CHAPTER FOUR

Choice from among Intentionally Selected Options

Patrick Shafto*,[1] and Elizabeth Bonawitz[§]
*Department of Mathematics and Computer Science, Rutgers University - Newark, Newark, NJ, USA
[§]Department of Psychology, Rutgers University - Newark, Newark, NJ, USA
[1]Corresponding author: E-mail: patrick.shafto@gmail.com

Contents

Abstract

How do people choose among a set of options? Previous work has provided either normative accounts of choice (suggesting that options maintain fixed utility) or subjective accounts (suggesting that utilities are context dependent). Neither account fully explains the systematicity and variability of people's choice behavior. We propose a novel factor, the intentional selection assumption. When people are provided with a set of options to choose among, they treat the set of options as an intentional selection by a person with a specific question in mind. By considering the likely relevant features of the options, the model shows how consideration of the goals and beliefs of the individual presenting the options can help resolve uncertainty about, and lead to variability in, the relative utilities of the options. We discuss how our model explains previous behavioral data and helps to bridge the normative and subjective accounts in the literature.

Psychology of Learning and Motivation, Volume 63
ISSN 0079-7421
http://dx.doi.org/10.1016/bs.plm.2015.03.006

1. INTRODUCTION

Given a set of options, how do people choose among the possibilities? For example, suppose you need a ride into the city and your neighbor points out that you have the choice between the train, a red bus, and a blue bus, and queries which you would prefer.[1] How might you decide which to choose? You may consider the relative costs of each mode of transportation. For example, the train may be faster or slower, or more or less expensive, than the buses.

Alternatively, you may have background knowledge that affects your interpretation of the question. For example, you may know that today is the day of a big sporting event in the city and that red and blue are the colors of the two competing teams. Given this knowledge, a simple calculation of the relative prices hardly seems adequate to answer your neighbor's question. Instead, it seems possible that your neighbor is asking about your stance toward the game: do you support the red team or the blue team or are you staying out of it? Under this scenario, if you support the red team, you might choose the red bus despite the fact that it takes longer and is more expensive.

This deceptively simple question—how one, when presented with a palate of options by another person, decides which option to choose—underpins a wide variety of disciplines. Psychology, economics, marketing, computer science and other fields have investigated choice, resulting in a variety of formal models of choice behavior (e.g., Luce, 1959; McFadden, 1977; Sutton & Barto, 1998; Thurstone, 1927; Yellot, 1977). Choice also plays a critical, but often overlooked, role in linking theoretical questions to empirical data. This is especially true in fields driven by behavioral experiments where the experimenter offers participants choices among different options (e.g., behavioral experiments and survey research).

In theory and practice, the options presented are treated as randomly selected. In theoretical treatments, this is manifest in explicit assumptions regarding the relationship between observed choices and unobserved but possible choices (Luce, 1959). In practice, this is manifest in the fact that people do not randomly select which options to present, and most often make no attempt to analyze for the effects of which options were (and were not)

[1] The example is adapted from the red bus/blue bus problem attributed to McFadden. As will become clear, the example here highlights a different phenomenon.

presented. Instead, options are often purposefully selected by the questioner to answer his or her question. Even when the options are not selected with intent, the chooser may perceive them as such. The implication is that the options presented to chooser may affect which choice he or she makes, which has implications for experimenters' inferences about utilities or preferences of the chooser.

Psychologists and economists have long known that the problem of inferring utilities from people's choices is non-trivial. Indeed, there is ample evidence that the options presented by a questioner affect which choices are made. One classic example is the Compromise Effect (Simonson & Tversky, 1992). In the task, participants chose between cameras that varied in price and quality. Comparing two cameras (one high in quality and price, and the other low in quality and price) led to equal choice between the two. However, when a third camera is added at either extreme (e.g., lower quality and lower price), then the intermediate camera is favored over the previously equally favored one.

This contextually flexible choice behavior is an instance of the broader argument against the tenets of classical economics, which assume that people have fixed and stable—objective—notions of utility that determine choice behavior. If fixed and stable utilities function as proposed by classical economics, then choice involves simply selecting the option with the highest utility. In the simplest variant, choosers might noisily maximize—choosing the option with highest utility, but occasionally making a mistake.

Adopting an axiomatic approach, Duncan Luce (1959) proposed the Luce choice rule, which varies from this noisy maximization form of choice in that choice is systematically probabilistic. The Luce choice rule suggests that the probability of each item is proportional to its utility relative to the other presently-available items; items are selected in proportion to their weight. That is, while objects may have some stable intrinsic utility, it is not necessarily the case that a chooser will always maximize by selecting the highest utility item. An option whose utility is twice as high as another will be selected twice as often as the other, however one-third of the time the option with less utility may still be selected. This proposal has influenced a great variety of subsequent work in psychology, economics, statistics, and computer science. The idea that people select options in proportion to their probability has also been proposed as a model of both adult and children's inductive inferences (e.g., see Bonawitz, Denison, Griffiths, and Gopnik (2014); Vul and Pashler (2008)). The convergence of these fields provides support for the idea that people's judgments reflect a probabilistic sample from normative utilities.

However, empirical research has cast doubt on normative accounts (e.g., Debreu, 1960; Tversky, 1972; Simonson & Tversky, 1992). For instance, consider the Compromise Effect. In this situation, people are not merely choosing in proportion to probabilities. If so, the addition of an option may affect the probability of selecting the high-quality, high-price camera and the low-quality, low-price camera, but it could not affect their *relative* probabilities. Based on this and numerous other phenomena, researchers have argued that choice is not merely probabilistic, but also that utility judgments are idiosyncratically affected by context and therefore not stable. These subjective theories, which assume utilities are context dependent, create a different set of challenges. One might propose that people do not have a stable concept of utility at all. However, if so, it is not clear how to explain the degree of systematicity observed in choice behavior.

It remains an open question how to explain the systematic variability of people's choices, especially as a consequence of contextual factors. It is possible that people are able to roughly approximate utilities. This would explain some of the systematicity of choice behavior. However, the approximation claim falls short of explaining how seemingly irrelevant aspects of context can reverse preferences, as in the Compromise Effect. An important challenge to explaining effects of context on people's choice thus remains.

Here we propose a novel factor that may influence people's choice behavior. This factor we call the intentional selection assumption. We suggest that when people are provided with a set of options to choose among, they treat the set of options as an *intentional selection* by a person with a *specific question* in mind. That is, choosers might consider the goals and beliefs of the individual presenting the options in order to help resolve uncertainty about the relative utilities of the options in a particular context. This proposal provides a novel application of recent research formalizing learning from others (Shafto, Goodman, & Frank, 2012), and we provide detail of how these social assumptions shape choice behavior.

The intentional selection assumption depends on a notion of the relevant features of interest in evaluating the utility of choices. Relevant features allow the utility of an option to change based on a chooser's inferences about the goals and beliefs of the person selecting the options, without moving to a fully subjective approach. That is, our relevant features assumption states that while utilities may be stable over features, the chooser may have uncertainty about which features are relevant in the particular context. The selected options can provide information as to which features are most relevant to consider in a given context. As a result, the overall utility of an item will depend on the relevant features that provide said utility, and the inferences about which features are most relevant will be context dependent. Our approach can therefore be seen as a middle ground between the normative and subjective approaches.

We begin with a review of the basic Luce choice model, which formalizes probabilistic choice given stable utilities, and discuss some results that are difficult to explain with that approach. We then discuss a recently proposed framework for reasoning about and learning from other people's actions (Shafto, Goodman, et al., 2012), and generalize the framework to apply to the choice behavior. We conclude by pointing to open questions and future directions.

2. THE LUCE CHOICE RULE

Consider the train/red bus/blue bus example discussed above. In the example, you are presented with three possible options: taking the train, taking the red bus, or taking the blue bus. For simplicity, we assume that each is equivalent in terms of time and price, and thus you are indifferent among the options a priori. The question is, which option should you choose? In this context, we review normative theory using the Luce choice rule.

Luce (1959) proposed two general principles. The first is that choice is probabilistic. The second is that choice probability should be independent of options that are not included. From these principles the Luce choice rule[2] was derived:

$$P_S(x) = \frac{u(x)}{\sum_{x' \in S} u(x')}. \tag{1}$$

This states that we choose an option x based on its utility, $u(x)$, relative to the utility of the other choices in the set of options S. Assuming stable,

[2] An alternative version of this is the Softmax rule (Sutton & Barto, 1998), where choice depends on a weighted exponential transform of the utility, $P_S(x) = \frac{exp(w \cdot u(x))}{\sum_{x' \in S} exp(w' \cdot u(x'))}$. This formalization allows connection to multinomial logit-based approaches (McFadden, 1977) and random-utility models (Train, 2003).

inherent utilities, the Luce choice rule provides a complete description of choice: choosers evaluate the utility of each option and choose an option that tends to maximize utility.

Let us consider the set $S = \{train, redbus, bluebus\}$ and $x_1 = train$, $x_2 = redbus$, and $x_3 = bluebus$. Our assumption of indifference among the choices implies that these utilities are identical, and thus the probability of choosing each example under Luce choice should be 1/3 as prescribed by Eqn (1).

If instead we assume that taking the train is twice as preferable as either of the other options, Luce choice rule gives a different answer. In this case, the utility of $x_1 = c$, and the utility of x_2 and x_2 is $\frac{1}{2}c$. Based on the Luce choice rule, the probability of choosing the train would be $\frac{c}{\frac{1}{2}c+\frac{1}{2}c+c} = \frac{c}{2c} = 1/2$ and the probability of choosing the red bus and the blue bus would be 1/4 each.

While Luce's approach formalizes choice elegantly, it generates predictions that are inconsistent with intuition. Let us return to the case where buses and trains all have identical utilities. If there were no sporting event and the choice between the red bus and blue bus were meaningless, then it seems clear that these options are not different (Debreu, 1960; McFadden, 1974). Because these two buses are perfect substitutes (the color adds no utility), choice boils down to a question of train versus bus. Thus, although we might be indifferent about the three options, we should not choose each $\frac{1}{3}$ of the time. Instead, the choice probabilities should be closer to 1/2 for the train, 1/4 for the blue bus, and 1/4 for the red bus. Indeed this prediction more closely matches human behavior (Tversky, 1972).

There may be various reasons for the failure of Luce choice rule to capture our intuitions about utilities in these cases (see Pleskac, 2013). Implicit in the Luce Choice model is the assumption of independence. This captures the chooser's assumptions about whether the options were provided independently of each other. It also captures the chooser's assumptions about whether the options that were presented are independent of those options that were not presented. Both of these assumptions may be incorrect when it comes to human judgments in social contexts, a point we return to in the next section.

3. EMPIRICAL ARGUMENTS AGAINST LUCE CHOICE

Empirical research has strenuously tested the assumption that choice is independent of not-present options. This is a very large literature (see

Rieskamp, Busemeyer, & Mellers, 2006; for a recent review), but for present purposes it is sufficient to cover a small number of illustrative results. We aim to span classic theoretical explanations by Tversky and colleagues (Elimination by Aspects and Componential Context Theory; Tversky (1972), Simonson and Tversky (1992)) to highlight some of the features of our approach.[3] We focus on three effects from the literature, including the Similarity Effect (Tversky, 1972), the Attraction Effect (Simonson & Tversky, 1992), and the Compromise Effect (Simonson & Tversky, 1992).

Debreu (1960) proposed a choice scenario that proved difficult for the Luce Choice rule to explain. Consider a music lover choosing between a recording of a Debussy piece or of one of two Beethoven pieces that are essentially the same (same orchestra, song, but different directors). If each recording is equally adored, any pairwise comparison of two of these recordings produces a 50/50 choice. It follows from Luce choice, then, that the choice between the three options should result in 1/3 probability for any of the three recordings. However, this is not how participants respond. Instead, people choose the Debussy piece with closer to 50% probability and then divide the remainder probability between the two Beethoven pieces. A different version of this problem was more recently proposed by McFadden (see also Train, 2003) involving red and blue buses (our introduction example and one we will return to later, albeit following different contextual assumptions). Debreu's example is one demonstration of the importance of context on choice behavior. People's responses depend on the options that are in the choice set.

A second example of context influencing choice is the Attraction Effect (Simonson & Tversky, 1992). In this example, participants were given a choice between a nice pen or 6 dollars. Participants generally favored the money over the pen. However, when a third option—a second, much less attractive pen—was offered in addition to the nice pen and money, the number of participants selecting the nice pen increased as compared to the first condition. That is, the addition of a seemingly irrelevant "distractor" pen increased participants' utility in the nicer pen as compared to the money.

[3] These examples also span types of violations of normative behavior, which are explainable by different kinds of models (Rieskamp et al., 2006).

A third example, also developed by Simonson and Tversky (1992) is the Compromise Effect, which we discussed in the introduction. In this case, when an additional camera is added to a pair previously rated as equal, the intermediately priced camera suddenly becomes more desirable. This example provides additional evidence for the role of contextual factors in influencing utility in choice tasks.

As these examples illustrate, there are multiple different demonstrations showing that context matters in an individual's choice behavior. We propose a formal account that depends on the choosers' intuitive psychological inferences based on the options presented, which leads to different features of relevance. Our goal is to provide a potentially unifying account for these three phenomena, which has proven challenging for theories of some models of choice. We begin with a brief background on the social influences on learning.

4. SOCIAL INFLUENCES ON LEARNING

The assumption of independence from non-present options can be viewed as an assumption that the options are randomly sampled. That is, one might believe that the options are generated without knowledge or relationship to the other options that may have already been drawn. In such a case, even if people did not precisely know the utility of all of the options, we would not expect systematic deviations from the predictions of Luce choice because any variance in the observed options should be unbiased. Nevertheless, the empirical literature shows that indeed, systematic deviations do occur.

Viewed as a random sampling, the assumption of independence of choice from not-present options is analogous to the assumption of random sampling omnipresent in the concept learning literature (e.g., Anderson, 1991; Fried & Holyoak, 1984; Kruschke, 1992; Love, Medin, & Gureckis, 2004; Medin & Schaffer, 1978; Nosofsky, 1984; Pothos & Chater, 2002). In these settings, the problem is typically for a learner to infer the correct concept, given a collection of labeled examples. Debates in the learning literature have focused on how concepts are represented (e.g., rule-based vs prototype vs exemplar), where the process by which the examples are selected is assumed to be ignorable.

Recent computational models and empirical research have suggested that people do not generically assume random sampling in all cases. For example, when drawing inferences, learners appear sensitive to whether examples are generated purposefully from within the concept, called Strong Sampling (Tenenbaum & Griffiths, 2001a; Xu & Tenenbaum, 2007a, 2007b). Learners show stronger inferences about representative examples drawn in pedagogical contexts; in these contexts teachers not only draw samples from within the concept but do so in order to maximize the chances that the learner will infer the correct concept (Bonawitz et al., 2011; Shafto & Goodman, 2008; Shafto, Goodman, & Griffiths, 2014; Tenenbaum & Griffiths, 2001b). For example, in the Bonawitz et al. (2011) study, children observed an experimenter act on a toy to bring about an effect and were then asked to play freely with the toy to figure out how it worked. The experimenter varied whether she generated the event accidentally (as in weak sampling) or whether she generated the event with the goal of teaching the child (promoting pedagogical inferences). Even though the evidence was identical—both groups observed the causal action on the toy—children in these different sampling contexts drew different inferences about the toy and consequently explored the toy in different ways. This shows that, given the very same data, people draw different inferences, a result that suggests learners are leveraging their knowledge about other people to facilitate learning.

Shafto, Goodman, et al. (2012) proposed a framework for formalizing these sorts of social effects on learning (see also Bonawitz et al., 2011 on children's exploratory play; Goodman, Baker, & Tenenbaum, 2009 on casual inference; Shafto, Eaves, et al., 2012 on epistemic trust; Frank & Goodman, 2012 on communication). An important contribution of this work is to focus on the inferential affordances provided to the learner by leveraging intuitive psychological reasoning. Specifically, because people's actions are goal directed, rather than random, we can reason about *why* they do what they do and this has implications for the kind and strength of inferences that can be drawn from an observation. In their framework, learners reason about hypotheses, h, given the observed data, d, selected by the individual,[4] and beliefs about the individual's knowledge, k, and goals, g.

The import of this framework is clear when comparing learning from actions selected by a knowledgeable person whose goal is to help versus

[4] In their paper, the authors differentiate data into actions and effects to facilitate discussion of causal reasoning. Here we simply refer to data.

learning from one who is naive. Formally, learners update their beliefs about hypotheses using Bayes' rule:

$$P(h|d,k,g) = \frac{P(d|g,k,h)P(g)P(k)P(h)}{\sum_{h'} P(d|g,k,h')P(g)P(k)P(h')}, \tag{2}$$

where the key modification is the idea that data, d, are potentially purposefully selected by a knowledgeable, goal-directed person, $P(d|g, k, h)$.

Consider learning from a naive informant; that is, someone who is neither knowledgeable nor necessarily goal directed. Therefore, we can conclude that the data we observe does not meaningfully depend on either their knowledge or their goal (they have neither). To capture this situation, we can reduce data selection from $P(d|g, k, h)$ to $P(d|h)$, and eliminate $P(g)$ and $P(k)$. This reflects the independence of the data from their beliefs or goals and is the social analog of random sampling.

Contrast learning from a naive informant with learning from a knowledgeable person whose goal is to help you learn. In this case, the sampling of the data, $P(d|g, k, h)$, is performed purposefully. Because the goal is to help the learner infer the correct hypothesis, the learner can replace $P(d|g, k, h)$ with $P(h|d, k, g)$.[5] This captures the idea that the person selecting the data is doing it so as to lead the learner to the correct answer.

Importantly, given the same hypothesis, one would expect very different data to be produced by a naive person and a knowledgeable and helpful person. While the naive person produces randomly sampled data, the knowledgeable and helpful person produces data that should disambiguate the correct hypothesis from other similar, but incorrect, hypotheses. Moreover, given the same data, the learner should draw very different inferences if he or she believes the person selecting it was naive, as opposed to knowledgeable and helpful.

For example, in cases where the person choosing the data is knowledgeable and helpful, the observed examples are not independent of the unobserved examples. This is most easily seen in the case of the Bonawitz et al. (2011) described above. The behavior of children in the direct instruction condition in this experiment might be explained by an appeal to making this inference. Children may have considered why the experimenter chose to show the particular evidence (squeaking the toy), given assumptions about the goals of the experimenter (that she was trying to help the

[5] We omit the normalizing constant, which involves summing over the possible data, for simplicity.

child learn about all the functions of the toy). Given the observation of only one function in the direct instruction context, children could infer that there was likely only one function by a simple counterfactual intuition: if there were more than one function, then the experimenter should have also shown those functions. Thus, in this case, lack of evidence becomes evidence of a lack. Under this model, when observing pedagogically sampled data, the learner draws an inference about why the observed data were presented, but also why the unobserved data were omitted. To explain why the teacher chose to demonstrate one function, but not two, the learner infers that there must not be a second function to be demonstrated. In contrast, these inferences would not be drawn in the accidental condition, in which one function is also demonstrated, but by "chance." In this case, the learner has no strong reason to suspect that there is only one function of the toy and may explore more broadly.

We propose that choice behavior is affected by similar social inferences. This reasoning about why options are chosen can modulate the likely relevance of different features. That is, different contrasting options may highlight the relevance of a particular feature and thus influence the perceived utility of the object. The overall implication being that assumptions about why presented (and omitted) options were chosen can affect inferences about features of relevance thus affecting perceived utility. We now turn to a formal description of this intuition.

5. A MODEL OF CHOOSING AMONG INTENTIONALLY SELECTED OPTIONS

Our model begins with the idea that an option, x is a composition of features. There are a potentially infinite set of features that might exist (e.g,. is red, is blue, and is \$15).[6] The utility, u_f, of each feature, f, contributes to the option's overall utility, $u(x)$, by simply summing over the utility of all features.

However, in addition to this unchanging vector of utilities over features, each feature is also weighed by two factors that depend on context: the commonality of the feature, ω, and the probability of relevance, $P(r)$. We return

[6] Throughout, we assume that features are binary, that is, either present or absent. Feature dimensions are modeled through common knowledge about the mutual relevance of, for example, different dollar amounts. That is, we assume that if the feature \$1 is relevant, then \$2, \$3, etc. are also relevant. This joins binary features to approximate continuous dimensions.

to the specification of these terms shortly. We begin by formalizing the contribution of these three factors to an option's overall utility in the following equation:

$$u(x) = \sum_{r} \sum_{f \in x,r} u_f \omega_f P(r). \tag{3}$$

In the case where all features are equally relevant, $P(r)$ is a constant, and the utility of an object, $u(x)$, is precisely the sum of the utilities of its features, as in normative theories of utility. When features are potentially differentially relevant, an option's overall utility is the sum over the weighted utilities of all possible features. As we will see, we will modify this equation slightly to capture how context influences uncertainty about the relevance. We now discuss the three contributing terms, the utility of each feature, u_f, the commonality of each feature, ω_f, and the probability of relevance of a set of features, $P(r)$.

5.1 Feature Utilities

The idea that options are evaluated in terms of the utilities of their features, here u_f, is not new (Restle, 1961). For example, the feature "tastiness" of an option "chocolate ice-cream from JP Licks" could have a large, and positive utility. The feature "price" could have a negative, smaller utility. In this way, if an option is weighed by the utilities of its features, then the total utility of a particular option would simply be the sum of these feature utilities, where positive features contribute to the larger utility and negative features subtract utility.

5.2 Commonality

In our model, options are by definition compositions of features. We propose that, given a set of options, S, one must account for the commonality of a feature across those options before computing that feature's contribution to the option's overall utility. The function of this term is to correct for the possibility of double counting a feature's utility. Consider two options, each composed of a single, different feature, $u_{x_1} = \{u_{f_1}\}$ and $u_{x_2} = \{u_{f_2}\}$. Imagine that the first option has vastly greater utility than the second, $u_{x_1} > u_{x_2}$, such that one would nearly always choose the first option (e.g., a \$1000 bill vs a \$1 bill). Now imagine adding many, many new options all of which are identical to u_{x_2}, that is, each has a single feature, f_2. If each option is considered relative to all other options, as in standard

Luce choice (Eqn (1)), this would lead to a reversal, where when the number of objects f_2 exceeded 1000, the chooser would be more likely to select a single dollar to 1000 dollars. More formally, according to Eqn (1), the probability of choosing the \$1000 bill from a set including m \$1 bills, assuming dollars translate directly to utility, is $\frac{1000}{1000+m}$. As $m \rightarrow \infty$ the probability of choosing the \$1000 bill goes to zero and the probability of choosing one of the \$1 bills goes to one, which is a clear violation of intuition.[7] We focus on features as the objects of choice, and therefore simply adding options does not necessarily change the structure of the problem. That is, f_2 should not gain simply because it is common.[8]

To ensure features remain the focus of choice, we add a commonality factor. Primarily, commonality depends on the number options in the set S that have a particular feature,

$$\omega_{f_S} = \frac{1}{n_{f_S}}, \tag{4}$$

where n_{f_S} is the number of times the feature appears among the options in the set. The commonality factor resolves problems related to the addition of identical objects. For example, in the monetary example described above, commonality weight of the \$1000 bill is $\frac{1}{1}$ and the weight of each \$1 bill is $\frac{1}{m}$, and the probability of selecting the \$1000 bill, assuming the simple Luce choice rule, is $\frac{1000}{1000+1}$. In sum, the commonality term simply ensures that choice is determined by the relative utilities of features, capturing the intuition that the choice should be unaffected by the addition of options that share one of those features.[9]

5.3 Feature Relevance

The main novel contribution of our model is in considering the feature relevance term, $P(r)$. The relevant features assumption proposes that an

[7] This issue is not isolated to cases where the exact same feature is replicated; cases where options have extremely similar features can lead to comparable problems. To handle these cases, one would have to generalize the notion of commonality, for example, by introducing distributions over similar features.

[8] Arguably the opposite is true—rare items tend to have greater utility—although discussion of choice with the possibility of resale is beyond the scope of this chapter.

[9] This formalization assumes that choosers choose exactly one option from among many. This equation would need to be generalized to account for choices of more than one option. Alternative formulations that capture the qualitative prediction—that the relative utility of a feature of a single object decreases as a function of the commonality of that feature across objects in the set of options—are possible.

item's utility is determined by the utility of its *relevant* features (as the name implies). The relevant features assumption allows for the possibility that certain features, such as the color of a bus, are not relevant to the calculation of utility in general and in particular. The relevant features assumption allows that there be variability in the utility of an item, such that it may depend on context in interesting ways. For example, in some contexts, the speed of a method of transportation may be most relevant, whereas in others the price may be more relevant. However, that features may be relevant or irrelevant does not tell us when context should affect inferences, only that it can.

Thus, an important challenge is in specifying how feature relevance is assessed. In our model, feature relevance depends on the full set of options selected as well as an inference about the intention of the questioner in providing the examples. That is, social choice proposes a two-part explanation for how context affects choice. The first is that the utility of an option is a function of the utilities of its features, each of which may or may not be relevant in a given context. The second is that the chooser assumes that the observed selection of options is chosen intentionally, with a specific question in mind.

Importantly, the context plays a role in helping the chooser infer which features f are relevant. Here we focus on the questioner's intentions as the key contextual factor that influences this relevance term.

5.3.1 The Intentional Selection Assumption

The intentional selection assumption states that the observed options are intentionally selected with a question—the relevant features—in mind. The chooser, observing the selected items, can reason about the intended question and use that inference to constrain the uncertainty about which features are relevant. The chooser will have uncertainty about the intentions of the questioner. Thus, the probability of relevance, $P(r)$, must take into account many possibilities. We can think about each of these possibilities, r, as a *hypothesis* that involves the set of features that should be relevant under that hypothesis, r.

Formally, for the set of observed options, S, to constrain the relevant features, the probability of relevance must depend on the set of observed options, S, $P(r|S)$. This results in a modification of Eqn (3):

$$u_S(x) = \sum_r \sum_{f \in x,r} u_f \omega_f P(r|S), \tag{5}$$

where the relevant features are inferred by the chooser based on the options presented by the questioner. Given uncertainty about the actual intended relevance hypothesis, we sum over all possible hypotheses, r.

As in previous work on social learning, the chooser infers $P(r|S)$ by assuming that the questioner has chosen the sample to help the chooser infer what features are relevant:

$$P(r|S) = \frac{P(S|r)P(r)}{\sum_{r'} P(S|r')P(r')}. \tag{6}$$

The probability of relevance, $P(r)$, reflects the chooser's a priori beliefs about the features the questioner is likely to find most relevant. The denominator of Eqn (6) is simply a normalizing constant, ensuring a proper probability distribution over the relevance hypotheses.

Intentional selection proposes that the options have been selected purposefully based on the relevant features. As noted, we use relevance, r, to capture a hypothesis that contains a particular limited set of features. For example, the relevance hypothesis could be "something about color is important" which might include all possible features of color (e.g., red, blue, and transparent). The probability of a particular set being chosen, given a particular hypotheses about relevance, r, is given by $P(S|r)$.

How might we evaluate the probability of observing a set of options given a particular hypothesis about relevance, $P(S|r)$? Intuitively, we might believe that in order to discern relevance, the ideal pair of options would *contrast* in utilities among those relevant features. For instance, if the feature *red* was relevant, to highlight this fact, a questioner would prefer to select an option that contrasts along this feature, leading to a set S containing options that are *red* and options that are *not red*. Similarly, in the case of a dimensional feature, such as price, a questioner should choose options that contrast (e.g., one high and one low) to emphasize the variability. Thus, a chooser should expect pairs of options to provide distributional information along the relevant features or dimensions in question.

Consider the case where the questioner selects *three* options with features that vary along an underlying dimension. As discussed above, a pair of contrastive options may be chosen to highlight distributional information—variability—along the feature or dimension of relevance. What additional information may be conveyed in a set of three options? While the two extreme options may indicate variability along a dimension, the third option may indicate the middle of the distribution. Consider cameras that vary along the dimensions of price and quality. One option may have features such as "low price" and "low quality." Another option may have features such as "high price" and "high quality." Because the dimensions along which these features contrast clearly trade-off in utility, the addition of a third option with the features "middle price" and "middle quality" can be expected to highlight the trade-off between these dimensions—the relevant features are the pair where the utilities of "price" and "quality" balance. In other words, the middle option can be expected to have features that are *representative* of the distributions over price and quality (Tenenbaum & Griffiths, 2001b).

Our model shows how choosers might make inferences about the set of options chosen for them. They can then use this information to inform their guesses about which features should be most relevant. This provides a context-dependent weight over features, affecting the overall utilities of options.

5.4 Choice among Options

Returning to Eqn (5), we can now see the contributions of each of the three components of the set of features for an option. Each option is considered with respect to all its features. Each feature involves the chooser's (stable) utility for that feature, u_f, weighed by our commonality term, ω_f, and the probability of relevance of that feature, given the set of options provided, $P(r|S)$.

Thus, we are left with a utility for a particular option, x, given the set of options provided, $u_S(x)$. The simplest model of choice follows from Luce, in which we choose an option x based on its utility relative the other choices in the set of options S. This gives us

$$P_S(x) = \frac{u_S(x)}{\sum_{x' \in S} u_S(x')}, \tag{7}$$

wherein choosers respond by selecting an option proportional to the utility of the other options.

Our model reflects stable utilities in the utility of each feature. It recovers normative notions of utility when all features are equally relevant. However, it also captures context dependence through the commonality and relevance terms. These terms depend on both the set of other options provided and the chooser's assumptions about the questioner's goals in providing the observed set of options. Thus, the model not only captures a systematicity in preference for certain options, but also allows for contextual factors (the set of other options provided) to influence the final choice. In this way, our model provides a middle ground between normative and subjective theories of choice.

6. EXAMPLES, REVISITED

Our model of feature relevance given intentional selection is a single account that can provide explanations for the Compromise, Similarity, and Attraction Effects. Rather than providing detailed derivations, we sketch how in principle each effect could be explained and how our approach relates to theoretical accounts by Tversky and colleagues (Simonson & Tversky, 1992; Tversky, 1972).

6.1 Compromise Effect

The Compromise Effect occurs when two items, such as cameras, vary on two attributes, such as price and quality. The addition of an option at either extreme, that is, higher price and quality or lower price and quality, often leads to choice of the middle option, regardless of the preference in a binary choice. Simonson and Tversky (1992) introduced Componential Context Theory to explain this effect. The key element of this theory is a distinction between the background context and the local context defined by the choice set, and shifts in choice are explained through the incorporation of loss aversion in the calculation of local and global utility.

In this case, the options demonstrate systematic variation across two features, where one option lies between the other two. How might the options have been intentionally selected? Recent research speaks directly to this question. Shafto et al. (2014; see also Tenenbaum & Griffiths, 2001b) investigated people's inferences from selections of three examples along a single dimension, where the examples were selected intentionally or randomly. Evidence suggests that people expected the examples to not be a random selection from the distribution, but that they represent the distribution. That is, the collection of examples should represent properties of the distribution: the variation and the middle.

Indeed, the relationship between the theoretical account offered by Shafto et al. (2014) and our model is implicit in the relevance computation. Relevance is guided by a learner's inference about the most likely set of examples, given a particular hypothesis about the relevant features. We can think of the different prices and qualities of the cameras as individual features of potential relevance. Given only two cameras, which one is most relevant remains unclear. However, the addition of a third camera highlights the middling features as the most relevant. The probability of the "middle" camera being the relevant price is higher, given two comparison price points (one above and one below), because it provides evidence for being a "representative" feature. Thus, the middle example can be expected to represent the optimum trade-off between the two features.

6.2 Similarity Effect

Recall the Debussy—Beethoven example in which a participant is presented with the choice between three recordings: Debussy or one of two Beethoven pieces that are essentially the same (same orchestra, song, different unfamiliar directors). The original theoretical account of this phenomena proposed by Tversky (1972) is Elimination by Aspects. Like our model, Elimination by Aspects proposed that alternatives are formalized as collections of features (or aspects), and each feature has an associated utility. However, Elimination by Aspects goes on to select an aspect proportional to its utility and eliminate options that do not have that aspect via an attention switching mechanism. The option that remains after eliminating all others is chosen. Because the two Beethoven recordings are nearly identical, the probability of selecting an aspect that is unique to one is small relative to the many unique aspects of the Debussy recording.

Our approach relies on commonality to explain the implications of extreme similarity between two of three options. In the Debussy—Beethoven example, the questioner has chosen options that vary along two main (meta) features: type of classical music and director. The commonality term ensures that the relative utilities of the composers are unaffected by the addition of a second Beethoven recording. Because our model posits *features* as the object of choice, the account avoids the problems that Luce choice encounters. Specifically, we propose that options are compositions of features, and as such a choice of two different options with an identical feature are weighed according to the commonality term. Thus, if given the choice between three options which vary on only two relevant features, our model predicts that the more unique option should receive approximately equal weight to the combined "identical" options, leading to approximately 50% weight for Debussy and 50% weight for the Beethovens (with each receiving 25% weight).

6.3 Attraction Effect

The Attraction Effect occurs when the introduction of a seemingly irrelevant option changes choice. For example, although people would choose

$6 over a nice Cross pen, the addition of a second, much less nice pen (that no one would choose) leads people to begin to choose the Cross pen over the $6 (Simonson & Tversky, 1992). Like similarity, one possibility is that Elimination by Aspects can explain Attraction Effects. The addition of the less nice pen has the effect of increasing attention to the quality of the pen, resulting in a shift in choice behavior.

Under our account, however, feature relevance is used to explain the intentional selection of examples. In the pairwise comparison of money and a pen, money is the clear winner because money is fungible and thus has greater utility (modulo the difficulty of obtaining a comparable pen). Adding in the less nice pen changes the intentional inference. Two very different pens highlights the relevance of the quality feature. In this context, the probability of selecting the nice pen should increase because this is the better selection along the most relevant dimension.

6.4 New Empirical Predictions

Our model also makes new empirical predictions. Intentional selection assumption should guide peoples' judgments about which features are likely present in a set of options. For example, imagine that a friend texts you that she is offering you the choice of three objects ("widgets") which vary on shape (triangle or square) and color (red or blue). You learn that Widget 1 is triangular and blue, Widget 2 is square and red, and Widget 3 is triangular…, but the text is cut off so you do not know what color this third triangular option is. Our model predicts that you should assume your friend was offering you a novel (representative) choice, and thus that the last widget is red, thus distinguishing it from the first option. If instead the options were produced by a machine that is randomly generating the object shapes and colors, you would have no such reason to presume the final object is red. This provides a first intuitive account of how intentional selection affects our intuitions about the features of objects presented to us.

A second test of this model could replicate classic findings, such as the Compromise Effect, but include an accidental version where options are presented by mistake. For example, consider Ariely's (2010) example in which participants are given a choice about which newspaper medium to purchase: an online subscription for $59 or a joint online and print subscription for $125. When presented with these two subscription options, people tend to choose the online subscription. However, if a third print-only subscription option is added for $125, people tend to choose the joint online and print subscription for the same price. What if participants were told that the website had accidentally not been updated so the print-only subscription option was not meant to be available (although they could still purchase it now if they wished)? In this "accidental" case, we break the intentional selection assumption and predict that we would not see the behavior suggestive of the Compromise Effect. Indeed, in ongoing work (Durkin et al., 2015) we replicated Ariely's (2010) past finding that people will be more likely to choose the joint print and online subscription option than the online option only if they are also presented with a print subscription of similar price. However, this was only true when all options were assumed to be intentionally selected. When the print-only option was presented accidentally (in our novel modification of the task), participants chose similarly to the control condition that only saw two options. These preliminary findings suggest that the Compromise Effect is actually a result of this intentional selection inference, as the effect evaporates when participants are told that options are not presented intentionally.

The studies described here provide a first qualitative test of the intentional selection assumption. However, future work should also explore quantitative predictions of our model. For example, we are currently investigating whether feature relevance changes as a function of the options presented in intentionally and accidentally sampled conditions. The goal of these quantitative assessments is to provide a more rigorous account of the factors that influence human decision in choice, as well as the assumptions that influence these factors.

7. DISCUSSION AND CONCLUSIONS

We have proposed a new account of choice behavior based on the intentional selection of examples along relevant features. We have argued that there are strong commonalities between choice and learning; in both cases, people must draw inferences based on observations and the process by which the observations were selected can facilitate this inference.

Our proposal is interestingly related to two theoretical accounts by Tversky and colleagues: Elimination by Aspects and Componential Context. Tversky (1972) proposed Elimination by Aspects to explain violations of independence from irrelevant alternatives. On this account, attentional shifts across aspects provide the explanation for choice behavior. Simonson and Tversky (1992) proposed Componential Context to explain Compromise Effects. On this account, differentiation between immediate and background context together with loss aversion provide the explanation for the preference for the middle example. Interestingly, as demonstrated by Rieskamp et al. (2006), Elimination by Aspects fails to account for Compromise Effects and Componential Context fails to explain Similarity Effects. Our approach integrates elements of both Elimination by Aspects and Componential Context in a single framework. Like Elimination by Aspects we focus on features, and specifically relevant features. Like Componential Context, we distinguish between the background context and the local context induced by the questioner's selection of options.

Our approach also differs from previous theory by Tversky and colleagues. The most notable deviation is in how we approach the problem. While Tversky's theories explain choice in terms of variables that refer only to the internal state of the chooser (e.g., attention and loss aversion), we explain choice in a social context. Thus, on our account, local context differs from global context because the questioner intentionally selected the options with a question in mind, and features are relevant or not based on inferences about the questioner's intended question. This shift in focus leads us to a different, and we hope clarifying, explanatory framework. In this framework, people are using the statistics of the world to draw sensible inferences about why other people act the way they do, and how they can best act in order to optimize mutual understanding. This is necessarily an inferential process, and thus results in natural variation and stability in choice behavior.

The domain has many, very influential findings and models, and our approach represents a very preliminary first step. One next step would be to derive more precise mathematical characterizations of when the model predicts phenomena to be observed (see Rieskamp et al., 2006). Clarifying connections to specific models and modeling frameworks is an important goal for future research. For example, we hope to draw connections to the multinomial logit family of models (e.g., McFadden, 1977; Train, 2003; Yellot, 1977). Despite some of these possible connections, one important difference between our approach and approaches from economics is the importance of semantic knowledge to our predictions. In this sense, our approach is marrying the more statistical approaches in economics with more psychological approaches pursued by Tversky and colleagues. We differ from both of these approaches in positing a role for reasoning about the questioner's intent, and thus a more clear articulation of the similarities and differences between our model and previous work will shed light on this unique aspect of our model.

Our approach has brought recent results in learning to bear on choice behavior, highlighting the fact that learning and choice have largely operated independently. In some ways, this is curious. For example, the majority of behavioral measures used to assess learning involve some form of choice. Among the most common is simply providing people with options to choose from. When these options are exhaustive, there is not likely to be any influence of social reasoning on choice. Oftentimes, to assess learning experimenters present a subset of the possible options and often these are chosen with respect to the theoretical questions of interest. While experimentalists treat responses in these contexts as unbiased representations of the learners' beliefs, our analysis suggests that choices may also reflect the learner's inferences about the experimenter. Indeed, this is not a novel proposal (e.g., Gonzalez, Shafto, Bonawitz, & Gopnik, 2012; Topal, Gergely, Miklosi, Erdohegyi, & Csibra, 2008), but our analysis provides a candidate computational account of why and what effects may be expected.

Considerable work remains for this account to be articulated at the level of specificity that is the standard for models of choice behavior. For example, our model makes assumptions about the kinds of possibilities learners consider for themselves and for others. The model assumes a space of candidate features, from which a relevant subset is sampled. Many in the learning literature have highlighted the nonindependence of modeling results and assumed set of features. Recent approaches of cross-categorization provide a formal framework that allows relevant subsets of a potentially infinite set of features (Mansinghka et al., in press; Shafto, Kemp, Mansinghka, & Tenenbaum, 2011). Similar formal tools may be useful here. However, the idea of relevant features is probably too narrow to fully characterize the ways in which we think about relationships among options. We have focused on this idea to highlight similarities and differences between our model and previous accounts, but considering more general approaches may be a useful direction for future research.

Similarly, our model assumes that learners reason about the kinds of things that are a priori relevant to others. This requires modeling others' beliefs and goals and how they relate to the kinds of questions one might ask. Recent research has made progress in formalizing models of intuitive Theory of Mind reasoning (Baker, Saxe, & Tenenbaum, 2011; Goodman et al., 2006), the relationship between intuitive psychological reasoning and learning (Shafto, Goodman, et al., 2012), and the relationship between knowledge, intent, and question asking (Gonzalez et al., 2012). For example, Gonzalez et al. (2012) investigated children's responses to neutral follow-up questions. For example, when children are asked supposedly neutral questions like "Is that your final answer?", children will often change their responses, but specifically when the questioner is assumed to be knowledgeable about the actual answer. These results suggest that children's responses to these questions depend on the epistemic state of the questioner: if the child believes the person knows the answer, children change their responses more than if they believe the person is ignorant of the answer. The account forwarded in that paper proposes that children are reacting to the relative difference between people labeling responses as correct (which is common) and people labeling responses as incorrect (which is less common). A similar account, based on the statistics of experience combined within a framework of intuitive psychological reasoning, may be fruitful here.

Although work remains to specify the model details and precisely test the model on existing empirical results, the model leads to several interesting, and to our knowledge, novel predictions. Most salient is the prediction that the intentional selection of examples ought to affect people's judgment. This suggests that, if one could manipulate social context compellingly, as has been done in the learning literature, the model would predict that the exact same set of options could lead to different choices. A second interesting prediction, which is arguably shared with other accounts based on feature subsets, pertains to existing (not uncontroversial; see Scheibehenne, Greifeneder, and Todd (2010)) empirical results regarding the difficulty of choice as it depends on the number of options (the paradox of choice). Given that increasing options often (but not always) leads to increasingly large sets of potentially relevant features, it may be possible to bring our account to bear on this controversy.

Our approach represents a middle ground between normative and subjective approaches to economics. Normative accounts suggest that the utility of an option is fixed and unchanging, whereas subjective accounts allow flexibility in assessments of utility. Each is challenged by experimental findings demonstrating both variability and stability in choice behavior. Our approach includes stable notions of utility that may vary as with changes in social context. This approach thus captures elements of stability of choice behavior, while positing an explanation for variability.

Given the apparent simplicity and pervasiveness of choice, it is not surprising that many, many researchers have found the topic a fruitful area of research. However, it is in some ways surprising that choice behavior has been so stubbornly resistant to characterization. We have sketched a novel model of choice based on social reasoning by the chooser about the questioner. It is too soon to say whether this approach will fare better than previous accounts, but given that choice is fundamental across so many domains of inquiry, explaining choice remains a fundamental problem in understanding human behavior.

REFERENCES

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review, 98*, 409–429.

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2011). Bayesian theory of mind: modeling joint belief-desire attribution. In *Proceedings of the thirty-third annual conference of the cognitive science society*.

Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends in Cognitive Sciences*, 497–500.

Bonawitz, E. B., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: teaching limits children's spontaneous exploration and discovery. *Cognition, 120*, 322–330.

Debreu, G. (1960). Review of R. D. Luce, individual choice behavior: a theoretical analysis. *American Economic Review, 50*, 186–188.

Durkin, K., Calgar, L. R., Bonawitz, E., & Shafto, P. (2015). *Explaining choice behavior: The intentional selection assumption*. *Proceedings of the 37th annual conference of the Cognitive Science Society*.

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science, 336*, 998.

Fried, L. S., & Holyoak, K. J. (1984). Induction of category distributions: a framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 234–257.

Gonzalez, A., Shafto, P., Bonawitz, E. B., & Gopnik, A. (2012). Is that your final answer? The effects of neutral queries on children's choices. In *Proceedings of the 34th annual conference of the cognitive science society*.

Goodman, N. D., Baker, C. L., Bonawitz, E. B., Mansinghka, V. K., Gopnik, A., Wellman, H., et al. (2006). Intuitive theories of mind: a rational approach to false belief. In *Proceedings of the twenty-eighth annual conference of the cognitive science society*.

Goodman, N. D., Baker, C. L., & Tenenbaum, J. B. (2009). Cause and intent: social reasoning in causal learning. In *Proceedings of the 31st annual meeting of the cognitive science society*.

Kruschke, J. K. (1992). Alcove: an exemplar-based connectionist model of category learning. *Psychological Review, 99*, 22–44.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). Sustain: a network model of category learning. *Psychological Review, 111*, 309–332.

Luce, R. D. (1959). *Individual choice behavior*. New York: John Wiley.

Mansinghka, V., Shafto, P., Jonas, E., Petschulat, C., Gasner, M., & Tenenbaum, J. B. *Crosscat: A fully bayesian nonparametric method for analyzing heterogenous, high dimensional data*, in press.

McFadden, D. L. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in econometrics* (pp. 105–142). New York: Academic Press.

McFadden, D. (1977). Quantal choice analysis: a survey. *Annals of Economic and Social Measurement, 5*, 363–390.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 100*, 254–278.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 104–114.
Pleskac, T. J. (2013). Decision and choice: Luce's choice axiom. In P. Bona (Ed.), *International encyclopedia of social and behavioral sciences* (2nd ed.).
Pothos, E. M., & Chater, N. (2002). A simplicity principle in unsupervised human categorization. *Cognitive Science, 26*, 303–343.
Restle, F. (1961). *Psychology of judgement and choice*. New York: Wiley.
Rieskamp, J., Busemeyer, J. R., & Mellers, B. A. (2006). Extending the bounds of rationality: evidence and theories of preferential choice. *Journal of Economic Literature, XLIV*, 631–661.
Scheibehenne, B., Greifeneder, R., & Todd, P. M. (2010). Can there ever be too many options? A meta-analytic review of choice overload. *Journal of Consumer Research, 37*, 409–425.
Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: modeling children's reasoning about others' knowledge and intent. *Developmental Science, 15*, 436–447.
Shafto, P., & Goodman, N. D. (2008). Teaching games: statistical sampling assumptions for pedagogical situations. In *Proceedings of the 30th annual conference of the cognitive science society*.
Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others: the consequences of psychological reasoning for human learning. *Perspectives on Psychological Science, 7*, 341–351.
Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: teaching by, and learning from, examples. *Cognitive Psychology, 71*, 55–89.
Shafto, P., Kemp, C., Mansinghka, V., & Tenenbaum, J. B. (2011). A probabilistic model of cross-categorization. *Cognition, 120*, 1–25.
Simonson, I., & Tversky, A. (1992). Choice in context: tradeoff contrast and extremeness aversion. *Journal of Marketing Research, 29*, 281–295.
Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
Tenenbaum, J. B., & Griffiths, T. L. (2001a). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences, 24*, 629–641.
Tenenbaum, J. B., & Griffiths, T. L. (2001b). The rational basis of representativeness. In *Proceedings of the 23rd annual conference of the cognitive science society* (pp. 1036–1041). Hillsdale, NJ: Erlbaum.
Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review, 34*(4), 273–286.
Topal, J., Gergely, G., Miklosi, A., Erdohegyi, A., & Csibra, G. (2008). Infants' perseverative search errors are induced by pragmatic misinterpretation. *Science, 321*, 1831–1834.
Train, K. (2003). *Discrete choice models with simulation*. Cambridge: Cambridge University Press.
Tversky, A. (1972). Elimination by aspects: a theory of choice. *Psychological Review, 79*, 281–299.
Vul, E., & Pashler, H. (2008). Measuring the crowd within probabilistic representations within individuals. *Psychological Science*, 645–647.
Xu, F., & Tenenbaum, J. (2007a). Sensitivity to sampling in Bayesian word learning. *Developmental Science, 10*, 288–297.
Xu, F., & Tenenbaum, J. B. (2007b). Word learning as Bayesian inference. *Psychological Review, 114*, 245–272.
Yellot, J. I. (1977). The relationship between luce's choice axiom, thurstone's theory of comparative judgement, and the double exponential distribution. *Journal of Mathematical Psychology, 15*, 109–144.