, 2014). Such assessments are extremely valuable from a theoretical standpoint, allowing researchers to distinguish between broad classes of theoretical accounts (see Luce, 2010), and also protecting them from excessively relying on phenomena cast at the aggregate level (for a discussion, see Regenwetter & Robinson, 2017). The goal of this section is to review some of this work and explicate some of the models used in it. As will become clear in a later discussion, there is enormous potential in the formal extension of these models, incorporating multiple kinds of variables and individual-level information. Among these is the potential for development of alternative/complementary explanatory accounts that are not necessarily causal or mechanistic (see Coulter, 1983; Hacker, 2013; Harré & Secord, 1973).

Hey and Knoll (2007) reported a study with the purpose of evaluating how far ahead people tend to plan. Participants dealt with trees in which DNs are interleaved with 50–50 CNs (see Fig. 2), each decision node

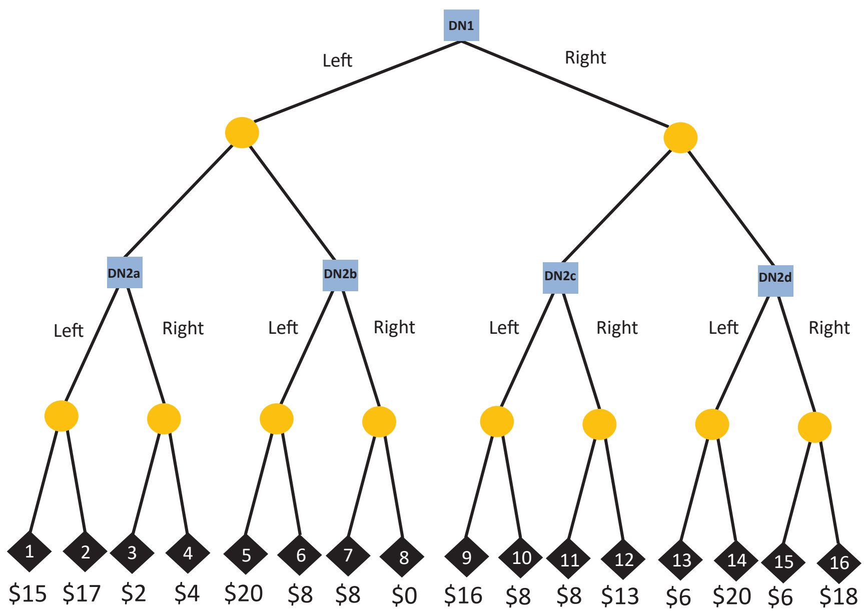

Fig. 2 Illustration of one of the kinds of decision trees used by Hey and Knoll (2007). For purposes of illustrative clarity, we suppressed the names of chance nodes (yellow circles), and in the case of outcome nodes (black diamonds) we only refer to their number. Please note that the two branches associated with each chance node have probability 0.50.

splitting into a Left and Right branch. The payoff structure across the different nodes was designed so that it would differentiate between three different kinds of individuals:

1. "people that plan ahead,"
2. "people that plan one step ahead," and
3. "people that do not plan."

To see how this differentiation takes place, we need to consider the dominance relationships that hold at each decision level. For simplicity, we will consider a tree with only two decision levels and 16 possible payoffs (see Fig. 2), although this tree only allows for a partial differentiation of types, namely between (1 or 2) and (3).

The possible payoffs associated with each of the first decision-level branches can be understood as equiprobable outcomes from two lotteries[b]:

$$\mathrm{L1} = (\mathrm{ON1}, \mathrm{ON2}, \mathrm{ON3}, \mathrm{ON4}, \mathrm{ON5}, \mathrm{ON6}, \mathrm{ON7}, \mathrm{ON8}),$$
$$\mathrm{R1} = (\mathrm{ON9}, \mathrm{ON10}, \mathrm{ON11}, \mathrm{ON12}, \mathrm{ON13}, \mathrm{ON14}, \mathrm{ON15}, \mathrm{ON16}),$$

Because of their possible outcomes, lottery R1 dominates L1, which means that a person that only focuses on the first decision level (DN1); i.e., someone that does not plan ahead, is expected to choose the Right branch.[c] At the second decision level (DN2a-DN2d), the payoffs associated with the Left and Right branches indicate that the former dominate the latter:

$$\mathrm{L21} = (\mathrm{ON1}, \mathrm{ON2})\ \mathrm{R21} = (\mathrm{ON3}, \mathrm{ON4}),$$
$$\mathrm{L22} = (\mathrm{ON5}, \mathrm{ON6})\ \mathrm{R22} = (\mathrm{ON7}, \mathrm{ON8}),$$
$$\mathrm{L23} = (\mathrm{ON9}, \mathrm{ON10})\ \mathrm{R23} = (\mathrm{ON11}, \mathrm{ON12}),$$
$$\mathrm{L24} = (\mathrm{ON13}, \mathrm{ON14})\ \mathrm{R24} = (\mathrm{ON15}, \mathrm{ON16}).$$

What these dominance relationships indicate is that someone who does not plan ahead should choose "Right" at the first level and "Left" at the second level. In contrast, someone that plans ahead (to some degree) is expected to begin by identifying their preferred second-level branches, and only then select the first level the branch that leads to them. In this specific example, the dominance of the "Left" branches at the second level establishes that this

[b] To avoid making things unnecessarily complex, we will not provide the actual outcomes unless they are essential for understanding the issue at stake.

[c] When trying to understand the concept of (stochastic) dominance, it is useful to refer to a simpler example, outside of the present decision tree problems. Let A = (2,5,8) and B = (2,6,9) denote 2 three-outcome lotteries with equiprobable outcomes. Lottery B dominates A because the probability of "gaining x or more" in B is always larger or equal than in B, for all x.

will be their second-level choice regardless of the first-level outcome. This means that the first-level decision can be represented as a choice between two reduced lotteries (with equiprobable outcomes):

$$L1^* = (ON1, ON2, ON5, ON6)\ R1^* = (ON9, ON10, ON13, ON14)$$

In this case, L1* dominates R1*, which means that someone who does plan ahead is expected to select "Left" at both levels. Larger trees with three decision levels (and 64 payoffs) were used to identify those that can only plan one step ahead. Hey and Knoll (2007) fit participants' choices with a number of true-and-error models (for a review, see Birnbaum, 2013) that can estimate the proportion of individuals in each of the three categories while accounting for the possibility that some of the responses incorrectly expressed participants' preferences (i.e., they were errors). As an example, consider the following characterization of choice probabilities based on a true-and-error model:

$$\begin{aligned} P(\text{Left1}, \text{Left2}) &= p^*(1-e1)^*(1-e2) + (1-p)^*e1^*(1-e2), \\ P(\text{Left1}, \text{Right2}) &= p^*(1-e1)^*e2 + (1-p)^*e1^*e2, \\ P(\text{Right1}, \text{Left2}) &= p^*e1^*(1-e2) + (1-p)^*(1-e1)^*(1-e2), \\ P(\text{Right1}, \text{Right2}) &= p^*(1-e1)^*(1-e2) + (1-p)^*(1-e1)^*e2, \end{aligned}$$

where p denotes the probability of someone planning (at least one step) ahead, and the complement $1-p$ being the probability that someone does not plan at all. In turn, parameters e1 and e2 denote the probability of an error response at the first and second levels, respectively. Note that these error parameters prevent the model from assuming that individuals make deterministic, error-free choices, without making any strong commitment to the source of said errors (attention failure, misunderstanding, motor-response error). In order to estimate these error probabilities, the same individuals need to make decisions across different decision trees. For an extensive review of true-and-error models, see Birnbaum (2013). The parameter estimates obtained by Hey and Knoll (2007), when considering people's choices for both two- and three-decision-level trees, indicated that 41% of the participants did not plan ahead, whereas 58% of them planned at least two steps ahead. Surprisingly, only 1% of the individuals were found to only plan one step ahead. This result indicates that participants can be sharply divided into two groups, with almost no one in between. This near-absence of intermediate cases suggests that individuals can be broadly understood as being of one of two kinds—naïve and sophisticated.

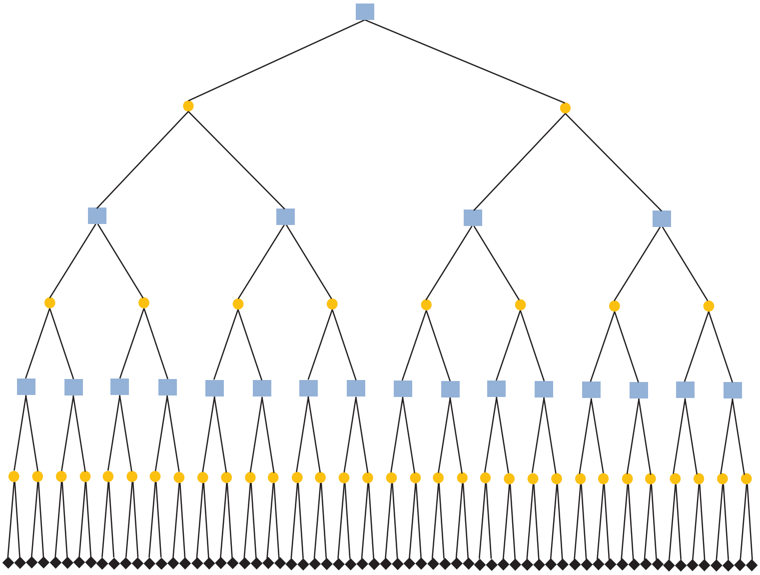

Fig. 3 Illustration of one of the kinds of decision trees used by Hey and Knoll (2011). For purposes of illustrative clarity, we suppressed the names of all nodes. Please note that the two branches associated with each chance node have probability 0.50.

In a follow-up study, Hey and Knoll (2011) investigated the reasoning processes that lead to people's planned choices in a task involving decision trees where three DNs were interleaved with three 50:50 CNs, for a total of 64 outcomes (see Fig. 3). One specific aspect of this study is that participants had to click on each final ON to learn about them. This was done in order to better track what kind of information was considered, and when. All of participants' mouse and keyboard activity was recorded for later analysis. Specifically, Hey and Knoll analyzed participants' verbal reports of how they were reasoning through different decision tree problems. These reports were recorded using a computerized notepad, along with their entire sequence of mouse and keyboard activity.

This activity was used to classify participants and allowed a distinction between the determination of a plan and its subsequent implementation. There is a clear dominance relationship in these trees, which means that there is a unique optimal solution. If participants indeed adhere to backward induction, a strategy that yields the optimal solution, then we should be able to observe this in their recorded activity: Specifically, they should begin by first evaluating the outcomes associated with the 16 third-level DNs. As before, these nodes involve a "Left" and "Right" alternative, each followed by a chance node. By comparing the pairs of outcomes associated with each decision, the participant should be able to identify the dominating decision. After making this determination, they can shift their attention to the second-level DNs, and consider which alternatives lead to the third-level dominating outcome pairs. Afterwards, the participant is left to determine which of the two first-level decision alternatives leads to their preferred second-level DNs. Based on their inspection of this "process and reasoning" data, Hey and Knoll (2011) determined that 31% of their participants conformed to backward induction, although there was some variation in the degree to which this strategy was successfully achieved: A portion of the backward-inducing participants (3% out of the 31%), designated "quasi-rationalists," failed to completely disregard the outcomes associated with some of the dominating options, which led to sub-optimal decisions at the first level. Another portion (16%), designated as "simplifiers," showed a sequence of forward and backward moves (or vice versa), in an attempt to deal with the complexity of the decision tree. Among the remaining 69% of participants whose behavior did not appear to conform to backward induction, 34% were classified as "forward workers" implementing a forward induction strategy in which they begin by selecting the first-level alternative that leads to their preferred set of 16 outcomes (out of four sets), and after that selecting their second-level alternative in terms of their preferred subset of four outcomes (again out of four sets). Finally, 24% of the participants were classified as "effort minimizers or ignorers" due to their neglect of relevant information, whereas the remaining 11% were classified as "strategy mixers" due to the inability to identify any consistent strategy across the different trees. An idealized illustration of some of these different strategies is provided in Fig. 4. This diversity in classifications shows that there is no one-size-fits-all account for people's failure to implement a backward induction strategy. Nevertheless, the results also show that this strategy is successfully implemented by a large proportion of individuals, even though the decision problems at hand are somewhat complex.

In another series of studies, Hey and Lotito (2009) and Hey and Panaccione (2011) focused on the way in which people handle their own dynamic inconsistencies when contrasting their planned and actual choices. Three hypotheses have been proposed: (1) people are "naïve," essentially ignoring the inconsistency, (2) "resolute," such that the realization of an inconsistency does not affect their choices, or (3) "sophisticated," acting in a way that resolves the inconsistency (see Hammond, 1988; Machina, 1989; McClennen, 1990). In order to observe cases of dynamic inconsistency,

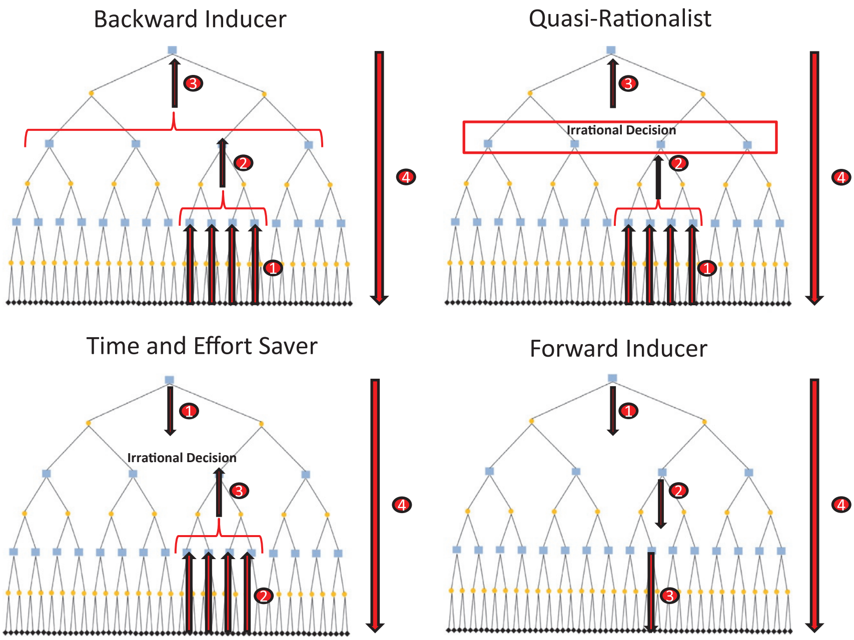

Fig. 4 Illustration of some of the strategies detected by Hey and Knoll (2011). Backward Inducers: Participants that began by evaluating the expected utilities of the branches of third-level decision nodes (step 1), and moved backwards to the first decision node (steps 2–3), and subsequently implemented their plan (step 4). Quasi-Rationalists: Participants that behaved just like the Backward-Inducer, but failed to disregard outcomes that would be eliminated at the second-level decision nodes, leading to asub- optimal plan. Time and Effort Savers: Participants that began by making a first-level decision in order to reduce the total number of outcome nodes that are to be considered. After this first decision, they worked backwards from this selected subset of outcome nodes. Forward Inducers: Participants that began by making a first-level decision that reduced the total number of outcomes from 64 to 16. A subsequent second-level decision further reduces this subset to only four outcome nodes.

Hey and Lotito (2009) relied on decision trees with outcomes and probabilities tailored to produce a preference pattern that violates expected utility theory (EUT) by producing a common ratio effect (see Cubitt, Starmer, & Sugden, 1998; Hey & Paradiso, 2006). Let us begin by considering the decision trees 1 and 2 illustrated in Fig. 5: According to EUT, people should prefer branch "left" over "right" in Tree 1's DN1 *if and only if* they prefer branch "right" over "left" in Tree 2. In other words, EUT cannot accommodate people that choose "right" or "left" in both trees. Based on previous empirical work, a reasonable number of people preferring "right" in both cases—which violates EUT—is expected (see Starmer, 2000).

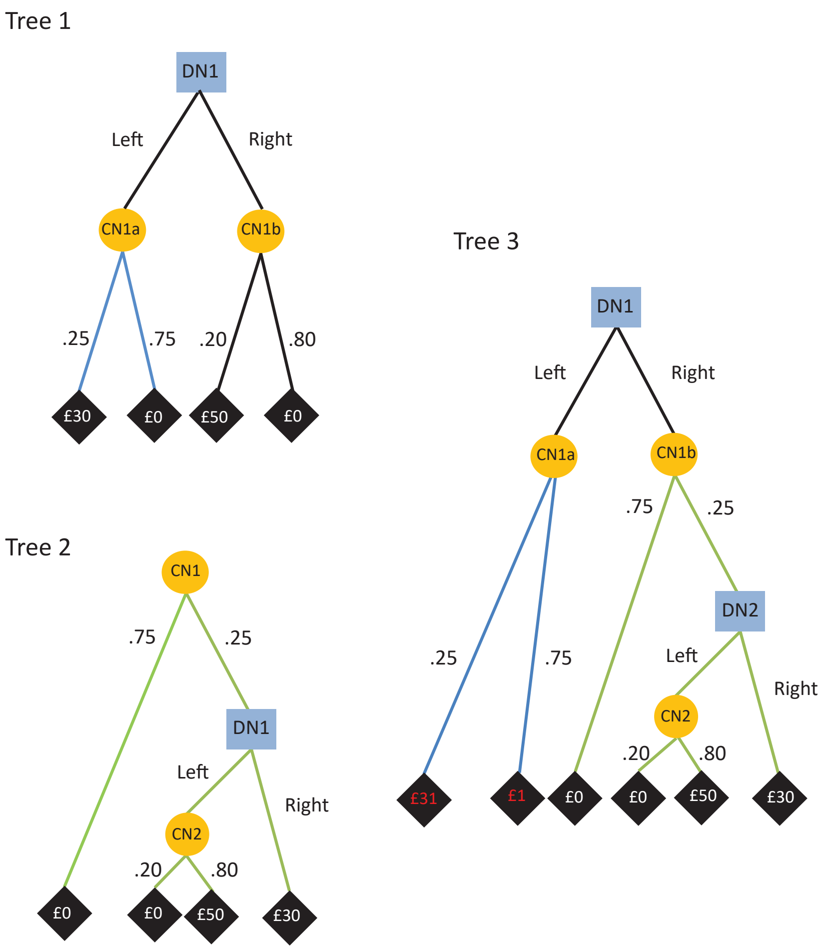

Fig. 5 Illustration of some of the sets of decision trees used by Hey and Lotito (2009). Some of the branches are colored differently (green and blue) in order to highlight the structural similarities between the trees 1/2 and 3. Some of the outcomes are colored differently (red) in order to highlight the dominance relationship between tree branches in Tree 1 and Tree 3.

Let us now turn to Tree 3: First, note that the expected utility of choosing "right" at DN1 and "left" at DN2 in Tree 3 is equivalent to choosing "right" in Tree 1's DN1 (note that the probability of earning £50 in both cases is $0.25 * 0.80 = 0.20$). Furthermore, note that the alternative of choosing "left" in Tree 3's DN1 dominates "left" in Tree 1's DN1. What this means is that if someone prefers "right" in Tree 3's DN1, it follows that they also prefer "right" in Tree 1's DN1. From the standpoint of DN1, choosing "right" at DN1 and "left" at DN2 is equivalent to choosing "right" in Tree 1. But from the standpoint of DN2, which the decision maker could eventually reach, the two options available are equivalent to the "left" and "right" options from Tree 2's DN1. And as stated earlier, people tend to prefer "right" over "left" in Tree 2. This situation introduces a potential conflict, in the sense that someone that plans to choose "right" and "left" in Tree 3's DN1 and DN2 respectively, is expected to not implement it later on, when actually reaching DN2. This situation is a case of dynamic inconsistency.

The three different kinds of people described above can be identified by the way that they value the different trees in terms of monetary bids (see Hey & Lee, 2005; Starmer & Sugden, 1991), and what they choose at each decision node. Let us consider the case of Tree 3: A naïve person will make their first choice at DN1 by evaluating the subjective utility of the two options. In contrast, a resolute person will essentially compare the expected utility of three options, (i) "right" in DN1 and "right" in DN2, (ii) "right" in DN1 and "left" in DN2, and (iii) "left" in DN1. But both naïve and resolute people agree in their evaluation of Tree 3, which is based on expected utilities of the three options above. Finally, a sophisticated person will engage in backward induction and first evaluate the expected utility of the two options at DN2. If they prefer "left" to "right" at DN2, then the choice at DN1 is ultimately between the expected utilities of "right-left" and "left." But if they prefer "right" at DN2, then the evaluation is based on a comparison between the expected utilities of "right-right" and "left." These differences will also determine the evaluation of Tree 3, which will be reflected in the bids made. People were classified into the three types using three different non-EUT models that instantiated the different ways in which "noisy" choices and evaluations are expected to take place under these types. The non-EUT models assumed a rank-dependent non-linear weighting of probabilities (e.g., Tversky & Kahneman, 1992) and individual classification was determined by the relative fit performance of these models (i.e., how likely were the data under each; for similar model-comparison approaches, see Glöckner, 2009; van den Berg, Awh, & Ma, 2014). Out of 50 participants, 25 were deemed to be (most likely) naïve, 20 were resolute, and only 5 were sophisticated. When focusing on the 28 participants whose probability weighting significantly differed from EUT (at 0.01 significance level), 15 were found to be naïve, 12 resolute, and only 1 sophisticated. A follow-up study by Hey and Panaccione (2011) used a somewhat different experimental procedure, and found that among classifiable individuals deviating from EUT, most of them were resolute (72%). But once again very few people were deemed to be sophisticated (9%). This discrepancy suggests some difficulties in classifying people in ways that generalize beyond specific experimental procedures. However, there is no indication that the lack of sophisticated individuals is somehow an experimental artifact.

Altogether, these results from Hey and colleagues provide a rich characterization of individual differences. We see that there is no one-size-fits-all account, with people widely varying in the way that they plan ahead, how they reason through decision tree problems, and in the way that they resolve potential conflicts between plans and immediate action. One of the merits of this approach is that the models used are essentially agnostic about the underlying processes. One can make this argument even in the cases of Hey and Lotito (2009) and Hey and Panaccione (2011), in which non-linear probability weighting functions were assumed. The reason being that such weightings can be emulated by processes such as the dynamic deliberation proposed by DFT (see Kellen et al., 2020). This agnosticism means that the resulting classifications do not depend on the success of any particular theory. Of course, this does not mean that these classifications cannot be revised, especially if a different, perhaps more fine grained-classification schemes (supported by the appropriate experimental designs) are developed.

3. Candidate explanations and future directions

At a glance, the above research paints a bleak picture of people's abilities to cope with the complexities of dynamic decision making. Numerous violations of dynamic inconsistency show that people cannot, or at least often fail to accurately predict their future choices. Likewise, demonstrations of consequential inconsistency show people's irrational reliance on past events to guide future actions. These findings comprehensively rule out the theory of backward induction. How then are we to reconcile these deficiencies with people's reasonably effective performance with both real world and laboratory tasks involving dynamic, multistage choice? If people lack the ability to properly implement normative strategies, what strategies do they use? In the following section, we review alternative accounts aimed at describing and explaining behavior in dynamic decision tasks.

3.1 Shifting points of reference

One possible explanation for dynamic and consequential inconsistencies relates to reference points. It is common in theories of decision making to interpret outcomes relative to some reference. That is, rather than judging an outcome in absolute terms, we compare it to a relevant value. In prospect theory (Kahneman & Tversky, 1979) the decision maker's current state of wealth is often used as a reference point, with outcomes that increase wealth (i.e., gains) being treated differently from outcomes that decrease wealth (i.e., losses). With this mechanism, Tversky and Kahneman (1981) accounted for framing effects, whereby preferences appear to change depending on whether outcomes are framed as gains or as losses.

Barkan and Busemeyer (1999, 2003) applied a similar logic to explain inconsistencies in dynamic decisions. To account for their pattern of results—in which decision makers became more risk-averse after experiencing a win, and more risk-seeking after experiencing a loss—they used a modified version of prospect theory to test two alternative explanations. According to the first, people do not consider the outcome of the fist gamble when making a prediction about the future (despite being asked to imagine winning or losing). However, after experiencing the outcome of the first gamble, they incorporate this into the value of the second gamble, and therefore experience a shift in reference point. For example, after receiving a gain of \$100 from the first gamble, this amount is added to outcomes of the second gamble. This amounts to a shift of reference point that moves the second gamble's gain further away, making it less attractive due to the diminishing marginal utility represented in the S-shaped utility function. The opposite is true when the first gamble is a failure. Receiving a loss of −\$100 shifts the reference point in the opposite direction, with a gain of \$200 now treated at one of only \$100. By moving the gain of the second gamble closer to reference point (where the utility function is more linear), the gamble becomes more attractive, and can have a higher subjective utility than was used to initially predict choice.

Barkan and Busemeyer (1999, 2003) compared this model to an alternative that draws from the *gambler's fallacy*, whereby a recent success increases the likelihood of future failure, and vice versa. That is, after experiencing the outcome of the first gamble, people's subjective probabilities change. Following a success, the likelihood of future success decreases, while it increases after a failure. Barkan and Busemeyer found that the reference point model was better able to explain their findings compared to the subjective probability model.

Johnson and Busemeyer (2001) proposed a model that also uses changes in psychological distance to explain dynamic inconsistency. Their explanation rests on the goal-gradient hypothesis instantiated in decision field theory (Busemeyer & Townsend, 1993), which in turn draws from the approach-avoidance conflict theory of Lewin (1935) and Miller, (1944). According to the model, choices depend on an approach tendency—which is determined by potential gains—and an avoidance tendency—which is determined by potential losses. Both components are dynamic and grow in strength as the distance between the current state and the final decision decreases. However, the rate of increase differs; at a distance, the approach tendency is greater than the avoidance tendency, but this reverses as the decision maker moves closer to the final DN. That is, when thinking ahead—with final consequences far in the future—people are primarily motivated to choose a sequence of actions that will maximize their eventual rewards. In contrast, when making a final decision, potential losses are now more salient than gains. As a result, people are primarily motivated to choose the action that minimizes their potential losses.

3.2 Mental simulation

Another common theme in psychological theories of planning and dynamic decision making is that of *mental simulation*. These theories posit that individuals deliberate by imagining the future consequences of their choices. Typically, they compare the outcomes resulting from each alternative to determine the best course of action. An early example comes from Kahneman and Tversky's (1982) *simulation heuristic*, which can be seen as a version of the availability heuristic, applied to multistage problems. According to the heuristic, when reasoning about counterfactuals, people mentally simulate alternative sequences of events, and determine the likelihood of an event based on how easy it is to imagine. Kahneman and Tversky used the simulation heuristic to explain why some scenarios were judged to produce more regret than others. They proposed that "near misses"—e.g., scenarios where small changes could have produced a better outcome—were judged (or predicted) to produce more regret because it was easier to "undo" events and actions, and image an alternative course of events.

In the domain of naturalistic decision making—where the aim is to study decision making in applied, rather than laboratory, settings—Klein and colleagues (Klein, 1993; Nemeth & Klein, 2010) have proposed the *recognition-primed decision* model to describe how people use mental simulations to make high-stakes decisions in dynamic environments under extreme time pressure. For example, a fireground commander will have experienced many urban fire scenarios in their career and will use this experience to decide how to fight a new fire. According to the model, they make split-second, high-risk decisions by matching the features of their current situation to patterns observed in previously experienced scenarios. Recognizing and matching these patterns activates memory for previous chosen actions, which the commander mentally simulates. If this simulation produces a satisfactory outcome, the choice is made without comparison to any alternatives.

There is also mounting neuroscientific evidence that mental simulation is key for certain kinds of decision making. For instance, Suzuki et al. (2012) conducted an fMRI study where participants were found to recruit their own mental processes to imagine and predict the choices of others. Animal studies—where rats learn to navigate branching mazes in search of food—have found evidence of mental simulation in hippocampal place cells. These cells preferentially code for location and form a neural map of the paths the rats have learned. Johnson and Redish (2007) and Dragoi and Tonegawa (2011) found that when rats paused at a junction in the maze, where they must choose which direction to move, they showed a pattern of "forward sweeping" neural activity whereby they imagined moving forward in one direction and then the other. The authors interpreted this as recruitment of hippocampal place cells to simulate the outcome of moving down one path vs. the other. Building on this work, Pezzulo and colleagues (Chersi & Pezzulo, 2012; Pezzulo, Rigoli, & Chersi, 2013) developed a reinforcement learning model that combines a model-free mechanism for habitual behavior with a model-based component that uses mental simulation to virtually "walk in the hippocampus" and to drive goal-directed choice.

In machine learning, simulation-based methods have been applied to various decision problems. For example, Monte-Carlo tree search is a heuristic method for deriving optimal choices by randomly simulating paths through multistage decision trees (Browne et al., 2012). These methods have been particularly successful in the domain of game playing (e.g., Lee, Müller, & Teytaud, 2010; Pepels, Winands, & Lanctot, 2014), where they often match, or even exceed, the performance of humans. For instance, Monte-Carlo tree search programs can compete with and beat champion Go players (see Silver et al., 2016), which shows that simulation can be the foundation of a powerfully effective strategy for making choices in complex dynamic decision environments.

In additional to research showing the importance of mental simulation in decision making and reasoning, there is also a growing body of evidence showing that people naturally think about complex choice scenarios using decision tree-like structures. This is especially true in causal reasoning, where individuals construct mental models representing the interdependence of events, actions, and outcomes (Sloman, 2005). People have also been found to use tree-like subjective representations across a diverse range of 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), learning (Gopnik et al., 2004; Tenenbaum, Griffiths, & Kemp, 2006), and the interpretation of causal language (Sloman, Barbey, & Hotaling, 2009; Walsh & Sloman, 2011). The effectiveness of this approach has even led to the prescriptions of so called fast-and-frugal decision trees for effective medical decision making (Gigerenzer & Kurzenhaeuser, 2005), and as a means of achieving good performance in signal detection tasks (Luan, Schooler, & Gigerenzer, 2011).

3.3 Decision field theory-planning

Hotaling (2020) observed that mental simulation can be seen as a natural extension of the process underlying many sequential sampling models (SSMs) of decision making. When deliberating over a single choice, SSMs posit that decision makers sample information relating to the consequences of choosing each alternative. For instance, models like DFT (Busemeyer & Townsend, 1993) and the leaky competing accumulator (Usher & McClelland, 2004) posit that individuals draw samples from the possible outcomes associated with each option. Outcomes are compared across options, and the process repeats until evidence for one option accumulates to a predetermined decision threshold. Essentially, these models explain choices in terms of a stochastic process whereby decision makers repeatedly imagine and compare the outcomes of each choice alternative.

Decision field theory-planning (DFT-P; Hotaling, 2020) applies the same logic to complex multistage decisions, positing that people use mental simulations to think through dynamic decision scenarios. As they deliberate over which action to take at present, they form a mental model of their situation, and imagine possible sequences of events. With each simulation, they trace a path through the decision tree from their current DN to a final ON. At each time-step in the decision process, the model compares the simulated outcome of each alternative in order to compute a momentary valence. These valences accumulate over time, until preference for one option exceed its decision threshold.

An important aspect of DFT-P pertains to the simulation of one's own decisions. While mentally tracing a path through a multistage tree, the decision maker must imagine the choice they will make upon reaching a future DN. The model therefore proposes that individuals plan future choices on-the-fly by repeatedly simulating their own future actions as they imagine and compare outcomes. Simulations differ from moment to moment as the decision maker attends to different possibilities. This contrasts sharply with the prescriptions of backward induction, wherein individuals work backward through the tree and commit to a choice at each future DN. Instead, DFT-P describes a process in which individuals do not make explicit plans for future DNs, but rather use repeated forward-looking mental simulations to imagine what they will do, resulting is probabilistic predictions of future choices.

Returning to the earlier example from Fig. 1, when Emma imagines the consequences of asking for a raise (CN1), she does so by simulating a sequence of events. She might first imagine that her boss will say "Yes" and will offer a choice between company stock and increased wages. Next, she simulates her future choice, and envisions choosing the stock option. Because she knows that the stock market was been strong recently, she imagines that this will yield a large payoff (e.g., ON3). Having now simulated a potential outcome of asking for a raise, she compares this outcome to the certain outcome of not asking for a raise (ON1). Assuming ON3 > ON1, this comparison produces evidence in favor of the riskier option and causes Emma's preference to evolve toward asking for a raise. However, at the next moment Emma may imagine a less fortuitous sequence of events. As her attention shifts to other features of the decision problem, she now imagines that her boss will reject her request, leaving her to decide whether to quit her job (DN2b). She imagines that she will quit and will then struggle to find work (ON7) due to a bad job market. Emma now compares this new simulated outcome to ON1. Assuming ON7 < ON1, this produces evidence pushing Emma's preference in the opposite direction, toward not asking for a raise. In this fashion, she repeatedly simulates outcomes according to her subjective beliefs about the likelihoods of various sequences of events.

In contrast to the normative approach of backward induction, DFT-P offers an explanatory theory of planning, wherein people form structured mental models of complex decision scenarios, and use mental simulation to predict, plan, and execute their choices. Whereas backward induction provides an optimal method for maximizing payoffs, given unlimited time and cognitive resources, DFT-P posits several important limitations. First, the decision maker is thought to have limited attention, and must therefore focus on the most important or salient outcomes. Second, the decision maker cannot perfectly predict their future actions, and must instead use mental simulation to imagine what they will likely choose in the future. Third, the decision maker does not exhaustively process all relevant information, but rather accumulates only enough evidence to meet their decision threshold. Thus, DFT-P provides a formal account of dynamic decision making in keeping with the principles of bounded rationality (Simon, 1982).

3.4 Extending model-based classifications

From a methodological perspective, there is considerable potential in extending the models used by Hey and colleagues. A careful look at these models, such as the true-and-error model used by Hey and Knoll (2007) reveals that they are members of the class of Multinomial Processing Tree (MPT) models (for a review, see Batchelder & Riefer, 1999). The class of MPT models has been extended in multiple ways such that one can now incorporate many kind of covariates (e.g., individuals' working-memory capacity) at the level of the probability parameters (e.g., parameters p, e1, and e2; see Klauer, 2010) as well as different kinds of outcome variables, such as response latencies and "process data" (e.g., mouse behavior; see Heck & Erdfelder, 2016; Heck, Erdfelder, & Kieslich, 2018; Klauer & Kellen, 2018). Given the modeling toolboxes currently available (e.g., Hartmann, Johannsen, & Klauer, 2020; Heck, Arnold, & Arnold, 2018; Singmann & Kellen, 2013), it would be straightforward to develop extended models that classified individuals in terms of their behavior (beyond choices) and relate these classifications with different kinds of individual-level covariates.

Also of note is the fact that these classifications can provide an important complement to ongoing efforts in the development of theoretical accounts such as DFT-P. For instance, the classifications could be used in the development of DFT-P accounts for the different kinds of individuals (see Hotaling & Busemeyer, 2012). Because these classifications are agnostic in the sense that they do not hinge on specific "processes" or "mechanisms," their value is preserved regardless of the theoretical account that is being considered (e.g., DFT-P or some future rival).

4. Conclusion

Dynamic decision problems are ubiquitous, from the mundane and every-day task of planning a route to work, to the complex and life-altering endeavor of planning one's career. In this article, we review two lines of research investigating how people handle the challenges of dynamic decision making. The first focuses on comparing human behavior to the prescriptions of a rational theory—*backward induction*—and uses deviations from this received view to motivate the development of precise cognitive process models to account for deviations. The second instead uses individual differences in dynamic decision making to make more coarse-grained and intuitive classifications of people into qualitatively distinct types (e.g., naïve vs. sophisticated). Altogether, these lines of research showcase the wealth of data that awaits researchers when stepping beyond static decision environments. We place dynamic decision making at center stage to promote and encourage the development of better, more complete, and encompassing theories. Among currently available accounts of decision making under risk, Decision Field Theory—and the dynamic mental simulation process that it postulates—offers a particularly promising candidate.

References

Anderson, P. L. (2009). The value of private businesses in the United States. *Business Economics*, *44*(2), 87–108. https://doi.org/10.1057/be.2009.4.

Barkan, R., & Busemeyer, J. R. (1999). Changing plans: Dynamic inconsistency and the effect of experience on the reference point. *Psychonomic Bulletin & Review*, *6*(4), 547–554. https://doi.org/10.3758/BF03212962.

Barkan, R., & Busemeyer, J. R. (2003). Modeling dynamic inconsistency with a changing reference point. *Journal of Behavioral Decision Making*, *16*(4), 235–255. https://doi.org/10.1002/bdm.444.

Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review*, *6*(1), 57–86. https://doi.org/10.3758/BF03210812.

Berkowitsch, N. A. J., Scheibehenne, B., & Rieskamp, J. (2014). Rigorously testing multialternative decision field theory against random utility models. *Journal of Experimental Psychology. General*, *143*(3), 1331–1348. https://doi.org/10.1037/a0035159.

Bertsekas, D. P. (2017). *Dynamic programming and optimal control*. Athena Scientific.

Birnbaum, M. H. (2013). True-and-error models violate independence and yet they are testable. *Judgment and Decision making*, *8*(6), 717–737.

Birnbaum, M. H., & Schmidt, U. (2010). Testing transitivity in choice under risk. *Theory and Decision*, *69*(4), 599–614. https://doi.org/10.1007/s11238-009-9147-1.

Bott, F. M., Kellen, D., & Klauer, K. C. (2021). Normative accounts of illusory correlations. *Psychological Review*, *128*(5), 856–878. https://doi.org/10.1037/rev0000273.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica*, *81*(3), 211–241. https://doi.org/10.1016/0001-6918(92)90019-A.

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

Busemeyer, J. R., & Goldstein, W. M. (1992). Linking together different measures of preference: A dynamic model of matching derived from decision field theory. *Organizational Behavior and Human Decision Processes*, *52*(3), 370–396. https://doi.org/10.1016/0749-5978(92)90026-4.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*(3), 432–459. https://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*(4), 530–545. https://doi.org/10.1037/0096-3445.129.4.530.

Chersi, F., & Pezzulo, G. (2012). Using hippocampal-striatal loops for spatial navigation and goal-directed decision-making. *Cognitive Processing*, *13*(Suppl. 1), 125–129. https://doi.org/10.1007/s10339-012-0475-7.

Costello, F., & Watts, P. (2017). Explaining high conjunction fallacy rates: The probability theory plus noise account. *Journal of Behavioral Decision Making*, *30*(2), 304–321. https://doi.org/10.1002/bdm.1936.

Coulter, J. (1983). *Rethinking cognitive theory*. St. Martin's Press.

Crupi, V., Fitelson, B., & Tentori, K. (2008). Probability, confirmation, and the conjunction fallacy. *Thinking & Reasoning*, *14*(2), 182–199. https://doi.org/10.1080/13546780701643406.

Cubitt, R., Starmer, C., & Sugden, R. (1998). Dynamic choice and the common ratio effect: An experimental investigation. *The Economic Journal*, *108*(450), 1362–1380. https://doi.org/10.1111/1468-0297.00346.

Cummins, R. (1983). *The nature of psychological explanation*. MIT Press.

Cummins, R. (2000). "How does it work?" vs. "What are the laws?" two conceptions of psychological explanation. In F. Keil, & R. Wilson (Eds.), *Explanation and cognition* (pp. 117–145). MIT Press.

DeGroot, M. H. (1970). *Optimal statistical decisions*. McGraw-Hill.

Diederich, A. (2003). MDFT account of decision making under time pressure. *Psychonomic Bulletin and Review*, *10*(1), 157–166. https://doi.org/10.3758/BF03196480.

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

Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, *469*(7330), 397–401. https://doi.org/10.1038/nature09633.

Fernbach, P. M., Darlow, A., & Sloman, S. A. (2010). Neglect of alternative causes in predictive but not diagnostic reasoning. *Psychological Science*, *21*(3), 329–336. https://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*(2), 168–185. https://doi.org/10.1037/a0022100.

Fudenberg, D., & Tirole, J. (1991). *Game theory*. MIT press.

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

Glöckner, A. (2009). Investigating intuitive and deliberate processes statistically: The multiple-measure maximum likelihood strategy classification method. *Judgment and Decision making*, *4*, 186–199.

Gonzalez, C., Fakhari, P., & Busemeyer, J. (2017). Dynamic decision making: Learning processes and new research directions. *Human Factors*, *59*(5), 713–721. https://doi.org/10.1177/0018720817710347.

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*(1), 3–32. https://doi.org/10.1037/0033-295X.111.1.3.
Hacker, P. M. S. (2013). *The intellectual powers: A study of human nature*. John Wiley & Sons.
Hammond, P. J. (1988). Consequentialist foundations for expected utility. *Theory and Decision*, *25*(1), 25–78. https://doi.org/10.1007/BF00129168.
Harré, R., & Secord, P. F. (1973). *The explanation of social behaviour*. Blackwell.
Hartmann, R., Johannsen, L., & Klauer, K. C. (2020). rtmpt: An R package for fitting response-time extended multinomial processing tree models. *Behavior Research Methods*, *52*, 1313–1338.
Heck, D. W., Arnold, N. R., & Arnold, D. (2018). TreeBUGS: An R package for hierarchical multinomial-processing-tree modeling. *Behavior Research Methods*, *50*, 264–284.
Heck, D. W., & Erdfelder, E. (2016). Extending multinomial processing tree models to measure the relative speed of cognitive processes. *Psychonomic Bulletin & Review*, *23*, 1440–1465.
Heck, D. W., Erdfelder, E., & Kieslich, P. J. (2018). Generalized processing tree models: Jointly modeling discrete and continuous variables. *Psychometrika*, *83*, 893–918.
Hey, J. D. (2005). Do people (want to) plan? *Scottish Journal of Political Economy*, *52*(1), 122–138. https://doi.org/10.1111/j.0036-9292.2005.00338.x.
Hey, J. D., & Knoll, J. A. (2007). How far ahead do people plan? *Economics Letters*, *96*(1), 8–13. https://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 behaviour. *Journal of Economic Psychology*, *32*(3), 399–409. https://doi.org/10.1016/j.joep.2011.02.011.
Hey, J. D., & Lee, J. (2005). Do subjects separate (or are they sophisticated)? *Experimental Economics*, *8*(3), 233–265. https://doi.org/10.1007/s10683-005-1465-8.
Hey, J. D., & Lotito, G. (2009). Naive, resolute or sophisticated? A study of dynamic decision making. *Journal of Risk and Uncertainty*, *38*(1), 1–25. https://doi.org/10.1007/s11166-008-9058-5.
Hey, J. D., & Panaccione, L. (2011). Dynamic decision making: What do people do? *Journal of Risk and Uncertainty*, *42*(2), 85–123. https://doi.org/10.1007/s11166-011-9114-4.
Hey, J. D., & Paradiso, M. (2006). Preferences over temporal frames in dynamic decision problems: An experimental investigation. *The Manchester School*, *74*(2), 123–137. https://doi.org/10.1111/j.1467-9957.2006.484_1.x.
Hilbig, B. E., & Moshagen, M. (2014). Generalized outcome-based strategy classification: Comparing deterministic and probabilistic choice models. *Psychonomic Bulletin and Review*, *21*(6), 1431–1443. https://doi.org/10.3758/s13423-014-0643-0.
Hotaling, J. M. (2020). Decision field theory-planning: A cognitive model of planning on the fly in multistage decision making. *Decision*, 7(1), 20–42. https://doi.org/10.1037/dec0000113.
Hotaling, J. M., & Busemeyer, J. R. (2012). DFT-D: A cognitive-dynamical model of dynamic decision making. *Synthese*, 67–80. https://doi.org/10.1007/s11229-012-0157-0.
Johnson, J. G., & Busemeyer, J. R. (2001). Multiple-stage decision-making: The effect of planning horizon length on dynamic consistency. *Theory and Decision*, *51*(2–4), 217–246. https://doi.org/10.1023/A:1015507110940.
Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, *27*(45), 12176–12189. https://doi.org/10.1523/JNEUROSCI.3761-07.2007.
Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–292. http://www.jstor.org/stable/1914185%5Cnhttp://www.jstor.org/%5Cnhttp://www.jstor.org/action/showPublisher?publisherCode=econosoc.%5Cnhttp://www.jstor.org.

Kahneman, D., & Tversky, A. (1982). The simulation heuristic. In *Judgment under uncertainty* (pp. 201–208). Cambridge University Press. https://doi.org/10.1017/CBO9780511809477.015.

Keeney, R. L., & Raiffa, H. (1976). *Decisions with multiple objectives: Preferences and value tradeoffs*. Wiley & Sons.

Kellen, D. (2019). A model hierarchy for psychological science. *Computational Brain & Behavior*, *2*(3–4), 160–165. https://doi.org/10.1007/s42113-019-00037-y.

Kellen, D. (in press). Behavioral decision research is not a Linda problem: Comment on Regenwetter, Robinson, and Wang (in press). Decision.

Kellen, D., Steiner, M. D., Davis-Stober, C. P., & Pappas, N. R. (2020). Modeling choice paradoxes under risk: From prospect theories to sampling-based accounts. *Cognitive Psychology*, *118*(November 2019), 101258. https://doi.org/10.1016/j.cogpsych.2019.101258.

Klauer, K. C. (2010). Hierarchical multinomial processing tree models: A latent-trait approach. *Psychometrika*, *75*, 70–98.

Klauer, K. C., & Kellen, D. (2018). RT-MPTs: Process models for response-time distributions based on multinomial processing trees with applications to recognition memory. *Journal of Mathematical Psychology*, *82*, 111–130.

Klein, G. A. (1993). A recognition-primed decision (RPD) model of rapid decision making. *Decision Making in Action: Models and Methods*, 138–147. https://doi.org/10.1002/bdm.3960080307.

Krynski, T. R., & Tenenbaum, J. B. (2007). The role of causality in judgment under uncertainty. *Journal of Experimental Psychology. General*, *136*(3), 430–450. https://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*(4), 225–228. https://doi.org/10.1109/TCIAIG.2010.2099154.

Lewin, K. (1935). *A dynamic theory of personality*. McGraw-Hill.

Ljungqvist, L., & Sargent, T. J. (2012). *Recursive macroeconomic theory*. MIT Press.

Loomes, G. (2010). Modeling choice and valuation in decision experiments. *Psychological Review*, *117*, 902–924. https://doi.org/10.1037/a0019807.

Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, *92*(368), 805. https://doi.org/10.2307/2232669.

Luan, S., Schooler, L. J., & Gigerenzer, G. (2011). A signal-detection analysis of fast-and-frugal trees. *Psychological Review*, *118*(2), 316–338. https://doi.org/10.1037/a0022684.

Luce, R. D. (2010). Behavioral assumptions for a class of utility theories: A program of experiments. *Journal of Risk and Uncertainty*, *41*(1), 19–37. https://doi.org/10.1007/s11166-010-9098-5.

Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, *27*(4), 1622–1668.

McClennen, E. F. (1990). Foundational explorations for a normative theory of political economy. *Constitutional Political Economy*, *1*(1), 67–99.

Meder, B., Hagmayer, Y., & Waldmann, M. R. (2009). The role of learning data in causal reasoning about observations and interventions. *Memory & Cognition*, *37*(3), 249–264. https://doi.org/10.3758/MC.37.3.249.

Miller, N. E. (1944). Experimental studies of conflict. In J. M. Hunt (Ed.), *1. Personality and the behavior disorders* (pp. 431–465). New York: The Ronald Press.

Navarro, D. J. (2019). Between the devil and the deep Blue Sea: Tensions between scientific judgement and statistical model selection. *Computational Brain & Behavior*, *2*(1), 28–34. https://doi.org/10.1007/s42113-018-0019-z.

Nemeth, C., & Klein, G. (2010). The naturalistic decision making perspective. In *Wiley encyclopedia of operations research and management science* (pp. 1–9). Wiley & Sons.
Oravecz, Z., Muth, C., & Vandekerckhove, J. (2016). Do people agree on what makes one feel loved? A cognitive psychometric approach to the consensus on felt love. *PLoS One*, *11*(4), e0152803. https://doi.org/10.1371/journal.pone.0152803.
Pepels, T., Winands, M. H. M., & Lanctot, M. (2014). Real-time Monte Carlo tree search in Ms Pac-Man. *IEEE Transactions on Computational Intelligence and AI in Games*, *6*(3), 245–257. https://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*(Mar), 1–15. https://doi.org/10.3389/fpsyg.2013.00092.
Raiffa, H. (1968). *Decision analysis: Introductory lectures on choices under uncertainty*. Addison-Wesley.
Regenwetter, M., & Robinson, M. M. (2017). The construct-behavior gap in behavioral decision research: A challenge beyond replicability. *Psychological Review*, *124*(5), 533–550. https://doi.org/10.1037/rev0000067.
Riefer, D. M., Knapp, B. R., Batchelder, W. H., Bamber, D., & Manifold, V. (2002). Cognitive psychometrics: Assessing storage and retrieval deficits in special populations with multinomial processing tree models. *Psychological Assessment*, *14*(2), 184–201. https://doi.org/10.1037/1040-3590.14.2.184.
Rottman, B. M., & Hastie, R. (2014). Reasoning about causal relationships: Inferences on causal networks. *Psychological Bulletin*, *140*(1), 109–139. https://doi.org/10.1037/a0031903.
Shiffrin, R. M., & Nobel, P. (1997). The art of model development and testing. *Behavior Research Methods, Instruments, & Computers*, *29*(1), 6–14. https://link.springer.com/content/pdf/10.3758%2FBF03200560.pdf.
Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., et al. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, *529*(7587), 484–489. https://doi.org/10.1038/nature16961.
Simon, H. A. (1982). *Models of bounded rationality*. MIT Press.
Singmann, H., & Kellen, D. (2013). MPTinR: Analysis of multinomial processing tree models in R. *Behavior Research Methods*, *45*, 560–575.
Sloman, S. (2005). *Causal models: How people think about the world and its alternatives*. Oxford University Press.
Sloman, S., Barbey, A. K., & Hotaling, J. M. (2009). A causal model theory of the meaning of cause, enable, and prevent. *Cognitive Science*, *33*(1), 21–50. https://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*(3), 303–333. https://doi.org/10.1016/j.cogsci.2003.11.001.
Soyland, A. J. (1994). *Psychology and metaphor*. Sage.
Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, *38*, 332–382.
Starmer, C., & Sugden, R. (1991). American economic association does the random-lottery incentive system elicit true preferences? An experimental investigation. *The American Economic Review*, *81*(4 (September)), 97. *The American Economic Review*, *81*(4), 971–978.
Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*(3). https://doi.org/10.1016/S0364-0213(03)00010-7.
Suzuki, S., Harasawa, N., Ueno, K., Gardner, J. L., Ichinohe, N., Haruno, M., et al. (2012). Learning to simulate others' decisions. *Neuron*, *74*(6), 1125–1137. https://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*(7), 309–318. https://doi.org/10.1016/j.tics.2006.05.009.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453–458.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, *90*(4), 293–315. https://doi.org/10.1037/0033-295X.90.4.293.

Tversky, A., & Kahneman, D. (1992). Advances in prospect-theory—Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. https://doi.org/10.1007/Bf00122574.

Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, *111*(3), 757–769. https://doi.org/10.1037/0033-295X.111.3.757.

van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological Review*, *121*, 124–149. https://doi.org/10.1037/a0035234.

van Ravenzwaaij, D., Donkin, C., & Vandekerckhove, J. (2017). The EZ diffusion model provides a powerful test of simple empirical effects. *Psychonomic Bulletin & Review*, *24*(2), 547–556. https://doi.org/10.3758/s13423-016-1081-y.

van Rooij, I., Wright, C. D., Kwisthout, J., & Wareham, T. (2018). Rational analysis, intractability, and the prospects of 'as if'-explanations. *Synthese*, *195*(2), 491–510. https://doi.org/10.1007/s11229-014-0532-0.

von Winterfeldt, D., & Edwards, W. (1986). *Decision analysis and behavioral research*. Cambridge University Press.

Wakker, P. P. (2010). *Prospect theory: For risk and ambiguity*. Cambridge University Press.

Walsh, C. R., & Sloman, S. A. (2011). The meaning of cause and prevent: The role of causal mechanism. *Mind & Language*, *26*(1), 21–52. https://doi.org/10.1111/j.1468-0017.2010.01409.x.