Decision
2015, Vol. 2, No. 4, 280–305

2325-9965/15/$12.00 http://dx.doi.org/10.1037/dec0000037

Transitivity in Context: A Rational Analysis of Intransitive Choice and Context-Sensitive Preference

Johannes Müller-Trede
University of California, San Diego

Shlomi Sher
Pomona College

Craig R. M. McKenzie
University of California, San Diego

Several apparent violations of transitivity have been reported in the literature on decision making. However, these effects have been shown to be compatible with random preference models, in which preferences are transitive at each point in time but vary at random over time. Such models imply that choice proportions will conform to a set of conditions called the triangle inequalities, and no clear triangle inequality violations have been empirically demonstrated to date. This article examines a broader class of choice models—"context-sensitive preference models"—in which the current and prior history of choice contexts can systematically influence decision makers' stochastic preferences. These models generate violations of the triangle inequalities even when preferences are always transitive. Furthermore, the article develops an analysis of decision making under incomplete information, in which rational decision makers draw inferences from the present choice context, but have limited memory for past contexts. It is shown that such decision makers can exhibit intransitive choice cycles of arbitrary magnitude as a result of context-dependent switching between transitive preference orders. Two experiments test the model's predictions, and clear violations of the triangle inequalities are observed.

Keywords: context effects, rationality, stochastic choice, transitivity, triangle inequality

Supplemental materials: http://dx.doi.org/10.1037/dec0000037.supp

Preferences are transitive if for any three options *A*, *B*, and *C*, when *A* is preferred to *B*, and *B* is preferred to *C*, then *A* is also preferred to *C*. The assumption that preferences are transitive plays a fundamental role in models of decision making and is widely accepted as a basic requirement of rational choice. Nonetheless, a large body of experimental research, building on the pioneering studies of May (1954) and Tversky (1969), appears to demonstrate robust intransitivities of preference.

Researchers studying transitivity can only observe the choices people make, however, and must infer underlying preferences from observed choices. This inference problem is nontrivial because people do not always choose consistently. Faced with multiple choices between *A* and *B*, for example, the same decision maker (DM) may choose *A* two thirds of the time, and choose *B* one third of the time. What conclusions can be drawn about the DM's "true" preferences? One solution to this problem is to infer preference from majority choice: If *A* is usually chosen over *B*, then *A* is preferred to *B* (Block & Marschak, 1960; Marschak, 1960). Accordingly, a DM's choices are said to satisfy *weak stochastic transitivity* if, whenever *A* is chosen at least as often as *B* in repeated pairwise choice, and *B* is chosen at least as often as *C*, then *A* is chosen at least as often as *C*. This

This article was published Online First May 11, 2015.

Johannes Müller-Trede, Rady School of Management, University of California, San Diego; Shlomi Sher, Department of Psychology, Pomona College; Craig R. M. McKenzie, Rady School of Management and Department of Psychology, University of California, San Diego.

The first two authors contributed equally to this work. This research was supported by National Science Foundation Grant SES-1062070. We thank Michael Birnbaum, Daniel Cavagnaro, Clintin Davis-Stober, Robin Hogarth, and Mike Regenwetter for helpful discussions and Lim Leong for valuable research assistance.

Correspondence concerning this article should be addressed to Johannes Müller-Trede, Rady School of Management, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093. E-mail: jmullertrede@ucsd.edu

was the standard adopted by Tversky (1969), who identified a subset of participants who violated weak stochastic transitivity (but see Iverson & Falmagne, 1985).

It has been pointed out, however, that if preferences are dynamic (i.e., a person may have different preference orders at different times), weak stochastic transitivity can be violated even by a DM who has a definite transitive preference order at each point in time (Loomes & Sugden, 1995; Regenwetter, Dana, & Davis-Stober, 2010). Regenwetter, Dana, and Davis-Stober (2011) recently analyzed a class of choice models called *random preference models*, in which a DM randomly draws a preference order from a distribution over transitive orders at each point in time.[1] In these models, the probability of selecting a preference order is unaffected by the choice context—that is, the set of options presented to the DM in the current choice problem. While random preference models allow for violations of weak stochastic transitivity, the choice behavior they induce must satisfy a set of conditions called the *triangle inequalities*, explained below. Regenwetter and colleagues analyzed new as well as previously published data on transitivity (including the data in Tversky, 1969), and argued that they are in fact consistent with the triangle inequalities, and therefore with random preference models (cf. Regenwetter et al., 2014). They conclude that there is little evidence for intransitivity of preference.[2]

In this article, we consider a broader class of stochastic choice models, in which the probability attached to a preference order can depend on present and past choice contexts. These *context-sensitive preference models* can lead to intransitive choice cycles which violate the triangle inequalities, even when underlying preferences are always transitive. The models are motivated by a range of experimental context effects, which suggest that people rely on the context of choice to "construct" their preferences online (Payne, Bettman, & Schkade, 1999; Slovic, 1995). Furthermore, we identify conditions under which context-sensitive preference is normatively justified. We show that the choices of a rational DM with limited prior knowledge, who constructs transitive preferences based on reasonable inferences under memory constraints, can exhibit intransitive cycles that violate the triangle inequalities. The choice problems we study are closely related to multiattribute choice problems in which May (1954) and Kivetz and Simonson (2000) obtained evidence for behavioral intransitivity. We show that, in extended repeated-choice designs, these choice problems lead to significant violations of both weak stochastic transitivity and the triangle inequalities at the within-subjects level. To our knowledge, this is the first clear demonstration of a violation of the triangle inequalities.

In the remainder of this article, we first briefly review the literature on (in)transitive choice and preference. We then examine situations where, we argue, rational choice can be captured by a context-sensitive preference model in which preferences are always transitive but choices can exhibit robust intransitive cycles. In two experiments, we show that people exhibit the choice patterns predicted by the model. The General Discussion considers normative and descriptive implications of this work for the study of (in)transitivity.

Interpreting Intransitive Choice

Transitivity of preference is a fundamental assumption of standard models of rational choice (but see Anand, 1993; Fishburn, 1991). Theoretically, stable cycles of intransitivity expose DMs to "money pumps," in which a DM repeatedly pays small premiums to exchange A for B, B for C, and C for A, losing money but gaining nothing in each cycle (Davidson, McKinsey, & Suppes, 1955). Empirically, people are sometimes disturbed to learn that their choices have been intransitive (Tversky, 1969) and may modify their choices to eliminate the intransitivity (Luce & Raiffa, 1957).

Nonetheless, a number of apparent violations of transitivity have been documented (e.g., Budescu & Weiss, 1987; Kivetz & Simonson,

[1] Regenwetter et al. (2010, 2011) refer to these models as "mixture models." The models are also equivalent to (distribution-free) "random utility models" (cf. Block & Marschak, 1960; Regenwetter et al., 2010). We adopt the terminology of "random preference models" to emphasize that the selection of a preference order is random with respect to the choice context.

[2] The issue remains a point of contention, however. See Birnbaum (2011); Birnbaum and Gutierrez (2007), and Cavagnaro and Davis-Stober (2014), who reach similar conclusions, as well as Myung, Karabatsos, and Iverson (2005) and Tsai and Böckenholt (2006), whose conclusions differ.

2000; Lindman & Lyons, 1978; Loomes, Starmer & Sugden, 1991; May, 1954; Montgomery, 1977; Ranyard, 1977; Tversky, 1969). Studies of transitivity usually examine choices between options which are characterized on multiple attributes or between risky gambles which vary in the probability and magnitude of a prize. Some studies employ between-subjects designs, in which transitivity is analyzed at the group level. For reasons discussed below, between-subjects tests of transitivity are problematic, however, and the most informative studies employ within-subjects designs. In within-subjects designs, each choice pair is presented multiple times to each participant, and repetitions of individual choice pairs are typically separated by filler items in order to minimize memory carryover and maximize cross-choice independence. Transitivity is then analyzed at the level of the individual participant. Some researchers have used initial screening procedures to select participants whose choices are likely to exhibit intransitivity, and subsequently verified this pattern in a more extensive series of choices (Ranyard, 1977; Tversky, 1969).

Regardless of whether transitivity is studied at the individual or group level, any empirical study of transitivity confronts the important inference problem described above. How can the "true" preferences of participants be inferred from fluid choice data, which may exhibit both systematic and random variability across time? We now examine the two aforementioned criteria, weak stochastic transitivity and the triangle inequalities, in greater detail.

Weak stochastic transitivity is the better-known criterion, and has been especially influential in tests of transitivity. Letting P_{xy} denote the probability that option x is chosen over option y, weak stochastic transitivity holds when, for all option triples $\{A, B, C\}$, if $P_{AB} \geq .5$ and $P_{BC} \geq .5$, then it is also the case that $P_{AC} \geq .5$. The criterion is problematic, however, as suggested above, because it can be violated by the choices of a DM with transitive preferences that change over time. To see how, consider a DM whose preferences fluctuate as follows. A random third of the time, the DM has a transitive preference order of $A \succ B \succ C$.[3] Another random third of the time, the DM has the transitive order $B \succ C \succ A$, and the rest of the time, the DM has the transitive order $C \succ A \succ B$. In a repeated-choice design, this DM will choose A over B two thirds of the time, choose B over C two thirds of the time, and choose C over A two thirds of the time. The DM would therefore violate weak stochastic transitivity—yet the DM's preferences are neither intransitive nor counternormative. Models of rational choice generally allow preferences to change over time (Anand, 1993; Bar-Hillel & Margalit, 1988; Regenwetter et al., 2010).

This aggregation problem substantially complicates the interpretation of violations of weak stochastic transitivity in within-subjects designs. The problem is only compounded in between-subjects designs, in which different participants' preference orders can combine to yield apparent intransitivity at the group level, even if each individual's preference order is both transitive and stable across time (Condorcet, 1785).

The triangle inequalities supply an alternative criterion which does not suffer from this problem. A DM's choices satisfy these inequalities if, for every option triple $\{A, B, C\}$, it is the case that

$$P_{AB} + P_{BC} - P_{AC} \leq 1. \tag{1}$$

Many models of transitive preference, including the random preference models advanced by Regenwetter et al. (2011), imply the triangle inequalities, which have long been recognized as a fundamental criterion for transitivity (Block & Marschak, 1960; Morrison, 1963).[4] Conversely, for choice sets of five or fewer options, if choice probabilities satisfy the triangle inequalities, there exists a random preference model that can induce these choice probabilities (Cohen & Falmagne, 1990; Falmagne, 1978; Fiorini, 2001).

The behavioral literature, however, has largely neglected the triangle inequalities, and

[3] We adopt the following standard notation. Given a set S of options, the binary relation $\succcurlyeq$ defined on S denotes *weak preference*. *Strict preference*, $x \succ y$, obtains if and only if $x \succcurlyeq y$ and not $y \succcurlyeq x$. *Indifference* ($x \sim y$) holds when both $x \succcurlyeq y$ and $y \succcurlyeq x$. Except where otherwise indicated, we assume that $\succcurlyeq$ is complete (i.e., for all $x, y \in S$, $x \succcurlyeq y$ or $y \succcurlyeq x$).

[4] To appreciate why random mixtures of transitive orders must satisfy the triangle inequalities, consider the example of a DM who prefers A to B 90% of the time, and also prefers B to C 90% of the time ($P_{AB} = P_{BC} = .9$). Suppose the times at which $A \succ B$ and the times at which $B \succ C$ have minimal overlap. In this case, they will overlap 80% of the time. In this 80% overlap, a transitive DM must also prefer A to C. That is, $P_{AC} \geq .8$. If there is greater overlap, the transitive DM will be constrained to prefer A to C even more often. Therefore, $P_{AB} + P_{BC} - P_{AC} \leq .9 + .9 - .8 = 1$.

rigorous statistical methods for testing the inequalities have only recently been developed (Davis-Stober, 2009). Applying these methods in a wide-ranging critique of the empirical literature, Regenwetter et al. (2010, 2011, 2014) examined old and new data sets and found little evidence for violations of the triangle inequalities. They concluded that there are no convincing demonstrations of intransitive preference, as existing data appear to be consistent with random preference models in which preferences vary over time but are always transitive.

Later, we report two experiments which, to our knowledge, demonstrate the first clear empirical violations of the triangle inequalities. The observed choice data therefore cannot be explained as resulting from transitive preferences that change randomly over time. Nonetheless, we argue that even these strong findings do not establish intransitivity of preference. Instead, we show that the observed violations of the triangle inequalities are consistent with transitive preferences that depend on the choice context and change systematically over time.

Context-Sensitive Preference Models

Conceptually, our analysis builds on the idea of mixtures of preference orders advocated by Regenwetter et al. (2010, 2011). Like a random preference model, a context-sensitive preference model assumes that DMs can be characterized by some fixed set of potential preference orders, and that at each point in time, one of these preference orders is selected. The models differ in how preference orders are selected, however. Random preference models posit a probability distribution over preference orders that is independent of the choice context. Independent draws from this distribution determine which order is selected, and choices are made according to the currently selected order. In a context-sensitive preference model, on the other hand, the probability that a given preference order is selected for a choice may depend on the choice context that the DM is currently facing, defined by the set of options under evaluation, as well as the history of choice contexts to which the DM has previously been exposed.

In light of a wide range of context effects that have been documented in the literature on decision making (see, e.g., Payne et al., 1999), context-sensitive preferences are psychologically plausible. Moreover, strong context effects can be normatively appropriate when prior knowledge is limited and the sample of available options comprises the DM's best evidence about the distribution of relevant attributes in the population of interest.

Sher and McKenzie (2014) recently developed a normative analysis of such choice situations. In their "options-as-information" model, rational DMs begin with an underspecified prior model of the distribution of relevant attributes. When a sample of options is encountered, DMs update their prior model to obtain a "posterior model" of the attribute space, which incorporates inferences drawn from the option sample. The attribute values of the available options are then normalized relative to the posterior model of the attribute distribution (e.g., to obtain z scores or percentiles).[5] Finally, they are combined using a weighted-additive rule, in which fixed weights are applied to the normalized attribute values, and the weighted normalized attribute values are summed to yield an overall evaluation of each option. This evaluation process is consistent with multiattribute utility theory, the predominant normative and prescriptive model of multiattribute choice (Edwards & Barron, 1994; Keeney & Raiffa, 1993). Importantly, in the options-as-information analysis, different option samples can lead to different models of the attribute space, and therefore to different evaluations and preferences. As a result, preferences are systematically context-dependent. Sher and McKenzie (2014) have applied this model to explain joint-separate reversals traditionally attributed to the evaluability hypothesis (Hsee, 1996; Hsee et al., 1999).

To see how context-dependent preferences can lead to an intransitive choice cycle, even when the DM's preferences are always complete and transitive, consider the sound systems in Table 1. Each sound system is defined on three (fictitious) attribute dimensions (harmonic

[5] A related process of normalization to context has been posited by other theories of choice and evaluation, including the Decision by Sampling model (Stewart, Chater, & Brown, 2006) and range-frequency theory (Parducci, 1965). The options-as-information model provides a normative analysis of normalization, in which attribute values are interpreted relative to an inferred model of the attribute space. It predicts that variability in the evaluation of options can be explained by variability in beliefs about the world (Sher & McKenzie, 2014).

Table 1
Stimuli for Experiment 1

	Sound system A	Sound system B	Sound system C
Harmonic range	10 mill	26 mill	16 mill
Sound depth	17 sones	12 sones	29 sones
Acoustic power	3.4 phons	2.0 phons	1.4 phons

range, sound depth, and acoustic power), each with its own (fictitious) unit of measurement. We consider a DM who may have some notion of the meaning of these attributes, and is aware that higher values are better on each dimension, but knows nothing about the distribution of each attribute in the larger population of sound systems. Thus, for example, the DM does not know whether 10 mill is a good or a poor value on the dimension of harmonic range.

In the options-as-information analysis, the DM first receives a sample of options—for example, $\{A, B\}$. On the basis of this sample, the DM draws inferences to form a posterior model of the distribution of each attribute. The DM then interprets each attribute value in light of this updated model. That is, each absolute attribute level (such as A's level of 10 mill on harmonic range) is transformed into a "normalized value" that reflects its attractiveness relative to the posterior model of the attribute. The weights assigned to the three attributes, on the other hand, are not affected by the option sample. Finally, the DM evaluates each option by computing a weighted sum of its model-normalized attribute values, and chooses the option with the highest value.

A formal analysis of this inference-and-evaluation process is provided in the Appendix. Here, we provide an intuitive sketch of the model and its implications. Because the DM has no prior knowledge of the attribute distributions (e.g., the mean value and the range of variation), the model assumes that similar inferences are drawn from all samples. A's normalized value for harmonic range, for example, relative to the model inferred from the option sample $\{A, B\}$, is assumed to be similar to C's normalized value for harmonic range relative to the model inferred from $\{B, C\}$, because both constitute the smaller of the two values for harmonic range in their respective two-option samples. Second, we assume that the DM's memory capacity is limited, so that, under the conditions tested in the experiment, the DM has forgotten about past choice pairs when a new choice pair is encountered. Under these assumptions, the options-as-information analysis is consistent with certain patterns of choice behavior but not others.

Consider, first, a DM who assigns roughly equal weights to the three attributes. On receiving the option sample $\{A, B\}$, the DM is likely to adopt a posterior model in which B is above and A is below average on harmonic range, while A enjoys analogous advantages on sound depth and acoustic power. The DM's posterior model is schematically depicted in the middle column of Figure 1. The figure shows how the

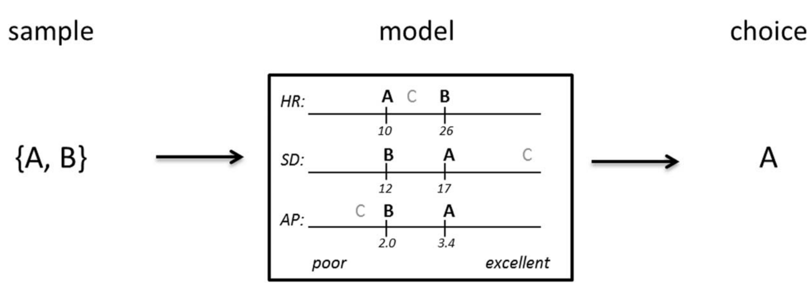

Figure 1. Schematic representation of an options-as-information analysis of the choice problem in Table 1. Horizontal lines represent posterior models of the distribution for the three attributes—harmonic range (HR), sound depth (SD), and acoustic power (AP)—when sound systems A and B are sampled. The position that would be occupied by the nonsampled option C, relative to the posterior model inferred from $\{A, B\}$, is also indicated. The symmetric form of the three models reflects the absence of prior knowledge about attribute distributions. The normative analysis predicts that A will be chosen if the weight the DM attaches to harmonic range does not exceed the sum of the weight attached to the other two attributes.

sampled options A and B (as well as the nonsampled option C) relate to this posterior model. Because the weights assigned to the attributes are about equal, B's model-normalized advantage on harmonic range will be offset by A's model-normalized advantage on the other two attributes, leading to a preference for A over B when $\{A, B\}$ is seen. The same process will lead the DM to adopt a model that favors B when $\{B, C\}$ is seen, and to adopt a model that favors C when $\{A, C\}$ is seen. The DM's choice behavior will therefore exhibit an intransitive cycle across the three contexts. Indeed, if attribute weights do not change over time, the DM will always choose A over B, always choose B over C, and always choose C over A in repeated pairwise choice.

But while the DM's choices exhibit a robust intransitive cycle, the DM's underlying preferences are never intransitive. To illustrate, consider a DM who computes z scores from option samples and assigns equal weights to the three attributes, and let $\succcurlyeq_S$ denote the preference order that results when sample S is received. As detailed in the Appendix, on encountering the sample $\{A, B\}$, the DM arrives at the transitive preference order $C \succ_{\{A,B\}} A \succ_{\{A,B\}} B$.[6] On sampling $\{B, C\}$, the DM draws different inferences, which result in the preference order $A \succ_{\{B,C\}} B \succ_{\{B,C\}} C$. Finally, on sampling $\{A, C\}$, the resulting order is $B \succ_{\{A,C\}} C \succ_{\{A,C\}} A$. That is, when attribute weights are equal in the normative model, the samples $\{A, B\}$, $\{B, C\}$, $\{A, C\}$ induce different preference orders which include the relations $A \succ_{\{A,B\}} B$, $B \succ_{\{B,C\}} C$, and $C \succ_{\{A,C\}} A$. This leads to an intransitive cycle of choices, even though no sample gives rise to intransitive preferences.

In what follows, we use the term *choice pattern* to refer to behavior in pairwise choice. In general, we designate by $xyzx$ an intransitive pattern of observed choices in which x is selected over y, y over z, and z over x. xyz refers to a transitive pattern of behavior in which x is selected over both y and z, and y is selected over z. For example, the intransitive choice pattern described in the previous paragraph is denoted by *ABCA*.

Table 2 shows that in addition to this intransitive choice pattern, there are also transitive patterns that are consistent with the inferential process described above. This is because a DM who places sufficiently strong weight on one attribute relative to the others will always choose the option that is favored on that attribute. For example, DMs who care strongly about harmonic range will choose B over C, C over A, and B over A. Regardless of the option sample they encounter and the inferences they draw, such DMs will choose according to the transitive pattern *BCA*. Likewise, DMs who care strongly about sound depth will exhibit a consistent *CAB* pattern, and those who place a great emphasis on acoustic power will exhibit an *ABC* pattern.

At the same time, some choice patterns are incompatible with the inferential model. DMs cannot exhibit a consistent *CBA* pattern, for example. Intuitively, a DM who chooses C over B when $\{B, C\}$ is sampled must greatly value sound depth, because it is the only attribute on which C outperforms B. Greatly valuing sound depth is inconsistent with choosing B over A when $\{A, B\}$ is sampled, however, because A outperforms B in terms of this attribute. In this fashion, the model rules out the transitive

Table 2
Normative Analysis of Choice Patterns for the Problem in Table 1

Choice pattern	Transitive?	Permitted?	Attribute weights
ABC	Y	Y	$w_a > w_h + w_s$
BCA	Y	Y	$w_h > w_s + w_a$
CAB	Y	Y	$w_s > w_a + w_h$
ACB	Y	N	NA
BAC	Y	N	NA
CBA	Y	N	NA
ABCA	N	Y	$\forall\, i, j, k, w_i < w_j + w_k$
ACBA	N	N	NA

Note. The table lists all possible transitive and intransitive choice patterns. For each choice pattern, the table indicates whether it is permitted by the normative context-sensitive model. For permitted choice patterns, the rightmost column gives the attribute weights that would lead a DM to exhibit the pattern. w_a, w_h, and w_s are the weights assigned to acoustic range, harmonic range, and sound depth, respectively. For the intransitive cycle *ABCA*, the weighting condition assumes that $i, j, k \in \{h, s, a\}$ and $i \neq j \neq k$. See Appendix for derivations.

[6] From the DM's point of view, C is a purely hypothetical or imaginary option, but it can nonetheless be evaluated in terms of the posterior model of the attribute space inferred from $\{A, B\}$. Furthermore, the posterior model, together with the DM's attribute weights, imply a valuation for a hypothetical option like C.

choice patterns *CBA*, *BAC*, and *ACB*, as well as the intransitive pattern *ACBA*.

This analysis is developed in more detail in the Appendix. The analysis rests on two critical assumptions—that sample-based inferences are scale-independent, and that the DM's memory is limited in capacity. The former assumption seems reasonable in situations where prior familiarity with the relevant attribute scales is minimal: The larger attribute value in the sample will be assigned a higher standing relative to the inferred posterior model, but the DM does not have prior knowledge as to whether the specific numerical difference between sampled values is large or small.

The assumption of limited-capacity memory is also essential to the analysis. If DMs retained perfect memory for previously sampled choice pairs, their posterior models for each attribute should stabilize over the course of a sequence of repeated choices. In the long run, this would lead to a consistent transitive choice pattern as the sequence proceeds. Human memory is far from perfect, however, and as noted above, experiments that examine transitivity at the individual level are deliberately designed to minimize memory carryover between repeated presentations of a choice problem (i.e., by including filler items and distributing choices across multiple experimental sessions). While these design features are motivated by the need to ensure that repeated pairwise choices are independent, they also minimize the likelihood of sustained learning over the course of the experiment. With each new sample, experimental participants may thus encounter the attribute values "as if for the first time." As a result, a normatively appropriate evaluation process can yield stable cycles of intransitive choice in typical within-subjects designs.

An options-as-information analysis of pairwise choices between the stimuli in Table 1 thus leads to a model of context-based inference and preference which is consistent with three transitive choice patterns and one intransitive pattern. The intransitive cycle, which arises in the simple case in which the three attributes are assigned similar weights, can be of arbitrary strength, and can therefore violate any criterion for transitivity, including weak stochastic transitivity and the triangle inequalities. In what follows, we show that this context-sensitive preference model is empirically testable, and we discuss how experimental tests of the model relate to tests of the two transitivity criteria.

Testing Context-Sensitive Preference Models

The recent mathematical advances of Davis-Stober (2009) have made statistical tests of weak stochastic transitivity and the triangle inequalities relatively straightforward, especially for small choice sets. These advances also allow us to test context-sensitive preference models. The approach is best illustrated in a geometric representation of the models (cf. Iverson & Falmagne, 1985; Regenwetter et al., 2010).

As before, let P_{xy} denote the probability that x is chosen over y in a pairwise choice between the two options. The unit cubes in Figure 2 depict probability spaces in which the three axes represent the probabilities P_{AB}, P_{BC}, and P_{AC}. The cubes' eight vertices correspond to the six transitive choice patterns and two intransitive patterns defined over the choice set $\{A, B, C\}$. In a within-subjects design, in which DMs make repeated pairwise choices between the stimuli, the choices of a perfectly consistent DM who always chooses the same option from each pair would lie on one of these vertices.

In this probability space, weak stochastic transitivity can be represented as a union of half-unit cubes, shown as the shaded volume in the top-left panel of Figure 2 (Iverson & Falmagne, 1985). The choices of a DM who satisfies weak stochastic transitivity would fall inside this shaded volume of the space. Similarly, random preference models can be represented in the probability space by a shape known as the linear ordering polytope, shown as the shaded volume in the top-right panel of Figure 1 (cf. Fiorini, 2001; Regenwetter et al., 2010). This shape is characterized by the triangle inequalities, so that the choice frequencies of a DM who satisfies the inequalities would fall inside the shaded volume. Note that the unshaded volumes in the two top panels of Figure 2 describe neighborhoods of the two vertices which correspond to the intransitive cycles—violations of both weak stochastic transitivity and of the triangle inequalities are associated with intransitive behavior.

The context-sensitive preference model developed in the previous section can also be represented in this probability space. In fact, the

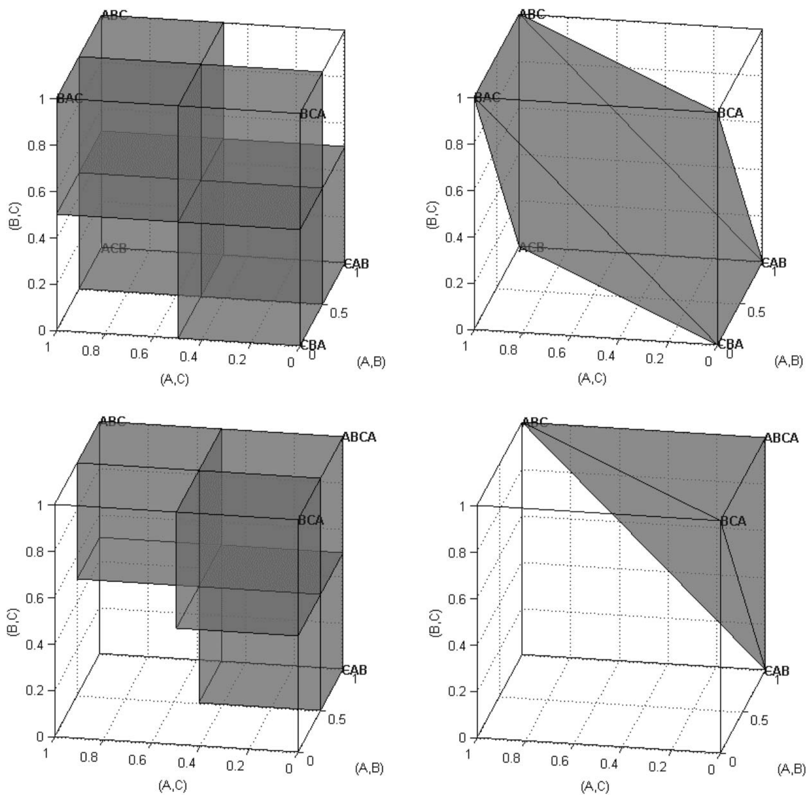

Figure 2. Geometric representations of weak stochastic transitivity (top-left), the triangle inequalities (top-right), and the "noisy response" and "flexible weight" specifications of the context-sensitive preference model (bottom-left and bottom-right, respectively) in a probability space whose axes represent choice probabilities in pairwise choices between *A*, *B*, and *C*. The individually labeled vertices correspond to the pure (i.e., perfectly consistent) choice patterns permitted by each model.

model has two plausible representations, corresponding to two ways in which within-subjects variability in choice behavior can be modeled. As emphasized above, behavior in repeated-choice experiments exhibits temporal inconsistency: The option selected from a given choice pair on one trial may be rejected on a subsequent trial involving the same pair. Traditionally, this temporal inconsistency has been modeled as arising from one of two sources—"noise" in the expression of preference, or preferences that change over time.

The first approach is implicit in the criterion of weak stochastic transitivity. According to this approach, even when "true" preferences are unique and stable, the expression of preference is inherently noisy. A DM who prefers *A* to *B*, for example, may nonetheless sometimes choose *B* over *A*. It is then reasonable to identify "true" preference with modal response (Block & Marschak, 1960; Marschak, 1960). This "noisy-response" approach can be adapted to tests of the context-sensitive preference model. Here, a given DM's attribute weights—and, therefore, the DM's preference order—are assumed to be unique and stable. If the expression of preference were not noisy, the DM's responses would then fall precisely on one of the four vertices corresponding to the four patterns identified as consistent with the context-sensitive preference model in the previous section (the transitive patterns *ABC*, *BCA*, and *CBA*, as well as the intransitive pattern *ABCA*). Allowing for the expression of preference to be noisy, on the other hand, and identifying preference with modal response, yields the probability space representation in the bottom-left panel of Figure 2.

The second approach to modeling within-subjects variability is implicit in the triangle inequality criterion. This approach assumes that the expression of preference is not subject to noise, but allows a DM's preferences to randomly vary. It, too, can be adapted for a test of the context-sensitive preference model. Here, the DM's attribute weights are allowed to randomly vary over time. This random variation arises from volatile tastes, and hence is assumed to be independent of the choice context. Thus, in this "flexible-weights" approach, the DM's dynamic construction of a preference order has both a context-independent component (fluctuating attribute weights) and a context-dependent component (inferences from option samples). At any given moment, behavior is determined both by the DM's current attribute weights and by the DM's current model of the attribute distributions. As a result, the DM's choice behavior, aggregated over time, will be a probabilistic mixture of the four choice patterns identified in the previous section as consistent with the model. This approach to modeling within-subjects variability yields the probability space representation in the bottom-right panel of Figure 2.

In empirically testing the context-sensitive preference model, we evaluate both models of variability described above—the "noisy response" representation (Figure 2, bottom-left) and the "flexible weights" representation (bottom-right). We also test the alternative models (weak stochastic transitivity and random preference models) depicted in the top row of Figure 2. For each participant, we test whether his or her choice data are consistent with each model by applying the statistical methodology developed by Davis-Stober (2009). These tests assess whether the participant's data lie far enough from the model's representation in probability space to count as a significant violation. Formally, all four models can be represented as convex polytopes or unions thereof, as in Figure 2; Davis-Stober (2009) derived likelihood ratio tests for multinomially distributed random variables (e.g., a participant's choice frequencies) under the inequality constraints imposed by such a representation. The key technical problem is that, when observed choice proportions lie outside the polytope, the asymptotic distribution of the log-likelihood ratio test is not a chi-square distribution. Davis-Stober (2009) provides a method for finding the chi-bar-square distribution that needs to be used instead. These methods are employed in testing for violations of the four models in Experiments 1 and 2 below. For additional technical details, we refer the reader to Davis-Stober (2009) and Regenwetter et al. (2010, 2014).

In interpreting tests of the four models, it is important to note that both specifications of the context-sensitive model are more restrictive than both of the alternative models. In Figure 2, the smaller the volume of a representation of a model, the more restrictive the model. Weak stochastic transitivity covers three quarters of the volume of the unit cube (top-left panel), and random preference models cover two thirds of its volume (top-right panel). The noisy-response specification of the context-sensitive preference model, on the other hand, covers half of the cube's volume (bottom-left), and the flexible-weights specification covers only a sixth of its volume (bottom-right): Both versions of the model make relatively specific predictions about the patterns of choice behavior that will be observed at the within-subjects level.[7]

A normative analysis of sample-based belief and preference updating with limited memory thus implies a context-sensitive preference model, in which preferences are always transitive, but violations of both weak stochastic transitivity and the triangle inequalities are possible. When a representation of response variability is specified, the context-sensitive model makes testable predictions in the sample space of relative choice frequencies. Next, we report two experiments designed to test the model and to probe for the triangle inequality violations that it predicts. The first experiment employs the stimuli in Table 1. A second experiment investigates a related multiattribute choice problem, described below, in which attribute information is sometimes missing. In both experiments, choice frequencies are compatible with the context-sensitive preference model (at least in the

[7] The context-sensitive preference model we test is more restrictive than the other models because it only permits underlying preferences that are consistent with the normative analysis of contextual inference. (This restriction outweighs the added flexibility that results from allowing for context-dependence.)

noisy-response specification). Furthermore, as predicted by the options-as-information analysis, triangle inequality violations are common at the individual-participant level.

Experiment 1

Method

Twenty-three undergraduate students at the University of California, San Diego (UCSD) participated in Experiment 1. Participants were recruited in weekly batches via advertisements posted around the campus and received $30 in compensation, and recruitment continued until we reached a sample size at least matching that in Regenwetter et al. (2011). Each participant attended four separate experimental sessions, usually within the same week. One participant was lost to attrition. The remaining 22 participants (59% female) ranged between 17 and 24 years of age ($M = 19.5$).

Participants, seated at individual computer stations, made 75 pairwise choices in each experimental session. Fifteen choice pairs (five each of $\{A, B\}$, $\{B, C\}$, $\{A, C\}$) involved the critical sound system stimuli in Table 1. The remaining 60 choices involved filler items, including other consumer products (digital cameras and microwaves) that were defined on multiple (nonfictitious) attributes, as well as gambles, defined in terms of visually represented probability wheels and hypothetical monetary outcomes. Each choice pair involving the critical stimuli was thus encountered 20 times by each participant across the four experimental sessions. The presentation order (left vs. right) of the options was randomly varied across trials.

At the beginning of each session, participants received brief instructions describing the different kinds of choices they would be making. For choices involving consumer goods, participants were asked to imagine that they were planning to make a purchase, and were told that they would receive some information about product attributes. Participants were not given any information regarding what counts as a good or bad value for any attribute; we assumed that participants would regard higher values on the attributes in Table 1 as better, but this information was not explicitly provided. In each choice problem, two stimuli from the same category were presented in a two-alternative forced choice task (i.e., with no indifference option). Several attention checks—choices between gambles in which one strictly dominated the other—were built into the choice sequence.[8] At the end of the last session, demographic information was elicited, and participants were asked to describe how they had made their decisions during the course of the experiment. The experiment did not include any measures or conditions that are not reported.

Results and Discussion

We begin with an aggregate summary of the pairwise choice data. Collapsing across all participants and all critical trials, there is evidence for the intransitive choice pattern predicted by the context-sensitive preference model. Participants chose A over B 62.5% of the time, B over C 67.7% of the time, and C over A 70.5% of the time. This aggregate effect, however, yields little insight into the choice proportions of individual participants.

These choice proportions are listed in Table 3. For each participant, the table also reports the results of four statistical tests—of weak stochastic transitivity, the triangle inequalities, and the two specifications of the context-sensitive model. We implemented the likelihood ratio tests described above (Davis-Stober, 2009) using the QTEST software (Regenwetter et al., 2014).

We find substantial evidence for violations of both transitivity criteria. The choices of seven of 22 participants descriptively violate weak stochastic transitivity, and four of these violations are significant at the .05-level. Most notably, the choices of these seven participants also descriptively violate the triangle inequalities, and five of these violations are significant. Thus 23% of participants exhibited individually significant triangle inequality violations. This clearly exceeds the test's expected Type I error rate.

By contrast, in its noisy-response specification (Figure 2, bottom-left panel), the context-sensitive preference model fits the data well. Only three participants violated this model de-

[8] Mean accuracy on the attention checks in Experiment 1 was 98%, and all participants had accuracy scores of at least 85%.

Table 3
Results From Experiment 1

Participant	Choice frequencies P_{AB}	P_{BC}	P_{AC}	Modal choice pattern	Weak stochastic transitivity G^2	p value	Random pref. model G^2	p value	Context-sens. pref. model (noisy resp.) G^2	p value	Context-sens. pref. model (flex. weight) G^2	p value
1	0.60	0.90	0.10	*ABCA*	0.81	0.19	6.1	**<.01**	0		0	
2	0.00	0.25	0.00	*CBA*	0		0		5.2	0.06	30.4	**<.01**
3	1.00	0.95	1.00	*ABC*	0		0		0		1.4	0.12
4	0.65	0.55	0.00	*ABCA*	0.2	0.55	1.6	0.10	0		0	
5	0.45	0.15	0.35	*CBA*	0		0		0.2	0.33	18.3	**<.01**
6	0.30	0.65	0.05	*BCA*	0		0		0		0.4	0.26
7	0.20	0.75	0.40	*BCA*	0		0		0		7.2	**<.01**
8	0.85	0.05	0.10	*CAB*	0		0		0		3.2	**0.04**
9	0.45	0.00	0.00	*CBA*	0		0		0.2	0.71	19.5	**<.01**
10	1.00	0.85	1.00	*ABC*	0		0		0		4.4	**0.02**
11	0.00	1.00	0.00	*BCA*	0		0		0		<.01	0.50
12	0.75	0.90	0.00	*ABCA*	5.2	**0.03**	18.3	**<.01**	0		0	
13	0.90	0.90	0.90	*ABC*	0		0		0		0.8	0.19
14	0.85	1.00	0.00	*ABCA*	10.8	**<.01**	33.8	**<.01**	0		0	
15	1.00	0.00	0.05	*CAB*	0		0		0		1.4	0.12
16	0.00	1.00	0.00	*BCA*	0		0		0		<.01	0.50
17	0.85	0.70	0.90	*ABC*	0		0		0		6.0	**<.01**
18	1.00	0.95	0.05	*ABCA*	19.8	**<.01**	35.3	**<.01**	0		0	
19	1.00	0.95	0.95	*ABC*	0		0		0		<.01	0.49
20	0.55	0.75	0.10	*ABCA*	0.2	0.33	1.4	0.12	0		0	
21	0.40	0.85	0.45	*BCA*	0		0		0		1.3	0.12
22	0.95	0.80	0.10	*ABCA*	7.7	**<.01**	15.9	**<.01**	0		0	

Note. pref. = preference; resp. = response; sens. = sensitive; flex. = flexible. P_{xy} refers to the proportion of trials in which a participant chose x from the pair $\{x, y\}$. For each of the four models, the log-likelihood ratio G^2 at the maximum-likelihood estimate is reported. When a model is not a perfect fit, the corresponding p value is indicated; significant violations are marked in bold.

scriptively, and none of these violations was significant at the individual-participant level (p values ranging from .06 to .71). The flexible-weights specification of the model (Figure 2, bottom-right panel), on the other hand, was descriptively violated by 15 participants. Seven of these violations were significant. Recall that this model is the most restrictive of those tested.

We now summarize the range of modal choice patterns in more detail. All participants who behaved intransitively exhibited the *ABCA* cycle that is predicted by the context-sensitive preference model when attribute weights are similar. Of the 15 participants who behaved transitively, 12 exhibited modal choice patterns predicted by the context-sensitive model when attributes have dissimilar weights: five participants exhibited an *ABC* pattern, five exhibited a *BCA* pattern, and two exhibited a *CAB* pattern. The modal responses of the three remaining participants followed a *CBA* pattern, though, as noted above, these departures from the model (in its noisy-response specification) were not significant. Overall, 19 of 22 participants made modal responses predicted by the context-sensitive preference model.

At the within-subjects level, then, the choice data of several participants violated both weak stochastic transitivity and the triangle inequalities, and hence are incompatible both with the noisy expression of a single transitive preference order, and with a random (i.e., context-independent) mixture of transitive preference orders. This finding is noteworthy in light of previous failures to demonstrate clear violations of the triangle inequalities (Regenwetter et al., 2011). The choice data of participants who behaved intransitively, and of those who behaved transitively, are consistent with a context-sensitive preference model in which the expression of preference is subject to noise.

Intransitivity and Missing Attribute Information

We now turn to a second type of choice problem in which reasonable contextual inferences can give rise to cycles of intransitive choice. The problem again involves unfamiliar attributes, but the DM's knowledge is even more limited, as values are selectively unavailable for some options and attributes. A problem of this type is presented in Table 4. In this problem, sound systems vary on the three (fictitious) attributes used earlier, and all options are equated on an additional (nonfictitious) dimension of harmonic distortion. However, each sound system lacks information for one attribute. For example, *A* is defined in terms of harmonic range, sound depth, and harmonic distortion, but information about acoustic power is missing. This problem is similar to those investigated by Kivetz and Simonson (2000), whose findings are discussed below.

An inferential analysis of this choice problem depends on how DMs interpret the missing values. We consider two cases. First, a DM may regard the absence of a value as a neutral fact. For such a DM, nothing new is learned about a product from discovering that information is unavailable on a particular attribute. Alternatively, a DM may see the absence of a value as suspicious, as suggested by a body of evidence indicating that missing attribute information is often interpreted as undesirable (e.g., Huber & McCann, 1982; Johnson & Levin, 1985).[9] For such a DM, omissions are potentially nonrandom and strategic.

Consider, first, a DM for whom learning that attribute information is missing is equivalent to not learning anything about the product. If the DM's prior knowledge of attribute distributions is minimal, inferences from option-samples will lead to an intransitive cycle in which the DM chooses *A* from {*A*, *B*}, *B* from {*B*, *C*}, and *C* from {*A*, *C*}. Suppose, for example, that such a DM is confronted with the set {*A*, *B*}. The DM has no reason to believe that *A*'s sound depth of 14 sones is better or worse than *B*'s omitted value on the same dimension. Likewise, the DM will not regard *B*'s value of 3.4 phons on acoustic power as better or worse than *A*'s omitted value on this dimension. However, relative to whatever model of harmonic range the DM infers from the sample {*A*, *B*}, *A* is plainly superior on this attribute. Consequently, the DM who encounters the sample {*A*, *B*} will construct a preference order in which *A* is preferred to *B*. This process is illustrated in Figure 3. The same process will lead the DM to select *B* over *C*, and to select *C* over *A*, in pairwise choice.

Next, we turn to a DM for whom missing values are suspicious. Relative to the posterior model inferred from an option sample, this DM treats missing values as presumptively below-average. On this interpretation, the missing values problem closely resembles the choice problem in Experiment 1, and is subject to a similar analysis. For example, on receiving the choice set {*A*, *B*}, *A* will be assigned a higher model-normalized value on harmonic range (e.g., *A* is above and *B* is below average relative to this posterior model) and on sound depth (*A* is about average and *B* is inferior), while *B* is assigned a higher model-normalized value on acoustic power (where *B* is about average and *A* is inferior). Thus, given {*A*, *B*}, *A* is favored over *B* on two model-normalized attributes; given {*B*, *C*}, *B* is favored on two attributes; and given {*A*, *C*}, *C* is favored on two attributes.

When missing values are viewed with distrust, the normative analysis of Experiment 1 generalizes, with minor modifications, to the choice problem in Table 4, and the same qualitative results hold (see Appendix for details). If the weights assigned to the attributes are not too discrepant, the model predicts the intransitive cycle *ABCA*. Alternatively, if one attribute is assigned markedly greater weight

Table 4
Stimuli for Experiment 2

	Sound system A	Sound system B	Sound system C
Harmonic range	26 mill	21 mill	—
Sound depth	14 sones	—	30 sones
Acoustic power	—	3.4 phons	2.1 phons
Harmonic distortion	<1%	<1%	<1%

[9] Under rare and specific circumstances (i.e., when the source of missing information is itself viewed negatively), missing attribute information can also be interpreted as favorable (Levin, Johnson, Russo, & Deldin, 1985). Because these circumstances do not apply in the present context, we do not consider this case.

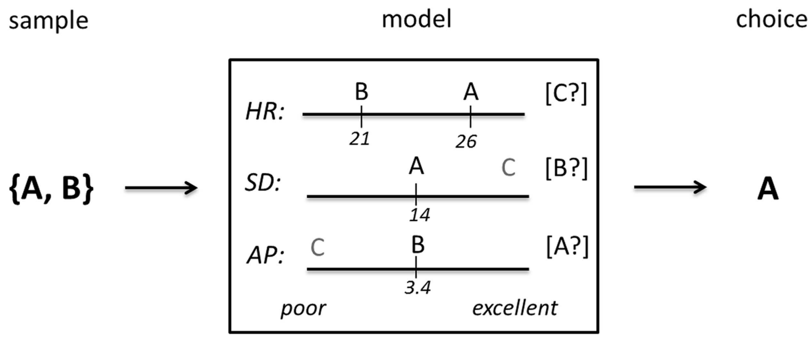

Figure 3. Schematic representation of an options-as-information analysis of the choice problem in Table 4. Bracketed options are missing values on the corresponding attribute. When missing values are seen as negative, *B* is equated with a point to the left of *A* on sound depth (SD), and *A* is equated with a point to the left of *B* on acoustic power (AP). The normative analysis then predicts a choice of *A* if the attribute weight attached to harmonic range (HR) does not exceed the sum of the weights attached to the other two attributes. When missing values are seen as uninformative, *A* and *B* are not differentiated relative to the posterior models of sound depth and acoustic power. The normative analysis then uniquely predicts that *A* will be chosen.

than the others, the three patterns *ABC*, *BCA*, and *CAB* are possible. As in Experiment 1, the model rules out the transitive patterns *CBA*, *BAC*, and *ACB*, as well as the intransitive cycle *ACBA*.

Kivetz and Simonson (2000) reported evidence for behavioral intransitivity in missing value problems similar to the one in Table 4. Using between-subjects designs, they found an intransitive *ABCA* pattern in the aggregate data. They also employed a limited within-subjects design, where each participant made each choice once; a substantial proportion of participants exhibited a cycle on this single triple of choices. Kivetz and Simonson did not demonstrate violations of the triangle inequalities, however, and so their results are, in principle, consistent with a random preference model in which preferences of all participants are always transitive but randomly fluctuate over time. Finally, Kivetz and Simonson reported a between-subjects experiment which they regarded as evidence against an inferential account of the intransitivity. We do not share their conclusion, but revisit their argument in more detail in the General Discussion.

Experiment 2 adapts the repeated-choice design of the first experiment to test weak stochastic transitivity, the triangle inequalities, and the two specifications (noisy response and flexible weights) of the context-sensitive preference model in a missing values problem. The analysis developed in this section suggests that, if anything, levels of behavioral intransitivity even higher than those seen in Experiment 1 may be expected. If missing values are regarded with suspicion, the intransitive *ABCA* pattern is one of four possible resulting choice patterns, as in Experiment 1. However, if missing values are seen as strictly uninformative, the intransitive choice pattern is *uniquely* predicted.

Experiment 2

Method

Twenty-four undergraduate students at UCSD participated in the experiment. Recruitment, compensation, and the determination of sample size were the same as in Experiment 1, and two participants were lost to attrition. One participant, whose postexperimental comments indicated that she had not been following instructions, was excluded from the analyses reported below. The remaining 21 participants (73% female) ranged in age from 18 to 26 years (M = 20.5).

The experimental procedure resembled that in Experiment 1, with the same number of sessions, number of choice problems per session, measures, and overall task design.[10] All choice problems involving consumer products—including the critical sound system stimuli in Table 4—were adjusted to include missing attribute information, however. Participants were instructed that they would receive information about different products, but that "this information may not always be complete and some information may be missing." As in Experiment 1, each pair of critical stimuli was encountered by each participant 20 times over the course of four sessions, and participants were not given any information regarding what constitutes a good or bad value for the product attributes. The experiment did not include any measures or conditions that are not reported.

Results and Discussion

The aggregate data again conformed to the predicted intransitive pattern. Collapsing across all participants and trials, *A* was selected over *B* 82.1% of the time, *B* over *C* 72.1% of the time, and *C* over *A* 70.0% of the time.

Is this pattern also seen at the individual-participant level? Table 5 presents the choice proportions for each of the 21 participants in pairwise choices between the critical stimuli. It also reports statistical tests of weak stochastic transitivity, the triangle inequalities, and the two specifications of the context-sensitive preference model.

As in Experiment 1, there is clear evidence for violations of both transitivity criteria. Eleven participants descriptively violated weak stochastic transitivity, and six of these violations are significant. Furthermore, 14 participants descriptively violated the triangle inequalities, and 12 of these violations are individually significant. That is, more than half of the participants in this study exhibited significant triangle inequality violations. By contrast, observed choice patterns were consistent with the context-sensitive model, especially in its noisy-response specification. Only one participant violated the noisy-response specification of the model descriptively, and this violation fell short of significance ($p = .08$). The more restrictive flexible-weights specification of the model was violated by seven participants descriptively, and four of these violations are significant.

We now turn to a summary of modal choice patterns in this experiment. All 11 participants with intransitive modal responses exhibited the *ABCA* cycle predicted by the context-sensitive preference model. Of the remaining 10 participants, the modal responses of nine were consistent with the model. Four participants exhibited an *ABC* pattern, one exhibited a *BCA* pattern, and two exhibited a *CAB* pattern. For two participants (marked with an asterisk in Table 5), the modal response is not uniquely identified, as *A* and *C* were selected with equal frequency. However, the two choice patterns that best describe their behavior (*ABC* and *ABCA*) are both consistent with the context-sensitive model. Only one participant chose according to a pattern (*CBA*) not permitted by the model. Overall, 20 of 21 participants chose according to modal patterns that are compatible with the context-sensitive preference model.

In summary, and echoing the results of Experiment 1, the noisy-response specification of the context-sensitive preference model successfully captures the data both of participants who behaved transitively and of participants who behaved intransitively in Experiment 2. Moreover, in Experiment 2, even the flexible-weights specification of the model, despite its considerably more restrictive predictions, fared better than both weak stochastic transitivity and random preference models. Remarkably, the majority of participants in this experiment exhibited individually significant violations of the triangle inequalities. When the information available to DMs is restricted, choice patterns incompatible with random mixtures over transitive orders can be abundant.

General Discussion

The transitivity axiom constrains the preferences of an ideal actor. In psychological research, however, preferences must be discerned through a veil of choices, and choice data often exhibit considerable within-subjects variability. Empirical tests of the transitivity axiom there-

[10] As in Experiment 1, we included attention checks involving dominated gambles. Mean accuracy on the attention checks in Experiment 2 was 91%. All participants had accuracy scores of at least 85%, with the exception of two participants (Participants 5 and 19 in Table 5) who consistently chose dominated gambles. Postexperimental comments of Participant 5 suggested that this resulted from a misreading of winning and losing sectors of the gambles' probability wheels, rather than from inattention. (Neither participant violated the triangle inequalities, so their inclusion does not increase the rate of violations in this experiment.)

Table 5
Results From Experiment 2

Participant	Choice frequencies			Modal choice pattern	Weak stochastic transitivity		Random pref. model		Context-sens. pref. model (noisy resp.)		Context-sens. pref. model (flex. weight)	
	P_{AB}	P_{BC}	P_{AC}		G^2	p value	G^2	p value	G^2	p value	G^2	p value
1	1.00	0.05	0.45	*CAB*	0		0		0		9.5	**<.01**
2	0.15	1.00	0.00	*BCA*	0		4.3	**0.02**	0		0	
3	1.00	0.75	0.35	*ABCA*	1.8	0.19	6.7	**0.01**	0		0	
4	1.00	0.80	0.40	*ABCA*	0.8	0.35	6.9	**<.01**	0		0	
5	1.00	0.65	0.85	*ABC*	0		0		0		2.2	0.07
6	0.85	0.85	0.00	*ABCA*	10.8	**<.01**	20.6	**<.01**	0		0	
7	0.95	0.75	0.00	*ABCA*	5.2	**0.03**	22.0	**<.01**	0		0	
8	0.30	0.15	0.00	*CBA*	0		0		3.3	0.08	13.2	**<.01**
9	1.00	1.00	0.75	*ABC*	0		7.5	**<.01**	0		0	
10	1.00	1.00	0.00	*ABCA*	27.7	**<.01**	48.7	**<.01**	0		0	
11	0.75	0.75	0.45	*ABCA*	0.2	0.32	0.1	0.39	0		0	
12	0.60	0.40	0.10	*CAB*	0		0		0		0.4	0.28
13	1.00	0.80	0.00	*ABCA*	7.7	**0.02**	30.7	**<.01**	0		0	
14	1.00	0.90	0.00	*ABCA*	14.7	**<.01**	37.4	**<.01**	0		0	
15	0.85	0.60	0.50	*	0		0		0		0.1	0.39
16	0.60	0.80	0.80	*ABC*	0		0		0		5.9	**0.01**
17	0.90	0.55	0.10	*ABCA*	0.2	0.33	4.8	**0.01**	0		0	
18	1.00	0.95	0.50	*	0		11.3	**<.01**	0		0	
19	0.55	0.85	0.75	*ABC*	0		0		0		4.5	**0.02**
20	1.00	0.95	0.05	*ABCA*	19.8	**<.01**	35.3	**<.01**	0		0	
21	0.75	0.60	0.25	*ABCA*	0.8	0.18	0.3	0.29	0		0	

Note. pref. = preference; resp. = response; sens. = sensitive; flex. = flexible. P_{xy} refers to the proportion of trials in which a participant chose x from the pair $\{x, y\}$. The asterisk * labels participants whose modal choice pattern is not uniquely identified. For each of the four models, the log-likelihood ratio G^2 at the maximum-likelihood estimate is reported. When a model is not a perfect fit, the corresponding p value is indicated; significant violations are marked in bold.

fore require assumptions about the relationship between preference and choice. Within-subjects variability in choice may reflect temporal fluctuations in "true" preferences (assuming they exist), noise in their expression or measurement, or a combination of these factors (Birnbaum & Bahra, 2012; Block & Marschak, 1960; Luce, 1995, 1997; Luce & Suppes, 1965; Regenwetter et al., 2011). An influential and formally elegant approach to the analysis of choice variability assumes that preferences fluctuate over time, but that all fluctuations in preference are random with respect to the choice context, and that preferences are expressed with perfect fidelity. Under these assumptions, transitivity of preference implies that the triangle inequalities must be satisfied. In their important critique of the transitivity literature, Regenwetter et al. (2011) argued that no clear violations of the triangle inequalities had been demonstrated. In doing so, they raised an implicit challenge, to which this article provides a response. In two repeated-choice experiments, strongly intransitive choice behavior was observed at the within-subjects level. Nearly a third of participants in Experiment 1, and more than half of participants in Experiment 2, exhibited individually significant violations of the triangle inequalities. These triangle inequality violations were demonstrated in typical multiattribute choice problems, similar to those in which May (1954) and Kivetz and Simonson (2000) previously reported more limited evidence of intransitivity.

However, while our experimental findings demonstrate that triangle inequality violations occur, our theoretical analysis shows that even these strongly intransitive choice patterns do not imply intransitivity—or irrationality—of preference. Instead, the observed choice patterns are consistent with a normative model in which transitive preferences are dynamically updated in a context-dependent manner, as inferences are drawn from sampled options. This article thus extends Regenwetter et al.'s (2011) analysis in two ways: First, it incorporates psychologically plausible context-dependence into stochastic models of transitive preference; and second, it shows how such context-dependence can be normatively appropriate.

This extended analysis allowed us to derive a testable model of context-sensitive preference. The individual choice data of participants in both experiments were consistent with the model, though the fit was sensitive to assumptions concerning the interpretation of within-subjects variability. A "flexible weights" specification of the model, in which expressions of preference are assumed to be noise-free, but attribute weights can vary randomly over time, could not account for all participants' choice data, though it was compatible with the responses of most participants in Experiment 2. A comprehensive fit to the data in both experiments was obtained for a "noisy response" specification of the model, in which attribute weights are assumed to be fixed, and observed choice is interpreted as a noisy expression of a "true" underlying preference signal.

The remainder of the discussion is organized as follows. First, we discuss psychological evidence for the inferential model, and compare it with potential heuristic accounts of the data. Second, we discuss applications to market settings. Third, we highlight parallels to the literature on intransitive choice in nonhuman animals. Next, we consider methodological limitations of the present experiments. We conclude by situating our model and findings in the broader context of research on constructed preference.

Inferences and Heuristics

The context-sensitive preference model is derived from a normative analysis of decision making under incomplete information and limited memory. It posits an evaluation process in which inferences are drawn as options are sampled, and options are evaluated in light of these inferences. Recent evidence suggests that this normative evaluation process is also psychologically plausible, and may explain some behavioral violations of procedure invariance and regularity.[11] Sher and McKenzie (2014) examined the role of sample-based inference in joint-separate reversals (Hsee et al., 1999), in which options that are superior on "difficult-to-evaluate" attributes receive poor evaluations when seen in isolation, but are prized when evaluated together with other options. Sher and McKenzie's participants drew markedly different inferences about the distribution of the unfamiliar attribute in different separate and joint evaluation contexts. Furthermore, when these inferences were supplied as background information to different participants, each of whom evaluated a single option, the effects of the inferences alone reproduced the standard joint-separate reversal. There is also evidence (Prelec, Wernerfelt, & Zettelmeyer, 1997; Sher, Müller-Trede, & McKenzie, 2015) that inferences from option-samples may play a major role in some well-known regularity violations, such as the attraction effect (Huber, Payne, & Puto, 1982; Simonson & Tversky, 1992) and the compromise effect (Simonson, 1989). Further research indicates that participants draw ecologically appropriate inferences from the way in which options are "framed" in the description of choice problems; these inferences can explain some important framing effects (McKenzie & Nelson, 2003; Sher & McKenzie, 2006, 2008; see also McKenzie, 2004; McKenzie, Liersch, & Finkelstein, 2006). Taken together, these findings suggest that participants in decision-making experiments routinely update beliefs and preferences in a coordinated and context-dependent fashion. When prior knowledge is limited, options are regarded as a source of information, and inferences from sampled options can have a large effect on evaluation and choice.

This process of belief- and preference-updating can account for the intransitive choice cycles exhibited by a number of participants in Experiments 1 and 2. However, heuristic accounts of these intransitive cycles are also possible. First, the intransitive cycles in both studies could be explained by a majority-rule heuristic, in which DMs tally up the number of attributes on which each option enjoys an advantage (cf. May, 1954; Zhang, Hsee, & Xiao, 2006). The majority rule heuristic is a special case of Huber's (1974, 1979) multidimensional

[11] The principle of *procedure invariance* states that different methods of eliciting preferences should yield the same ordering of options. *Regularity* states that adding new options to a choice set should not increase the popularity of any option in the initial set.

preference model, in which the attributes favoring each option in a choice pair are weighted, and the total weights are compared. Alternatively, in missing-value problems like those examined in Experiment 2, DMs may selectively underweight attribute dimensions with missing information, as proposed by Kivetz and Simonson (2000). As noted earlier, in defending this proposal, Kivetz and Simonson also argued against an inferential model of decision making in these problems. In what follows, we first address their argument, and then weigh the inferential model against the heuristic accounts.

Kivetz and Simonson reported a between-subjects missing-value experiment in which participants were presented with all three options, but were asked to choose among only two of them, the third being listed as "out of stock." (In addition, Kivetz and Simonson provided parenthetical background information about attribute ranges in this and other studies.) Although participants in all conditions had been exposed to all three options, the modal responses across conditions were collectively intransitive, leading the authors to conclude that the effect could not be explained by sample-based inferences. However, we note that this finding can be as readily explained by absence of attention as by absence of inference. Inferences can operate on attended and remembered information only. If participants employed a strategy of restricting attention to available options, inferences from attended items would generate the intransitive choice pattern. Furthermore, as Sher and McKenzie (2014) noted, although even pallid information about attribute ranges can strongly influence evaluations of individual items, these effects can disappear when the foreground evaluation set is enlarged, perhaps because background information is less likely to be attended to under these conditions. Barring independent measures of attention, we do not think that insensitivity to parenthetical information and unavailable options provides clear evidence against an inferential account.

The present experiments lack the process measures that would be needed to arbitrate conclusively between the inferential and heuristic accounts. Instead, the goal of these studies was to establish the occurrence of unambiguous triangle inequality violations (a finding that is of considerable interest even if a heuristic model is favored), and to clarify their normative interpretation. Importantly, however, the normative model provides a unified account of the behavior of all participants in both experiments, and explains why more intransitive cycles are observed in Experiment 2 than in Experiment 1. In contrast, Kivetz and Simonson's (2000) weighting heuristic can explain the intransitive cycles in Experiment 2, but it cannot explain the behavior of participants who chose transitively in that experiment, and it does not apply at all to Experiment 1. A majority rule heuristic can explain the intransitive cycles in both experiments, but it, too, cannot account for transitive behavior in either study (which includes a majority of participants in Experiment 1) or the difference between the two experiments (see also Birnbaum & Diecidue, in press).

Huber's (1979) multidimensional preference model, on the other hand, coincides with the inferential model in the choice patterns it predicts in these problems. This equivalence arises because, when prior knowledge of attribute distributions is minimal, sample-based inferences are assumed to be symmetric, and model-normalized scores on an attribute are simply determined by rank on that attribute. However, we note that the two models make diverging predictions in the same choice problems when DMs have relevant prior knowledge about attribute distributions. For example, a DM may know that *A*'s advantage over *B* on acoustic power is substantial, while *B*'s advantage over *C* on the same dimension is practically negligible. In the normative model, such asymmetries in interpretation can lead to any transitive choice pattern, while Huber's model of weighted ordinal comparisons can only generate the permitted choice patterns in Table 2.[12] According to the inferential model, then, the choice predictions in Table 2 should hold when prior knowledge is limited, but they may fall apart when additional knowledge is acquired. Without further amendments, the multidimensional preference model (Huber, 1979) suggests

[12] For example, consider a DM who places more weight on acoustic power than on sound depth and ignores harmonic range. If the DM knows that *B* and *C* in Table 1 are negligibly different on acoustic power while *A* is clearly superior, and that *A* and *B* are negligibly different on sound depth while *C* is clearly superior, the inferential model predicts that the DM will exhibit an *ACB* choice pattern. This pattern cannot be generated in Huber's model.

that these predictions should be robust to increasing knowledge.

Ultimately, however, one need not regard the normative and heuristic approaches as mutually exclusive. The inferences posited by the normative model are elementary, and one can readily substitute a qualitative comparison of attribute weights for the quantitative inequalities in the model. The normative analysis would then provide a rational basis for the DM's selection of an appropriate heuristic: When prior knowledge is minimal and absolute attribute values are uninterpretable, choose based on a weighted sum of attribute ranks. When prior knowledge allows for a more nuanced interpretation of attribute differences, employ a richer model that explicitly takes this information into account. From this perspective, the inferential model is not in irreconcilable conflict with the heuristic accounts, but instead may describe the intuitive logic that determines which heuristics are used by which participants for which choice problems (cf. Gigerenzer et al., 1999; Payne, Bettman, & Johnson, 1993).

Intransitive Choice in the Market

Our normative model considers a DM who has minimal knowledge of the distribution of relevant attributes. This assumption is appropriate for the choice problems we have examined, and for many others in the literature, in which attributes are often fictitious, ambiguous, or unfamiliar, and in which measures are often taken to prevent learning in within-subjects designs. Outside of the laboratory, levels of background knowledge are more variable and opportunities for learning are less constrained. A consumer shopping for mattress covers, for example, may set out with a more or less fuzzy concept of *thread count* and a more or less hazy sense of what counts as a good or poor value. In addition, processing of sampled options may be deeper and interference weaker in the shopping aisle than in the laboratory, resulting in better (though still limited) memory for previously sampled options.

When background knowledge is complete or learning is perfect, the normative model predicts that intransitive choice cycles should disappear. In the intermediate range of partial knowledge and imperfect learning, patterns of predicted behavior depend on how confident consumers are in their beliefs and on how steeply they update these beliefs when new options are sampled. Studies of context effects seldom measure inferences explicitly; however, the existing evidence suggests that context-based inferences can be sizable and consequential in laboratory settings (Prelec et al., 1997; Sher & McKenzie, 2014). Explicitly measuring inferences in more naturalistic environments is an important task for future research. The normative model implies that intransitive choice patterns will be found in market settings when (but only when) DMs draw transient sample-based inferences about the interpretation of attribute values.

As noted earlier, a well-known argument for the normative status of transitivity appeals to the fact that, in idealized market settings, DMs with intransitive preferences are vulnerable to money pumps (Davidson, McKinsey, & Suppes, 1955). Is the DM in the inferential model, whose preferences are transitive in each context but whose choices can be strongly intransitive across contexts, similarly exposed to money pumps? If we assume that the DM carries forward no information from prior choices, he or she would indeed be vulnerable to money pumps. The money pump then dramatizes the cost of forgetting rather than the cost of irrationality (i.e., those who know more can exploit those who know less, and the effect can be repeated if the ignorant are also forgetful). However, the above assumption—that no information is retained—becomes dubious in a money pump situation, both from a formal and from a psychological perspective. A basic feature of the money pump situation is that the DM carries the selected good forward to later trades. The DM therefore carries forward the information that the current possession was previously chosen, and thus that it was superior relative to a past context. A normative analysis of repeated choice with retained selections would need to take this retained information into account. Furthermore, from a psychological perspective, it is plausible that DMs will learn more, and more durably, from real choices with prolonged ownership of chosen goods than from evanescent choices in a hypothetical forced-choice setting. These normative and psychological considerations may limit the cost of forgetting (i.e., the danger of repeated exploitation of ignorance) in market environments involving repeated trades.

Intransitive Choice in the Animal Kingdom

Demonstrations of behavioral intransitivity are not restricted to human decision making. Patterns of intransitive choice have been documented in animal species ranging from honeybees (Shafir, 1994) to birds (Waite, 2001). Apparent violations of regularity have also been reported in nonhuman animals (Shafir, Waite, & Smith, 2002). These findings are seen as a challenge to optimality analyses of animal behavior, and as a window into the cognitive limitations of our evolutionary cousins. Thus, the title of one study trumpets the discovery of "Irrational Choices in Hummingbird Foraging Behavior" (Bateson, Healy, & Hurly, 2002).

However, some students of animal behavior have argued that observed violations of transitivity and regularity are consistent with optimal foraging theory. Their arguments bear striking parallels to the normative models of human decision making examined here. In our discussion of the potential rationality of intransitive choice patterns, we have highlighted the dependence of preference on the variable state of the DM (random preference models) and the information contained in the choice set (context-sensitive preference models). Optimal foraging models of animal choice have placed a similar emphasis on both state- and context-dependence.

Schuck-Paim, Pompilio, and Kacelnik (2004) argued that some regularity violations result from state-dependent preference. In research on context effects in animal behavior, context manipulations are sometimes confounded with the energetic state of the subject. This confound arises because animals need extensive exposure to a context before choices can be measured, and this training can alter the state, and hence the "preferences," of the organism. Schuck-Paim et al. (2004) found that a standard regularity violation disappeared when they carefully controlled the energetic state of their subjects (European starlings).

Houston, McNamara, and Steer (2007) showed that optimal context-dependence of animal foraging behavior can lead to intransitive choice. This model has been extended, and applied to regularity violations, by Trimmer (2013) and McNamara, Trimmer, and Houston (2014). The principle that motivates these models is the temporal autocorrelation of typical natural environments. That is, foraging options available now are differentially likely to be available in the near future. The current context thus provides information about the future environment. Because optimal foraging is partly determined by expectations about future resources, choices need not be consistent across contexts, and a fitness-maximizing decision rule for foraging behavior can generate intransitive choice patterns (cf. Fawcett et al., 2014). While the environment of a bird foraging for nectar differs considerably from that of a human buying a sound system, contexts can be informative in both settings. When this is the case, optimal preferences are context-sensitive and optimal choice patterns can be cyclical.

Methodological Limitations

The present studies employ a two-alternative forced choice (2AFC) task, in which the participant must express a preference. The 2AFC task has been the standard paradigm in studies of intransitivity going back to Tversky (1969), and it was used by Regenwetter et al. (2011) in testing the triangle inequalities (but see Regenwetter & Davis-Stober, 2012). It is partly motivated by the difficulty of designing a no-preference option which, from the participant's point of view, is clearly differentiated from a weak-preference or vague-preference option. The reliance on 2AFC tasks, however, precludes sensitivity to indifference and incomplete preference. Indifference ($A \sim B$) occurs when two options are regarded as precisely equated in value (cf. Footnote 3). A preference order is incomplete if it is not defined for some pairs of options (i.e., it is not the case that $A \succ B$, $B \succ A$, or $A \sim B$).

In the domain of multiattribute choice, we do not regard the neglect of indifference as a major problem. True indifference is a state in which options are balanced on a subjective "knife's edge," and thus should be unstable in the face of small variations in attribute weights. The neglect of incomplete preferences is potentially more significant. While decision-theoretic models usually assume complete preferences, a number of economists (e.g., Aumann, 1962; Eliaz & Ok, 2006; Mandler, 2001) and philosophers (e.g., Raz, 1985) have convincingly argued that completeness is not a requirement of rationality. Furthermore, in a theoretical analysis, Mandler (2005) has shown that intransitive, but rational, loops can occur in sequential choice when preferences are incomplete. Though largely neglected in psychological studies of decision making, incomplete preferences are both normatively and descriptively plausible, and we think that new empirical and conceptual paradigms that allow for incompleteness could enrich our understanding of the psychology of choice. Nonetheless, the present theoretical analysis shows that strongly intransitive cycles are compatible with rational choice under limited memory, even if preferences are assumed to be complete at each point in time.

In our analysis of within-subjects variability in choice, we considered two sources—fluctuations in the DM's underlying preferences, and noise in their expression. In testing the model, we examined each potential source of variability separately. We found that the model could fully account for the data with response noise alone (i.e., the model accounted for the modal responses of virtually all participants) but not with variable attribute weights alone (i.e., some participants' choice data fall outside the polytope that spans the permitted choice patterns). The differential success of these two pure models should be interpreted with caution, however. First, the two models occupy different volumes of choice probability space. Second, the correct model of choice variability need not be pure; a more realistic conception of variability might seek to simultaneously accommodate both sources and assess their relative contributions. We agree with Regenwetter et al. (2011, p. 54) that the development of testable hybrid models of choice variability is an important task for future research.[13]

Inference in Constructed Preference

A major lesson of recent psychological research is the pervasive context-dependence of human decision making. People seem to construct their preferences on the fly, consulting the context of available options as they go (Payne et al., 1999; Slovic, 1995). This context-sensitive instability of preference and choice is widely seen as a challenge to models of rational choice.

In the domain of learning, on the other hand, context-sensitive instability of belief is a hallmark of rationality. Building on this observation, this article aims to contribute to a rational analysis of constructed preference. It is guided by two assumptions: first, that learning is ubiquitous, even in artificial laboratory environments; and second, that inferences and preferences are inseparable, both in ideal and in real decision making. Accordingly, the present research both affirms and qualifies the theoretical framework of constructed preference. Constructed preference has a significant normative dimension that complements, and may sometimes illuminate, its varied empirical manifestations. In contrast to theories of preference formation that posit inconsistent weighting of attributes (e.g., Bhatia, 2013; Bordalo, Gennaioli, & Shleifer, 2013; Tversky & Simonson, 1993), the present model highlights the plasticity of the DM's beliefs about the choice environment. Our model and findings complement the rational analysis of apparent flaws in reasoning and biases in judgment (e.g., Klayman & Ha, 1987; Le Mens & Denrell, 2011; McKenzie, 2003; McKenzie & Mikkelsen, 2007; Oaksford & Chater, 1994; Tenenbaum & Griffiths, 2001).

This research also amplifies the concerns raised by Regenwetter et al. (2010, 2011) and Birnbaum (2011, 2012) about the distance that separates intransitive choice data from theoretical conclusions about intransitivity of preference. When preferences track inferences, and inferences depend on context, behavioral violations of weak stochastic transitivity—and the triangle inequalities—do not imply true violations of transitivity. To understand the structure of preference, it is valuable to model the dynamics of inference.

[13] For a step in this direction using an alternative modeling approach, see Birnbaum and Gutierrez (2007) and Birnbaum and Bahra (2012).

References

Anand, P. (1993). The philosophy of intransitive preference. *The Economic Journal, 103,* 337–346. http://dx.doi.org/10.2307/2234772

Aumann, R. J. (1962). Utility theory without the completeness axiom. *Econometrica, 30,* 445–462. http://dx.doi.org/10.2307/1909888

Bar-Hillel, M., & Margalit, A. (1988). How vicious are cycles of intransitive preference? *Theory and Decision, 24,* 119–145. http://dx.doi.org/10.1007/BF00132458

Bateson, M., Healy, S. D., & Hurly, T. A. (2002). Irrational choices in hummingbird foraging behavior. *Animal Behaviour, 63,* 587–596. http://dx.doi.org/10.1006/anbe.2001.1925

Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review, 120,* 522–543. http://dx.doi.org/10.1037/a0032457

Birnbaum, M. H. (2011). Testing mixture models of transitive preference: Comment on Regenwetter, Dana, and Davis-Stober (2011). *Psychological Review, 118,* 675–683. http://dx.doi.org/10.1037/a0023852

Birnbaum, M. H. (2012). A statistical test of the assumption that repeated choices are independently and identically distributed. *Judgment and Decision Making, 7,* 97–109.

Birnbaum, M. H., & Bahra, J. P. (2012). Testing transitivity of preferences using linked designs. *Judgment and Decision Making, 7,* 524–567.

Birnbaum, M. H., & Diecidue, E. (in press).Testing a class of models that includes majority rule and regret theories: Transitivity, recycling, and restricted branch independence. *Decision*.

Birnbaum, M. H., & Gutierrez, R. (2007). Testing for intransitivity of preferences predicted by a lexicographic semiorder. *Organizational Behavior and Human Decision Processes, 104,* 96–112. http://dx.doi.org/10.1016/j.obhdp.2007.02.001

Block, H. D., & Marschak, J. (1960). Random orderings and stochastic theories of responses. In I. Olkin, S. Ghurye, H. Hoeffding, W. Madow, & H. Mann (Eds.), *Contributions to probability and statistics* (pp. 97–132). Stanford, CA: Stanford University Press.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy, 121,* 803–843. http://dx.doi.org/10.1086/673885

Budescu, D. V., & Weiss, W. (1987). Reflection of transitive and intransitive preferences: A test of prospect theory. *Organizational Behavior and Human Decision Processes, 39,* 184–202. http://dx.doi.org/10.1016/0749-5978(87)90037-9

Cavagnaro, D. R., & Davis-Stober, C. P. (2014). Transitive in our preferences, but transitive in different ways: An analysis of choice variability. *Decision, 1,* 102–122. http://dx.doi.org/10.1037/dec0000011

Cohen, M., & Falmagne, J.-C. (1990). Random utility representation of binary choice probabilities: A new class of necessary conditions. *Journal of Mathematical Psychology, 34,* 88–94. http://dx.doi.org/10.1016/0022-2496(90)90014-Z

Condorcet, M. (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix* [Essay on the application of the probabilistic analysis of majority vote decisions]. Paris, France: Imprimerie Royale.

Davidson, D., McKinsey, J., & Suppes, P. (1955). Outlines of a formal theory of value, I. *Philosophy of Science, 22,* 140–160. http://dx.doi.org/10.1086/287412

Davis-Stober, C. P. (2009). Analysis of multinomial models under inequality constraints: Applications to measurement theory. *Journal of Mathematical Psychology, 53,* 1–13. http://dx.doi.org/10.1016/j.jmp.2008.08.003

Edwards, W., & Barron, F. H. (1994). SMARTS and SMARTER: Improved simple methods for multiattribute utility measurement. *Organizational Behavior and Human Decision Processes, 60,* 306–325. http://dx.doi.org/10.1006/obhd.1994.1087

Eliaz, K., & Ok, E. A. (2006). Indifference or indecisiveness? Choice-theoretic foundations of incomplete preferences. *Games and Economic Behavior, 56,* 61–86. http://dx.doi.org/10.1016/j.geb.2005.06.007

Falmagne, J.-C. (1978). A representation theorem for finite random scale systems. *Journal of Mathematical Psychology, 18,* 52–72. http://dx.doi.org/10.1016/0022-2496(78)90048-2

Fawcett, T. W., Fallenstein, B., Higginson, A. D., Houston, A. I., Mallpress, D. E. W., Trimmer, P. C., & McNamara, J. M. (2014). The evolution of decision rules in complex environments. *Trends in Cognitive Sciences, 18,* 153–161. http://dx.doi.org/10.1016/j.tics.2013.12.012

Fiorini, S. (2001). Determining the automorphism group of the linear ordering polytope. *Discrete Applied Mathematics, 112,* 121–128. http://dx.doi.org/10.1016/S0166-218X(00)00312-7

Fishburn, P. C. (1991). Nontransitive preferences in decision theory. *Journal of Risk and Uncertainty, 4,* 113–134. http://dx.doi.org/10.1007/BF00056121

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York, NY: Oxford University Press.

Houston, A. I., McNamara, J. M., & Steer, M. D. (2007). Violations of transitivity under fitness maximization. *Biology Letters, 3,* 365–367. http://dx.doi.org/10.1098/rsbl.2007.0111

Hsee, C. K. (1996). The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations of alternatives. *Organizational Behavior and Human Decision Processes, 67,* 247–257. http://dx.doi.org/10.1006/obhd.1996.0077

Hsee, C. K., Loewenstein, G. F., Blount, S., & Bazerman, M. H. (1999). Preference reversals between joint and separate evaluations of options: A review and theoretical analysis. *Psychological Bulletin, 125,* 576–590. http://dx.doi.org/10.1037/0033-2909.125.5.576

Huber, J., & McCann, J. (1982). The impact of inferential beliefs on product evaluations. *Journal of*

Marketing Research, 19, 324–333. http://dx.doi.org/10.2307/3151566
Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and similarity. *Journal of Consumer Research, 9,* 90–98. http://dx.doi.org/10.1086/208899
Huber, O. (1974). An axiomatic system for multidimensional preferences. *Theory and Decision, 5,* 161–184. http://dx.doi.org/10.1007/BF00143246
Huber, O. (1979). Nontransitive multidimensional preferences: Theoretical analysis of a model. *Theory and Decision, 10,* 147–165. http://dx.doi.org/10.1007/BF00126336
Iverson, G. J., & Falmagne, J.-C. (1985). Statistical issues in measurement. *Mathematical Social Sciences, 10,* 131–153. http://dx.doi.org/10.1016/0165-4896(85)90031-9
Johnson, R. D., & Levin, I. P. (1985). More than meets the eye: The effect of missing information on purchase evaluations. *Journal of Consumer Research, 12,* 169–177. http://dx.doi.org/10.1086/208505
Keeney, R. L., & Raiffa, H. (1993). *Decisions with multiple objectives: Preferences and value trade-offs*. New York, NY: Cambridge University Press. http://dx.doi.org/10.1017/CBO9781139174084
Kivetz, R., & Simonson, I. (2000). The effects of incomplete information on consumer choice. *Journal of Marketing Research, 37,* 427–448. http://dx.doi.org/10.1509/jmkr.37.4.427.18796
Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review, 94,* 211–228. http://dx.doi.org/10.1037/0033-295X.94.2.211
Le Mens, G., & Denrell, J. (2011). Rational learning and information sampling: On the "naivety" assumption in sampling explanations of judgment biases. *Psychological Review, 118,* 379–392. http://dx.doi.org/10.1037/a0023010
Levin, I. P., Johnson, R. D., Russo, C. P., & Deldin, P. J. (1985). Framing effects in judgment tasks with varying amounts of information. *Organizational Behavior and Human Decision Processes, 36,* 362–377. http://dx.doi.org/10.1016/0749-5978(85)90005-6
Lindman, H. R., & Lyons, J. (1978). Stimulus complexity and choice inconsistency among gambles. *Organizational Behavior and Human Performance, 21,* 146–159. http://dx.doi.org/10.1016/0030-5073(78)90046-6
Loomes, G., Starmer, C., & Sugden, R. (1991). Observing violations of transitivity by experimental methods. *Econometrica, 59,* 425–439. http://dx.doi.org/10.2307/2938263
Loomes, G., & Sugden, R. (1995). Incorporating a stochastic element into decision theories. *European Economic Review, 39,* 641–648. http://dx.doi.org/10.1016/0014-2921(94)00071-7
Luce, R. D. (1995). Four tensions concerning mathematical modeling in psychology. *Annual Review of Psychology, 46,* 1–26. http://dx.doi.org/10.1146/annurev.ps.46.020195.000245
Luce, R. D. (1997). Several unresolved conceptual problems of mathematical psychology. *Journal of Mathematical Psychology, 41,* 79–87. http://dx.doi.org/10.1006/jmps.1997.1150
Luce, R. D., & Raiffa, H. (1957). *Games and decisions*. New York, NY: Wiley.
Luce, R. D., & Suppes, P. (1965). Preference, utility and subjective probability. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 3, pp. 249–410). New York, NY: Wiley.
Mandler, M. (2001). A difficult choice in preference theory: Rationality implies completeness or transitivity but not both. In E. Millgram (Ed.), *Varieties of practical reasoning* (pp. 373–402). Cambridge, MA: MIT Press.
Mandler, M. (2005). Incomplete preferences and rational intransitivity of choice. *Games and Economic Behavior, 50,* 255–277. http://dx.doi.org/10.1016/j.geb.2004.02.007
Marschak, J. (1960). Binary-choice constraints and random utility indicators. In K. J. Arrow, S. Karlin, & P. Suppes (Eds.), *Proceedings of the first Stanford symposium on mathematical methods in the social sciences, 1959* (pp. 312–329). Stanford, CA: Stanford University Press.
May, K. O. (1954). Intransitivity, utility, and the aggregation of preference patterns. *Econometrica, 22,* 1–13. http://dx.doi.org/10.2307/1909827
McKenzie, C. R. M. (2003). Rational models as theories—not standards—of behavior. *Trends in Cognitive Sciences, 7,* 403–406. http://dx.doi.org/10.1016/S1364-6613(03)00196-7
McKenzie, C. R. M. (2004). Framing effects in inference tasks—And why they are normatively defensible. *Memory & Cognition, 32,* 874–885. http://dx.doi.org/10.3758/BF03196866
McKenzie, C. R. M., Liersch, M. J., & Finkelstein, S. R. (2006). Recommendations implicit in policy defaults. *Psychological Science, 17,* 414–420. http://dx.doi.org/10.1111/j.1467-9280.2006.01721.x
McKenzie, C. R. M., & Mikkelsen, L. A. (2007). A Bayesian view of covariation assessment. *Cognitive Psychology, 54,* 33–61. http://dx.doi.org/10.1016/j.cogpsych.2006.04.004
McKenzie, C. R. M., & Nelson, J. D. (2003). What a speaker's choice of frame reveals: Reference points, frame selection, and framing effects. *Psychonomic Bulletin & Review, 10,* 596–602. http://dx.doi.org/10.3758/BF03196520

McNamara, J. M., Trimmer, P. C., & Houston, A. I. (2014). Natural selection can favour 'irrational' behaviour. *Biology Letters, 10,* 20130935. http://dx.doi.org/10.1098/rsbl.2013.0935

Montgomery, H. (1977). A study of intransitive preferences using a think-aloud procedure. In H. Jungerman & G. de Zeeuw (Eds.), *Decision-making and change in human affairs* (pp. 347–362). Dordrecht, the Netherlands: Reidel. http://dx.doi.org/10.1007/978-94-010-1276-8_22

Morrison, H. W. (1963). Testable conditions for triads of paired comparison choices. *Psychometrika, 28,* 369–390. http://dx.doi.org/10.1007/BF02289558

Myung, J., Karabatsos, G., & Iverson, G. (2005). A Bayesian approach to testing decision making axioms. *Journal of Mathematical Psychology, 49,* 205–225. http://dx.doi.org/10.1016/j.jmp.2005.02.004

Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review, 101,* 608–630. http://dx.doi.org/10.1037/0033-295X.101.4.608

Parducci, A. (1965). Category judgment: A range-frequency model. *Psychological Review, 72,* 407–418. http://dx.doi.org/10.1037/h0022602

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York, NY: Cambridge University Press. http://dx.doi.org/10.1017/CBO9781139173933

Payne, J. W., Bettman, J. R., & Schkade, D. A. (1999). Measuring constructed preference: Towards a building code. *Journal of Risk and Uncertainty, 19,* 243–270. http://dx.doi.org/10.1023/A:1007843931054

Prelec, D., Wernerfelt, B., & Zettelmeyer, F. (1997). The role of inferences in context effects: Inferring what you want from what is available. *Journal of Consumer Research, 24,* 118–126. http://dx.doi.org/10.1086/209498

Ranyard, R. H. (1977). Risky decisions which violate transitivity and double cancellation. *Acta Psychologica, 41,* 449–459. http://dx.doi.org/10.1016/0001-6918(77)90003-8

Raz, J. (1985). Value incommensurability: Some preliminaries. *Proceedings of the Aristotelian Society, 86,* 117–134.

Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2010). Testing transitivity of preferences on two-alternative forced choice data. *Frontiers in Psychology, 1,* 148. http://dx.doi.org/10.3389/fpsyg.2010.00148

Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2011). Transitivity of preferences. *Psychological Review, 118,* 42–56. http://dx.doi.org/10.1037/a0021150

Regenwetter, M., & Davis-Stober, C. P. (2012). Behavioral variability of choices versus structural inconsistency of preferences. *Psychological Review, 119,* 408–416. http://dx.doi.org/10.1037/a0027372

Regenwetter, M., Davis-Stober, C. P., Lim, S. H., Guo, Y., Popova, A., Zwilling, C., . . . Messner, W. (2014). QTEST: Quantitative testing of theories of binary choice. *Decision, 1,* 2–34. http://dx.doi.org/10.1037/dec0000007

Schuck-Paim, C., Pompilio, L., & Kacelnik, A. (2004). State-dependent decisions cause apparent violations of rationality in animal choice. *PLoS Biology, 2,* e402. http://dx.doi.org/10.1371/journal.pbio.0020402

Shafir, S. (1994). Intransitivity of preferences in honey bees: Support for "comparative" evaluation of foraging options. *Animal Behaviour, 48,* 55–67. http://dx.doi.org/10.1006/anbe.1994.1211

Shafir, S., Waite, T. A., & Smith, B. H. (2002). Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behavioral Ecology and Sociobiology, 51,* 180–187. http://dx.doi.org/10.1007/s00265-001-0420-8

Sher, S., & McKenzie, C. R. M. (2006). Information leakage from logically equivalent frames. *Cognition, 101,* 467–494. http://dx.doi.org/10.1016/j.cognition.2005.11.001

Sher, S., & McKenzie, C. R. M. (2008). Framing effects and rationality. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 79–96). Oxford, UK: Oxford University Press. http://dx.doi.org/10.1093/acprof:oso/9780199216093.003.0004

Sher, S., & McKenzie, C. R. M. (2014). Options as information: Rational reversals of evaluation and preference. *Journal of Experimental Psychology: General, 143,* 1127–1143. http://dx.doi.org/10.1037/a0035128

Sher, S., Müller-Trede, J., & McKenzie, C. R. M. (2015). *Asymmetric inferences in asymmetric dominance*. Manuscript in preparation.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research, 16,* 158–174.

Simonson, I., & Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research, 29,* 281–295. http://dx.doi.org/10.2307/3172740

Slovic, P. (1995). The construction of preference. *American Psychologist, 50,* 364–371. http://dx.doi.org/10.1037/0003-066X.50.5.364

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology, 53,* 1–26. http://dx.doi.org/10.1016/j.cogpsych.2005.10.003

Tenenbaum, J., & Griffiths, T. (2001). The rational basis of representativeness. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd annual conference of the cognitive science society* (pp.

1036–1041). Edinburgh, UK: Cognitive Science Society.
Trimmer, P. C. (2013). Optimal behaviour can violate the principle of regularity. *Proceedings of the Royal Society B, 280*. http://dx.doi.org/10.1098/rspb.2013.0858
Tsai, R.-C., & Böckenholt, U. (2006). Modelling intransitive preferences: A random-effects approach. *Journal of Mathematical Psychology, 50,* 1–14. http://dx.doi.org/10.1016/j.jmp.2005.11.004
Tversky, A. (1969). Intransitivity of preferences. *Psychological Review, 76,* 31–48. http://dx.doi.org/10.1037/h0026750
Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science, 39,* 1179–1189. http://dx.doi.org/10.1287/mnsc.39.10.1179
Waite, T. A. (2001). Intransitive preferences in hoarding gray jays (*Perisoreus canadensis*). *Behavioral Ecology and Sociobiology, 50,* 116–121. http://dx.doi.org/10.1007/s002650100346
Zhang, J., Hsee, C. K., & Xiao, Z. (2006). The majority rule in individual decision making. *Organizational Behavior and Human Decision Processes, 99,* 102–111. http://dx.doi.org/10.1016/j.obhdp.2005.06.004

Appendix

A Context-Sensitive Preference Model

Consider the two-alternative forced-choice task in Experiment 1, in which participants make repeated pairwise choices from the triple of options $\{A, B, C\}$. As discussed in the text, a participant's choices can be represented as a point in the sample space of relative choice frequencies shown in Figure 2. The sample space's eight vertices are associated with the six transitive patterns and the two intransitive cycles that arise in a two-alternative forced-choice design, and in what follows, we prove that, under the assumption of symmetric sample-based inferences, the context-sensitive preference model defined in Equation (1) below is consistent with four of these patterns, but is not consistent with the other four. Like most previous research on transitivity (i.e., Regenwetter et al., 2011; Tversky, 1969), our analysis does not consider weak preference relations that allow for indifference between options (see General Discussion).

Let w_h, w_s, and w_a denote a set of non-negative weights that the DM assigns to the attributes harmonic range, sound depth, and acoustic power, respectively. Although the DM is assumed to have minimal prior knowledge about an attribute distribution, these weights need not be identical, because attribute labels may provide suggestive but relevant information about their meaning. The weights are not affected by the values seen in a particular option sample, however.

The DM's posterior models of the attribute distributions that describe the overall population of sound systems, on the other hand, are affected by the values seen in a particular option sample. Each numerical attribute value is normalized relative to a posterior model. For a given option x, we denote its normalized attribute value for harmonic range, relative to a model inferred from option sample S, by $h_{x|S}$. Similarly, x's normalized attribute values for sound depth and acoustic power, relative to the model inferred from S, are denoted by $s_{x|S}$ and $a_{x|S}$, respectively.

(*Appendix continues*)

The overall valuation v of an option x relative to an option sample S, is denoted by $v_{x|S}$, and can be computed as a weighted sum of the normalized attribute values. Thus, for example, the valuation assigned to A when $\{A, B\}$ is sampled is given by

$$v_{A|\{A,B\}} = w_h h_{A|\{A,B\}} + w_s s_{A|\{A,B\}} + w_a a_{A|\{A,B\}}. \quad (1)$$

A central assumption in our analysis, discussed in greater detail in the text, is that, in the absence of relevant prior information, DMs draw symmetric inferences from the option samples, so that the model-normalized values of the smaller and the greater of two values in a two-option sample do not depend on the particular values encountered, or on the attribute scales. This implies that, for a given two-option sample, the differences between two options' normalized attribute values differ only in sign. When $\{A, B\}$ is sampled, for example,

$$-(h_{A|\{A,B\}} - h_{B|\{A,B\}}) = (s_{A|\{A,B\}} - s_{B|\{A,B\}}) = (a_{A|\{A,B\}} - a_{B|\{A,B\}}). \quad (2)$$

Next, we show that, under this assumption, the cycle *ABCA* and the patterns *ABC*, *BCA*, and *CAB* are consistent with Equation (1), but the cycle *ACBA* and the patterns *ACB*, *BAC*, and *CBA* are not.

1. The Cycle *ABCA* is consistent with the Context-Sensitive Model.

A is chosen over B if and only if

$$\begin{aligned} & w_h h_{A|\{A,B\}} + w_s s_{A|\{A,B\}} + w_a a_{A|\{A,B\}} \\ & \quad > w_h h_{B|\{A,B\}} + w_s s_{B|\{A,B\}} + w_a a_{B|\{A,B\}} \\ \Leftrightarrow\ & w_h(h_{A|\{A,B\}} - h_{B|\{A,B\}}) + w_s(s_{A|\{A,B\}} - s_{B|\{A,B\}}) \\ & \quad + w_a(a_{A|\{A,B\}} - a_{B|\{A,B\}}) > 0 \\ \Leftrightarrow\ & -w_h + w_s + w_a > 0. \end{aligned}$$

Similarly B is chosen over C if and only if $w_h - w_s + w_a > 0$, and C is chosen over A if and only if $w_h + w_s - w_a > 0$. This system of inequalities states that the weight assigned to each of the attributes is less than the sum of the weights assigned to the other two attributes. Any such weight assignment thus gives rise to the intransitive cycle *ABCA*.

2. The patterns *ABC*, *BCA*, and *CAB* are consistent with the model.

A is chosen over C if and only if $-w_h - w_s + w_a > 0$. But this implies that $-w_h + w_s + w_a > 0$, and hence the selection of A over B; and that $w_h - w_s + w_a > 0$, and hence that B is selected over C. Thus if the weight assigned to acoustic power is greater than the sum of the weights assigned to the other two attributes, the DM will exhibit the pattern *ABC*.

Analogously, any set of weights that satisfies the inequality $w_h - w_s - w_a > 0$ leads to the pattern *BCA*, and any set of weights that satisfies the inequality $-w_h + w_s - w_a > 0$ leads to the pattern *CAB*.

3. The remaining patterns (*ACB*, *BAC*, *CBA*, and *ACBA*) are inconsistent with the model.

As noted above, the selection of A over C implies both that A is chosen over B and that B is chosen over C. This rules out the patterns *ACB*, *BAC*, and *ACBA*. Similarly, C is selected over B if and only if $-w_h + w_s - w_a > 0$, which implies that $-w_h + w_s + w_a > 0$; hence A is chosen over B, ruling out the pattern *CBA*.

Example: Even when inferences lead to intransitive choice cycles across contexts, they generate transitive preferences within contexts. To illustrate this point, we consider a DM who (a) assigns the unit weights $w_h = w_s = w_a = 1$ to the three attributes, and (b) computes z-scores for the model-normalized attribute values $h_{x|S}$, $s_{x|S}$, and $a_{x|S}$.

From the option sample $\{A, B\}$, the DM computes maximum likelihood estimates of the mean and standard deviation of each attribute (e.g., 18 and 8, respectively, for harmonic range). This leads to overall valuations of $v_{A|\{A,B\}} = 1$ and $v_{B|\{A,B\}} = -1$ for the two options in the sample. Even though the DM has not encountered C, its valuation, too, can be computed relative to the posterior models inferred from the option sample $\{A, B\}$. In particular, C's model-normalized attribute values

(*Appendix continues*)

are given by $h_{C|\{A,B\}} = (16 - 18)/8 = -0.25$, $s_{C|\{A,B\}} = 5.8$, and $a_{C|\{A,B\}} = -1.86$, and its overall valuation is $v_{C|\{A,B\}} = 3.69$. Upon seeing the option sample $\{A, B\}$, the DM's decision rule thus generates the (transitive) preference order $C >_{\{A,B\}} A >_{\{A,B\}} B$.

The same decision rule generates different preference orders when the DM is exposed to other option samples. Relative to the models inferred from the option sample $\{B, C\}$, for example, the options' overall valuations are $v_{A|\{B,C\}} = 3.05$, $v_{B|\{B,C\}} = 1$, and $v_{C|\{B,C\}} = -1$, resulting in the order $A >_{\{B,C\}} B >_{\{B,C\}} C$. Relative to the models inferred from $\{C, A\}$, $v_{A|\{A,C\}} = -1$, $v_{B|\{A,C\}} = 2.1$, and $v_{C|\{A,C\}} = 1$, which implies $B >_{\{A,C\}} C >_{\{A,C\}} A$. Because the three context-dependent preference orders include the relations $A >_{\{A,B\}} B$, $B >_{\{B,C\}} C$, and $C >_{\{A,C\}} A$, the DM will choose in accordance with the intransitive cycle *ABCA*.

Finally, in the context of the stimuli employed in Experiment 2, we consider two cases—that missing values are seen as negative (and so, e.g., $s_{A|\{A,B\}} - s_{B|\{A,B\}} > 0$) and that they are seen as uninformative (and so $s_{A|\{A,B\}} - s_{B|\{A,B\}} = 0$). In modeling the first case, it is natural to relax the assumption of inferential symmetry from the previous choice problem; instead, we assume that inferences are only symmetric for the two dimensions which exhibit missing values. In other words, the perceived (model-normalized) difference between a missing value and a known value need not be the same as the perceived difference between two known values. This assumption can be modeled by introducing a scaling factor α, as follows:

$$\alpha(h_{A|\{A,B\}} - h_{B|\{A,B\}}) = (s_{A|\{A,B\}} - s_{B|\{A,B\}}) = -(a_{A|\{A,B\}} - a_{B|\{A,B\}}). \quad (3)$$

It is straightforward to show that, for $\alpha \leq 1$, the qualitative results of the analysis in (1) – (3) above remain unchanged: The intransitive cycle *ABCA*, and the transitive patterns *ABC*, *BCA*, and *CAB* are consistent with the model, and the other four patterns are not. A scaling factor $\alpha \leq 1$ means that the perceived difference between a known and an unknown value is not larger than the perceived difference between two known values. That is, learning that an option is worse than a second option has a stronger negative effect on its valuation than not knowing how an option compares to a second option. This assumption seems highly compelling, normatively as well as psychologically. (We note, however, that, if missing information about an option did have a larger negative effect on its valuation than learning that it is worse than a second option, the transitive patterns *ACB*, *BAC*, and *CBA* would no longer be inconsistent with the model, because, for $\alpha > 1$, there exist solutions to each of the corresponding systems of inequalities.)

When missing values are uninformative, α in Equation (3) is set to 0, and pairwise choice is determined by the difference on the unique attribute with common information, resulting in the intransitive choice pattern *ABCA*.

Received January 14, 2015
Revision received February 24, 2015
Accepted March 20, 2015 ■