

A Rational Model for Individual Differences in Preference Choice

Sheeraz Ahmad (sahmad@cs.ucsd.edu)

Department of Computer Science and Engineering
University of California, San Diego

Angela J. Yu (ajyu@ucsd.edu)

Department of Cognitive Science
University of California, San Diego

Abstract

Human preference choice suffers curious contextual effects: the relative preference between two multi-attribute options (e.g. cars with differing safety and economy ratings) can dramatically shift, depending on the presence/absence of additional options. This phenomenon defies any simple utility-based account of choice, and has been taken to imply irrationalities/sub-optimality in human decision-making, or reflect idiosyncrasies in neural processing. Recently, we used a Bayesian model to show that these contextual effects are normative consequences of observers using options to learn about the “market”. However, it had an unsavory implication that all decision-makers asymptotically converge to the same beliefs/behavior. Here, we propose a new model that uses both market and personal utilities to make choices. This model still captures the contextual effects, while also allowing asymptotic differences in individual preferences and providing a general framework for explaining how consumption informs one’s beliefs and preferences.

Keywords: decision making; preference choice; multi-attribute; contextual effects; individual differences; Bayesian learning

Introduction

Humans are regularly faced with decisions involving a choice between options with multiple attributes. For example, one may have to choose between a car that has a higher safety rating but lower mileage and another that has a lower safety rating but higher mileage; or one may have to choose between a PhD applicant who has better grades but worse letters and another that has worse grades but better letters. There may not be a universal or obvious way to make these decisions, and indeed humans often exhibit inconsistencies in their choices. One particular class of peculiarities in human preference choice has garnered special attention in psychology research, namely a type of contextual effects whereby an individual’s relative preference between two options can be altered, or even reversed, when a third option (a ‘decoy’) is introduced into the choice-set (Kahneman & Tversky, 1979; Kahneman et al., 1982; Tversky & Simonson, 1993).

The discovery of these contextual effects has caused difficulty for traditional, utility-based accounts of preference-based decision-making in humans. If each option has an associated scalar utility for an individual, and the probability of choosing option A monotonically depends on the utility of option A in comparison to the utility of option B, then the relative preference between A and B should not change when a third option, C, is added or removed (Luce, 1959; Thurstone, 1954; Tversky, 1972). This contradiction has contributed to the school of thought that human decision-making

is irrational and sub-optimal (Tversky, 1972; Kahneman et al., 1982). More recently, it has been proposed that these peculiarities may arise from specific idiosyncrasies in neural architecture or dynamics in brain areas that support preference choice (Busemeyer & Townsend, 1993; Roe et al., 2001; Usher & McClelland, 2004; Trueblood, 2012).

In contrast to a purely utility-based account of preference choice, we recently proposed an alternative normative model of decision-making (Shenoy & Yu, 2013), which assumed that humans do not have fixed, perfectly known utility values assigned to options, but instead may suffer uncertainties about how to assign utilities both within an attribution dimension and also *jointly* for a combination of attribute values in a multi-attribute scenario. Consequently, observers use available options not only to make choices, as assumed by previous utility-based models of preference choice, but also to *learn* about the range and distribution of attribute values generally available in the “market”, as well as some market-based sense for how the attributes ought to be valued against each other. By this account, the addition of a third option confers extra information about the “market” and may therefore influence the relative preference between the two original options, with the effect expected to be particularly strong when the decision-maker has relatively little experience with the particular “market.” For example, in the PhD applicant example, suppose a professor is evaluating applications of students from a foreign country (whose academic structure is not well known to him), and he must choose between an applicant A who has a test score of 290 and a grade of 90, and an applicant B with a test score of 300 and a grade of 80. Based on this data alone, he may not have a strong preference between the applicants, but if a third applicant C comes in with a test score of 290 and a grade of 130, the professor may strongly shift his *relative* preference between A and B toward B (though he may prefer C over both A and B), since C’s grade of 130 shows that a grade of 90 over 80 is not really much of an advantage at all. Indeed, we showed that this inferential model was able to account for all three classical contextual effects: attraction, compromise, and similarity (Shenoy & Yu, 2013).

This earlier model provided a normative explanation for why contextual effects arise in rational decision-makers. It also provided a means for modeling individual differences, by allowing different individuals to have different prior beliefs about the distribution of attribute values in different dimensions and how a combination of attribute values from multiple

dimensions jointly determine the utility function. However, it makes the odd prediction that all decision-makers will eventually converge to the same beliefs about the market, given sufficient exposure to the various options available in the market; in Bayesian parlance, this is because the iteratively updated posterior distribution will converge to the same (delta) function regardless of the prior distribution. This prediction flies in the face of empirical data (Malaviya & Sivakumar, 1998; Müller et al., 2012) and casual observations that different individuals often exhibit systematic and long-lasting differences in their preference choices.

In this work, we propose an alternate Bayesian account of preference choice, which captures the notion that each person entertains beliefs about both the market-based utility function (learned from exposure to available options, or “window shopping”) and a personal utility function (learned from choosing/consuming specific options), and combines the two in making preference choices. The introduction of this personal utility function allows individuals to have persistent diversity in their preference choices. In the following, we first describe the three classical contextual effects, then the new Bayesian model, followed by simulation results that compare our model against classical contextual effects and more subtle individual differences, as well as against the previous model (Shenoy & Yu, 2013) in the context of asymptotic learning; we conclude with a discussion of the broader implications of this work, relationship to related work, and potential directions for future research.

Three Contextual Effects

Three classical contextual effects have been observed when starting from two originally equally preferable options (say X and Y), each with two attributes: (1) attraction effect (Fig. 1A) (Huber & Payne, 1982; Heath & Chatterjee, 1995), where introduction of an option Z, that is close to Y and is dominated by it on both attributes, leads to an increased preference of Y over X. (2) compromise effect (Fig. 1B) (Simonson, 1989), where introduction of an extreme option Z, that is better than Y on one attribute and inferior on the other, as well as is farther from X, again leads to an increased preference of Y over X. (3) similarity effect (Fig. 1C) (Tversky, 1972), where the introduction of an option Z, that is almost comparable to Y on both attributes, leads to an increased preference of X over Y.




Figure 1: Three classical contextual effects. Options X and Y are equi-preferable in the absence of a third option. Arrow denotes the direction of the preference shift when a decoy, Z, is added. (A) Attraction effect. (B) Compromise effect. (C) Similarity effect.

Consumption Model

We consider an option space with two attributes, which are combined using a Cobb-Douglas utility function (Douglas, 1976), parametrized by γ , which specifies the relative importance of the two attributes to the joint utility function. The utility or value of an option (x, y) is $v = x^\gamma y^{1-\gamma}$. Note that the two attributes nonlinearly combine to determine the utility function: this captures diminishing marginal utility, or the idea that differential change in an already abundant attribute contributes less to the overall utility than a similar change in a scarce attribute (Hicks, 1932).

In our model, there are two parameters that contribute to the attribute trade-off, in turn affecting the multi-attribute utility function. The market utility, $v_m = x^\gamma y^{1-\gamma}$, is parametrized by market tradeoff parameter γ , and the personal utility, $v_p = x^\lambda y^{1-\lambda}$, is parametrized by the personal tradeoff parameter λ . We propose the following way to combine these towards a net utility: $v = v_m^w v_p^{1-w}$, since it leads to a simplified form: $v = x^\zeta y^{1-\zeta}$ ($\zeta \triangleq w\gamma + (1-w)\lambda$), and intuitively provides an “average” of the two utilities (in this case, a generalized geometric mean). The parameter w can be interpreted as a personality trait dictating how much an individual values uniqueness as opposed to conforming to the market. For example, consider a consumer buying a car; even though she may be more price-conscious herself, she might still want to buy a trendy, more expensive car because of external considerations like status symbol, peer pressure or resale value ($w > 0.5$). A more rebellious consumer may value uniqueness more and make a decision primarily based on her own preference, giving little consideration to market preferences ($w < 0.5$).

The generative model (Fig. 2A) also assumes that depending on whether the option is consumed or not ($c_i \in \{0, 1\}$), the individual gets different levels of satisfaction. If the option is not consumed ($c_i = 0$), a small level of stochastic satisfaction is received, which can be interpreted to be resulting from mental simulation (thinking about the consumption of the option) but which yields no information about the hidden preference λ . On the other hand, if the option is consumed, the satisfaction derived is a noisy version of the personal utility, $v_p = x^\lambda y^{1-\lambda}$ for an option (x, y) , which in turn is informative about λ . For example, after accepting a few of the PhD applicants, the professor may decide, after all, that a student’s grades are a better predictive of success for conducting research in her lab than recommendation letters; while another professor may decide just the opposite.

For inference, when the options are not consumed ($c_i = 0 \forall i$), the node s_i does not depend on o_i or λ , thus the directed edges from these to s_i are effectively removed and the model reduces to a simplified version of our previous model (Shenoy & Yu, 2013) (for $w = 1$). More generally, the posterior belief about λ can be updated based on the observed options ($\mathbf{o} = \{o_i\}$) and satisfaction ($\mathbf{o} = \{o_i\}$) as:




Figure 2: (A) Graphical model showing how options (o_i) are generated in the market, based on the hidden value (v_i), market preference (γ), individual preference (λ), weight (w) and the parameter θ_i . Also describes the *satisfaction* as a noisy representation of the personal value, depending on the individual preference (λ) and whether the option is consumed (c_i). (B) For a particular setting of the attribute trade-off parameter (ζ), all options lying on the same curve have the same value/utility as denoted on the curve; different curves signify different values. (C) For a fixed value (v_i), all options lying on the same curve have the same attribute trade-off (ζ), as denoted on the curve; different curves signify different trade-offs.

$$\begin{aligned}
P(\lambda|\mathbf{o}, \mathbf{s}) &\propto P(\mathbf{s}|\mathbf{o}, \lambda)P(\mathbf{o}|\lambda)P(\lambda) \\
&= \left\{ \prod_i P(s_i|o_i, \lambda; c_i) \right\} \left\{ \int_{\gamma, \theta, \mathbf{v}} P(\gamma)P(\theta)P(\mathbf{v})P(\mathbf{o}|\mathbf{v}, \gamma, \theta, \lambda; w) \right. \\
&\quad \left. d\gamma d\theta d\mathbf{v} \right\} P(\lambda) \\
&= \left\{ \prod_i P(s_i|o_i, \lambda; c_i) \right\} \\
&\quad \left\{ \int_{\gamma} P(\gamma) \prod_i \left[\int_{\theta_i, v_i} P(\theta_i)P(v_i)P(o_i|v_i, \gamma, \theta_i, \lambda; w) d\theta_i dv_i \right] d\gamma \right\} P(\lambda) \tag{1}
\end{aligned}$$

Similarly, the joint posterior on the hidden values can be updated as:

$$\begin{aligned}
P(\mathbf{v}|\mathbf{o}, \mathbf{s}) &\propto \left\{ \int_{\gamma, \lambda} P(\gamma)P(\lambda) \prod_i \left[\int_{\theta_i} P(\theta_i)P(o_i|v_i, \gamma, \theta_i, \lambda; w) d\theta_i \right] \right. \\
&\quad \left. [P(s_i|o_i, \lambda; c_i)] d\gamma d\lambda \right\} \left\{ \prod_i P(v_i) \right\} \tag{2}
\end{aligned}$$

As in our previous model, we assume that the decision policy involves sampling from the posterior $P(\mathbf{v}|\mathbf{o})$; the sample $\hat{\mathbf{v}}$ is then used to choose an option: $\pi(\hat{\mathbf{v}}) = \text{argmax}_j \hat{v}_j$. Thus, stochasticity in choice upon presentation of the same set of

options is expected because of the residual uncertainty in the posterior distributions of the option values (Debreu, 1958; Blavatsky, 2008).

Results

In this section, we apply the proposed models to different multi-attribute choice tasks, in which sometimes subjects have to choose among just two options, and sometimes among three options whereby a “decoy” is added to the two original options (see Fig. 1). We first show how our model accounts for the three classical contextual effects; we then use the model to capture several observed individual differences in existing experimental literature.

For all simulations, the market parameters for the preference γ are $a = 2$ and $b = 2$, and for the utility/value v_i , $k = 20$ and $\mu = 50$. Other parameters used are, $\sigma_\theta = 20$ and observation noise $\sigma_0 = 2$. Lastly, the options are $X = (40, 60)$, $Y = (60, 40)$, $Z = (50, 30)$ for attraction, $Z = (80, 20)$ for compromise, and $Z = (65, 35)$ for similarity effect. Since no simple closed form expressions exist for the different posteriors (e.g. Eq. 1) and approximations based on discretization of continuous variables are inexact and inefficient, we use a program called JAGS (Plummer, 2003) that uses Gibbs sampling (Geman & Geman, 1984) to generate samples from the posterior distribution of the desired model parameters (\mathbf{v} , γ and λ in our case).

For the simulations in Fig. 3, we assume that there is no consumption ($c_i = 0$), and that the individual relies solely on the information from market options ($w = 1$). With these settings, our model reproduces all three contextual effects,




Figure 3: Proportion of choices for different effects and their explanation. (A) X and Y are equi-preferable. (B) Adding an option Z inferior to Y increases the preference for Y relative to X. (C) Adding an option Z even more extreme than Y, in relation to X, increases the preference for Y. (D) Adding an option Z very similar to Y, but not clearly more or less preferable to Y, decreases the preference for Y relative to X. (E) X and Y lie on the equi-preference curve which represents a fixed value; any option above this curve would be considered more valuable, and below would be considered less valuable. (F) The attraction decoy makes Y appear better on average. (G) The compromise decoy makes Y appear better on average. (H) The valuation of Y and Z are highly correlated, such that they tend to be considered both better than X or both worse than X, for different settings of the hidden variable γ . When they are both worse than X, X is chosen; when they are both better than X, Y is chosen half the time (Z the other half). Thus, on average, X is chosen more often than Y, even though they are on average considered about equally valuable.

so that even though the two options are equi-preferable by themselves (Fig. 3A), the preference shifts towards Y when an attraction or compromise decoy is added to the choice set (Fig. 3B and C), and towards X when a similarity decoy is added (Fig. 3D). Next, in order to understand how the model works for all the contextual effects, we show the resulting equi-preference curves, the locus of options that all have the same value for a given preference γ (in other words, all the points on the curve $x^\gamma y^{1-\gamma} = v$). Clearly, there are infinitely many equi-preference curves, one corresponding to each value of v , but we only plot the ones, for visual simplicity, that always pass through the option X. When there are only two options, such an equi-preference curve with γ set to its posterior mean also passes through Y, making the two options equally attractive (Fig. 3E). For the attraction effect, option Y lies above such an equi-preference curve, making it appear relatively more lucrative (Fig. 3F); however option X is still selected owing to the stochasticity in the option values. Similar explanation holds for the compromise effect (Fig. 3G). For the similarity effect, the equi-preference curve with γ set to its posterior mean is not particularly informative (Fig. 3H, blue curve), since this effect arises due to the close correlation between the valuation of the option Y and Z in the model. As can be seen, the two options are likely to appear better or worse than X together (green and red curves respectively). Therefore, when they appear better, the choice gets split between them; when they appear worse, X gets chosen,

leading to overall higher frequency of choosing X than Y.

We also investigate scenarios where individual exhibit different behavior based on their previous experience. Experiments have shown that contextual effects are not always robust, with individual differences emerging when subjects value attributes differently (Malaviya & Sivakumar, 1998; Müller et al., 2012). To show that our model can capture such deviations, we simulate the scenario where an individual prefers attribute y more over attribute x ($\lambda = 0.35$). Furthermore, we assume that the individual relies equally on self and market preference ($w = 0.5$). With these settings, we observe that compromise effect becomes insignificant (Fig. 4A), which is what has been observed for consumers who are less quality conscious (attribute x) and more price conscious (attribute y) when making brand choices (Müller et al., 2012).

Lastly, we show how consuming the options can help an individual discover self-preference, and how the process can lead to divergence in the choice behavior of two individuals who learn their preference (from consumption) along with the market preference (from options). In the model proposed earlier (Shenoy & Yu, 2013), the only way the individuals could differ is if they have different priors over the market preference parameter (γ). However, with increase experience, the individual beliefs would converge, consequently leading to the same choice for all individuals who are experienced. In the simulations, we consider two individuals starting with different priors on γ , $Beta(2,3)$ and $Beta(3,2)$,




Figure 4: (A) Individual difference in contextual effects. Compromise effect becomes insignificant in the consumption model for individuals that value attribute 1 less ($\lambda = 0.35$). (B) & (C) Evolution of belief about market preference with experience for two individuals starting with different priors, using the previous simple model and our consumption model respectively. The translucent bounded region shows the standard deviation around the mean belief.

respectively. When both of these individuals experience the same set of options, we note that their beliefs eventually converge (Fig. 4B). This contradicts the every day observation, where individuals with very similar experience with a product category, can still converge on buying different brands. Next, we consider the same two individuals and simulate their beliefs under our consumption model. We additionally note that individual 1 prefers attribute 2 more ($\lambda = 0.3$), whereas individual 2 prefers attribute 1 more ($\lambda = 0.7$), both individuals giving equal weight to market-preference and self-preference ($w = 0.5$). The extra parameters are $k_s = 100$ and $\mu_s = 5$. Now the individual beliefs about joint preference diverge, in accordance with their hidden preference (Fig. 4C). Thus, the model provides an explanation for why two individuals with the same experience may end up preferring different options, and more generally a framework for explaining different framing effects, as well as variations thereof, based on the individual’s experience, personality, and other causes of individualized preferences.

Discussion

In this paper, we proposed a Bayesian model that takes both individual preference as well as effects of consuming options into account. This model reduces to a simplified version of our previous proposed model (Shenoy & Yu, 2013) in the absence of consuming chosen options (and experiencing corresponding satisfaction levels), and can explain not only all three contextual effects but also observed individual differences arising from experience and personal preference. Furthermore, the model relaxes an assumption we made in our prior work that the observer needs to infer both the option values as well as the “fair market value”. Instead, it assumes that the observer only need to infer the relative utility (value) of the available options, and not what constitute “fair” in the market place in absolute terms. Therefore, our current framework simplifies the previous model, as well as provides a general framework for explaining how consumption combines with “window-shopping” to inform one’s beliefs and

preferences, thus leading to diversification of individual preferences.

However, there are some experimental findings that still prove challenging for the new model. For example, the attraction effect has sometimes been observed to diminish for consumers with a low level of experience (Malaviya & Sivakumar, 1998), contradicting a straight-forward prediction of our model. Another curious phenomenon is the *phantom decoy effect* (Pratkanis & Farquhar, 1992; Pettibone & Wedell, 2007), which is very similar to the attraction effect, except the decoy Z is slightly better than one of the options, say Y (rather than worse, as in the attraction effect) and that the subject is informed that this decoy is not actually available as a choice; in that case, human observers reliably shift their preference toward Y instead of away from it, as our model would currently predict. These more nuanced cases require further attention and provide fruitful avenues for future research.

In the future, we also wish to investigate whether an individual has a fixed relative weight for personal and market preferences, or whether this trade-off can change either with experience or context. One possibility is, as the individual gains more experience, the trade-off starts to favor self-preference, thus requiring a more sophisticated model where the trade-off parameter (w) is dependent on experience (perhaps tied to internal uncertainty/confidence). Another direction is to extend the model to allow for “vicarious satisfaction” so that the high demand of an option, with say attribute 1 as the larger attribute, would signal that attribute to be more preferable. Such a model could provide a computational explanation to the *phantom decoy effect*. Lastly, humans may actively seek which options to consume, in order to figure out their self-preference, e.g. trying a Vietnamese restaurant after having tried a Thai restaurant to see if a slightly different spice level would be more satisfying, and our model can be extended to incorporate this active decision making component. This can potentially be achieved by choosing options based on a more sophisticated criteria that takes into account not only the immediate values or satisfaction (as is

done in the current formulation), but also the long term value and informational goals as well as the cost of consuming an option. Insights from the field of active learning (see (Settles, 2009) for a survey) can provide the foundation for such a pursuit. In summary, our work provides a novel computational framework to account for individual differences in a variety of observed preference choice behavior, and opens up venues for future investigations into more sophisticated models of preference-based decision making.

Acknowledgments

We thank Laurence Aitchison and Peter Dayan for stimulating discussions on ideas related to the current paper. This material is based upon work supported in part by the U. S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number W911NF1110391.

References

- Blavatsky, P. R. (2008). Stochastic utility theorem. *Journal of Mathematical Economics*, 44(11), 1049–1056.
- Busemeyer, J., & Townsend, J. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, 100(3), 432–459.
- Debreu, G. (1958). Stochastic choice and cardinal utility. *Econometrica*, 26, 440–444.
- Douglas, P. (1976). The Cobb-Douglas production function once again: its history, its testing, and some new empirical values. *The Journal of Political Economy*.
- Geman, S., & Geman, D. (1984). Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*(6), 721–741.
- Heath, T., & Chatterjee, S. (1995). Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research*, 22, 268–284.
- Hicks, J. R. (1932). *The theory of wages*.
- Huber, J., & Payne, J. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgement under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica*, 47, 313–27.
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.
- Malaviya, P., & Sivakumar, K. (1998). The moderating effect of product category knowledge and attribute importance on the attraction effect. *Marketing Letters*, 9(1), 93–106.
- Müller, H., Vogt, B., & Kroll, E. B. (2012). To be or not to be price conscious: a segment-based analysis of compromise effects in market-like framings. *Psychology & Marketing*, 29(2), 107–116.
- Pettibone, J. C., & Wedell, D. H. (2007). Testing alternative explanations of phantom decoy effects. *Journal of Behavioral Decision Making*, 20(3), 323–341.
- Plummer, M. (2003). *Jags: A program for analysis of bayesian graphical models using gibbs sampling*.
- Pratkanis, A. R., & Farquhar, P. H. (1992). A brief history of research on phantom alternatives: Evidence for seven empirical generalizations about phantoms. *Basic and Applied Social Psychology*, 13(1), 103–122.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological review*, 108(2), 370–392.
- Settles, B. (2009). *Active learning literature survey* (Computer Science Technical Report No. 1648). University of Wisconsin–Madison.
- Shenoy, P., & Yu, A. (2013). Rational preference shifts in multi-attribute choice: What is fair? In *Proceedings of the 35th annual conference of the cognitive science society*.
- Simonson, I. (1989). Choice based on reasons: the case of attraction and compromise effects. *Journal of Consumer Research*, 16(158-74).
- Thurstone, L. (1954). The measurement of values. *Psychological Review*.
- Trueblood, J. S. (2012, October). Multialternative context effects obtained using an inference task. *Psychonomic bulletin & review*, 19(5), 962–8. doi: 10.3758/s13423-012-0288-9
- Tversky, A. (1972). Elimination by aspects: a theory of choice. *Psychological review*, 79(4).
- Tversky, A., & Simonson, I. (1993). Context dependent preferences. *Management Science*, 39, 1179–89.
- Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological review*, 111(3), 757.