Journal of Behavioral Decision Making, *J. Behav. Dec. Making* (2016)
Published online in Wiley Online Library (wileyonlinelibrary.com) **DOI:** 10.1002/bdm.1964

Decision Making in Environments with Non-Independent Dimensions

SUDEEP BHATIA*
[1]*University of Pennsylvania, Philadelphia, PA USA*

ABSTRACT

This paper tests whether the dimensions involved in preferential choice tasks are evaluated independently from one another. Common decision heuristics satisfy dimensional independence, and multi-strategy models that assume that decision makers use a repertoire of these heuristics predict that they are unable to represent and respond to dimensional dependencies in the decision environment. In contrast, some single-strategy models are able to violate dimensional independence, and subsequently adapt to environments that feature interacting dimensions. Across five experiments, this paper documents systematic violations of the assumption of dimensional independence. This suggests that decision makers are able to modify their behavior to respond to dimensional dependencies in their environment, and in turn those models that are unable to do this do not provide a full account of human strategy selection and behavior change. This paper ends with a discussion of ways in which some existing models can be modified to incorporate violations of dimensional independence. Copyright © 2016 John Wiley & Sons, Ltd.

KEY WORDS multi-attribute choice; decision making; independence; common-consequence effect

INTRODUCTION

How do people make decisions, and how do these decisions vary across environments with differing statistical structures and reward characteristics? This difficult but important problem has been the topic of considerable inquiry in decision-making research. Many scholars in this field have argued that decision makers rely on a repertoire of distinct strategies, with different strategies performing differently in different environments (Gigerenzer, Todd, and the ABC Research Group, 1999; Payne, Bettman, & Johnson, 1993; also Newell & Simon, 1972; Rieskamp & Otto, 2006). These multi-strategy approaches propose that decision makers are often able to select strategies that lead to the highest accuracy with the lowest costs. Thus, for example, when the attributes, or other sources of information, involved in a particular domain are highly correlated, simple strategies that utilize only one attribute yield accurate, robust, and relatively effortless decisions and subsequently are most likely to be used by decision makers (e.g., Davis-Stober, Dana, & Budescu, 2010).

Some scholars in decision-making research have, in contrast, argued that decision makers utilize only one general process (Bhatia, 2013, in press; Glöckner & Betsch, 2008; Lee & Cummins, 2004; Newell, 2005; Roe, Busemeyer, & Townsend, 2001). Unlike the aforementioned multi-strategy approaches, this single process does not vary substantially across different environments. Rather, its parameters are assumed to be able to adjust and adapt to the statistical structures and reward characteristics involved in a given domain.

The existence of different environments with different characteristics generates the research goal of understanding the (meta-decision) processes that people use to choose between strategies, in the case of multi-strategy models, or adjust parameters in different environments, in the case of single-strategy models (see Marewski & Link, 2014, for a discussion). This strategy selection problem, which is the topic of this journal's special issue, is key to developing a full characterization of human decision making. Without specifying how people select different strategies or adapt a single process to best respond to the environment in consideration, both multi- and single-strategy decision-making models will remain incomplete. Understanding these processes is also necessary for comparing multi-strategy models with single-strategy models. These two approaches propose radically different ways of modifying behavior to suit the characteristics of the environment, and evaluating the predictions of these approaches involves comparing how well the strategy selection processes assumed by these approaches are able to respond to changes in environmental structure, and in turn, how well these approaches are able to predict behaviors across different types of decision environments.

This paper attempts to shed light on strategy selection processes, and in turn the relative desirability of different types of decision models, by examining preferential choice environments with dependencies between the dimensions that characterize choice alternatives (Fishburn & Wakker, 1995; Keeney & Raiffa, 1993). In these environments, changing some dimensions of a choice alternative can affect how other dimensions of that alternative are evaluated, and subsequently how likely the decision maker is to choose the alternative. Consider, for example, a late-night snack consisting of cookies and milk. Cookies are particularly tasty when eaten with milk, and replacing milk with a different beverage (e.g., beer) can affect how desirable the cookies appear to the decision maker, and subsequently the decision maker's evaluation of the late-night snack.

These types of dependencies appear to be quite intuitive and are, in fact, an important component of the economic modeling of preferential choice (e.g., Rubinstein, 2012; Varian, 1992). However, multi-strategy models of decision making, which typically utilize heuristic rules, are unable to generate these dependencies. These heuristics include the lexicographic rule (Fishburn, 1974; Tversky, 1969), in which decisions are made on the basis of only one dimension; the

*Correspondence to: Sudeep Bhatia, Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA. E-mail: bhatiasu@sas.upenn.edu

equal weights rule (Dawes, 1979), in which underlying dimensions are aggregated without differential weighting; the tallying rule, in which binary pairwise comparisons on different dimensions determine the decision maker's response (Alba & Marmorstein, 1987; Russo & Dosher, 1983); and the weighted-additive rule, in which decision makers make choices based on the weighted values of the alternatives on their various dimensions (Keeney & Raiffa, 1993). All of these individual strategies assume that dimensions are evaluated separately from each other and in turn predict that preferences satisfy dimensional independence. As a result, multi-strategy decision-making approaches, which propose that decision makers use one or the other of these strategies (Gigerenzer & Todd, 1999; Payne et al., 1993; Rieskamp & Otto, 2006), also predict that preferences satisfy dimensional independence.

Single-strategy models of decision making, such as sequential sampling models (Bhatia, 2013; Diederich, 1997; Lee & Cummins, 2004; Newell, 2005; Roe et al., 2001; Usher & McClelland, 2004), usually make similar predictions. These models often assume the independent evaluation of underlying dimensions, which makes it impossible for these models to capture interdependencies such as the ones discussed earlier. One exception is the multi-attribute linear ballistic accumulator (MLBA; Trueblood, Brown, & Heathcote, 2014), which does not make this assumption and thus can accommodate such dependencies. There are also other single-strategy models, such as those relying on neural networks (Glöckner & Betsch, 2008; Holyoak & Simon, 1999), those utilizing configural weighting (Birnbaum, Parducci, & Gifford, 1971; Birnbaum, 1974, 2008), and those based on exemplar memory (Bröder, Newell, & Platzer, 2010; Juslin & Persson, 2002; Juslin, Karlsson, & Olsson, 2008), that are able to violate dimensional independence. These models involve recurrent connections between the nodes representing decision alternatives and their underlying dimensions (in the case of neural network models), context-based additive terms in an otherwise linear model (in the case of configural weighting), or else similarity-based calculations with previously experienced alternatives (in the case of exemplar models), and these assumptions allow the models to learn and respond to dependencies between the various dimensions.

Are decision makers capable of aggregating the dimensions that characterize available choice alternatives in a non-independent manner? A negative answer to this question would imply that the processes involved in strategy selection (i.e., choosing decision strategies or adjusting model parameters) are limited, in that there are some types of reward functions that decision makers are unable represent, and thus settings in which strategy selection processes fail to modify behavior to best adapt to the reward structure of the environment. However, these limitations would be consistent with most multi-strategy (and some single-strategy) decision-making approaches. A positive answer to this question would, in contrast, suggest that individuals can make adaptive decisions in complex environments with non-independent dimensions, and in turn, that the strategies in the multi-strategy repertoire, and the processes involved in selecting these strategies, do not provide a complete account of human behavior. Such an answer would also provide support for single-strategy approaches that are able to adjust their parameters to appropriately handle dimensional dependencies.

In this paper, we test for violations of independence in choices between bundles composed of different objects (Studies 1 and 2), and real and artificial objects composed of different attributes (Studies 3–5). If the dimensional values of these alternatives do not alter how other dimensions are processed, then changing values on a dimension that is common across all alternatives should not affect choice. In Studies 1, 2, and 3, we use this insight to design binary choice problems in which two bundles contain the same amount of some object, or two objects contain the same amount of some attribute. We vary this common object or attribute across choice problems and find that this affects choice proportions, violating dimensional independence. In Studies 4 and 5, we test for violations of independence with artificial choice alternatives, for which non-independent attribute–reward relationships are learnt through experience. We find that a majority of participants alter their choices as common attributes are changed, again demonstrating violations of independence. These studies also show that decision makers are able to make adaptive decisions in environments that require the aggregation of dimensions in a non-independent manner. Note that in order to simplify terminology, we use the term *strategy selection* to refer more generally to behavior change across environments. Thus, in this paper, this term refers to both the selection of different strategies by multi-strategy models and changes to parameter values in single-strategy models.

INDEPENDENCE IN DECISION MODELING

In this paper, we will limit our attention to riskless preferential choice tasks involving decisions between bundles composed of different objects or decisions between objects composed of different attributes. Both decisions between bundles and decisions between individual objects can be seen as involving choice alternatives whose components are dimensions in a multi-dimensional choice space. These types of decisions have been a topic of considerable research, and many multi-strategy and single-strategy decision models have been developed to describe the ways in which decision makers behave when offered these types of decisions.

One important property in multi-dimensional decisions is dimensional independence. Dimensional independence formally states that a pair of dimensions x and y are independent if for all x_1, y_1, x_2, and y_2 the alternative (x_1, y_1) is chosen over the alternative (x_2, y_1) if and only if (x_1, y_2) is chosen over (x_2, y_2) (Keeney & Raiffa, 1993). Intuitively, dimensions on which the alternatives are identical should be cancelled out and should not influence preferences. If one set of common dimensions is replaced by another set of common dimensions, relative preferences and choice proportions should remain the same.

DOI: 10.1002/bdm

Dimensional independence in multi-strategy models

Do people satisfy dimensional independence? There are a large range of different multi-strategy approaches currently studied in decision-making research. The ones that could be applied to the types of preferential choice tasks discussed here include the Payne et al. (1993) adaptive decision maker framework, Gigerenzer and Todd's (1999) fast and frugal framework, and Rieskamp and Otto's (2006) strategy selection learning theory. These approaches involve different meta-decision processes for strategy selection, such as cost–benefit analysis for the adaptive decision maker framework, ecological feasibility for the fast and frugal framework, and reinforcement learning for the strategy selection learning theory. However, when applied to preferential choice problems with alternatives defined on multiple dimensions, they propose that decision makers use a similar repertoire of strategies. These strategies include lexicographic decision rules, equal weights rules, tallying rules, and the weighted-additive rule, and all of these individual strategies satisfy dimensional independence.

Consider, for example, the lexicographic rule (Fishburn, 1974). This rule states that that decision makers first consider their favorite dimension. If the two available alternatives differ on this dimension, they select the alternative with the highest value on the dimension. If two alternatives are identical on this dimension, they move on to a second dimension and repeat the process again. As this rule assumes dimension-wise comparisons, dimensions are considered separately. Thus, decision makers do not think about the second dimension when evaluating the alternatives on the first dimension, and subsequently, the second dimension cannot influence their preference for the first dimension. If the first dimension is common to both alternatives, then decision makers consider the second dimension but do not base their evaluations of the second dimension on the values of the first. Indeed, it is fairly easy to verify that the dimensional independence condition specified earlier is necessarily guaranteed for the lexicographic rule for any set of preferences within and between dimensions.

The equal weight, tallying, and weighted-additive rules (Alba & Marmorstein, 1987; Dawes, 1979; Keeney & Raiffa, 1993; Russo & Dosher, 1983) also satisfy independence. This is because the preferences specified by all three of these rules can be written as weighted sums of some dimension–value transformations. For the equal weights rule, the preference for an alternative is simply $x+y$. For the tallying rule, the preference for an alternative is $V(x)+V(y)$, where $V(x) = 1$ if the alternative has a higher amount of attribute x, and $V(x)=0$ otherwise (and likewise for $V(y)$). For the weighted-additive rule, the preference for an alternative is $w_x x+w_y y$, where w_x and w_y are the weights on the two dimensions. For all three rules, the decision maker chooses one alternative if its weighted sum is higher than the weighted sum of the other alternative. This implies that dimensions that are common across both alternatives are cancelled out and thus do not influence the decision. This, as discussed earlier, generates dimensional independence.

The aforementioned four rules are the best studied in decision-making research. There are, in addition to these, variants of these rules. These variants also satisfy dimensional independence. For example, lexicographic semi-orders (Tversky, 1969), which use a difference threshold on each dimension in order to evaluate whether one alternative is better than the other on that dimension, and the CONF rule (Karelaia, 2006), which considers multiple discriminating dimensions before making a decision, are both similar to the lexicographic rule in that they consider dimensions sequentially. These two heuristics will also thus never violate dimensional independence. This is also the case for elimination-based heuristics, such as elimination by aspects and elimination by least attractive (Svenson, 1979; Tversky, 1972). Fast and frugal trees, which combine both lexicographic and elimination rules, also feature this property, as they also consider dimensions sequentially (Dhami, 2003; Luan, Schooler, & Gigerenzer, 2011).

Ultimately, as all different component strategies that could be assumed by multi-strategy decision models satisfy independence, the multi-strategy models do so as well. Thus, the adaptive decision maker framework (Payne et al., 1993) would predict that decision makers never violate independence no matter how the costs and benefits to violating independence are structured. Likewise, the fast and frugal framework (Gigerenzer & Todd, 1999) would suggest that dimensional independence should be satisfied regardless of the ecological characteristics of the environments decision makers are exposed to. Similarly, strategy selection learning theory (Rieskamp & Otto, 2006), and related reinforcement learning-based approaches, would never generate violations of independence no matter what rewards the decision maker experiences. Despite the fact that all of these approaches propose different strategy selection mechanisms, these mechanisms are all unable to recognize and adapt to environments in which choice dimensions are non-independent.

That said, it is important to note that the approaches described here are not necessarily limited to the few strategies we have outlined. These strategies are knowledge-based, in that they utilize information about the features of the choice alternatives, particularly their underlying dimensions, to differentiate between the alternatives. Some of the aforementioned multi-strategy approaches, such as the fast and frugal framework, also however involve memory-based heuristics, such as heuristics, which choose alternatives based on whether or not the alternatives are recognized by the decision maker. Relatedly, Marewski and Schooler's (2011) cognitive niche approach implements many of these memory and knowledge-based strategies and formalizes the strategy selection problem, in an ACT-R model (and in doing so, blurs the line between a multi-strategy model and a single-strategy model). There is no doubt that memory-based heuristics are at play in many decision-making scenarios. However, as this paper is limited to the simple setting with choice alternatives defined on quantifiable dimensions, it will only consider the aforementioned knowledge-based heuristics, and subsequently, it will judge the descriptive accuracy of the various multi-strategy approaches based only on the predictions their component knowledge-based heuristics make.

DOI: 10.1002/bdm

Dimensional independence in single-strategy models

Dimensional independence is a feature not only of multi-strategy approaches to decision making. Single-strategy models such as sequential sampling models (Bhatia, 2013; Diederich, 1997; Lee & Cummins, 2004; Newell, 2005; Roe et al., 2001; Usher & McClelland, 2004) also satisfy this property. Most sequential sampling models, for example, assume that decision makers attend to dimensions sequentially and aggregate the values of the dimensions into their preferences for the various alternatives dynamically over the course of the decision. If a dimension that is common to both alternatives is sampled, the preferences for the two alternatives change equally. This does not affect the modal choice predictions made by these models. This is also the case for the heuristic accumulator model proposed by Bhatia (in press), which attempts to generate behavior resembling heuristic decision making using the accumulator framework (also Lee & Cummins, 2004). This model can be seen as integrating a multi-strategy framework within a single process model, and like the two classes of models it integrates, the heuristic accumulator model satisfies independence. Importantly, the MLBA (Trueblood et al., 2014), which assumes preference accumulation without sequential attribute sampling, does not satisfy dimensional independence. Essentially, this model combines a non-linear transformation of dimensional values with a non-linear attention weighting function for dimensions, to violate dimensional independence in some specific settings.

There are some other single-strategy approaches that are also able to accommodate dimensional dependencies. One of these approaches involves exemplar-based decision making. Although models based on exemplar memory are fairly uncommon in preferential choice research, there have been many attempts to use these models in related domains, such as multi-cue judgment (Bröder et al., 2010; Juslin & Persson, 2002; Juslin et al., 2008; Olsson, Enkvist, & Juslin, 2006; Persson & Rieskamp, 2009; Platzer & Bröder, 2013; Von Helversen & Rieskamp, 2008, 2009). In this domain, decision makers are assumed to store mental representations of cue combinations and associated criterion values that have been experienced previously. Judged criterion values for new cue combinations are based on their similarity with previously experienced cue combinations. As this approach does not assume an independent structure (or really any structure) for the function that describes the relationship between the cues and the criterion, it is able to display violations of independence. Exemplar models provide a good account of quantitative judgments when cue–criterion relationships are non-linear, which suggests that they may be able to describe choice behavior with non-independent reward functions as well (see also Medin & Schaffer, 1978, and Nosofsky, 1984, for related models).

Another type of model that permits dimensional dependencies is the configural weight model. Configural weight models have typically been applied to study perceptual judgments (Birnbaum et al., 1971) impression formation in social psychology (Birnbaum, 1974; Mellers, Richards, & Birnbaum, 1992), risky choice (Birnbaum, 2008), and certain types of multi-cue judgments (Mellers, 1980; also Garcia-Retamero, Hoffrage, Dieckmann, & Ramos, 2007). These models typically take a linear form but allow the weighting on a particular dimension for an alternative to depend on the various features of the set of alternatives offered to the decision maker, such as the rank order of its scale value among the dimensions to be integrated, or the overall range of the dimensions in consideration. These dependencies, although embedded in a typically linear or additive aggregate rule, nonetheless cause the models to generate dimensional dependencies. For this reason, these types of models are able to accurately capture the dependence of likeability judgments on different types of adjectives used to describe decision makers, certain paradoxes in risky choice, and cue aggregation rules that use compound cues, which are cues that aggregate pairs of cues in a potentially non-independent manner. Thus, it may be possible that these types of models are able to describe multi-dimensional choice behavior with non-independent reward functions as well.

It is also possible to model dimensional dependencies using two-layered bidirectional neural networks (Glöckner & Betsch, 2008; Holyoak & Simon, 1999; Thagard, 1989). These networks are used to explain coherence shifts in choice, according to which the weighting or activation of dimensions changes over the time course of the decision, to cohere with emerging preference (e.g., DeKay, Patiño-Echeverri, & Fischbeck, 2009; Russo, Meloy, & Medvec, 1998; Simon, Krawczyk, & Holyoak, 2004; Glöckner, Betsch, & Schindler, 2010). In order to capture these coherence effects, these bidirectional networks typically represent dimensions in a bottom layer and choice alternatives in a top layer and additionally assume that there is recurrent connectivity between the various nodes within and between layers. This means that activation does not only spread from the bottom layer to the top (i.e., from dimensions to choice alternatives). It also spreads from the top layer to the bottom layer (from choice alternatives to dimensions), and additionally within the respective node of the bottom layer and the top layer (between different alternatives and between different dimensions). Critically for the purposes of this paper, this implies that individual nodes representing certain dimensions affect the activation of other nodes representing other dimensions, generating violations of independence.

It is useful to note that these bidirectional networks make principled predictions regarding the connections within the bottom, dimensional layer. Particularly, they assume that dimensions that are coherent with each other have positive connections with each other and dimensions that are incoherent with each other have negative connections with each other. If we interpret coherence as a form of complementarity, then these networks would predict that nodes corresponding to objects that complement each other (e.g., cookies and milk, in the aforementioned example) would have positive connections with each other. Activating one of these nodes would increase the activation of the other, which in turn would influence choices between various bundles. Likewise, nodes corresponding to attributes that complement each other would have positive connections with each other and would in turn influence choices between various objects composed of these attributes.

DOI: 10.1002/bdm

Strategy selection and dimensional independence

The fact that all of the strategies typically within the repertoires of multi-strategy approaches to decision making satisfy independence implies that the meta-decision processes assumed to guide strategy selection in these approaches will not be able to change behavior so as to violate independence, even in settings where these violations are beneficial for the decision maker. In contrast, approaches utilizing a single strategy could potentially satisfy dimensional independence in settings when doing so is desirable, and violate dimensional independence when it is not. These changes in behavior would be driven by changes to the underlying parameters of these models, in the case of neural networks or configural weight models (as well as MLBA, in settings where it is able to generate dimensional dependencies), or changes to learnt representations, in the case of exemplar models. They would not however involve substantial changes to the processes that characterize the decision, that is, they would not involve the selection of any fundamentally different strategy.

In this sense, tests of dimensional independence provide a valuable approach to comparing multi-strategy and single-strategy models, and in turn comparing the strategy selection processes involved in adapting behavior to a given environment. If there exist environments in which dimensional independence is violated, then it would be clear that multi-strategy approaches (but not the single-strategy models discussed earlier) provide an incomplete account of behavior. This would necessitate the addition of individual strategies capable of violating independence to the repertoires of multi-strategy models, as well as the development of strategy selection mechanisms that are capable of recognizing environments in which non-independent choice is desirable. The goal of this paper is thus to examine these environments and rigorously test the possibility of dimensional independence being violated.

It is important to note that there has been some related work regarding non-linearity in multi-cue judgment, which is a domain in which the multi- and single-strategy models we consider are often assumed to apply (Bröder et al., 2010; Garcia-Retamero, Hoffrage, & Dieckmann, 2007; Garcia-Retamero, Hoffrage, Dieckmann, et al., 2007; Hoffrage, Garcia-Retamero, & Czienskowski, 2008; Juslin & Persson, 2002; Juslin et al., 2008; Mellers, 1980; Olsson et al., 2006; Persson & Rieskamp, 2009; Platzer & Bröder, 2013). This research examines the mechanisms that decision makers use to make judgments based on a number of continuous or non-continuous cues. It typically finds that decision makers are able to learn cue–criterion relationships when underlying functions are non-additive and non-linear, even though this is typically harder than learning similar additive linear cue–criterion relationships. This work is related to research on function learning that finds that decision makers can learn quadratic and exponential relationships between a criterion and a single cue (DeLosh, Busemeyer, & McDaniel, 1997).

Research on multi-cue judgment does suggest that decision makers may be able to make preferential choices using non-independent dimensional evaluations. However, multi-cue judgment and preferential choice are two very different domains, and it is not necessarily the case that insights from one domain can be easily transferred to the other. First, there are important psychological differences between making judgments based on beliefs about cue–criterion relationships and making choices based on values and preferences over dimensions: Mechanisms that are active in one setting may not be active in another. More importantly, however, although non-independence implies non-additivity and non-linearity, it is not the case that non-additivity or non-linearity necessarily implies non-independence. In fact, the non-additive and non-linear cue–criterion functions used in Mellers (1980), the work of Juslin and coauthors (e.g., Juslin et al., 2008), and other papers (e.g., Von Helversen & Rieskamp, 2008, 2009, which are concerned with quantitative estimation), are unable to generate violations of dimensional independence. If these functions were used to model preferences, they would always satisfy the dimensional independence property described in the aforementioned section. Thus, the results of the aforementioned papers do not provide evidence for individuals violating independence in multi-cue judgment. One exception to this is the work of Garcia-Retamero, Hoffrage, Dieckmann, et al. (2007) and Hoffrage et al. (2008), which we discuss in detail later on in this paper.

There is one domain in which violations of independence have played a prominent role in empirical and theoretical research on preferential decision making. This is risky choice. The so-called rational model of risky choice, expected utility theory, assumes that gamble outcomes are weighted by their probabilities and aggregated into a single gamble utility. This type of weighted aggregation is similar to that specified by the weighted-additive rule, and like the weighted-additive rule, results in gamble outcomes being evaluated separately. Subsequently outcomes that are common across gamble are ignored. The famous common-consequence effect (Allais, 1953; Kahneman & Tversky, 1979) shows that this prediction is violated. Changing outcomes that are common across gambles can lead to reversals in preference between the gambles.

Risky choice and multi-dimensional preferential choice involve some differences. For example, the gambles that are used in risky choice typically involve only one component: A probabilistic monetary payoff. Additionally, the probabilities used to weigh and aggregate payoffs are objective and are provided to the decision maker as part of the choice problem. Multi-dimensional preferential choice, in contrast, involves alternatives with numerous rich, interacting components, whose weights are almost always subjective. Nonetheless, there are a large number of similarities between these two domains, and some effects, such as context effects, have been shown to emerge in both settings (Huber, Payne, & Puto, 1982; Wedell, 1991). It thus may be reasonable to use insights from experimental work on independence violations in risky choice to test for similar violations in riskless choice. Particularly, we can design pairs of binary choice problems, which vary common dimensions in much the same way that the choice problems in the common-consequence effect vary common outcomes. If decision makers satisfy dimensional independence, then

DOI: 10.1002/bdm

changing common dimensions in pairs of choice alternatives should not change decision makers' preferences between these alternatives.

STUDY 1

One setting where independence can be violated involves choices between bundles of objects. Here the bundles are the choice alternatives, and the individual objects are the dimensions on which these alternatives are defined. Different objects can be complements or substitutes, implying that adding or removing a particular object from a bundle can increase or decrease the desirability of other objects in that bundle. The lexicographic heuristic, the tallying heuristic, the equal weights heuristic, and the weighted-additive heuristic are all unable to accommodate these changes to preferences. This implies that multi-strategy models of decision making, which involve the selection and use of one of these strategies, are also in turn unable to accommodate these changes to preferences. This is also the case for single-strategy models that satisfy independence, such as most accumulator models. In contrast, exemplar models accommodate dependencies between the objects in these bundles if prior experiences with bundles composed of complementary objects led to higher rewards than experiences with bundles with non-complementary objects. Likewise configural weight models can accommodate dependencies if the weighting function permits terms that allow for scaling-based or range-based comparisons between the objects. MLBA is also able to generate dependencies based on how it allocates attention weights and values to the various dimensions. Finally, recurrent neural network models would generate dependencies between objects if nodes corresponding to complementary objects share positive recurrent connections and nodes corresponding to non-complementary objects share negative recurrent connections.

The goal of Study 1 is to test whether we can observe violations of independence in choices between bundles. It uses two different types of bundles composed of various household objects and varies the common objects across these bundles. Violations of independence would be observed if choice proportions for the bundles change as these common objects are varied.

Method

Participants

One hundred participants in Study 1 (mean age = 30, 29% female), recruited through Amazon Mechanical Turk (MTurk), completed the study online.

Materials and procedure

Study 1 used two different decision problems, which are as follows:
Problem 1: Imagine that you have won a prize at your nearby electronics store. You have a choice between the following two options. Which do you prefer?

A. One Blue-ray disk player and a set of 10 Blue-ray movies
B. One color printer and a set of 10 Blue-ray movies
A′. One Blue-ray disk player and a set of five color print cartridges
B′. One color printer and a set of five color print cartridges

Problem 2: Imagine that you have won a prize at your nearby household appliances and furniture store. You have a choice between the following two options. Which do you prefer?

A. One brown three-seat couch and one brown two-seat couch
B. One washing machine and one brown two-seat couch
A′. One brown three-seat couch and one dryer
B′. One washing machine and one dryer

All participants answered both problems sequentially. This was a between-participants design, and for each decision problem, each participant was given either a choice between alternatives A and B or a choice between alternatives A′ and B′. This was randomized across participants and across problems. Note that A and B have the same common object (e.g., Blue-ray movies). This is replaced by a different common object in alternatives A′ and B′ (e.g., print cartridges). If people evaluate objects independently, then these common objects should not affect preferences between the choice alternatives. Subsequently, a participant who prefers A over B should also prefer A′ over B′. Changes in observed choice proportions between the two versions of the decision problems would indicate violations of dimensional independence.

Results and discussion

In the first decision problem of Study 1, 81% of participants chose A over B, but only 37% chose A′ over B′. This is a significant difference ($\beta = 2.02$, $z = 4.31$, 95% CI = [1.10, 2.93], $p < .01$) for a logistic regression testing the effect of the common object on choice. Similarly, in the second problem of Study 1, 66% of participants chose A over B, but only 16% chose A′ over B′. This too is a significant difference ($\beta = 2.32$, $z = 4.76$, 95% CI = [1.36, 3.28], $p < .01$). In aggregate, A was chosen over B 68% of the time and A′ was chosen over B′ 26% of the time, which is a significant difference ($\beta = 2.27$, $z = 5.40$, 95% CI = [1.45, 3.10], $p < .01$) for a logistic regression testing the effect of the common object on choice, with fixed effects for the decision problem, as well as random effects on the participant level. These results are summarized in Table 1 and are illustrated in Figure 1.

Both the individual comparisons for the decision problems and the pooled comparisons for the decision problems show that choice proportions are reversing between the A

Table 1. Number of participants selecting A over B and A′ over B′ for each of the three decision problems used in Studies 1 and 2

Prob.	A	B	Total	A′	B′	Total
Exp. 1—P1	39	9	48	19	33	52
Exp. 1—P2	33	17	50	8	42	50
Exp. 2	27	73	100	9	91	100

DOI: 10.1002/bdm

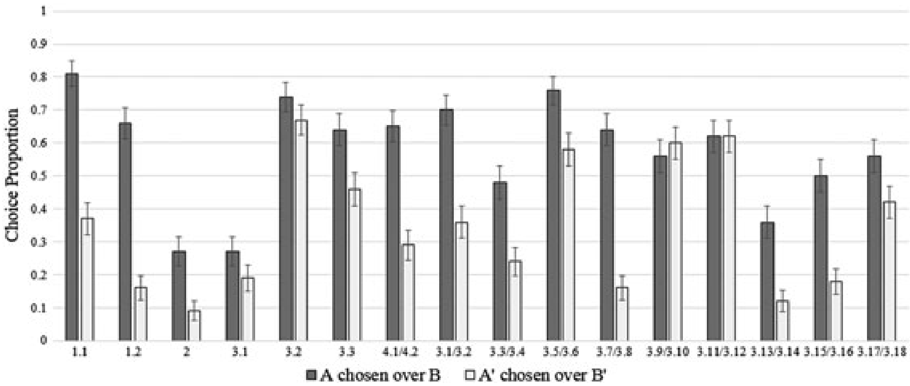

Figure 1. Choice proportions for all relevant decision problems in this paper. Here, dark grey bars correspond to choice proportions of A over B, and light grey bars correspond to choice proportions of A′ over B′. In all of these problems, models that assume dimensional independence predict that choice that these two sets of choice proportions are identical. Error bars depict one standard error above and below the mean

and B problems and the A′ and B′ problems. This suggests that decision makers can violate independence in preferential choice between bundles of objects, contradicting the predictions of many existing multi-strategy models. These models, which are composed of heuristics incapable of violating independence, would predict that decision makers would never choose both A over B and B′ over A′, as observed earlier. This prediction holds regardless of which of the specific heuristics these multi-strategy models apply.

STUDY 2

Study 1 involved a between-subjects design in which each participant was given either alternatives A and B or A′ and B′ to choose between. Study 2 tested whether we could observe significant violations of independence even in a highly transparent setting where the two choice problems were given to the same participant, one after the other. As with Study 1, it considered choices between bundles composed of various objects.

Method

Participants

One hundred participants (mean age = 32, 21% female), recruited through Amazon MTurk, completed the study online.

Materials and procedure

Study 2 was identical to Study 1. However, it tested for independence violations within participants. Study 2 involved the following choice:

Imagine that you want to eat a snack, and you have a choice between the following two equally priced options. Which do you prefer?

A. One muffin and one coffee
B. One sandwich and one coffee
A′. One muffin and one Coke
B′. One sandwich and one Coke

All participants answered this problem twice: Either with only A and B as the available alternatives or with only A′ and B′ as the available alternatives. The two versions of the problem were answered one after the other in a random order, with no intervening gap or break. Note that the drink in both versions is common to the two alternatives. If participants satisfy independence, the specific drink available should not matter. Thus, if participants choose A over B, they should also choose A′ over B′ (and vice versa).

Results and discussion

Of the participants, 27% chose A over B, but only 9% chose A′ over B′. This is a statistically significant difference ($\beta = 1.31$, $z = 3.17$, 95% CI = [0.50, 2.13], $p < .01$) for a logistic regression with random effects on the participant level. These results suggest that independence can be violated even when the choice problems testing for these violations are presented in a highly transparent manner to the participants. Again, this contradicts the predictions of many existing multi-strategy models, which are composed of heuristics incapable of violating independence, as well as single-strategy models such as most accumulator models, which involve the independent evaluation of attributes. These models predict that the relative preferences for A over B and for A′ over B′ would be identical. The results from this study are summarized in Table 1 and are illustrated in Figure 1.

STUDY 3

The choice problems used in Studies 1 and 2 differ in one important way from those that are typically studied in decision research. Most decision-making studies use choice alternatives involving single objects defined on multiple attributes, such as individual apartments, cars, or laptops. In contrast, the choice alternatives used in Studies 1 and 2 involve bundles of objects, such as printers and print cartridges. Both bundles of objects and single objects are composed of multiple dimensions, and multi-strategy decision models predict independence in both the settings. The results of Studies 1 and 2 thus provide strong evidence against these predictions. That said, in order to convincingly demonstrate violations of independence, it is useful to consider choice problems involving the types of multi-attribute single objects that make

DOI: 10.1002/bdm

up most decision-making studies. This was the goal of Study 3. Study 3 was also performed offline, in a laboratory setting, in order to fully control the experimental environment.

Method

Participants

One hundred fifty participants (mean age = 22, 46% female) took part in the study for course credit and monetary compensation. The data were collected in a behavioral laboratory at a university in England and in a behavioral laboratory at a university in Portugal.

Materials and procedure

Study 3 used an experimental setup similar to that of Studies 1 and 2. Unlike these studies, however, each choice alternative was a single good, and information about the attributes of this good was presented in a table format. Particularly, participants had to choose between two apartments, which were defined in terms of their distance to work, their distance to public transport, and their size (problem 1); two cars, which were defined in terms of their horse power, maximum speed, seating capacity, and safety rating (problem 2); and two laptops, which were defined in terms of processor speed, RAM, screen size, and warranty length (problem 3). The specific attribute values used are presented in Table 2.

We predicted that the relationships between the different attributes in the decision problems would affect people's choices and lead to violations of independence. Particularly, people should be willing to accept a larger distance to public transportation if their work place is nearby, but not if it is far away. Likewise, people should prefer less powerful cars with high safety ratings if the cars have seven seats (and are suitable for family travel) but prefer more powerful cars with lower safety ratings if the cars have only two seats. Finally, people should prefer portable laptops with smaller screens if the laptops are powerful and can be used for a variety of applications but should prefer less portable laptops with big screens if they are not powerful and only be used for watching videos and other media. For all these three decision problems, these predictions specify a preference for A over B but B′ over A′.

Participants were presented with the three problems sequentially, and for each problem, they were presented with a choice of between options A and B or between options A′ and B′. This was randomized across participants and across decision problems. Overall, higher choice frequencies of A over B compared with A′ over B′, across the three problems, would constitute violations of independence.

Table 2. The choice alternatives and number of participants (out of 150) selecting the choice alternatives, in Study 3

Alternative	Work Dist. (km)	Pub. Tran. Dist. (km)	Size (ft^2)		# Choosing
1-A	0.5	7	500		40
1-B	0.5	1	410		110
1-A′	2	7	500		28
1-B′	2	1	410		122
Alternative	**Seats**	**Horse power**	**Max. speed (km/hour)**	**Safety**	**# Choosing**
2-A	7	135	160	9/10	111
2-B	7	340	310	4/10	39
2-A′	2	135	160	9/10	100
2-B′	2	340	310	4/10	50
Alternative	**Processor (GHz)**	**RAM (GB)**	**Warranty (years)**	**Screen size (in.)**	**# Choosing**
3-A	1.7	4	2	12.4	96
3-B	1.7	4	2	17.3	54
3-A′	3.6	16	1	12.4	69
3-B′	3.6	16	1	17.3	81

Results and discussion

We observed changes in choice frequencies of A over B and A′ over B′ in our predicted direction, in all three decision problems. However, unlike our prior studies, these differences were not all individually significant. Particularly, 27% of participants chose A over B and 19% of participants chose A′ over B′ in decision problem 1. This is not a statistically significant difference ($\beta = .47$, $z = 1.20$, 95% CI = [−.30, 1.25], $p = .23$) for a logistic regression testing for the effect of the common attribute on choice. Likewise, 74% of participants chose A over B, and 67% of participants chose A′ over B′ in decision problem 2. This is also not a statistically significant difference ($\beta = .29$, $z = 0.79$, 95% CI = [−0.43, 1.00], $p = .43$). Finally, 64% of participants chose A over B, and 46% of participants chose A′ over B′ in decision problem 3. This is a statistically significant difference ($\beta = .72$, $z = 2.17$, 95% CI = [0.07, 1.38], $p < .05$). These results are summarized in Table 2 and are illustrated in Figure 1.

These results suggest that violations of independence are emerging in problem 3, but that these violations are either very weak or non-existent in problems 1 and 2. Although the results for problem 3 are sufficient to demonstrate that people can violate independence in multi-attribute choice, it is useful to pool the choice proportions across all problems and look at aggregate choices. This would, in a manner similar to that of a meta-analysis, indicate whether violations of independence are strong enough to emerge even when data from different problems are considered together. Note that for all three problems, we made a priori predictions regarding changes in choice proportions. Particularly, we predicted that the preference for A over B would be higher than that for A′ over B′, in all three problems. In this sense, the relative choice share of A over B can be seen as corresponding to the main dependent variable from one experimental condition, and the choice share of A′ over B′ can be seen as corresponding to the dependent variable from a second experimental condition. Subsequently, averaging the choice share of A over B and the choice share of A′ over B′, across all three problems, would give us choice proportions that could be used to test for independence violations on the aggregate level.

DOI: 10.1002/bdm

Overall, 53% of participants chose A over B, but 46% of participants chose A′ over B′. This is a statistically significant difference ($\beta = .54$, $z = 2.41$, 95% CI = [0.11, 0.98], $p < .05$), for a logistic regression testing for the effect of the common attribute on choice, with fixed effects for the decision problem and random effects on the participant level. This indicates that behavior is changing even when we consider all three decision problems together. Even though the violations of independence are stronger in problem 3, and weaker in problems 1 and 2, they are nonetheless strong enough to lead to significant differences in choice proportions on the aggregate level.

STUDY 4

Studies 1, 2, and 3 examine independence in settings where information is explicitly provided to decision makers. Although these studies provide very strong evidence for violations of independence, they nonetheless use realistic choice options whose various dimensions are often featured in everyday choices. It may be the case that different participants have had different experiences with these dimensions and thus evaluate and aggregate these dimensions differently. Study 4 controlled for this issue by testing for independence in binary choices between artificial choice alternatives (which share common attributes). Attribute–reward relationships in this study were learnt through experience.

By controlling the reward structure of the environment that decision makers were offered, Study 4 provides a powerful test of the strategy selection mechanisms proposed by various multi-strategy and single-strategy models. In order to maximize rewards, participants need to aggregate dimensions non-independently. Multi-strategy models, which are composed of heuristics that satisfy independence, predict that decision makers would not be able to learn and adapt in this type of environment and thus would have relatively low rewards. In contrast, single-strategy models, which allow for dimensional dependencies, such as exemplar models, configural weight models, MLBA, and recurrent neural networks, would be able to adapt to these dependencies and generate relatively high rewards, by either learning examples of high-reward alternatives or modifying their parameters to adequately capture the dependencies that characterize these alternatives. In this sense, Study 4, unlike Studies 1–3, can be used to test not only how well various multi- and single-strategy models describe choice but also how well the strategy selection mechanisms proposed by these approaches perform in complex environments.

Finally, note that the design for this experiment is closely related to the types of experiments used in function learning research (DeLosh et al., 1997). Function learning experiments however do not involve violations of independence and additionally often involve continuous judgments, and relations between only two variables (rather than multiple variables). In this sense, the results of Study 4 would be novel.

Method

Participants

Fifty participants (mean age = 30, 20% female), recruited through Amazon MTurk, completed the study online. Participants were paid based on their responses (refer to succeeding discussions). One of these participants failed to answer all the questions in the study and is excluded from the analysis.

Materials and procedure

All participants were given six decision problems, which involved choices between pairs of alternatives defined on three attributes: x, y, and z. Each alternative was associated with a number of points, P, which was specified by the following deterministic reward function:

$$P = 2xy + 5z$$

The total number of points that participants earned was the sum of the points of the six alternatives that they chose in the six decision problems. Each point was worth \$0.01, and participants were paid the monetary value of points that they earned in the study.

The point values of the available alternatives were not shown to participants in the six decision problems. Instead, participants learnt about the points associated with different compositions of attributes before making their choices. In this learning task, 60 alternatives with varying amounts of x, y, and z were presented to the participants alongside their corresponding point values. These alternatives were different from the alternatives used in the choice task. Participants did not make any choices in the learning task.

The alternatives used in this study contained between 0 and 5 units of each attribute in each alternative. Additionally, for simplicity, alternatives that contained nonzero amounts of attribute x did not contain any attribute z, and vice versa. There are 72 different alternatives that can be composed of this structure. Twelve of these alternatives were used in the six decision problems. These are presented in Table 3. The remaining 60 alternatives were used in the learning task. The order of the alternatives within the decision problems (i.e., whether these alternatives were presented as “A” or “B”), the order of the six decision problems, and the order of the 60 alternatives in the learning task were all randomized across participants.

Table 3. Decision problems and number of participants selecting A over B (out of 49), in Study 4

	Alternative A				Alternative B				
Prob.	x	y	z	Points	x	y	z	Points	# Choosing A
1	3	5	0	30	0	5	3	15	32
2	3	0	0	0	0	0	3	15	14
3	2	3	0	12	0	3	2	10	27
4	1	2	0	4	0	2	1	5	15
5	5	4	0	40	0	4	5	25	32
6	4	1	0	8	0	1	4	20	14

A is the optimal response in odd-numbered decision problems, and B is the optimal response in even-numbered decision problems.

DOI: 10.1002/bdm

Note that the reward function presented earlier aggregates attributes non-independently. As a result of this, the optimal choice between a pair of alternatives with a common attribute can reverse if the common attribute is varied. This insight was used to design decision problems 1 and 2, which were the focus of interest in this study. Alternatives A and B in problem 1 both contain 5 units of *y*. In problem 2, these 5 units of *y* are removed from both these alternatives. If people evaluate attributes independently, then this change should not affect observed choices: A participant who prefers A over B in problem 1 should also prefer A over B in problem 2. The reward function however generates more points for A in problem 1, and more points for B in problem 2, and thus the optimal choice involves a switch from choosing A in problem 1 to choosing B in problem 2. If participants are able to make optimal choices in this study, then they will violate independence.

Decision problems 4–6 largely served as fillers. These problems were selected so that the optimal choice involved the selection of the alternative with more *x* (alternative A in Table 3) or the alternative with more *z* (alternative B in Table 3), with equal frequency. Although non-linear multi-attribute models do not put any strong restrictions on choices in these decision problems, models that assume linear aggregation (e.g., most heuristic choice models) predict that decision makers should either always select the alternative with more *x* or always select the alternative with more *z*. Decision makers who aggregate attributes optimally will violate these predictions.

Results and discussion

We observed strong violations of independence in decision problems 1 and 2. Particularly, 65% of participants chose alternative A in problem 1, but only 29% of participants chose alternative A in problem 2. This is statistically significant ($\beta = 1.54$, $z = 3.55$, 95% CI = [0.69, 2.40], $p < .01$) for a logistic regression testing the effect of the decision problem on the choice of A, with participant-level random effects. This shift indicates that the majority of participants can choose adaptively when doing so requires the non-independent aggregation of attributes, and that non-independence may be quite common in experienced-based choice.

More generally, we found that the optimal response was chosen 66% of the time across the six decision problems, and that participants were more likely to select an alternative with more *x* (alternative A in Table 3) when it was the optimal response than when it was not the optimal response ($\beta = 1.67$, $z = 5.70$, 95% CI = [1.10, 2.26], $p < .01$, for a logistic regression testing the effect of the optimality of A on the choice of A, with participant-level random effects). This violates the predictions of all linear models of attribute aggregation and provides further evidence that decision makers can learn non-independent attribute–reward relationships. Table 3 displays observed choice proportions for the six decision problems. Choice proportions for the first two decision problems are illustrated in Figure 1.

As with the results of Studies 1–3, this contradicts the predictions of many existing multi-strategy models, as well as single-strategy models such as sequential sampling models that involve the independent evaluation of attributes. These models predict that the relative preference for A over B and would not change across decision problems 1 and 2. In contrast, the difference in the choice proportions for A over B in these two problems can be accommodated by single-strategy models such as exemplar models, configural weight models, MLBA, and recurrent neural networks.

Unlike Studies 1–3, Study 4 also shows that decision makers possess strategy selection mechanisms that are able to learn and adapt to dimensional dependencies in the environment. As a result, they are able to make better decisions than those predicted by all heuristics, and subsequently, by multi-strategy models of decision making. For example, by choosing A over B in problem 1 but B over A in problem 2 (which was the modal choice pattern in the experiment), decision makers can earn a total of 45 points. In contrast, models that do not violate independence, and thus predict that decision makers choose A over B or B over A in both problems 1 and 2, can only earn a total of 30 points.

STUDY 5

Study 4 has shown that people are able to learn to make adaptive decisions even if these decisions involve the non-independent aggregation of attribute values. Its design however has a few limitations. First, the learning task used in this study did not involve choice. Rather, it explicitly presented the point values of different attribute combinations. Typically, the reward structure in a preferential choice domain is learnt through choice itself. Related work has found that the type of feedback offered to decision makers in multi-cue judgment affects learning and strategy use (Bröder, Glöckner, Betsch, Link, & Ettlin, 2013; Pachur & Olsson, 2012). Thus, in order to establish the generality of the results of Study 4, it is necessary to consider a similar choice-based learning task with feedback. Additionally, the reward function used in Study 4 is very simplistic. Even though the fact that decision makers were able to learn this function indicates that they can violate independence, it may nonetheless be useful to attempt to replicate the results of Study 4 with a more complex non-independent reward function. This is the goal of Study 5.

Note that by controlling the reward structure of the environment that decision makers were offered, Study 5, like Study 4 (but unlike Studies 1–3), is able to provide a powerful test of how various strategy selection mechanisms perform in complex environments. Once again, in order to maximize rewards, participants need to aggregate dimensions non-independently, something that multi-strategy models predict participants are unable to do.

Method

Participants

Fifty participants (mean age = 21, 67% female) took part in this study for monetary compensation. Participants were given a show-up fee and additionally paid based on their

DOI: 10.1002/bdm

responses (refer to succeeding discussions). The data were collected in a behavioral laboratory at a university in England.

Materials and procedure

Study 5 used a setup similar to that of Study 4. Particularly, all participants made choices between pairs of alternatives defined on three attributes: x, y, and z. Each alternative was associated with a number of points, P, and decision makers were required to select the alternative with the highest number of points. The attributes contained between 0 and 4 units each, and the point value of each alternative was generated by the following deterministic reward function:

$$P = \begin{cases} 5, & x = y \\ z, & x \neq y \end{cases}$$

In this reward function, alternatives with equal amounts of x and y have the highest possible value, 5. Alternatives with non-equal amounts of x and y have a value that is equal to the amount of z that they contain (which is a maximum of 4). It is possible to violate independence by giving decision makers pairs of choice problems, in which changes to the common attributes affect whether or not the amount of x in the alternatives is equal to the amount of y. Note that the reward function used here violates not only independence but monotonicity as well. It is possible to reduce the point value of an alternative by increasing the amount of x or y.

Study 5 involved a different type of learning task compared with Study 4. In this task, decision makers were required to make choices between pairs of alternatives. After making their choice, they were shown the point values associated with the two alternatives. This choice-based learning task involved 50 binary choices, generating feedback for a total of 100 different choice alternatives.

After the learning task, participants made 18 choices between two alternatives. There was no feedback in this part of the study. These 18 choices corresponded to nine pairs of choice problems designed to test for violations of independence by varying the common attribute. The 18 decision problems used in the second part of the study are shown in Table 4. Note, for example, that both Alternatives A and B in decision problem 1 consist of 4 units of y. These alternatives are identical to the ones used in decision problem 2, except that the latter consist of 2 units of y. Thus, attribute y is common across the alternatives in both decision problem 1 and decision problem 2. Participants who prefer A over B in problem 1 should also prefer A over B in problem 2, if they are unable to learn non-independent reward functions. However, as the reward function used in this study generates more points for A in problem 1 (where $x = y$), and more points for B in problem 2 (where $x \neq y$ and z in A is less than z in B), the optimal choice involves a switch from choosing A in problem 1 to choosing B in problem 2. This logic also applies to all other pairs of choice problems in Table 4. These problems are organized so that A is the optimal choice in odd-numbered problems and B is the optimal choice in even-numbered problems. Pairs of choice problems that are identical, except for the common attribute, are listed sequentially (e.g., 1 and 2, 3 and 4, and 5 and 6). If participants are able to make optimal choices in this study, then they will switch between A and B in each pair of these choice problems, violating independence.

Table 4. Decision problems and number of participants selecting A over B (out of 50), in Study 5

	Alternative A				Alternative B				
Prob.	x	y	z	Points	x	y	z	Points	# Choosing A
1	4	4	1	5	1	4	2	2	35
2	4	2	1	1	1	2	2	2	18
3	3	3	2	5	2	3	3	3	24
4	3	1	2	2	2	1	3	3	12
5	3	3	3	5	1	3	4	4	38
6	3	2	3	3	1	2	4	4	29
7	1	1	1	5	3	1	4	4	32
8	1	4	1	1	3	4	4	4	8
9	4	4	2	5	1	4	4	4	28
10	4	2	2	2	1	2	4	4	30
11	3	3	3	5	2	3	4	4	31
12	3	1	3	3	2	1	4	4	31
13	2	2	1	5	3	2	3	3	18
14	2	4	1	1	3	4	3	3	6
15	1	1	2	5	2	1	3	3	25
16	1	3	2	2	2	3	3	3	9
17	4	4	3	5	2	4	4	4	28
18	4	1	3	3	2	1	4	4	21

A is the optimal response in odd-numbered decision problems, and B is the optimal response in even-numbered decision problems.

Study 5 also differed from Study 4 in terms of incentives. Particularly, participants were not paid based on the point values of their chosen alternatives. Rather, one participant in each session was chosen randomly, and one of this participant's 18 choice problems was also chosen randomly. If they chose the alternative with the highest point value in their randomly chosen decision problem, they were compensated with a £10 Amazon.co.uk gift card. Finally, note that the order of the alternatives within the decision problems (i.e., whether these alternatives were presented as "A" or "B"), the order of choices in the learning task, and the order of the choices in the main portion of the experiment were all randomized.

Results and discussion

Once again, we observed strong violations of independence. Particularly, 58% of participants chose alternative A in odd-numbered problems (where A was the optimal response), but only 36% of participants chose alternative A in even-numbered problems (where B was the optimal response). This is a statistically significant difference ($\beta = 1.01$, $z = 6.78$, 95% CI = [0.72, 1.30], $p < .01$) for a logistic regression testing the effect of the optimality of A (i.e., whether it was in an odd-numbered or even-numbered decision problem) on the choice of A, with participant-level random effects. On a question level, participants were significantly ($p < .05$) more likely to select A when it was optimal than not, in six out of nine pairs of decision problems; marginally significantly ($p < .1$) more likely to select A when it was optimal

DOI: 10.1002/bdm

than not, in one out of nine pairs of decision problems; completely nonsignificantly ($p > .1$) less likely to select A when it was optimal than not, in one out of nine pairs of decision problems; and equally likely to select A when it was optimal and when it was not, in one out of nine pairs of decision problems, with each comparison controlling for participant-level random effects. More generally, participants were able to choose the optional response an average of 61% of the time, which is statistically higher than random ($\beta = .44$, $z = 5.56$, 95% CI = [0.28, 0.60], $p < .01$), controlling for participant-level random effects. All of these results are summarized in Table 4 and are illustrated in Figure 1.

These results again indicate that the majority of participants can choose adaptively when doing so requires the non-independent aggregation of attributes, and that non-independence may be quite common in experienced-based choice even when reward values are learnt in through a choice task, and the functions underlying these reward values are highly complex. As with Study 4, these results show that decision makers possess strategy selection mechanisms that are able to learn and adapt to dimensional dependencies in the environment, and are able to make better decisions than those predicted by all heuristic models, and subsequently, by multi-strategy models of decision making. The modal choices of participants in this study, for example, generate total rewards of 72 points. The highest possible rewards using a decision strategy that satisfies independence, in contrast, generates only 63 points. Applying the optimal strategy in every question would generate 76 points.

GENERAL DISCUSSION

Dimensional independence is the assumption that the desirability of a relevant dimension in a choice alternative is not affected by the values of the alternative on other dimensions. This implies that changing objects that are common across bundles, or changing attributes that are common across different objects, does not alter people's choices. In five studies, we tested the assumption of attribute independence by varying common dimensions across available alternatives. Our choice problems were analogous to the common-consequence effect questions used to test the independence assumption in risky choice (Allais, 1953; Kahneman & Tversky, 1979). These types of questions provide critical tests of independence, regardless of the specific quantitative or qualitative assumptions underlying different theories. As with the common-consequence effect, we observed that choice proportions changed significantly as common dimensions were varied. This suggests that the amount of a dimension in an alternative can affect how other dimensions in that alternative are evaluated, implying violations of the assumption of dimensional independence.

By now, there are many theoretical approaches that attempt to explain both choice behavior and its cognitive underpinnings. Many of these models assume that decision makers have a repertoire of strategies, or heuristic rules, with which they make decisions (e.g., Gigerenzer & Todd, 1999; Payne et al., 1993; Rieskamp & Otto, 2006). Different environments call for the selection of different strategies, and decision makers are typically able to choose strategies that give adaptive responses for the environment in which the decision is being made. While multi-strategy models provide a valuable perspective on preferential choice, none of the strategies typically assumed to be part of the strategic repertoires of these models are able to account for the violations of dimensional independence documented in this paper.

This is also the case for some single-strategy models of decision making, which propose that preferential choice is the product of a single process (Bhatia, 2013; Lee & Cummins, 2004; Newell, 2005; Roe et al., 2001). The processes proposed by these models satisfy dimensional independence. Subsequently, none of these models are able to predict the types of behaviors observed in our studies.

In contrast to the aforementioned approaches, single-strategy models such as neural networks, MLBA, configural weight models, and exemplar models are able to predict the violations observed in this paper. Recurrent neural networks, for example, typically involve two layers of nodes, with nodes in the bottom layer corresponding to dimensions and nodes in the top layer corresponding to alternatives (Glöckner & Betsch, 2008; Holyoak & Simon, 1999; Thagard, 1989). There are recurrent connections both between the two layers and within the layers. These recurrent connections imply that changing the values on some dimensions can affect node activations for other dimensions. This can generate violations of dimensional independence. Indeed, the connection strengths between the dimensions have an intuitive interpretation: They correspond to the complementarities between the various objects or attributes in the choice task. Again, consider the late-night snack example presented at the start of this paper. In this example, dimensions like milk and cookies would have positive connections with each other as milk and cookies taste particularly good when eaten together. In contrast, dimensions like beer and cookies would not have positive connections (and perhaps even have negative connections).

Recurrent neural networks provide unique insights for understanding strategy selection, particularly the ways in which decision makers can behave independently when doing so is adaptive, but non-independently when doing so is not. These networks learn through changes in their connection weights, and it is these changes that represent the fundamental mechanisms at play in strategy selection for these models. There has been much research on building recurrent neural networks that are able to learn optimal connection weights for a given environment (Elman, 1990; Kosko, 1988), and it seems that these results could be easily applied to model the strategy selection problem in non-independent decision making.

Configural weight models are a different type of model that are also able to accommodate dimensional dependencies (Birnbaum et al., 1971; Birnbaum, 1974, 2008). These models adopt a linear rule in the aggregation of dimensions and also permit terms that combine or contrast the values on different dimensions. These models also, in some settings, allow the weight on a particular dimension to depend on the values of other dimensions, thereby permitting dimensional interactions. Unlike neural networks, these models are seldom applied in a

DOI: 10.1002/bdm

learning context (but see Garcia-Retamero, Hoffrage, Dieckmann, et al. (2007) and Hoffrage et al. (2008), which we discuss later). Thus, it is difficult to state exactly how they can model the strategy selection problem. That said, it is likely that such an application of these models would involve experience-based changes in parameters to permit dimensional dependencies when doing so is adaptive and avoid dimensional dependence when doing so is not.

Exemplar memory models, in contrast, do not learn specific parameter values or specific rules for making decisions. Rather, they store prior experiences and use similarity to these experiences to predict reward values (Bröder et al., 2010; Juslin & Persson, 2002; Juslin et al., 2008). Thus, these models are able to indirectly learn non-independent dimension–reward relationships by associating certain types of objects with higher rewards. In many ways, because of this, these models do not face a strategy selection problem at all. The exact same mechanisms (with the same parameters) can be used in simple environments as well as more complex environments. It is not the parameters or strategies that change across environments in exemplar models but rather the memories that these models utilize.

Of course, the fact that we were able to document violations of independence does not mean that they are present in all decision environments. In fact, it is likely that many environments are simple enough that decision makers can use the heuristics of multi-strategy models or single-strategy models such as those based on sequential sampling, to adaptively decide between various alternatives. However, dimensional independence does seem likely in more complex preferential choice settings, which involve multiple complementary objects (e.g., cookies and milk) or multiple complementary attributes (e.g., sweet and sour), and it is here where models that are unable to accommodate dimensional dependencies will have a hard time predicting behavior. Future work should attempt to understand which environments generate dependencies and test the aforementioned set of models in both independent environments and dependent environments. Such tests should also involve quantitative comparisons (rather than qualitative comparisons, as in this paper), as this would help establish which models are better overall. It may, after all, be the case that models unable to violate independence still account for a larger proportion of the variance in human behavior. This can only be tested with quantitative comparisons.

Additionally, it would be valuable to examine how multi-strategy approaches could be tweaked to handle rewards interdependencies, and this should also be the focus of future research. There are many different ways to approach this. One possibility is to use memory-based heuristics, such as the recognition heuristic, that are sometimes assumed to accompany the knowledge-based heuristics studied in this paper (Gigerenzer & Todd, 1999; Gigerenzer & Goldstein, 1996). It is not however clear how recognition, by itself, could be used to make non-independent decisions in settings where choice alternatives characterized by non-independent rewards do not differ in terms of their ability to be recognized.

An alternative approach could involve making the strategy section problem more complex. For example, most multi-strategy approaches currently select strategies based on general environmental characteristics. As such, the specific strategies that are applied from one choice problem to another in a given environment do not necessarily differ. It could be possible to build on this approach to select strategies to apply based on not only the environment but also the composition of the specific alternatives in the choice set. For example, if decision makers given the cookie and milk problem described earlier use a lexicographic rule with cookies as the most important dimension, when cookies are presented with milk, but use a weighted-additive rule when cookies are presented with beer, then their decisions could violate independence in the manner observed in the aforementioned studies. This could, for example, be accomplished within a reinforcement learning framework (Rieskamp & Otto, 2006) with context-based weights for different strategies (e.g., Sutton & Barto, 1998).

Another possibility would be to retain the current, simpler, approach to strategy selection and instead transform the dimensions presented in the choice problem in a more complex manner. Thus, in the cookies and milk problem, decision makers could learn to map the two dimensions, cookies and milk, to a single latent dimension, before applying various heuristic rules and strategy selection processes to this latent dimension. This would allow for these models to violate independence with standard heuristics, as latent dimensions would vary between alternatives, even if common dimensions do not. This approach could also be applied to accumulator models, so that these models are able to violate independence.

It is important to note that one extension of the fast and frugal framework that adopts a variant of this approach involves the configural take of the best rule proposed by Hoffrage et al. (2008; also Garcia-Retamero, Hoffrage, and Dieckmann, 2007, and Garcia-Retamero, Hoffrage, Dieckmann, et al., 2007). This heuristic does consider not only individual cues but also compound cues, which are cues that aggregate pairs of cues in a potentially non-independent manner (e.g., with the exclusive or XOR disjunction), based on the decision maker's causal knowledge of the domain at hand. Although this rule is used by Garcia-Retamero, Hoffrage, Dieckmann, et al. (2007) and Hoffrage et al. (2008) to study logical rules in multi-cue judgment and category learning (also Shepard, Hovland, & Jenkins, 1961), it could be easily extended to a multi-attribute setting and combined with other heuristic rules used by multi-strategy approaches to accommodate the findings discussed in this paper. Indeed, Hoffrage et al. (2008) also specify the environments in which these types of heuristics are adaptive, thereby presenting valuable insights regarding the selection of these heuristics and resulting deviations from the assumption of attribute independence.

Yet another approach would involve incorporating new strategies that aggregate dimensions non-independently. These non-independent rules could be based on the types of utility functions used in economics research to model object complementarities (e.g., Varian, 1992). Note that our results are also closely related to well-known findings in

DOI: 10.1002/bdm

risky choice research, such as the common-consequence effect, which demonstrates violations of the key assumptions of expected utility theory (Kahneman & Tversky, 1979). There are heuristics that are able to account for these interdependencies (Brandstätter, Gigerenzer, & Hertwig, 2006; Katsikopoulos & Gigerenzer, 2008), indicating that it may be possible to use a variant of these heuristics to capture dimensional interdependencies in riskless choice.

Of course, it is also possible to integrate insights from current multi-strategy and single-strategy models to create hybrids, which can violate independence. One possibility involves using Bhatia's (in press) accumulator model, which is able to instantiate many of the heuristic rules proposed by multi-strategy approaches, in a single neural network model. This neural network model does not violate independence because, unlike recurrent approaches, it does not allow for connections between different dimensions. However, because it relies on a neural network instantiation, these connections could easily be built in. The resulting models would generate behaviors resembling various heuristic rules, while also allowing for values on one dimension to influence others. It is possible that other types of hybrid models, such as those relying on an ACT-R framework rather than a neural network framework, could also be built using these insights (Marewski & Schooler, 2011). However, this requires additional research.

Ultimately, the dimensions of choice alternatives often have complex interrelationships, which strongly affect the reward values of choice alternatives. Decision makers are able to process these relationships and make choices that violate independence, when doing so is beneficial. Without being able to explain the emergence of dimensional dependencies, both single- and multi-strategy models of decision making will remain incomplete.

ACKNOWLEDGEMENTS

Funding was received from the Economic and Social Research Council grant ES/K002201/1.

REFERENCES

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: Journal of the Econometric Society*, *21*(4), 503–546.

Alba, J. W., & Marmorstein, H. (1987). The effects of frequency knowledge on consumer decision making. *Journal of Consumer Research*, *14*(1), 14–25.

Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review*, *120*(3), 522–543.

Bhatia, S. (In Press). Choice rules and accumulator networks. *Decision*.

Birnbaum, M. H. (1974). The nonadditivity of personality impressions. *Journal of Experimental Psychology*, *102*(3), 543–561.

Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review*, *115*(2), 463–501.

Birnbaum, M. H., Parducci, A., & Gifford, R. K. (1971). Contextual effects in information integration. *Journal of Experimental Psychology*, *88*(2), 158–170.

Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: making choices without trade-offs. *Psychological Review*, *113*(2), 409.

Bröder, A., Glöckner, A., Betsch, T., Link, D., & Ettlin, F. (2013). Do people learn option or strategy routines in multi-attribute decisions? The answer depends on subtle factors. *Acta Psychologica*, *143*(2), 200–209.

Bröder, A., Newell, B. R., & Platzer, C. (2010). Cue integration vs exemplar-based reasoning in multi-attribute decisions from memory: A matter of cue representation. *Judgment and Decision Making*, *5*, 326–338.

Davis-Stober, C. P., Dana, J., & Budescu, D. V. (2010). Why recognition is rational: Optimality results on single-variable decision rules. *Judgment and Decision Making*, *5*(4), 216–229.

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, *34*, 571–582.

DeKay, M. L., Patiño-Echeverri, D., & Fischbeck, P. S. (2009). Distortion of probability and outcome information in risky decisions. *Organizational Behavior and Human Decision Processes*, *109*(1), 79–92.

DeLosh, E. L., Busemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: The sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(4), 968.

Dhami, M. K. (2003). Psychological models of professional decision making. *Psychological Science*, *14*(2), 175–180.

Diederich, A. (1997). Dynamic stochastic models for decision making with time constraints. *Journal of Mathematical Psychology*, *41*, 260–274.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*(2), 179–211.

Fishburn, P. C. (1974). Lexicographic orders, utilities and decision rules: A survey. *Management Science*, *20*(11), 1442–1471.

Fishburn, P., & Wakker, P. (1995). The invention of the independence condition for preferences. *Management Science*, *41*(7), 1130–1144.

Garcia-Retamero, R., Hoffrage, U., & Dieckmann, A. (2007). When one cue is not enough: Combining fast and frugal heuristics with compound cue processing. *The Quarterly Journal of Experimental Psychology*, *60*(9), 1197–12.

Garcia-Retamero, R., Hoffrage, U., Dieckmann, A., & Ramos, M. (2007). Compound cue processing within the fast and frugal heuristics approach in nonlinearly separable environments. *Learning and Motivation*, *38*(1), 16–34.15.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast-and-frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.

Gigerenzer, G., & Todd, P. M. (1999). Simple heuristics that make us smart. New York: Oxford University Press.

Glöckner, A., & Betsch, T. (2008). Modeling option and strategy choices with connectionist networks: Towards an integrative model of automatic and deliberate decision making. *Judgment and Decision Making*, *3*(3), 215–228.

Glöckner, A., Betsch, T., & Schindler, N. (2010). Coherence shifts in probabilistic inference tasks. *Journal of Behavioral Decision Making*, *23*(5), 439–462.

Hoffrage, U., Garcia-Retamero, R., & Czienskowski, U. (2008). Compound cue processing in linearly and nonlinearly separable environments. *Psychological Record*, *58*(2), 10.

Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, *128*, 3–31.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, *9*(1), 90–98.

Juslin, P., Karlsson, L., & Olsson, H. (2008). Information integration in multiple cue judgment: A division of labor hypothesis. *Cognition*, *106*(1), 259–298.

Juslin, P., & Persson, M. (2002). PROBabilities from EXemplars (PROBEX): A "lazy" algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, *26*(5), 563–607.

DOI: 10.1002/bdm

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.
Karelaia, N. (2006). Thirst for confirmation in multi-attribute choice: Does search for consistency impair decision performance? *Organizational Behavior and Human Decision Processes*, *100*(1), 128–143.
Katsikopoulos, K. V., & Gigerenzer, G. (2008). One-reason decision-making: Modeling violations of expected utility theory. *Journal of Risk and Uncertainty*, *37*(1), 35–56.
Keeney, R. L., & Raiffa, H. (1993). Decisions with multiple objectives: Preferences and value trade-offs. New York: Cambridge University Press.
Kosko, B. (1988). Bidirectional associative memories. *IEEE Transactions on Systems, Man, and Cybernetics*, *18*(1), 49–60.
Lee, M. D., & Cummins, T. D. (2004). Evidence accumulation in decision making: Unifying the "take the best" and the "rational" models. *Psychonomic Bulletin & Review*, *11*(2), 343–352.
Luan, S., Schooler, L. J., & Gigerenzer, G. (2011). A signal-detection analysis of fast-and-frugal trees. *Psychological Review*, *118*(2), 316.
Marewski, J. N., & Link, D. (2014). Strategy selection: An introduction to the modeling challenge. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*(1), 39–59.
Marewski, J. N., & Schooler, L. J. (2011). Cognitive niches: An ecological model of strategy selection. *Psychological Review*, *118*(3), 393.
Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*(3), 207.
Mellers, B. A. (1980). Configurality in multiple-cue probability learning. *The American Journal of Psychology*, *93*(3), 429–443.
Mellers, B. A., Richards, V., & Birnbaum, M. H. (1992). Distributional theories of impression formation. *Organizational Behavior and Human Decision Processes*, *51*(3), 313–343.
Newell, B. R. (2005). Re-visions of rationality? *Trends in Cognitive Sciences*, *9*(1), 11–15.
Newell, A., & Simon, H. A. (1972). Human problem solving (Vol. 104, No. 9). Englewood Cliffs, NJ: Prentice-Hall.
Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*(1), 104.
Olsson, A. C., Enkvist, T., & Juslin, P. (2006). Go with the flow: How to master a nonlinear multiple-cue judgment task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(6), 1371.
Pachur, T., & Olsson, H. (2012). Type of learning task impacts performance and strategy selection in decision making. *Cognitive Psychology*, *65*(2), 207–240.
Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). The adaptive decision maker. New York: Cambridge University Press.
Persson, M., & Rieskamp, J. (2009). Inferences from memory: Strategy- and exemplar-based models compared. *Acta Psychologica*, *130*, 25–37.
Platzer, C., & Bröder, A. (2013). When the rule is ruled out: Exemplars and rules in decisions from memory. *Journal of Behavioral Decision Making*, *26*(5), 429–441.
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.
Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*(2), 207.
Rubinstein, A. (2012). Lecture notes in microeconomic theory: The economic agent. Princeton, NJ: Princeton University Press.
Russo, J. E., & Dosher, B. A. (1983). Strategies for multiattribute binary choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*(4), 676.
Russo, J. E., Meloy, M. G., & Medvec, V. H. (1998). Predecisional distortion of product information. *Journal of Marketing Research*, *35*(4), 438–452.
Shepard, R. N., Hovland, C. L., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, *75*, 1–42.
Simon, D., Krawczyk, D. C., & Holyoak, K. J. (2004). Construction of preferences by constraint satisfaction. *Psychological Science*, *15*(5), 331–336.
Sutton, R. S., & Barto, A. G. (1998). Reinforcement learning: An introduction. Cambridge, MA: MIT press.
Svenson, O. (1979). Process descriptions of decision making. *Organizational Behavior and Human Performance*, *23*(1), 86–112.
Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, *12*(03), 435–467.
Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review*, *121*(2), 179–205.
Tversky, A. (1969). Intransitivity of preferences. *Psychological Review*, *76*, 31–48.
Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*(4), 281.
Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, *111*(3), 757–769.
Varian, H. R. (1992). Microeconomic analysis. New York: W. W. Norton & Company.
von Helversen, B., & Rieskamp, J. (2008). The mapping model: A cognitive theory of quantitative estimation. *Journal of Experimental Psychology: General*, *137*(1), 73.
von Helversen, B., & Rieskamp, J. (2009). Models of quantitative estimations: Rule-based and exemplar-based processes compared. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(4), 867.
Wedell, D. H. (1991). Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*(4), 767.

Authors' biography:

Sudeep Bhatia is an Assistant Professor in the Department of Psychology at the University of Pennsylvania. His research focuses on judgment and decision making, with an emphasis on modeling the cognitive processes at play in judgments and decisions.

Authors' address:

Sudeep Bhatia, University of Pennsylvania, Philadelphia, PA USA.

DOI: 10.1002/bdm