Investigating Consumer Decision Strategies With Systems Factorial Technology

Gavin J. Cooper[a,*], Guy E. Hawkins[a]

[a] *School of Psychology, University of Newcastle, Callaghan, NSW, 2308, Australia*

Abstract

People routinely make multi-attribute decisions about consumer items, such as products and services. The potentially complex decision strategies underlying such consumer decisions have recently been investigated in detail, with most researchers restricting their focus to a relatively small subset of heuristics so as to retain tractability in analyses and identifiability of parameterized cognitive models. Many of these heuristics can be conceived as special cases of a smaller number of overarching dimensions: processing all or a subset of the attribute information describing the consumer options, and processing that attribute information in series or in parallel. These higher-level dimensions correspond to two latent factors of focus in Systems Factorial Technology (SFT), a non-parametric modeling technique that aims to uncover the mental architectures that generate observed decision behavior. Here, we develop a simplified consumer decision scenario and report proof-of-concept evidence regarding the ability of SFT to discriminate between mental architectures, and as a consequence whole classes of decision strategies, in the newly developed consumer task. Our results suggest that most people make decisions prior to processing all available product information. Furthermore, people appear to process numerically presented attribute information in serial, and pictorially presented attribute information in parallel. This extension of SFT beyond its classic domain of application in

*Corresponding author
Email address: `Gavin.Cooper@uon.edu.au` (Gavin J. Cooper)
URL: `www.gavincooper.net` (Gavin J. Cooper)

perceptual processing provides a relatively simple, non-parametric approach to investigating consumer decision strategies.

Keywords: Consumer choice, Decision strategy, Mental architecture, Information processing, Systems Factorial Technology (SFT)

1. Introduction

Consumers routinely make decisions about products and services. Some are largely inconsequential; what to eat for breakfast or which socks to purchase. Others are consequential with substantial implications for financial, physical, and mental well-being; should I make this investment or should I take this course of chemotherapy. Businesses, non-profits and governments invest great effort persuading consumers to select one option over others. Central to these efforts is the desire to understand how consumers think and behave so as to improve products and services, more effectively target potential customers, and create competitive advantage. Understanding consumer thought and behavior involves understanding the strategies consumers use to process information about products and services.

The strategies people use to select between consumer options have been studied for decades (for a seminal review, see Payne et al., 1993). These decision strategies, often referred to as heuristics, aim to describe how a decision maker shifts from an initial state of knowledge to a final state of knowledge at which point they feel confident that a decision has been made. The cognitive steps that fall between the initial and final state of knowledge generally consider how people process information, when they stop searching for information, and ultimately how they select an option.

In this paper, we first provide an overview of consumer decision strategies, and then highlight previously unidentified links between features of these decision strategies and the mental architecture of information processing systems – statements about how people process and integrate multiple sources of information to inform decisions – as studied in the domains of perception and cognition. Our aim is to classify consumer decision strategies into a smaller subset of higher-level mental architectures that are common to multi-attribute decisions in the consumer and perceptual domains; this approach assumes that the consumer decision strategies (outlined below) are specific model implementations derived from general mental architectures. Table 1 provides a first-pass at this classification, and is outlined in detail later. With this classification in place, we then present two experimental tests of a consumer-based decision task. Owing to our classification of decision strategies into a smaller set of mental architectures, we draw upon a suite of analysis tools known as Systems Factorial Technology (SFT; Townsend & Nozawa, 1995) to rule in or rule out different mental architectures in different experimental contexts. This allows us to determine the viability of whole classes of consumer decision strategies in the contexts under investigation.

Table 1: Fifteen consumer decision strategies previously proposed in the literature (column 2; for review, see Pfeiffer, 2012), with our assessment of their corresponding processing architectures (column 3). Through the methods of Systems Factorial Technology (SFT), each processing architecture predicts a unique pair of the Mean Interaction Contrast (MIC) and Survivor Interaction Contract (SIC) as computed from response time data (columns 4 and 5).

Set	Decision Strategy	Possible Architectures	MIC Predictions	SIC Predictions
1	Lexicographic heuristic Minimum difference lexicographic rule Elimination by aspects	Self-terminating, serial	MIC $= 0$	$SIC(t)$ is flat
2	Compatibility test Conjunctive strategy Satisficing heuristic Satisficing-plus heuristic	Self terminating, (serial or parallel)	$MIC = 0$ or $MIC > 0$	$SIC(t)$ is flat or $SIC(t)$ is positive
3	Disjunctive strategy Dominance strategy Simple majority decision rule	Exhaustive, (serial or parallel)	$MIC = 0$ or $MIC < 0$	$SIC(t)$ is negative $\rightarrow$ positive with equal sized deflections or $SIC(t)$ is negative
4	Equal weight heuristic Weighted additive rule Additive difference strategy Majority of confirming dimensions heuristic Frequency of good and/or bad features heuristic	Coactive	$MIC > 0$	$SIC(t)$ is negative $\rightarrow$ positive with larger positive deflection.

1.1. Consumer Decision Strategies

Here, we provide a non-exhaustive characterization of decision strategies that are relevant to our thesis; for more detailed coverage we refer the reader to previous reviews of the topic (Payne et al., 1993; Pfeiffer, 2012; Riedl et al., 2008). Throughout, we suppose a decision maker is considering products that are described by two dimensions or *attributes*, such as price and quality, though the issues we consider can be generalized to cases with more than two attributes. We refer to consumer products with the more generic term *option*, and to introduce specific decision strategies we assume the decision maker is presented with an exemplar option that is high in both price and quality; high price is a negative feature, high quality is a positive feature. We also note that across different consumer contexts options might be presented in sets of two or more available options where the task is to select the most (or least) preferred option, or they might be presented one option at a time where the task is to accept or reject the option from further consideration according to some criterion/criteria.

The various decision strategies proposed in the literature differ with respect to their characteristic assumptions about how people process the attribute information that defines each option. Some strategies propose that decision makers compare an option's attributes – the specific price and quality rating on offer – to an *aspiration level*, the threshold or 'cutoff' for what is deemed an acceptable level of the attribute. A simple example of an aspiration level is the maximum price a decision maker will consider for a given product category; generally, once the aspiration level is met (not met) the corresponding attribute is deemed positive (negative). Strategies that propose an option with a negative value on one or more attribute/s cannot be 'made up' or compensated by positive values on other attributes are known as *non-compensatory*; in our example, a decision maker might reject the option because the price is too high, independent of whether the quality is excellent or terrible.

Many, though not all, influential non-compensatory strategies implicitly or explicitly assume the use of aspiration levels. For instance, the satisficing heuristic (Simon, 1955) states that the attributes of an option are compared against their corresponding aspiration levels. As soon as the decision maker encounters an attribute that does not meet its aspiration level, the option is rejected and the next option is considered; a choice is made once all attributes of a single option exceed their aspiration levels, and no further options are considered (i.e., the satisficing heuristic simply considers options in the order they appear and does not necessarily consider all available options). Similar to the satisficing heuristic, though considering all available options prior to choice commitment, some strategies will reject an option once k attributes violate their corresponding aspiration levels where $k \leq m$ (compatibility test; Beach & Mitchell, 1987), m is the number of attributes, and k set by the decision maker; the degenerate cases are also considered different strategies, where $k = 1$ (conjunctive strategy; Coombs & Kao, 1955) or $k = m$ (disjunctive strategy; Coombs & Kao, 1955).

Aspiration levels are also used in *compensatory* strategies, which assume that a negative value on one attribute is traded off against, or can be compensated by, a positive value on another attribute; in our example, a decision maker might be willing to bear the high price of the option *because* they believe they are receiving an option of high quality. The motivating idea behind compensatory strategies is that the decision maker forms an overall impression of an option's attributes, rather than a focus on any attribute in particular. This 'overall impression' might be derived through a relatively simple cognitive process such as using aspiration levels to generate a simple tally of the frequency of 'good' (met the aspiration level) and 'bad' (did not meet the aspiration level) attributes of an option (Alba & Marmorstein, 1987; Montgomery & Svenson, 1976); an attribute might also be 'neutral' (has no effect on the decision). Different decision rules can then be applied to the tallies such as selecting the option with the largest number of good attributes, the fewest bad attributes, or some maximization operation over a combination of good and bad attributes.

Aspiration levels are not common to all decision strategies, though; the trade-off embodied in compensatory strategies could take place without any reference to aspiration levels. For example, the additive model (Fishburn, 1970) and weighted additive models (Tversky, 1969; also known as multi-attribute utility theory, Montgomery & Svenson, 1976) suppose that the 'overall impression' of an option is derived through a relatively demanding cognitive process where a subjective value or *utility* is assigned to each of an option's attributes, summing those utilities across attributes, and comparing an option's summed utility to an internal or external criterion and/or the utility of other available options. Similarly, not all non-compensatory strategies assume aspiration levels. For example, the lexicographic heuristic (Tversky, 1969) assumes that attributes are evaluated sequentially in the order of their subjective importance or *weight*; attributes with high subjective weight bear stronger influence on choices than attributes with lower subjective weight. Each attribute comparison can terminate the decision process if there is a single option that is superior to all others (for similar, see minimum difference lexicographic rule; Montgomery & Svenson, 1976).

1.2. Discriminating Between Consumer Decision Strategies

Discriminating between the various decision strategies has been a persistent challenge in the literature, which has hindered theoretical consensus. The challenge is due to the problem of reverse inference: different strategies can predict similar decisions, so given a set of observed decisions how are we to infer the specific strategy that was in use? This can be made all the more challenging since the strategies in use depend on (typically) unobserved properties of the decision maker. The reverse inference problem afflicts many investigations into the strategies underlying multi-attribute decisions, and is even a well-known challenge for simpler, single-attribute decisions (e.g., perceptual decision making, Ratcliff & Smith, 2004).

A range of methods have been proposed that aim to discriminate between the various consumer decision strategies in particular and the reverse inference problem more generally. These methods generally fall into one of two classes: process-oriented and outcome-oriented methods. Process-oriented methods aim to understand decision strategies by tracing the processes or sequence of steps people use to acquire decision-relevant information, where the decision that people ultimately make is typically of less interest; common examples include Mouselab (Payne et al., 1988), eye tracking, Active Information Search (Engländer & Tyszka, 1980), and many more (e.g., Reisen et al., 2008). As noted by others (Reisen et al., 2008), process-oriented methods provide rich information about the search for information but are generally silent about how people integrate the information to inform their choice. For example, they can provide rich data on the order in which decision makers investigate attributes and options though those data are typically non-informative with respect to other questions of strategy, like whether the attribute information was or was not compared to aspiration levels.

In contrast, outcome-oriented methods tend to emphasize the choices people ultimately make and infer the most likely decision strategy on the basis of the chosen option. The outcome-oriented approach typically includes mathematical or structural modeling of the relationship between the inputs (composition of the choice alternatives) and outputs (chosen options), typically with regression-based methods or discrete choice models (Train, 2009). As with the process-oriented methods, the outcome-oriented approaches provide insight into some aspects of decision strategy, including the subjective weight that was assigned to different attributes, though are silent on others, such as the order in which attribute information was considered.

There have been some attempts to integrate the process- and outcome-oriented methods, such as DecisionTracer (Riedl et al., 2008), which represent an exciting avenue to potentially move beyond the shortcomings of either method when considered in isolation. Nevertheless, such unified approaches are not without their own shortcomings. For example, intervening on the search process (via process-tracing methods such as Mouselab or Active Information Search) can interfere with the decision strategies used in naturalistic environments where information is typically not concealed-until-searched, and unintrusive process-tracing methods such as eye tracking can be prohibitive in some contexts (e.g., experimental environments required to accurately track eye movements, or equipment cost; Riedl et al., 2008).

Here, we investigate a novel application of an outcome-based method to provide unique process-based insight into consumer decision strategies – Systems Factorial Technology (SFT; Townsend & Nozawa, 1995). SFT is a non-parametric, response-time based framework that permits conclusions about the mental architecture of information processing systems – how people process and integrate multiple sources of information to inform decisions; we provide further explanation in the next section, including a classification of decision strategies into classes of mental architectures. An advantage of SFT in this context is its capacity to predict unique signatures in data for different mental architectures. This means that observing a particular signature in data provides evidence both in favor of the observed architecture and against alternative architectures, which in turn allows us to determine the viability of whole classes of consumer decision strategies. With this approach, SFT has been foundational to understanding mental architectures in various domains including perceptual identification and classification (Eidels et al., 2010; Fifić & Townsend, 2010), multidimensional categorization (Fific et al., 2008; Little et al., 2011), selective attention (Chang et al., 2016), face perception (Cheng et al., 2018), and even social phenomena including the other race effect (Yang et al., 2018) and clinical conditions including autism spectrum disorders (Johnson et al., 2010). For an overview of SFT we refer the reader to Harding et al. (2016). For a detailed background with example applications we recommend Algom et al. (2015) or Little et al. (2017).

1.3. Consumer Decision Strategies as Mental Architectures

The strategies people use to select between multi-attribute options – whether those options represent consumer products that vary in price and quality or nondescript objects that vary in shape and brightness – can be conceptualized as *mental architectures*. A mental architecture is a statement about how a system deals with multiple sources of information: how it processes each information source, stops searching for new information, integrates information, and ultimately performs an action on the basis of the information. Mental architectures are interpreted at a higher level of abstraction than the decision strategies introduced earlier, in the sense that multiple strategies can adhere to a single mental architecture but no two mental architectures can describe a single decision strategy. A consequence is that if a single decision strategy appears to fall within the purview of more than one mental architecture then that strategy has not been sufficiently specified, in the sense that there is more than one valid computational implementation of the sequence of steps that encode the decision strategy. Throughout, we consider five mental architectures that have been investigated in great detail in the literature (for a review of the properties of the architectures, see Algom et al., 2015).

When presented with options described by multiple sources of information, the decision maker can take one of two approaches: they can process all of the available information, or only a subset of the information, before committing to a decision; this is known as the *stopping rule*. In the context of multi-attribute and multi-alternative choice, the stopping rule could refer to partial or complete processing of the available (i) attribute information, or (ii) options. We refer to the former in our explanation below and test in our experiments, though we note the methods can also apply to the latter. If all sources of attribute information are independently processed prior to making a decision, the stopping rule is said to be *exhaustive*. This would be like ensuring that both the price and quality are processed to completion – say, compared to their corresponding aspiration levels – before deciding to accept or reject an option. For example, the disjunctive strategy states that an option is only rejected once it is determined that *all* of the option's attributes do not meet their corresponding aspiration levels. An exhaustive stopping rule is also used in the dominance strategy (Lee, 1971) which proposes that an option will only be chosen if it is equal to or better than the available alternatives across all attributes, and the simple majority rule (Arrow, 2012) that assumes a similar approach to the dominance strategy though ensures that a choice is always made – the option that is best across the largest number of attributes. This classification of strategies into elements of mental architectures is shown as Set 3 in Table 1. We note that the classification presented in Table 1 is merely a proposal, which is open to further debate.

Nevertheless, the (tentative) classification serves our aims in this article.

Similar to exhaustive stopping rules, some strategies assume that all sources of information are used to inform a decision except that each source is not processed independently. Rather, the sources are integrated into a single representation reflecting the value of an option, and the decision is then made on the basis of the integrated value. This is known as a *coactive* architecture, and it appears in some of the most commonly studied decision strategies; in Table 1 these strategies are categorized as Set 4. For example, the equal weight and weighted additive strategies assume that a subjective utility is assigned to each of an option's attributes and those attribute utilities are summed to form an overall utility for each option. The decision maker selects the option with the highest total utility; the individual attribute values do not influence the decision outcome directly, only indirectly via their influence on the integrated utility of the option.

In contrast to exhaustive or coactive processing of all available attribute information, people might make a decision after processing only a subset of it, known as a *self-terminating* or minimum time stopping rule. A self-terminating stopping rule is in place when a decision maker accepts or rejects an option once they have processed some, though not all, of the attribute information; for example, price *or* quality. Self-terminating stopping rules are common to a number of decision strategies, shown in Table 1 as Sets 1 and 2. For example, the lexicographic heuristic only considers as many attributes as required to identify a single option that is superior to alternatives. In a similar vein, the satisficing heuristic will reject an option as soon as an attribute is encountered that violates its corresponding aspiration level.

Orthogonal to stopping rules, another key property of a mental architecture is whether multiple sources of information are processed in serial or parallel. *Serial* processing occurs when a decision maker sequentially processes each source of attribute information – one at a time. This would be like first considering the quality (or price) of an option, perhaps comparing it to the corresponding aspiration level or assigning a subjective utility to the attribute value, and only then processing the price (quality). Serial processing is a key feature of the lexicographic heuristic and elimination by aspects (Tversky, 1972), where attributes are assessed one at a time in the order of the subjective importance (weight). In contrast, a decision maker might process the different information sources simultaneously – all information at the same time, known as *parallel* processing. In our consumer example, this would not require one-after-the-other processing of the two attributes; the comparison to the aspiration levels, or assignment of subjective utilities, happens at the same time for both attributes.

With just a few exceptions, it is clear from Table 1 that there have been few theoretical commitments regarding serial or parallel processing in previously proposed decision strategies; strategies we have classified as serial *or* parallel (column 2) have not been sufficiently specified in the literature to identify the processing style of the strategy. This places some limitations on our approach to testing mental architectures in consumer decisions, such as the capacity to discriminate smaller subsets of strategies (e.g., compatibility test from the conjunctive strategy). Nevertheless, we argue that our approach to testing mental architectures (classes of decision strategies) rather than pairs or small sets of specific decision strategies is potentially a new avenue for investigation in the literature. We also argue that the current lack of theoretical commitment to properties of the mental architectures is an opportunity: to gain greater theoretical insight, we ought to commit our psychological theories to sequences of computational rules that have unambiguous interpretations. In doing so, we will have greater capacity to discriminate between those theories in data (for extensive discussion of this argument, see Lewandowsky & Farrell, 2011). We now provide an overview of SFT and the unique predictions it generates for mental architectures.

1.4. Systems Factorial Technology And The Double Factorial Paradigm

Here, we provide proof of concept of the capacity for SFT to provide insight into the mental architecture underlying consumer decisions. SFT comprises a suite of non-parametric analyses that, when combined with the *double factorial* *paradigm* (Townsend & Nozawa, 1995, described below), permit process-based inferences from outcome-based response time data. SFT can address questions including: Do people process *all available information* prior to making a decision (i.e., exhaustive processing) or a *subset of information* (i.e., self-terminating)? Do people integrate disparate sources of information into a single representation (i.e., coactive)? Are multiple sources of information processed *one source at a time* (i.e., in serial) or *multiple sources simultaneously* (i.e., in parallel)? We note that whilst SFT has been extended to investigate mental architectures in tasks involving $m > 2$ attributes (Yang et al., 2014), most literature to date has focussed on the study of $m = 2$ sources of information, a convention we also follow here.

We first provide an overview of the double factorial design that is critical to SFT followed by our modification appropriate for the study of consumer-like choices. The double factorial design is most simply described in the context of perceptual detection. Suppose a participant is presented with two sources of information such as two lights, where a single light is located at the left and the right of a display. On each trial, each source of information can have a target present (e.g., a light switched on) or absent (the light is switched off). The double factorial paradigm is typically administered with one of two types of decision rules. If the participant is instructed to indicate whether *any* target is present then a self-terminating stopping rule (sometimes called an *OR* rule) is an efficient strategy, since the presence of just one target is sufficient to respond accurately. In contrast, if the participant is instructed to indicate whether *both* targets are present, then an exhaustive stopping rule (sometimes called an *AND* rule) is required to accurately respond to the task. This 2 (target: present, absent) $\times$ 2 (source: left, right) design can be thought of as the 'first' factorial component of the 'double factorial' paradigm.

The 'second' factorial component of the 'double factorial' paradigm corresponds to a 2×2 manipulation embedded within the cells of the first factorial component. This second factorial component is a manipulation of the salience of the information presented at each source, when a target is present. In the perceptual detection example a high salience target might be a bright light and a low salience target might be a dim light, where the left and right lights can each appear with high or low salience. These four cells of the embedded design are, by convention, labeled by their level of salience (H=High, L=Low). Thus, for a low salience stimulus in the first source of information and a high salience stimulus in the second, we get a trial label of LH; the corresponding labels for all four cells of interest within the design are HH, HL, LH, and LL.

An effective salience manipulation is one where a high (low) salience target speeds (slows) processing of the corresponding source of information, and has no impact on the speed of processing for the other source of information. That is, if the light at the left location is changed from High (H) to Low (L) salience then it only affects the speed of processing of the stimulus appearing at the left location; it has no effect on the stimulus appearing at the right location. This independent manipulation of the processing speed for each source of information is known as *selective influence* and it is integral to the interpretability of the signatures derived from SFT.

We can test in data whether the assumption of selective influence was upheld by examining the distribution of response times in each of the four double target cells (HH, HL, LH, and LL). We use the survivor function, which represents the probability that an event has not occurred by time t; for a random variable X the survivor function is $S_X(t) = Pr\{X > t\}$, the complement of the more familiar cumulative distribution function. Thus for a particular cell from our double factorial design, say, when the first source of information has 'Low' salience and the second source has 'High' salience, or LH, we refer to the survivor function of the response times as $S_{\mathrm{LH}}(t)$. An effective salience manipulation implies that four cells of primary interest to our design satisfy the ordering $S_{HH} < \{S_{HL}, S_{LH}\} < S_{LL}$. That is, decisions are fastest when both sources of information have high salience (HH), intermediate when one source is of high salience and the other is low salience (HL, LH), and slowest with two low salience sources (LL).

With an effective salience manipulation in place, we use the four survivor functions to generate the Survivor Interaction Contrast (SIC),

$$\mathrm{SIC}(t) = \big[S_{\mathrm{LL}}(t) - S_{\mathrm{LH}}(t)\big] - \big[S_{\mathrm{HL}}(t) - S_{\mathrm{HH}}(t)\big].$$

Figure 1 illustrates the SFT-generated SIC signatures corresponding to the five mental architectures introduced earlier (for a deeper intuition behind the shape of SICs, we refer the reader to Townsend & Nozawa, 1995). The upper row shows that the SIC for a coactive architecture has an early negative deflection that is quite small followed by a later positive deflection that is considerably larger. Parallel architectures exhibit only a negative deflections (exhaustive) or positive deflections (self-terminating), whereas serial architectures have both negative followed by positive deflections of equal size (exhaustive) or no deflection from the baseline at all (self-terminating). The SIC predictions for the different mental architectures, and as a consequence the various decision strategies, are shown in the rightmost column of Table 1.

Discriminating between serial exhaustive and coactive architectures in data based on the SIC alone is challenging, as they share the qualitative features of a negative deflection followed by a positive deflection from the baseline. The issue is resolved by measuring area between the SIC(t) and baseline for all t; that is, integrating the SIC with respect to t. Since responses time are a strictly positive random variable, the integration of the SIC reduces to an interaction contrast of the mean response times for each cell (Houpt et al., 2014), the Mean Interaction Contrast (MIC),

$$\mathrm{MIC} = \big[M_{\mathrm{LL}} - M_{\mathrm{LH}}\big] - \big[M_{\mathrm{HL}} - M_{\mathrm{HH}}\big].$$

The MIC can discriminate between a serial exhaustive architecture (MIC $= 0$, since the early negative and later positive deflections are of equal size) and a coactive architecture (MIC > 0, since the early negative deflection is smaller than the later positive deflection). MICs for the five architectures are illustrated in Figure 1, and in the fourth column of Table 1.

Our aim is to identify the MIC and SIC signatures in data from a consumer-based SFT task. Given the classifications in Table 1, this will allow us to report

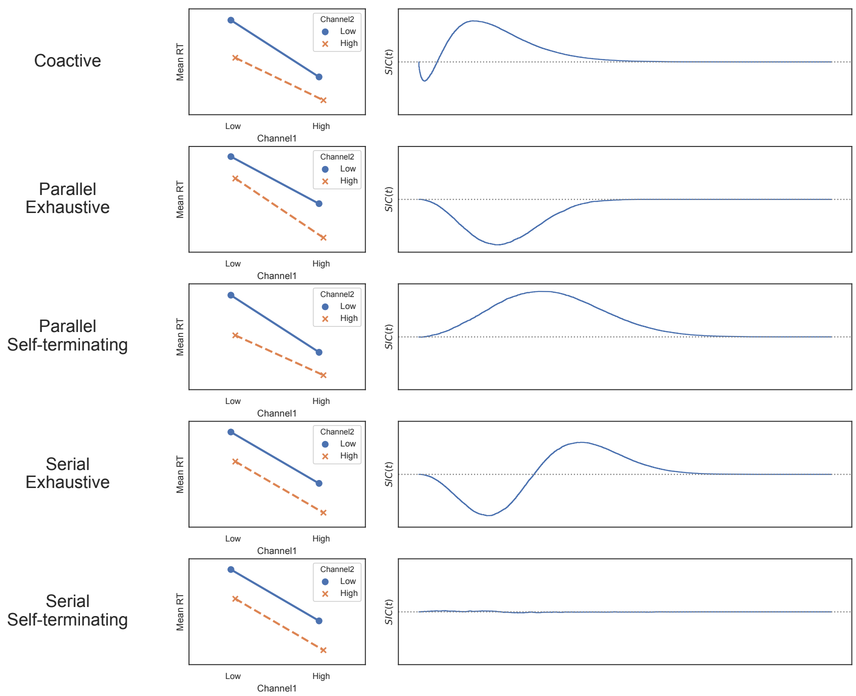

Figure 1: Five mental architectures (rows) and their Mean Interaction Contract (MIC; left column) and Survivor Interaction Contrast (SIC; right column) as predicted by SFT.

evidence in favor of or against different mental architectures, and as a consequence decision strategies. For example, if we were to observe MICs of 0 and SICs with early negative and later positive deflections of equal size – representing a serial exhaustive architecture – then we can conclude that participants made decisions in a manner that were consistent with the disjunctive strategy, dominance strategy, and/or the simple majority decision rule, shown as Set 3 in Table 1. Critically, such a diagnosis of the architecture would allow us to conclude that the observed decisions *were not* consistent with any of the other decision strategies.

1.5. Systems Factorial Technology In Consumer Decisions

We modified the stimulus display of the double factorial design so as to test hypotheses about mental architectures in consumer decisions. To describe this modification, we begin with a description of the standard outcome-based method for examining consumer preferences in the applied choice literature – the discrete choice experiment (DCE). DCEs are choice scenarios where participants select between hypothetical products or services. DCEs can be composed of an arbitrary number of alternatives each of which can be defined by an arbitrary number of attributes.

Figure 2A provides a schematic overview of the generic DCE structure: a number of options are presented in a choice set, where each option in the set is described along several *attributes*. Figure 2B provides an example of four such attributes that might be of relevance for consumers selecting a hotel when traveling for business: cost per night, review rating, distance to the central business district, and gym availability. Each option takes on a particular value for each attribute – these are *attribute levels*; for example, the cost of the hotel might be $180 or $330 per night, or the hotel may or may not have a gym. Given this setup, DCEs ask participants to state their most-preferred option from the set of hypothetical options. Participants complete a number of such hypothetical choice sets, each of which contains a different configuration of attribute levels for each option. Through the attributes and levels structure of the choice sets and the pattern of choices across those sets, detailed inferences can be drawn regarding the utilities (subjective value) of the options that were most likely to generate participants' expressed preferences (for detailed overview, see Louviere et al., 2000; Train, 2009). Although DCEs and their associated analyses have proven useful in accounting for and predicting choices, they are silent on the decision strategies people used to arrive at their choices.

We modified the appearance of the typical DCE choice setup so as to adhere to the requirements of the double factorial design, and accordingly the SFT analysis tools. As a first test, this involved the simplest possible multi-attribute choice scenario: a single option defined by two attributes. The participant was asked to accept or reject a single presented option according to a pair of aspiration levels, one for each attribute, which were specified by the experimenter. For example, the sample stimulus shown in Figure 2C might be accompanied with

(A) Schematic discrete choice experiment

☐	☐		☐
Option 1	**Option 2**	**…**	**Option *N***
Attribute 1	Attribute 1		Attribute 1
Attribute 2	Attribute 2		Attribute 2
…	…		…
Attribute *M*	Attribute *M*		Attribute *M*

(B) Example discrete choice experiment with hotels

	☐	☐	☐
Attributes	**Hotel 1**	**Hotel 2**	**Hotel 3**
Cost/night	$250	$330	$180
Review rating	75%	89%	68%
Distance to CBD	2km	.5km	2km
Gym available	Yes	Yes	No

(C) Double factorial task with hotels

$250

75%

Figure 2: (A) Schematic discrete choice experiment (DCE) to study consumer preference, and an example with hotel choices (B). Several options are displayed. Options are defined by attributes, and each attribute can take on a varying number of levels. The participant's task is to select their most preferred option. (C) An adaptation of the DCE stimulus in (B) to the double factorial design studied in this paper. A single option is shown that is defined by a set of attributes and levels. The participant's task is to indicate if they are willing to accept or reject the option.

the aspiration levels "accept options that are cheaper than $300 and have review scores greater than 60%". In this case, the correct response would be to accept the option since it met the aspiration level on both attributes. If the aspiration levels had instead been described as "accept options that are cheaper than $200 with review scores greater than 60%", the correct response would be to reject the option since it did not meet the aspiration level of the price attribute. The experimenter-specified aspiration levels can be thought of as representing a fixed referent option such that each trial presented a novel multi-attribute option that was to be judged as better than the fixed multi-attribute referent, in which case the correct response was to accept the presented option, or worse than the referent, and therefore to reject the presented option.

Our consumer SFT task mirrors the typical stimulus setup used in SFT-based studies of, say, perceptual detection: each trial presents an option that is defined by two attributes, where the decision is determined by comparing those attributes to their aspiration levels. We note upfront that our task considerably simplifies the DCE. One important difference is a switch from preferential choice (as in DCEs) to veridical choice (i.e., there is a correct answer, as in our task). Our use of experimenter-specified aspiration levels introduces a level of artificiality to this consumer context, since in naturalistic contexts consumers determine their own aspiration levels, or might not use aspiration levels at all. Nevertheless, we deemed the experimental control gained through this design choice to be a necessary starting point in providing proof-of-concept evidence that SFT, via the double factorial design, is a promising tool to study the architecture underlying high-level cognition such as consumer decisions. We return to the issue of experimenter- vs consumer-specified aspiration levels in the General Discussion. We consider the work presented here to be a starting point for expanding the scope of SFT for the study of the mental architectures in consumer decisions; for example, generalization to an arbitrary number of attributes (Yang et al., 2014), choice sets with more than one displayed option, or user-specified aspiration levels.

Here, we present two experiments that provide a starting point for investigating classes of decisions strategies that can be ruled in or ruled out given individual participant data. Experiment 1 examined decision making strategies when attribute information was presented in a numeric fashion (i.e., dollar values for price, % ratings for quality), which is quite common in the consumer literature. Experiment 2 generalized the experimental design to examine decision strategies when attribute information was presented in a symbolic manner (i.e., dollar symbols for price, star ratings for quality), which is an increasingly common method to convey product information in online product and service providers. As we show below, a benefit that emerged from presenting information in the symbolic format was a decrease in overall participant response times, which we suspect was due to the simpler processing requirements when comparing the features of the presented option to the referent option, in turn increasing the precision of the conclusions drawn from the SFT analysis.

2. Experiment 1

The aim of Experiment 1 was to examine the mental architecture of consumer decisions with a SFT-based analysis of a simplified consumer choice task. At the outset, participants were given aspiration levels in the form of a decision rule; this was a dollar value and a quality rating to use as a 'response threshold'. They were asked to apply these aspiration levels to a series of decisions about consumer-like options, determining whether each option should be accepted (i.e., met the aspiration levels) or rejected (i.e., did not meet one or both of the aspiration levels).

2.1. Method

2.1.1. Participants

Sixty four undergraduate psychology students participated in exchange for course credit. All participants had normal or corrected-to-normal vision and provided informed consent prior to participation.

2.1.2. Design and Materials

The experiment involved a hypothetical consumer choice scenario where participants were asked to review candidate hotels for an executive's upcoming business trip; Figure 2C provides an example of the stimulus display. The executive was willing to stay in any hotel that was cheaper than $500 per night *and* had a review rating better than 50%; these were the experimenter-defined aspiration levels, reflecting an *AND* rule in the SFT nomenclature. Participants were presented with a single candidate hotel at a time and were instructed to *accept* any hotel that met both aspiration levels (i.e., the hotel was cheaper than

$500 and had a rating better than 50%) and to reject all other hotels. That is, they were to reject hotels that did not meet either aspiration level (i.e., more expensive than $500 *or* poorer than 50% rating) or both criteria (i.e., more expensive than $500 *and* poorer than 50% rating).

Consistent with the structure of the double factorial paradigm, we manipulated the difficulty of the hotel judgments by selectively increasing or decreasing the numeric distance between an option's attribute values and the aspiration levels. This manipulation assumes that it is more challenging to discriminate numbers that are numerically closer to a criterion than numbers that are numerically distant. This has been shown for two-digit numbers by Hinrichs et al. (1981) and in three-digit numbers in Hinrichs et al. (1982). We grouped the numeric distance manipulation into two levels that represent *low* and *high* salience, reflecting values that exceeded the aspiration levels by a small and large margin, respectively. The attribute values of each hotel were then composed of a randomly sampled value from one of three distributions, separately for each attribute – the *low* or *high* non-overlapping salience distributions, or a *distractor* distribution. The latter was introduced to ensure that some stimuli required a reject decision (i.e., too expensive, poor rating) to prevent automated 'accept' responding in the absence of stimulus processing. Figure 3 illustrates the three distributions of the salience manipulation, separately for both attributes. Values for the high salience, low salience and distractor cells were randomly sampled from normal distributions with, for the price attribute, means of $200, $400 and $700, respectively, and standard deviation $25, and for the rating attribute, means of 80%, 60% and 30%, respectively, and standard deviation 2.5%. The mean of the distractor distributions was set to the midpoint between the low and high salience distributions yet located on the opposite side of the aspiration level (i.e., greater than $500, poorer than 50%). The distributions ensured that the price was always three digits and rating was always two digits.

Factorially crossing the low salience (L), high salience (H) and distractor (D) levels for each attribute produced a design with nine cells, which we represent with the two letter coding system introduced earlier; the first and second letters

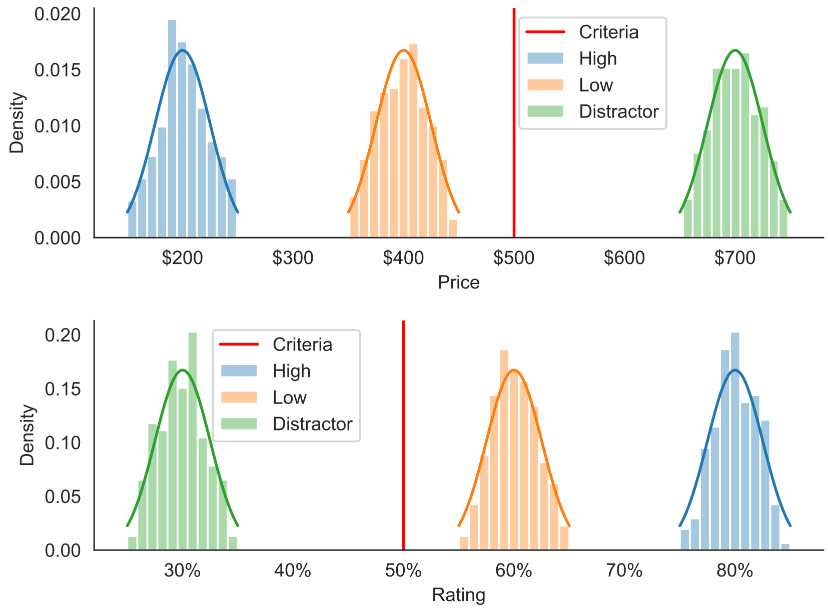

Figure 3: The distributions of values for the price and rating attributes for the Accept condition of the experiment. Attribute values for each hotel were randomly sampled from one of three distributions: high salience, low salience or distractor (blue, orange and green distributions, respectively), according to the contingencies described in the main text. Density curves show the data-generating distributions and histograms show the sampled values for a representative participant. The vertical red line illustrates the experimenter-defined decision criterion placed on each attribute.

correspond to the level of the price and rating attributes, respectively. Four cells of the design are critical to our SFT analysis: the presence of two attribute values that met the aspiration levels (i.e., HH, HL, LH, LL), where the correct response was to accept the hotel; these are known as double targets. As these cells reflect the key data source for SFT we over-represented them in the design to maximize data for analysis: 66% of all trials were from these four critical cells. The remaining five cells comprised trials where only one of the attributes met its aspiration level (DL, DH, LD, HD; single targets) or neither attribute (DD; double distractor). In these five cells the correct response was to reject the hotel and they were evenly split across the remaining 33% of trials. The $\frac{2}{3}$

vs $\frac{1}{3}$ split ensured an appropriate balance between response types throughout the experiment (i.e., accepting or rejecting the hotel).

We refer to the task as described thus far as the *Accept* condition: the *AND* rule specified an aspiration level for each attribute that was required in order to accept a hotel, while all other hotels (i.e., single targets, double distractors) were to be rejected. We also tested a *Reject* condition where the *AND* rule specified an aspiration level for each attribute that was required in order to reject a hotel, while all other hotels were to be accepted. All details of the Reject condition were as described above except for the following: the instructions were described as an executive that was *not* willing to stay in any hotel that was more expensive than \$500 per night *and* had a review rating less than 50%, and for the price and rating attributes the H, L and D distributions had means of \$800, \$600 and \$300, and 20%, 40% and 70%, respectively. The focus of the task – the Accept or Reject condition – was manipulated between subjects in a pseudo-random manner to ensure equal sample sizes.

2.1.3. Procedure

The participant was introduced to a choice task that involved making decisions about hotels. They received instructions that described the aspiration levels they were to use to guide their decisions, which differed depending on the participant's focus condition (Accept, Reject). The rule was described at the start of the task and provided as a reminder during the rest breaks following each block.

The trial timeline was as follows. A centered fixation cross was shown for 1 second, which was then replaced with the hotel stimulus for the trial. Participants were free to respond at any time post-stimulus onset. Once a response was registered, the stimulus was removed from the display, and then the fixation cross for the next trial was shown. The stimulus timed out after a display time of 4.5 seconds, after which it was removed from the display and no response was recorded for the trial. If the participant did not respond before the trial time out, or if they responded within 300ms of stimulus onset, they received an additional 2 second penalty prior to commencement of the next trial. Participants completed a practice block of 30 trials with correct/error feedback to familiarize them with the task. The main task consisted of 4 blocks of 210 trials for a total of 840 trials. Correct/error feedback was not provided in the main task.

We counterbalanced across participants the display order of the two attributes (price above rating as in Figure 2C, or rating above price; both attributes were always presented simultaneously), and whether the left or right hand ('Z' or '/' keys) corresponded to accept or reject responses. The task was presented in `Expyriment (version 0.90)`, a Python library for cognitive experiments (Krause & Lindemann, 2014), on Windows 7 PCs with 23-inch monitors and a 60Hz refresh rate. The stimuli were presented in white text on a black background at approximately 35mm in width and height and subtended a visual angle of 2.7° at a viewing distance of 75cm.

2.2. Results

We took a stepped analysis approach to ensure that the data from our consumer-like version of the double factorial task met the assumptions of SFT. We first screened for particularly fast or slow responses followed by an assessment of accuracy across key cells of the design, and then assessed the selective influence assumption of the salience manipulations; the selective influence assumption, critical to SFT, states that correct responses to high salience stimuli must be faster than correct responses to low salience stimuli. We assessed the number of participants who met a standard of performance at each step of the analysis, where participants who fell below this standard (e.g., low accuracy) or violated selective influence were excluded prior to commencing the next step of the analysis.

2.2.1. Fast, Slow, And Inaccurate Responses

We removed all trials with no response (.5% of trials), responses faster than 300ms on the assumption they were fast guesses (.1% of trials), and responses slower than the .95 quantile of each individual's response time distribution to reduce the potential influence of very slow responses in the tail of the distribution (5% of trials, by definition). We then classified each of the nine cells of the design into four categories so as to assess mean accuracy: *double targets* where both the price and rating met the aspiration levels; *price single target* where the price met the aspiration level but rating did not; *rating single target* where the rating met the aspiration level but price did not; and *double distractors* where neither the price or rating met the aspiration level.

Figure 4 displays response accuracy across the four categories separately for each participant. Around one third of participants had quite low accuracy, hovering around 60-70% for the two single target categories. We used a performance standard of at least 80% accuracy in the *least*-accurate of the four categories. This criterion, marked by a vertical line in Figure 4, led to the removal of 27 of the 64 participants (42%). Although this number may seem high relative to standard cognitive psychology experiments, we believe it is warranted in this context given it is the first application of SFT to decisions in a consumer-like domain. With this strict criterion we ensured that we worked with the best data so as to provide the best chance for a successful proof-of-concept test of the approach; phrased differently, if SFT does not provide useful insights with this restricted data set, it is highly unlikely to provide useful insights when considering a more inclusive data set. Once we confirm that SFT is viable for consumer-like decisions, we can develop the methodology so as to improve performance across participants, and hence obtain more inclusive data sets for analysis. At a practical level, maintaining participants who performed with relatively low accuracy would likely introduce two potential problems, since SFT operates on the distribution of correct response times: there would be insufficient data to analyze distributional properties, and there would be a greater chance of including participants who did not perform the task as instructed. Interestingly, significantly more participants were excluded from the Reject condition (21) than the Accept condition (6), $\chi^2(1) = 8.3$, $p = .004$. This suggests that an *AND* rule in the Reject-focused condition was more challenging for participants to implement than an *AND* rule in the Accept-focused condition, possibly because it conflicts with the natural decision rule people use when considering consumer options – deciding which product to buy rather than rejecting the potentially numerous options that will be rejected.

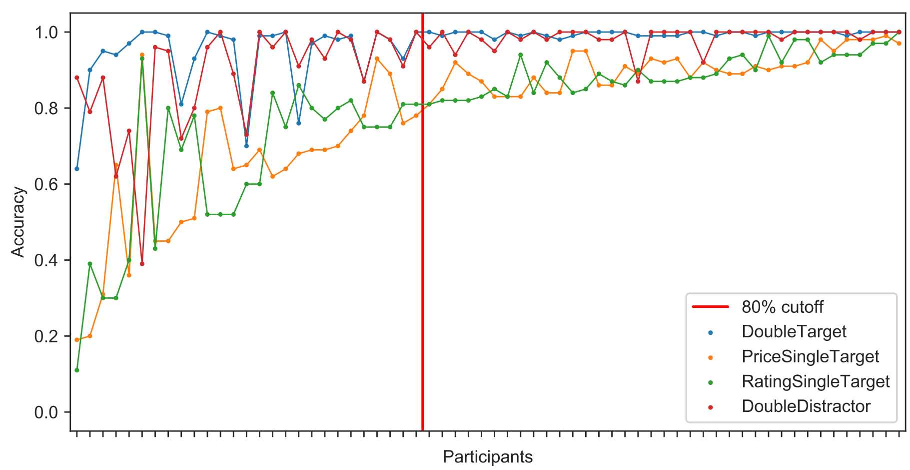

Figure 4: Accuracy across the four categories – double target, single target price, single target rating and double distractor – separately for the 64 participants in Experiment 1, sorted by minimum accuracy across the four trial types. The vertical red line indicates the 80% criterion used as the performance standard.

2.2.2. *Selective Influence*

We next assessed whether our salience manipulation produced data consistent with the assumption of selective influence. Selective influence was tested by comparing individual participant response time distributions from the four cells of interest (HH, HL, LH, LL) with one-sided Kolmogorov-Smirnov tests (Houpt & Townsend, 2010), assuming $\alpha = 0.15$ (Johnson et al., 2010). A violation of selective influence was considered to be any statistically significant ordering of response time distributions such that $S_{LL} < \{S_{HL}, S_{LH}\} < S_{HH}$; that is, where the LL cell was faster than either HL or LH, and/or that HL or LH were faster than HH. There were 7 significant violations of selective influence; these participants were not analyzed further. Visual inspection of the survivor functions indicated an additional 12 participants with incorrect ordering of the response time distributions in at least one component of the distribution. It is likely that these orderings were not so extreme as to be detected by the significance test, though we nevertheless removed these participants from further analysis; their survivor functions are displayed in supplementary material.

Figure 5 shows the response time survivor functions from the four critical cells for those participants that passed the criteria of no statistically significant violations as well as a visual inspection of the response time distribution ordering. Of these 18 participants (48.6% of the accurate responders sub-sample), two showed statistically significant ordering of *all* response time distributions such that $S_{HH} < \{S_{HL}, S_{LH}\} < S_{LL}$; that is, the HH cell was faster than LH/HL, and LH/HL were faster than LL. The remaining participants showed a subset of the significant effects while all demonstrated the expected visual ordering of functions. Figure 5 makes it clear that there was considerable across-participant variability in the response patterns. For example, some participants showed very strong differentiation (i.e., HH was faster than HL and LH which in turn were faster than LL; e.g., participants 12 and 53), while others showed very little differentiation between the four cells across all components of the distribution (e.g., participant 19).

2.3. Survivor Interaction Contrast

Figure 6 shows the SIC curves corresponding to the survivor functions in Figure 5. We assessed deviations from 0 in the SICs with the D statistic (Houpt & Townsend, 2010). Following Fox & Houpt (2016), we assumed a lenient significance criterion ($\alpha = 0.33$) since a conservative p-value can bias the test toward flat SICs (i.e., $SIC(t) = 0$ for all t).

As shown in Figure 6 the majority of participants made decisions consistent with a serial self-terminating architecture (10 of 18, flat SIC with no statistically significant deviation from baseline). The remaining participants made decisions consistent with a parallel processing architecture, with an even split between a self-terminating stopping rule (4 of 18, positive-only statistically significant deflection in SIC) and an exhaustive stopping rule (4 of 18, negative-

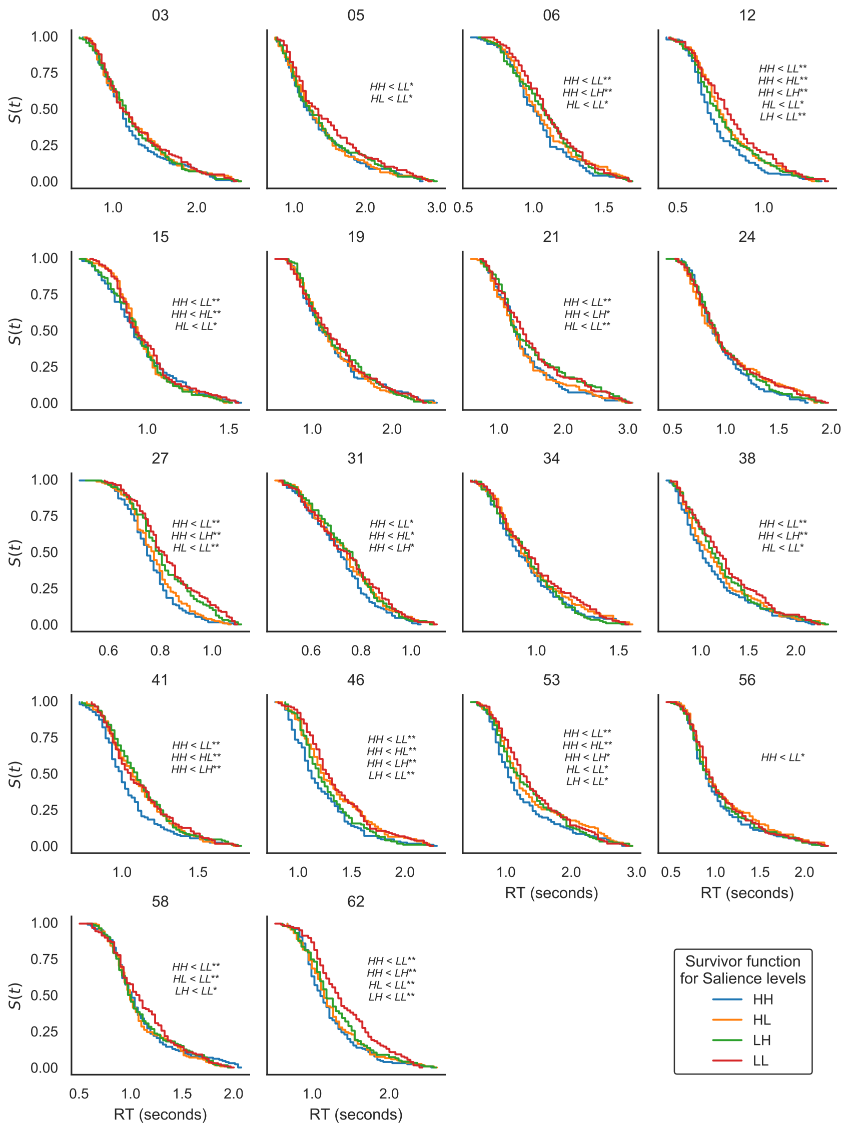

Figure 5: Response time survivor functions for participants included in the primary analysis of Experiment 1. Survivor functions are shown separately for each of the four double target cells (HH, HL, LH, LL). Panels are marked with the results of individual-participant Kolmogorov-Smirnov tests between cumulative distribution functions (complement of the survivor function). p-value indicators for the effects are shown in superscript (** $p < 0.05$, * $p < .15$). Non-significant effects are not shown.

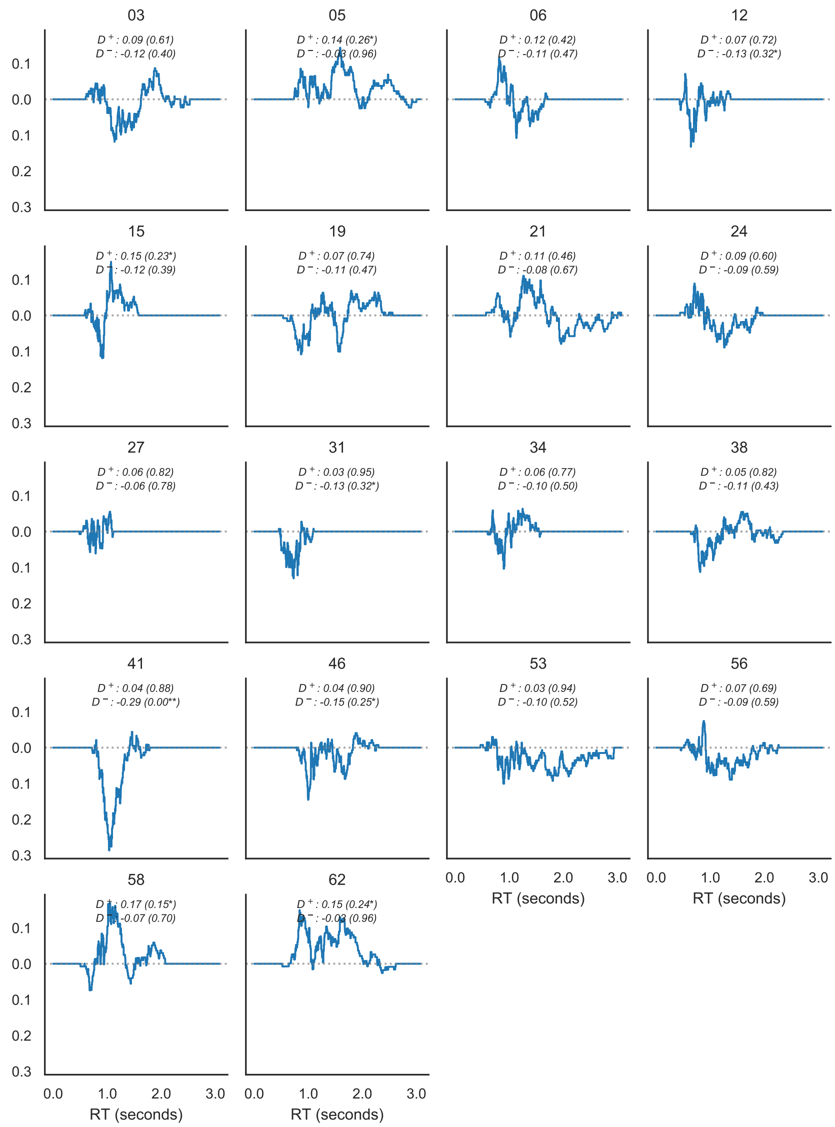

Figure 6: SIC curves for participants included in the primary analysis of Experiment 1. Panels are marked up with the results of individual-participant D statistic tests. p-value indicators for the effects are shown in superscript (** $p < 0.05$, *$p < .33$).

only statistically significant deflection in SIC). We also note that the Accept- and Reject-focused conditions were exactly even split across the three architectures (i.e., 5:2:2 vs 5:2:2). Table 2 summarizes the SIC results in terms of the sets of decision strategies outlined in Table 1, which we discuss in greater detail below.

Table 2: The decision strategies of each participant, as determined by the shape of the SIC curves, in Experiment 1. MIC predictions are omitted as they were not pertinent to the results of Experiment 1. Decision strategy sets are as described in Table 1.

Participant(s)	SIC	Architecture	Decision Strategies
03, 06, 19, 21, 24, 27, 34, 38, 53, 56	flat	Serial Self-terminating	Sets 1 and 2
-	negative $\rightarrow$ positive	Serial Exhaustive	Set 3
05, 15, 58, 62	positive	Parallel Self-terminating	Set 2
12, 31, 41, 46	negative	Parallel Exhaustive	Set 3
-	negative $\rightarrow$ positive	Coactive	Set 4

An interesting result from Experiment 1 is that we did not observe any SICs consistent with the remaining two mental architectures considered in SFT: serial exhaustive and coactive. This is convenient from an analysis perspective because the latter two architectures cannot be unambiguously discriminated on the basis of their predicted negative-then-positive SIC; one must analyze the MIC, which has a different prediction for the two architectures, to discriminate them. Since an analysis of participant MICs would not provide new evidence, we do not report it here.

2.4. Discussion

There was considerable variability in the mental architectures that were most consistent with participant responses, leading to a reasonably wide range of possible decision strategies that participants might have used in our consumer-based SFT task (Table 2). Nevertheless, there were some commonalities that allow for interesting conclusions.

2.4.1. Decision Strategy

The decisions of many participants (78%, 14 of 18) were consistent with a self-terminating stopping rule. For the majority of this subset (10 of 14), product attribute information was processed in serial; the hotel's price was processed first and only then was the quality considered, or the reverse order –

quality was processed first and then price. Such a mental architecture is consistent with well-studied strategies that assume a sequential processing approach where the participant is willing to commit to a decision before all available attribute information has been processed, including elimination by aspects and the lexicographic heuristic (Set 1, Table 1). These participants' data were also consistent with decision strategies that do not specify whether processing occurs in serial or parallel (Set 2, Table 1), as are the remaining minority of the subset of participants who used a self-terminating stopping rule (4 of 14; with evidence for a parallel processing architecture). These strategies include the satisficing heuristic, compatibility test and the conjunctive strategy. A minority of participants (4 of 18) used a parallel exhaustive architecture, consistent with the simple majority rule, and the dominance and disjunctive strategies.

It is interesting that most participants made decisions consistent with a self-terminating stopping rule given that the nature of the decision task was such that both pieces of attribute information (price and quality) were necessary to provide a correct response to double targets (i.e., trials where both attributes met their aspiration levels). Intuitively, one might expect the *AND* rule that participants were instructed to use would most naturally map to coactive or exhaustive processing architectures, rather than a self-terminating rule. One possibility is that the simplified and low-stakes nature of the decision task allowed participants to engage in heuristic-based minimal processing (such as a self-terminating stopping rule) with minimal consequences on decision outcomes (Hoyer, 1984; Wright, 1975). We return to the point of simplifying decision strategies in the General Discussion.

One of the strongest findings was that no participants made decisions consistent with a coactive decision architecture. Given the strategy classification in Table 1, this allows us to exclude as candidate explanations a complete set of decision strategies that includes some of the simpler heuristics, such as the frequency of good and/or bad features and the majority of confirming dimensions heuristic. It also allows us to exclude more detailed coactive strategies such as the equal weights heuristic and its generalization, the weighted additive rule. The latter two strategies have been extensively studied in the literature (e.g., Carpenter et al., 2016; Riedl et al., 2008; Van de Calseyde et al., 2014), so our SFT-based classification that allows us to exclude this class of strategies as potential explanations of decision behavior in this context is quite powerful.

2.4.2. Participant Performance

Our stepped analysis procedure led to the removal of more participants than is typical in psychological studies. Although a small proportion of participants may not have sufficiently engaged with the task, as is possible in any psychological experiment, the majority of participants reported that they understood the goal of the task and attempted to follow the task instructions. We suggest that one possible cause of the low accuracy rates that led to the exclusion of approximately 1/3 of the sample could be the result of some participants erroneously applying a *disjunctive* choice rule instead of the instructed *conjunctive* choice rule. In the Accept-focused condition, such an approach would reject a hotel when the price *and* quality both failed to meet the aspiration levels (disjunctive), as opposed to when just one attribute did not meet its aspiration level; the reverse holds for the Reject-focused condition. The outcome of this effect can be seen in Figure 4: some participants had low accuracy in *price single target* or *rating single target* trials, or both, but still performed with high accuracy on the *double distractor* trials.

Another aspect of the design that may have contributed to the relatively high exclusion rate is that the salience manipulation had quite a small effect, in the sense that there was little differentiation between the response time distributions across the four key cells of the double factorial design. The effect of a weak salience manipulation is such that it becomes relatively easy to observe violations of selective influence, as even small fluctuations in the response times of a few trials can lead to an incorrect ordering between two distributions. In addition, the numeric presentation of attribute information may have induced a mixture of cognitive processes in transforming the presented digits to internal magnitudes, for subsequent comparison to the aspiration levels. This might have occured as a holistic process when comparing two digit values to an internal referent (Dehaene et al., 1990; Hinrichs et al., 1981; Zhang & Wang, 2005), like comparing the hotel's two-digit quality rating to its aspiration level. In contrast, values represented with three or more digits tend to be processed in a stepped manner, sequentially from most to least significant digit (Hinrichs et al., 1982; Poltrock & Schwartz, 1984); this process might have been used in comparing the hotel's three-digit price to its aspiration level. This potential convolution of processes (holistic and stepped) to represent the two attribute components of the stimulus may have introduced complexity to the response time distributions, shrouding the effect of the critical salience manipulation.

3. Experiment 2

In Experiment 2 we aimed to address the issues of a weak salience manipulation and differential stimulus representation processes that may have contributed to participant performance in Experiment 1. We investigated this aim by testing whether the decision strategies observed in Experiment 1 were driven by the underlying structure of the decision task or the way in which product information was presented to the decision maker, by manipulating the mode of information presentation. In Experiment 1, attribute information was presented as a text-based description – for example, a hotel might cost $432 with a 67% quality rating. Although this is the status quo presentation mode for applied DCE decision research (e.g., Louviere et al., 2000, 2015), it differs to a rapidly growing presentation mode in online purchasing contexts – perceptual or *symbolic* attribute information. For example, `yelp` presents price and quality ratings for restaurants as dollar signs and stars, respectively; `booking.com` makes use of star ratings for all accommodation offerings; and insurance comparison sites such as `iSelect` and `CompareTheMarket` present information as a mixture of text, stars, and tick boxes. An underlying motivation of symbolic presentation modes is to allow the user a means to rapidly gain an impression of products in isolation and allow for simpler comparison across product offerings.

In Experiment 2, we therefore used the same underlying task structure as in Experiment 1, including the use of experimenter-specified aspiration levels, except that attribute information was presented symbolically – for example, a hotel might have a price represented with 3 (from a maximum of 5) dollar signs and a quality rating of 4 (of 5) stars. This manipulation of presentation mode allowed us to investigate whether the manner in which product information was delivered to the decision maker influenced the way they processed that product information, by comparing results between Experiments 1 and 2. It also draws a closer to the existing literature in perceptual processing that has been studied with SFT (e.g., Fific et al., 2008, 2010; Little et al., 2011).

3.1. Method

3.1.1. Participants

Thirty-four undergraduate psychology students participated in exchange for course credit. All participants had normal or corrected-to-normal vision and provided informed consent prior to participation.

3.1.2. Design and Materials

The design and materials of Experiment 2 were identical to the hypothetical consumer choice scenario in Experiment 1 except where noted. The primary difference was that attribute information was presented as horizontally sequenced perceptual symbols rather than digits: hotel prices were displayed as a number of dollar signs from a total of five – a cheap hotel as $ through to an expensive hotel as $$$$$ – and hotel quality was displayed as a 5 star rating system – a poor rating as ★ through to an excellent rating as ★★★★★. In the Accept-focused condition, the instructions stated the executive was willing to stay in any hotel that was cheaper than $$$ *and* had a rating better than ★★★; these were the experimenter-defined aspiration levels, again reflecting an *AND* decision rule. In the Reject-focused condition, the executive was not willing to stay in any hotel that was more expensive than $$$ *and* had a rating less than ★★★.

The Accept- and Reject-focused instructions were again manipulation between subjects.

The salience of the hotel price and quality attributes was manipulated by increasing or decreasing the number of symbols for each attribute relative to the aspiration levels; this manipulation is similar in spirit to the numeric distance effect underlying the salience manipulation in Experiment 1. Attribute values that exceeded the aspiration levels by one symbol were defined as *low* salience, and by two symbols as *high* salience. Table 3 outlines this manipulation in terms of the numbers of each attribute symbol for each salience level, separately for the Accept- and Reject-focused conditions, all of which were judged relative to the aspiration levels: $$$ and ★★★.

We manipulated within subjects the symbolic display across two levels, intended to represent different formats that consumers may come across when making online purchase decisions. The display modes differed with respect to whether the dollar signs and ratings were shown in a relative sense (i.e., in the context of the full range of their scale, as a 'score' out of the maximum possible score) or an absolute sense (with no reference to the range of the scale). In the *Relative* condition, the rating stars were shown in bright yellow, $rgb(254, 233, 0)$, and the dollar signs in bright green, $rgb(24, 255, 0)$, with the remaining (unfilled) stars or dollar signs shown in gray, $rgb(127, 127, 127)$. For example, if a hotel had 4 stars (in yellow) and 2 dollar signs (in green) it would be shown alongside 1 additional gray star and 3 additional gray dollar signs. In the *Absolute* condition, price and rating symbols were shown in the same colors as the Relative condition but the grayed (unused) symbols were not shown. In both conditions, stimuli appeared on a black background.

In both display conditions there were placeholders for 5 stars and 5 dollar signs that were horizontally and vertically centered in the display; the placeholders were gray in the relative condition and invisible in the absolute condition. Trials with fewer than 5 stars/dollar signs were filled from the leftmost placeholder toward the right, giving the display a 'left-aligned' appearance. This display setup mirrors the way symbolic attribute information is typically en-

Table 3: The high, low and distractor levels of the salience manipulation of the price and quality attributes in Experiment 2. Salience levels are shown as symbolic stimuli, similar to their presentation in the experiment, for the Accept-focused (upper) and Reject-focused (lower) conditions.

Accept

Price	High	Low	Distractor
Quality			
High	$ ★★★★★	$$ ★★★★★	$$$$ ★★★★★
Low	$ ★★★★	$$ ★★★★	$$$$ ★★★★
Distractor	$ ★★	$$ ★★	$$$$ ★★

Reject

Price	High	Low	Distractor
Quality			
High	$$$$$ ★	$$$$ ★	$$ ★
Low	$$$$$ ★★	$$$$ ★★	$$ ★★
Distractor	$$$$$ ★★★★	$$$$ ★★★★	$$ ★★★★

countered online, and also facilitates rapid processing of the stimulus via a reduction of information: participants knew they could always begin viewing from one position (the leftmost placeholder).

3.1.3. Procedure

The procedure closely followed Experiment 1 in terms of the instructions (with minor adjustment to account for the symbolic stimulus presentation), trial timeline, and counterbalancing display order of the two attributes (price above rating, or rating above price). The main task consisted of 2 blocks of 210 trials in the Relative display condition and 2 blocks of 210 trials in the Absolute display condition, with order of the two display types counterbalanced across participants: Relative blocks first and Absolute blocks second, or the reverse. Participants completed two 30-trial practice blocks with correct/error feedback: one before each of the Relative and Absolute conditions, to familiarize them with the different stimulus appearance in the two conditions. Participants did not receive correct/error feedback in the main task.

3.2. Results

We followed the same stepped analysis approach as Experiment 1 with the addition that all analyses were conducted separately for the *Relative* and *Absolute* display conditions.

3.2.1. Fast, Slow, And Inaccurate Responses

We removed all trials with no response (.2% of trials), responses faster than 300ms (1% of trials), and responses slower than the .95 quantile of each individual's response time distribution. Figure 7 shows that response accuracy across the four categories – double target, price single target, rating single target, double distractor – was comparable to Experiment 1: 19 and 17 participants met the 80% threshold for the Absolute (56% of participants) and Relative condition (50%), respectively; 15 participants (44%) met the threshold in both conditions. Despite the similarity across experiments in the proportion of participants who met the accuracy threshold, of the participants that exceeded the threshold there appeared to be a more rapid rise toward ceiling performance across all 4 categories in Experiment 2, particularly in the Absolute condition.

3.2.2. Selective Influence

Figure 8 shows the response time survivor functions for the Absolute and Relative conditions. There were fewer statistically significant violations of the

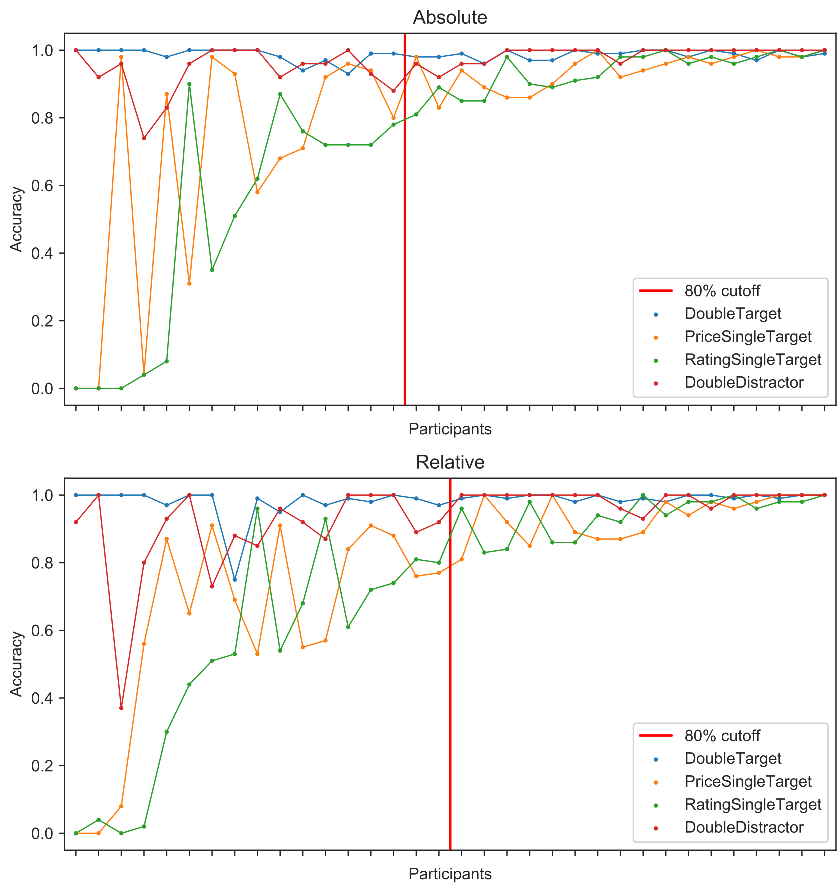

Figure 7: Accuracy across four trial types – double target, price single target, rating single target and double distractor – separately for the 34 participants in Experiment 2, split by the two display conditions (Absolute in the upper panel, Relative in the lower panel). The vertical red line indicates the 80% criterion used as the performance standard.

expected ordering of response time survivor functions in Experiment 2: 4 participants in the Absolute condition and 1 participant in the Relative condition. Nevertheless, visual inspection indicated an additional 3 participants in the Absolute condition and 7 in the Relative condition with incorrect ordering of the response time distributions in at least one component of their distribution. As in Experiment 1, participants who did not show the expected ordering of the HH, HL, LH and LL cells were removed from further analysis; their survivor functions are shown in the supplementary material. Of the participants that remained, the effect of the salience manipulation – in terms of the visual differentiation between response time distributions of the four cells – appeared to be larger in Experiment 2 than in Experiment 1 (compare Figure 8 with Figure 5).

3.2.3. Survivor Interaction Contrast

Figure 9 shows the SIC curves corresponding to the survivor functions shown in Figure 8. Unlike Experiment 1, there was an approximately even split between serial self-terminating (flat SIC) and parallel self-terminating (positive SIC) architectures, with just a couple of participants characterized as a parallel exhaustive architecture (negative SIC), when considered across the two display conditions. This pattern differs to Experiment 1, where there were more than twice as many serial vs. parallel self-terminating architectures. The difference across experiments appears to have been driven by the Relative display condition: slightly more parallel than serial processing architectures, whereas the distribution of the two architectures in the Absolute condition is not too dissimilar to that observed in Experiment 1. As in Experiment 1, we did not observe any data consistent with serial exhaustive or coactive architectures. Table 4 summarizes the SIC results into the sets of decision strategies outlined in Table 1.

Table 4: The decision strategies of each participant, as determined by the shape of the SIC curves, separately for the Absolute and Relative display modes in Experiment 2. MIC predictions are omitted as they were not pertinent to the results of Experiment 2. Decision strategy sets are as described in Table 1.

Display Mode		SIC	Architecture	Decision Strategies
Absolute	Relative			
05, 09, 10, 11, 15, 18	04, 05, 09, 15	flat	Serial Self-terminating	Sets 1 and 2
-	-	negative → positive	Serial Exhaustive	Set 3
03, 13, 29, 30	11, 14, 21, 22, 29	positive	Parallel Self-terminating	Set 2
06, 22	-	negative	Parallel Exhaustive	Set 3
-	-	negative → positive	Coactive	Set 4

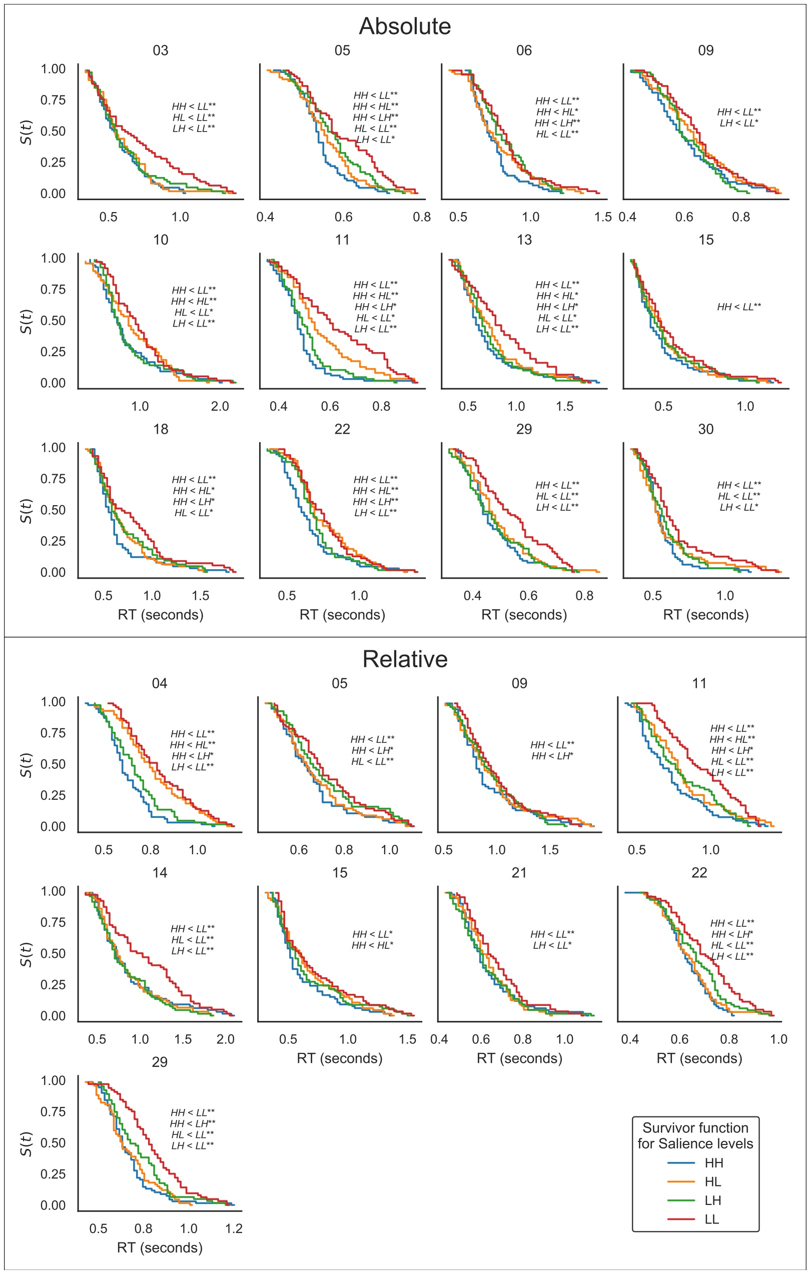

Figure 8: Response time survivor functions for the participants included in the primary analysis of the Absolute (upper) and Relative (lower) display conditions of Experiment 2. All other details are as described in Figure 5.

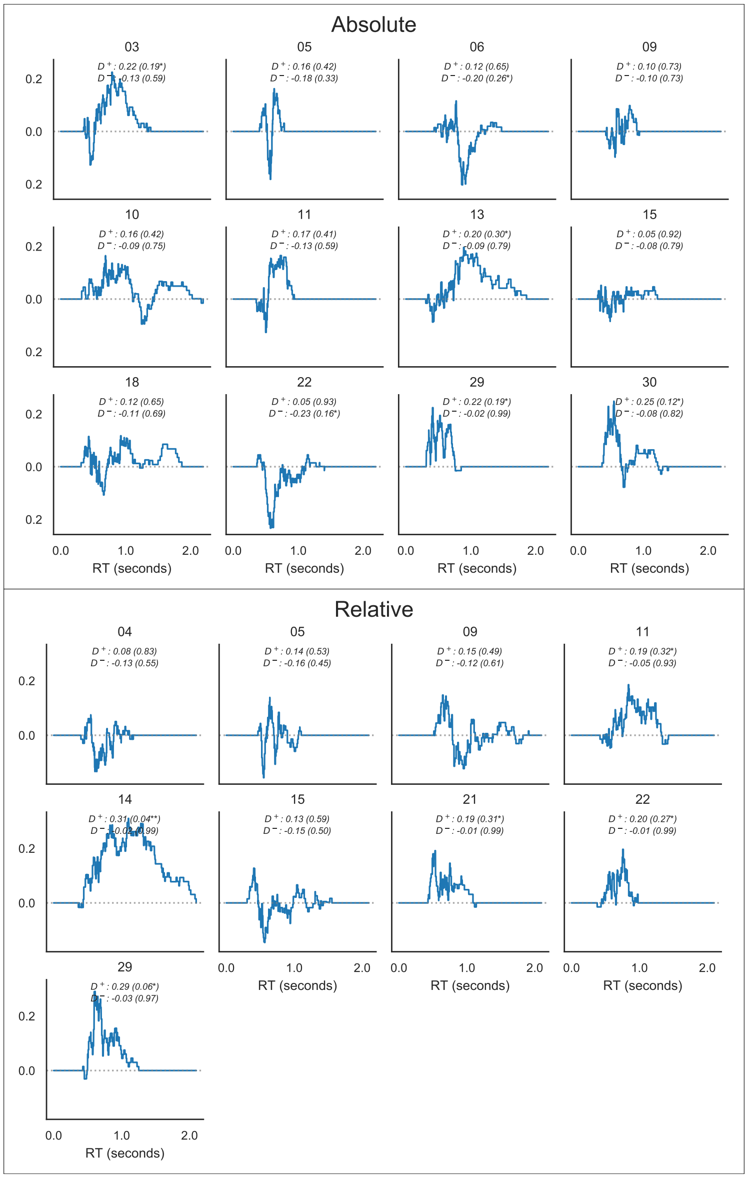

Figure 9: SIC curves for participants included in the primary analysis of Experiment 2, separately for the Absolute (upper panels) and Relative (lower panels) display conditions. All other details are as described in Figure 6.

3.3. Discussion

As in Experiment 1, the vast majority of participants (90%, 19 of 21) used a self-terminating stopping rule (Sets 1 & 2 from Table 1), consistent with simpler rule-based strategies such as elimination by aspects, the lexicographic heuristic and the satisficing heuristic. Interestingly, the majority of this subset of participants processed attribute information in serial when it was presented in an Absolute manner – that is, star or price rating shown without reference to its scale – but in parallel when it was presented relative to the maximum score possible on each attribute scale.

The shift toward more parallel processing in Experiment 2 might be due to the simplification of the stimulus display relative to the numerical representation used in Experiment 1. Indeed, the symbolically presented attribute information of Experiment 2 appeared to more strongly differentiate processing than numerically presented attribute information, as seen in the separation of the response time distributions across key cells of the design (compare Figure 8 with Figure 5). One interpretation of this result is that the manipulation of high and low salience attribute information, critical to SFT, was indeed more visually salient when represented symbolically (as a smaller or larger sequence of objects on screen) than when represented by digit value (where the same number of digits was always present).

The symbolic representation might also have allowed for different processing strategies. For example, the attribute information might be interpreted similarly to a perceptual stimulus – like two lines, representing price and quality – from which a length judgment can be made. It is possible that such a length-judgment approach was easier to apply with the added context of the maximum score possible on each attribute scale (Relative display condition), though this is speculative. With such a strategy, it is plausible that a greater number of participants recruited the more efficient parallel processing architecture in the symbolic task compared to less efficient serial processing in the numerical task.

One might argue that such symbolic displays no longer represent 'consumer'-style decisions. We do not believe this is necessarily the case. Consumers often encounter symbolically represented information in, for example, online purchase contexts. We are also not the first to represent preferential choice stimuli in symbolic displays; for example, others have investigated decision processes through the use of pie charts that represent the tradeoff between option value and risk (e.g., Guo et al., 2017; Ordonez & Benson III, 1997), visual counters to represent probabilities or outcome values (e.g., Dambacher et al., 2016), or even bar charts in personnel selection scenarios (e.g., Tversky, 1969).

As in Experiment 1, we observed no evidence for coactive architectures, again allowing us to rule out strategies such as the weighted additive rule. Unlike Experiment 1, however, in Experiment 2 we also observed no evidence (Relative) or only a small minority of participants (2 of 10, Absolute) using an exhaustive stopping rule, suggesting this form of decision rule was particularly unlikely to be used to process the symbolic display that was manipulated in Experiment 2. This allows us to effectively rule out additional decision strategies as potential explanations of performance: the disjunctive and dominance strategies, and the simple majority rule (Table 1).

4. General Discussion

We have provided proof-of-concept evidence for the utility of a new approach to discriminate between cognitive strategies in consumer decisions. Over decades of research, the consumer decision literature has used a range of techniques in an attempt to uncover the cognitive processes involved in searching for relevant information, deciding when enough information has been evaluated, and ultimately making a choice. Traditionally, these methods have tended to probe one component of the decision process more heavily than others: most commonly, focused on the acquisition of information or the ultimate product choice. Here, we investigated whether a suite of analysis tools that can address both components, and has proven insightful in the context of multi-attribute perceptual detection and classification – known as Systems Factorial Technology (SFT; Townsend & Nozawa, 1995) – might also provide unique insight into consumer decision strategies. Across two experiments, we found there is potential for the use of SFT in understanding decision strategies in the consumer domain. In the decision contexts we studied, there was converging evidence across both experiments to 'rule in' classes of decision strategies that involve sequential, rule-based processing, such as the lexicographic heuristic and elimination by aspects. Both experiments also allowed us to 'rule out' classes of strategies that summate all available product information prior to processing the set of available options, such as the weighted additive model. We discuss these issues in detail below, along with some limitations to the SFT approach.

A major benefit of an SFT application to the consumer choice context is the ability to reduce the large set of candidate decision strategies to a smaller set of higher-level mental architectures. This dimension reduction in turn permits us to simultaneously test the plausibility of a large number of decision strategies in a single experimental paradigm, such that we can now investigate the mental architectures in use across different contexts. With this approach, finding evidence in favor of a particular mental architecture allows a collection of strategies defined by a few common elements to be ruled feasible in the investigated context, and failing to find evidence for other mental architectures allows a different collection of strategies to be ruled infeasible in that context.

In the decision tasks we investigated, involving a sequence of independently presented two-attribute products that were each compared to a fixed referent, we observed consistent evidence for mental architectures with self-terminating stopping rules. The 'stopping rule' addresses the psychological question: when have I seen enough information to make a decision? A self-terminating rule suggests a decision is made as soon as a minimally sufficient amount of information has been processed. This differs to an exhaustive rule where a decision is only made once all available information has been processed, which might be considered a more comprehensive decision style in the context we investigated. The widespread use of the self-terminating stopping rule was interesting in our context because participants were explicitly instructed to use experimenter-defined aspiration levels that required the use of information from both attributes to specify the appropriate response. We suggest that the observation of a self-terminating rules, consistent with rule-based heuristic strategies, might be the result of an effort-minimization approach to the task: from the participant's perspective, achieving a sufficient level of accuracy with the least effort (for a similar explanation in the context of perceptual decision making, see Hawkins et al., 2012).

Such an effort-minimization approach may have been a consequence of making repeated low stakes consumer decisions, though it might also be an adaptive strategy in certain real-life situations. For example, in online and real-life supermarket shopping, we make repeated product choices, most of which are low stakes (e.g., selecting potato chip brand A or B typically has little consequence). In those cases, particular products might be selected on the basis of their most salient features rather than all available information (e.g., Benn et al., 2015). This is not necessarily a negative feature. It might be taken as a demonstration that people will adapt their cognitive strategies to the contexts they find themselves in. In our experiments, considering the decision accuracy across all cells of the double factorial paradigm and the SFT classification to mental architectures, the most plausible task strategy appears to be one of responding on the basis of just one attribute at a time (price, or quality; self-terminating rule) and occasionally switching the focal attribute across trials. This process would give rise to the pattern of accuracy observed in Figures 4 and 7; a similar explanation was covered in the Discussion of Experiment 1. It would be interesting to test whether incorporating a performance-contingent reward structure in the same decision context might push people to adopt a more demanding exhaustive strategy.

Although it appears that people used a common stopping rule across experiments, we found evidence that the way that product information is presented can influence how that decision-relevant information is processed. In particular, we found stronger evidence for serial processing when attribute information was presented numerically yet parallel processing when it was presented symbolically; response times were also considerably faster for the symbolically presented information. One implication of this result is that the proficiency of decision making in situations involving time constraints might be facilitated by symbolic information presentation, which might promote more efficient parallel processing architectures without impacting the "correctness" of the choices made. Furthermore, a natural extension of the SFT consumer decision task is the presentation of attribute information in a *mixed* mode. This follows from a common trend seen in online shopping where the price of a product is presented as a precise numeric value while the quality is presented as a star rating system (e.g., `Amazon`). An investigation of the processing architecture and stopping rules in play when dealing with this kind of context switching within a single choice could be beneficial to the companies that regularly present information in this way.

One shortcoming of the tasks we developed is that the preferential choice process typically involved in consumer contexts was reduced to a veridical choice process: the aspiration levels were defined by the experimenter (not the decision maker), which meant each trial had an objectively 'correct' choice. Of course, in most consumer contexts the individual sets their own aspiration levels based on the collection of personal experiences they bring to each decision scenario. We deemed the reduction from a preferential to veridical choice process as necessary to provide proof-of-concept testing of the possibilities of SFT in a consumer domain. Nevertheless, we suspect that this reduction might have at least partially contributed to our high participant exclusion rates: in our tasks there were clearly defined correct responses, and a substantial proportion of participants were removed on the basis of low accuracy. In preferential choice there are no 'correct' choices, and as such the issue of high exclusion rates based on response accuracy is no longer relevant.

Given the emerging evidence we present here, we suggest that future applications of SFT to consumer decisions take more flexible approaches. One such extension might be through a two-phase paradigm. Phase 1 might be a typical free-choice DCE scenario, used to identify individual participant aspiration levels. In Phase 2 we can specify low and high salience attribute values relative to each individual's aspiration levels identified in Phase 1, after which the product options could be presented in much the same way as the task we developed here. Although this approach allows the participant to determine their own aspiration levels, it requires the assumption that those aspiration levels are stationary across phases of the task.

Yet another approach might be to present product information across a considerable range of attribute levels, in a similar format to Experiment 1, allowing the participant to determine their own aspiration levels on acceptable and unacceptable products. Then, during the analysis phase, an optimization process can be used to estimate the aspiration levels that provided the best account of the set of accept/reject decisions that were observed. With those optimized aspiration levels, one can subset trials into sets of high and low salience trials for each attribute. This would represent a more naturalistic decision process as participants freely select their own aspiration levels, and it would not require a stationarity assumption, though it may be challenging to apply in practice particularly with respect to adequate trial numbers for analysis.

Overall, we find promise for the use of SFT in consumer choice contexts, given careful experimental design and analysis. We conclude with two recommendations for future research in this domain. Firstly, when tasked with a complex decision rule, such as a conjunction, correct use of the rule may increase if participants are provided with feedback throughout the task. This may prevent participants from erroneously slipping into a simpler choice rule, such as focusing on a single product attribute. Furthermore, we suggest future research explores preferential rather than veridical choice rules, to enhance the ecological validity of the approach. Secondly, symbolically presented attribute information appears to be more rapidly processed than numerically presented attribute information, and as a result produces clearer salience effects. It also has a clear mapping to some purchasing contexts, such as quality and value ratings for online purchases.

Acknowledgments

We thank Keelin Garry for assistance with data collection. GEH was supported by an Australian Research Council Discovery Early Career Researcher Award (DE170100177). This research was supported by an Australian Government Research Training Program (RTP) Scholarship.

References

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

Algom, D., Eidels, A., Hawkins, R. X., Jefferson, B., & Townsend, J. T. (2015). Features of response times: Identification of cognitive mechanisms through mathematical modeling. *Oxford library of psychology. The Oxford handbook of computational and mathematical psychology*, (pp. 63–98).

Arrow, K. J. (2012). *Social choice and individual values* volume 12. Yale university press.

Beach, L. R., & Mitchell, T. R. (1987). Image theory: Principles, goals, and plans in decision making. *Acta Psychologica*, *66*, 201–220.

Benn, Y., Webb, T. L., Chang, B. P., & Reidy, J. (2015). What information do consumers consider, and how do they look for it, when shopping for groceries online? *Appetite*, *89*, 265–273.

Van de Calseyde, P. P., Keren, G., & Zeelenberg, M. (2014). Decision time as information in judgment and choice. *Organizational Behavior and Human Decision Processes*, *125*, 113–122.

Carpenter, S. M., Yates, J. F., Preston, S. D., & Chen, L. (2016). Regulating emotions during difficult multiattribute decision making: the role of pre-decisional coherence shifting. *PloS one*, *11*, e0150873.

Chang, T.-Y., Little, D. R., & Yang, C.-T. (2016). Selective attention modulates the effect of target location probability on redundant signal processing. *Attention, Perception, & Psychophysics*, *78*, 1603–1624.

Cheng, X. J., McCarthy, C. J., Wang, T. S., Palmeri, T. J., & Little, D. R. (2018). Composite faces are not (necessarily) processed coactively: A test using systems factorial technology and logical-rule models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*, 833.

Coombs, C., & Kao, C. (1955). *Nonmetric Factor Analysis*. Bulletin (University of Michigan. Engineering Research Institute). University of Mich. Engineering Research Institute.

Dambacher, M., Haffke, P., Groß, D., & Hübner, R. (2016). Graphs versus numbers: How information format affects risk aversion in gambling, .

Dehaene, S., Dupoux, E., & Mehler, J. (1990). Is numerical comparison digital? analogical and symbolic effects in two-digit number comparison. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 626.

Eidels, A., Townsend, J. T., & Algom, D. (2010). Comparing perception of stroop stimuli in focused versus divided attention paradigms: Evidence for dramatic processing differences. *Cognition*, *114*, 129–150.

Engländer, T., & Tyszka, T. (1980). Information seeking in open decision situations. *Acta Psychologica*, *45*, 169–176.

Fific, M., Little, D. R., & Nosofsky, R. M. (2010). Logical-rule models of classification response times: A synthesis of mental-architecture, random-walk, and decision-bound approaches. *Psychological Review*, *117*, 309.

Fific, M., Nosofsky, R. M., & Townsend, J. T. (2008). Information-processing architectures in multidimensional classification: A validation test of the systems factorial technology. *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 356.

Fifić, M., & Townsend, J. T. (2010). Information-processing alternatives to holistic perception: Identifying the mechanisms of secondary-level holism within a categorization paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*, 1290.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. John Wiley and Sons, New York.

Fox, E. L., & Houpt, J. W. (2016). The perceptual processing of fused multispectral imagery. *Cognitive research: principles and implications*, *1*, 31.

Guo, L., Trueblood, J. S., & Diederich, A. (2017). Thinking fast increases framing effects in risky decision making. *Psychological science*, *28*, 530–543.

Harding, B., Goulet, M.-A., Jolin, S., Tremblay, C., Villeneuve, S.-P., & Durand, G. (2016). Systems factorial technology explained to humans. *Tutorials in Quantitative Methods for Psychology*, *12*, 39–56.

Hawkins, G. E., Brown, S. D., Steyvers, M., & Wagenmakers, E.-J. (2012). An optimal adjustment procedure to minimize experiment time in decisions with multiple alternatives. *Psychonomic Bulletin & Review*, *19*, 339–348.

Hinrichs, J. V., Berie, J. L., & Mosell, M. K. (1982). Place information in multidigit number comparison. *Memory & Cognition*, *10*, 487–495.

Hinrichs, J. V., Yurko, D. S., & Hu, J.-M. (1981). Two-digit number comparison: Use of place information. *Journal of Experimental Psychology: Human Perception and Performance*, *7*, 890.

Houpt, J. W., Blaha, L. M., McIntire, J. P., Havig, P. R., & Townsend, J. T. (2014). Systems factorial technology with R. *Behavior Research Methods*, *46*, 307–330. doi:10.3758/s13428-013-0377-3.

Houpt, J. W., & Townsend, J. T. (2010). The statistical properties of the survivor interaction contrast. *Journal of Mathematical Psychology*, *54*, 446–453.

Hoyer, W. D. (1984). An examination of consumer decision making for a common repeat purchase product. *Journal of consumer research*, *11*, 822–829.

Johnson, S. A., Blaha, L. M., Houpt, J. W., & Townsend, J. T. (2010). Systems factorial technology provides new insights on global–local information processing in autism spectrum disorders. *Journal of Mathematical Psychology*, *54*, 53–72.

Krause, F., & Lindemann, O. (2014). Expyriment: A Python library for cognitive and neuroscientific experiments. *Behavior Research Methods*, *46*, 416–428. URL: http://link.springer.com/10.3758/s13428-013-0390-6. doi:10.3758/s13428-013-0390-6.

Lee, W. (1971). *Decision theory and human behavior*. Wiley.

Lewandowsky, S., & Farrell, S. (2011). *Computational modeling in cognition*. California: Sage.

Little, D. R., Altieri, N. E., Fifić, M. E., & Yang, C.-T. E. (2017). *Systems factorial technology: A theory driven methodology for the identification of perceptual and cognitive mechanisms*. Elsevier Academic Press.

Little, D. R., Nosofsky, R. M., & Denton, S. E. (2011). Response-time tests of logical-rule models of categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*, 1.

Louviere, J. J., Flynn, T. N., & Marley, A. A. J. (2015). *Best–Worst Scaling: Theory, Methods and Applications*. Cambridge, UK: Cambridge University Press.

Louviere, J. J., Hensher, D. A., & Swait, J. D. (2000). *Stated Choice Methods: Analysis and Applications*. Cambridge, UK: Cambridge University Press.

Montgomery, H., & Svenson, O. (1976). On decision rules and information processing strategies for choices among multiattribute alternatives. *Scandinavian Journal of Psychology*, *17*, 283–291.

Ordonez, L., & Benson III, L. (1997). Decisions under time pressure: How time constraint affects risky decision making. *Organizational Behavior and Human Decision Processes*, *71*, 121–140.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 534.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.

Pfeiffer, J. (2012). Fundamentals on decision-making behavior. In *Interactive Decision Aids in E-Commerce* (pp. 15–45). Springer.

Poltrock, S. E., & Schwartz, D. R. (1984). Comparative judgments of multi-digit numbers. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 32.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two–choice reaction time. *Psychological Review*, *111*, 333–367.

Reisen, N., Hoffrage, U., & Mast, F. W. (2008). Identifying decision strategies in a consumer choice situation. *Judgment and decision making*, *3*, 641.

Riedl, R., Brandstätter, E., & Roithmayr, F. (2008). Identifying decision strategies: A process-and outcome-based classification method. *Behavior Research Methods*, *40*, 795–807.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*, 99–118.

Townsend, J. T., & Nozawa, G. (1995). Spatio-temporal properties of elementary perception: An investigation of parallel, serial, and coactive theories. *Journal of Mathematical Psychology*, *39*, 321–359.

Train, K. (2009). *Discrete choice methods with simulation*. (2nd ed.). Cambridge: Cambridge University Press.

Tversky, A. (1969). Intransitivity of preferences. *Psychological review*, *76*, 31.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*, 281–299.

Wright, P. (1975). Consumer choice strategies: Simplifying vs. optimizing. *Journal of marketing research*, (pp. 60–67).

Yang, C.-T., Fifić, M., Chang, T.-Y., & Little, D. R. (2018). Systems factorial technology provides new insights on the other-race effect. *Psychonomic bulletin & review*, *25*, 596–604.

Yang, H., Fific, M., & Townsend, J. T. (2014). Survivor interaction contrast wiggle predictions of parallel and serial models for an arbitrary number of processes. *Journal of Mathematical Psychology*, *58*, 21–32.

Zhang, J., & Wang, H. (2005). The effect of external representations on numeric tasks. *The Quarterly Journal of Experimental Psychology Section A*, *58*, 817–838.