A comparison of quantum and multinomial processing tree models of the interference effect

Christopher R. Fisher (christopher.fisher.27.ctr@us.af.mil)
Cubic Defense
Beavercreek, OH 45324 USA

Lorraine Borghetti (lorraine.borghetti.1.ctr@us.af.mil)
Air Force Research Laboratory
Wright Patterson AFB, OH USA

Joseph W. Houpt (joseph.houpt@utsa.edu)
University of Texas at San Antonio
San Antonio, TX 78249

Leslie M. Blaha (leslie.blaha@us.af.mil)
Air Force Research Laboratory
Wright Patterson AFB, OH USA

Christopher Stevens (christopher.stevens.28@us.af.mil)
Air Force Research Laboratory
Wright Patterson AFB, OH USA

Abstract

We compare the qualitative predictions of an existing quantum model and a novel multinomial processing tree (MPT) model of the interference effect using parameter space partitioning (PSP). An interference effect occurs when categorizing a stimulus changes the marginal probability of a subsequent decision, leading to a violation of the LOTP. The PSP analysis revealed that our MPT model can produce the same qualitative patterns as the quantum model. Further analysis, however, revealed that the models differ in several important ways. First, a larger volume of the MPT model's parameter space produces a smaller number of interference effects compared to the quantum model. Second, the distribution of volume across patterns is more diffuse for the MPT model, indicating it is more flexible than the quantum model. We discuss limitations and future directions.

Keywords: Multinomial processing trees; Quantum cognition; Interference effects; Categorization; Model flexibility

Introduction

An interference effect occurs when an action or judgment changes the marginal probability of a subsequent decision (Wang & Busemeyer, 2016; Busemeyer et al., 2011). One reason interference effects are interesting from a theoretical perspective is that they violate a law of classical probability theory (CPT) called the *law of total probability* (LOTP). Adherence to the LOTP means that for decision D and set of categories $\{C_i\}_{i \in I}$, the marginal distribution of D is given by $\Pr(D) = \sum_{i=1}^{n} \Pr(D \mid C_i) \Pr(C_i)$. Previous research has demonstrated that categorizing face interferes with the subsequent decision to attack, such that $\Pr(D) \neq \sum_{i=1}^{n} \Pr(D \mid C_i) \Pr(C_i)$ (Wang & Busemeyer, 2016; Busemeyer et al., 2011).

Interference effects present a challenge for many models that are based on CPT because they violate the LOTP. For example, two models based on CPT—a Markov model and a signal detection model—are unable to account for the entire pattern of interference effects that have been observed empirically (Wang & Busemeyer, 2016). By contrast, a quantum model called the belief-action entanglement (BAE) model provides an account of the interference effect (Wang & Busemeyer, 2016). The reason that the BAE model is successful in accounting for interference effects is that the less restrictive axioms of quantum probability theory allow for the violation of the total law of probability.

Our primary goal is to demonstrate as a proof of concept that a model based on CPT can produce interference effects. Specifically, we show that a multinomial processing tree (MPT; Riefer & Batchelder, 1988) composed of a categorization process, a category revision process, and a decision process is sufficient to produce interference effects. Our second goal is to compare the qualitative patterns of interference effects that the new model and the BAE model can produce. Understanding the prediction space of a model is important for understanding its behavior, assessing flexibility, and identifying diverging predictions between different models. An overly flexible model provides a less persuasive account of the data than a less flexible model (Roberts & Pashler, 2000).

The remainder of this article is organized as follows. In the next section, we describe the categorization-decision paradigm used to study interference effects. Next, we provide a brief overview of the BAE quantum model of interference effects. We then introduce a new MPT model which can also produce the empirical pattern of interference effects. We compare the qualitative patterns of interference effects each model can produce using a method called parameter space partitioning (Pitt et al., 2006). We conclude with a discussion of the limitations of the proposed model and the need for a unified account of interference effects, order effects and other phenomena based on CPT.

Categorization-Decision Paradigm

One popular paradigm for studying interference effects is the categorization-decision sequential choice paradigm (Wang & Busemeyer, 2016). Prior research with this paradigm has demonstrated that inclusion of an explicit categorization stage interferes with subsequent decision making (Wang & Busemeyer, 2016). On each trial, subjects are presented with a face and must decide whether to attack or withdraw. Each face is either a good guy, who is likely to be friendly, or a bad guy who is likely to be hostile. Although subjects do not know the category associated with each face (good vs. bad), they can use facial features, such as width, as cues to aid in the decision process. For simplicity, we define type-b and type-g faces as faces most likely to be in the bad or good category, respectively. The extended paradigm involves three conditions (Wang & Busemeyer, 2016). In the decision-only condition (*d*), subjects make a single decision: to attack or withdraw from each face. In the categorize and decide condition (*cd*), subjects categorize each face as good or bad before proceeding to the attack/withdraw decision. In the third condition (*xd*), subjects are given the true category of each face prior to making a decision.

According to many models based on CPT, the marginal probability of attacking (irrespective of category membership) should be equal in each condition as required by the LOTP, which states:

$$\begin{aligned} \Pr_{\mathrm{d}}(A = a \mid F = x) = \\ \Pr_{\mathrm{cd}}(A = a \mid F = x, C = g)\Pr_{\mathrm{cd}}(C = g \mid F = x) + \\ \Pr_{\mathrm{cd}}(A = a \mid F = x, C = b)\Pr_{\mathrm{cd}}(C = b \mid F = x) \end{aligned} \tag{1}$$

where random variables A, F and C represent the action, facial feature, and category, respectively. Possible actions are a for attack and w for withdrawal; possible values for facial feature are tb for type-b and tg for type-g, and possible categories are b for bad and g for good. Each probability statement is subscripted by its condition; for example, *cd* is the categorize and decide condition. The left-hand-side represents the case in which no category judgment is made, and the right-hand-side represents two possible cases—one in which the face is categorized as bad, and another in which the face is categorized as good. Because good and bad are mutually exclusive and exhaustive states of the world, the probability of each state should sum to the probability in which neither state is known. If this equation is true, the LOTP holds, and no interference effect occurs. However, if the LOTP does not hold, it follows that the act of categorizing the face interferes with the subsequent decision.

An example of a typical interference effect pattern can be found in Table 1. The pattern is typified by interference effects of approximately equal magnitude but opposite direction in the *xd* condition, a positive interference effect for type-b faces in the *cd* condition, and the absence of an interference effect for type-g faces in *cd* condition. This asymmetrical pattern in *cd* has been challenging for CPT models, such as signal detection and Markov models, to predict (Wang & Busemeyer, 2016).

Table 1: Interference effects reported in Experiment 2 of Wang & Busemeyer (2016). Values are computed as the difference of the left and right hand side of Equation 1.

xd		cd	
type-b	type-g	type-b	type-g
0.03	-0.03	0.04	0.00

Belief-Action Entanglement Model

The belief-action entanglement (BAE) model is a quantum model of interference effects (Wang & Busemeyer, 2016). Importantly, the axioms on which quantum models are based allow for the violation of certain rules in classical probability, such as the LOTP. In the BAE, beliefs are represented by four orthornormal basis vectors corresponding to the four combinations of category (good vs. bad) crossed with action (attack vs. withdraw). Prior to making a decision, a person is in an indefinite state called a superposition, which is a linear combination of the four basis vectors.

During the deliberation process, a person's indefinite state evolves according to a wave function with different potentials to attack or withdraw. The decision dynamics are governed by four utility parameters which represent the utilities of attacking under different conditions. For example, $\mu_{\mathrm{tg},b}$ is the utility of attacking a type-g face that has been categorized as bad. The utility parameters are are assumed to be symmetric for type-g faces: $\mu_{\mathrm{tg},g} = -\mu_{\mathrm{tg},b}$. However, for type-b faces, the utilities can be asymmetrical, which allows interference to occur. In the *d* and *cd* conditions, entanglement aligns beliefs and actions to be consistent with each other. A parameter, γ, controls the degree of entanglement as well as its direction. Importantly, the entanglement and the utility parameters interact to produce interference effects. An interference effect will occur whenever the entanglement parameter is nonzero and the utility parameters for a given feature type (e.g., type-b) are asymmetrical (e.g., $\mu_{bg,b} \neq -\mu_{tb,g}$). The BAE also includes a parameter j, which represents the probability of categorizing a face into its most likely category (e.g., type-g categorized as good).

Judgment Revision Model

We developed a novel multinomial processing tree (MPT) model of the categorize-decide paradigm called the Judgment Revision model (JRM). Although the JRM is based on CPT, it can produce interference effects under specific conditions. A MPT characterizes how latent cognitive processes map onto categorical responses which follow a multinomial distribution (Riefer & Batchelder, 1988). As the name implies, MPTs are organized as a tree-like structure in which nodes represent cognitive states or processes and branches that connect nodes represent the transition from one cognitive state or process to another. Each branch is associated with a parameter representing a transition probability between cognitive states or processes. A series of transitions ultimately terminates at a response node representing a specific response category. The probability of following a specific path to a response node (i.e., a series of connected branches) is computed as the product of transition probabilities. In a MPT, several paths can terminate at a response node representing the same response category; in this case, the marginal probability of a specific response is the sum of all path probabilities linked to the response category.

The JRM assumes interference effects emerge from the interaction of three cognitive processes. The first cognitive process is the decision to attack a face, which is represented by parameter a. The probability of attacking depends on both the face type and the category of the face, leading to the use of two indices: (1) the first index represents the feature type (tg for type-g and tb for type-b), and (2) the second index represents the category (g for good and b for bad). The second cognitive process is the categorization of a face as good or bad. The parameter j represents the probability of categorizing a face into its most likely category (e.g., type-g as good). The third cognitive process is the decision to continue with the initial category judgment or to revise it, which is captured by parameter c. With probability c, a person is certain in the initial category judgment and continues to the decision process without revising the initial category. With probability $1-c$, a person is uncertain in the initial category judgment and revises it from good to bad (or vice versa) before continuing to the decision process. As we detail later, if one can assume that certainty in the categorization (i.e., c) can vary across some conditions, the JRM can produce the observed interference effect pattern.

Predictions

Category Given Condition In the *xd* condition, subjects are given both the feature and the category cues prior to making a decision to attack or withdraw. Parameters j and c play no role in this condition because the correct category information is provided, thus leading to simplified equations. The probability of attacking a type-b face in category b is:

$$\Pr_{\mathrm{xd}}(A=a \mid F=\mathrm{tb}, C=b) = a_{\mathrm{tb},b}.$$

The probability of attacking a type-b face in category g is:

$$\Pr_{\mathrm{xd}}(A=a \mid F=\mathrm{tb}, C=g) = a_{\mathrm{tb},g}$$

The probability of attacking a type-g face in category b is:

$$\Pr_{\mathrm{xd}}(A=a \mid F=\mathrm{tg}, C=b) = a_{\mathrm{tg},b}$$

The probability of attacking a type-g face in category g is:

$$\Pr_{\mathrm{xd}}(A=a \mid F=\mathrm{tg}, C=g) = a_{\mathrm{tg},g}$$

To compute the marginal probability of attacking in the *xd* condition, it is necessary to multiply the conditional attack probabilities by the objective category probabilities, p. For this, we assume that p is the same for both b and g faces; thus p is the probability that a face belongs to the most probable category (e.g., type-g is in category g). Formally,

$$p = \Pr(C=g \mid F=\mathrm{tg}) = \Pr(C=b \mid F=\mathrm{tb}).$$

The marginal probability of attacking a type-g face in the *xd* condition is:

$$\begin{aligned}
&\Pr_{\mathrm{xd}}(A=a \mid F=\mathrm{tg}) \\
&\quad = \sum_{n\in\{g,b\}} \Pr_{\mathrm{xd}}(A=a, C=n \mid F=\mathrm{tg}) \\
&\quad = \sum_{n\in\{g,b\}} \Pr_{\mathrm{xd}}(A=a \mid C=n, F=\mathrm{tg}) \Pr(C=n \mid F=\mathrm{tg}) \\
&\quad = p \cdot a_{\mathrm{tg},g} + (1-p) \cdot a_{\mathrm{tg},b}.
\end{aligned}$$

Similarly, the marginal probability of attacking a type-b face in the *xd* condition is:

$$\Pr_{\mathrm{xd}}(A=a \mid F=\mathrm{tb}) = p \cdot a_{\mathrm{tb},b} + (1-p) \cdot a_{\mathrm{tb},g}$$

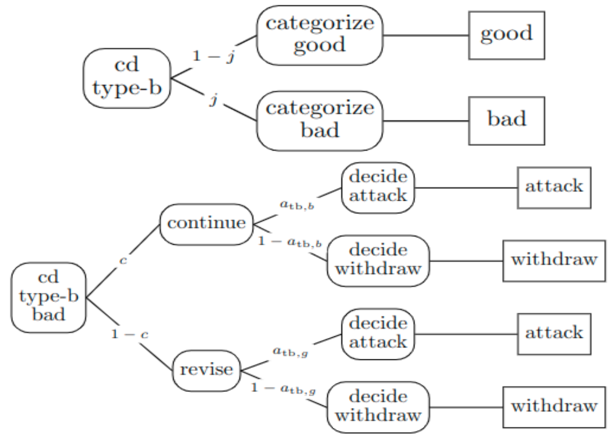

Figure 1: Example trees for the *cd* condition. The top tree represents the categorization process for a type-b stimulus in the *cd* condition. The bottom tree represents the decision process for a type-b stimulus categorized as bad

Categorize and Decide Condition In the *cd* condition, subjects are instructed to categorize the face before deciding whether to attack or withdraw. The first tree in Figure 1 illustrates the categorization process for a type-b face. In the first branch, a type-b face is categorized as good with probability $1-j$. In the second branch, a type-b face is categorized as bad with the complementary probability j. The probability of categorizing a type-b face as good is given by:

$$\Pr_{\mathrm{cd}}(C=g \mid F=\mathrm{tb}) = 1-j$$

The probability of categorizing a type-g face as good is:

$$\Pr_{\mathrm{cd}}(C=g \mid F=\mathrm{tg}) = j$$

After categorizing the face, a person must decide to attack or withdraw. As shown in Figure 1, there are two paths leading to a decision to attack. In the first path, a person is certain with probability c and continues with the initial category judgment of bad. The face is then attacked with probability $a_{\mathrm{tb},b}$. In the second path, a person is uncertain with probability $1-c$ and revises the initial category judgment from bad to good. Next, the face is attacked with probability $a_{\text{tb},g}$. This process can be represented mathematically with the following equation:

$$\Pr_{\text{cd}}(A = a \mid F = \text{tb}, C = b) = c \cdot a_{\text{tb},b} + (1 - c) \cdot a_{\text{tb},g}$$

One important point to note is that the JRM does not require certainty in category judgments to be equal in all conditions. In particular, we assume that c is higher in the *cd* condition in which a type-b face is categorized as good. The c parameter in this condition is denoted as c_k to distinguish it from c in the other conditions. Importantly, when $c_k > c$, the JRM can produce a positive interference effect for type-b faces in the *cd* condition. Without this assumption, the JRM can only produce interference effects in the *xd* conditions. Justification for this assumption can be found in Table 2 where certainty is measured as the degree to which conditional attack probabilities are close to the boundaries 0 or 1. As expected, we tend to see more certainty in *xd* because all information is provided. However, this pattern is reversed for type-b face categorized as good in the *cd* condition. Thus, we assume $c_k > c$. The probability of attacking a type-b face categorized as good is:

$$\Pr_{\text{cd}}(A = a \mid F = \text{tb}, C = g) = c_k \cdot a_{\text{tb},g} + (1 - c_k) \cdot a_{\text{tb},b}.$$

The probability of attacking a type-g face categorized as bad is given by:

$$\Pr_{\text{cd}}(A = a \mid F = \text{tg}, C = b) = c \cdot a_{\text{tg},b} + (1 - c) \cdot a_{\text{tg},g}.$$

The probability of attacking a type-g face categorized as good is given by:

$$\Pr_{\text{cd}}(A = a \mid F = \text{tg}, C = g) = c \cdot a_{\text{tg},g} + (1 - c) \cdot a_{\text{tg},b}.$$

Table 2: Conditional attack probabilities reported in Wang & Busemeyer (2016) Experiment 2 .

	Good		Bad	
	type-g	type-b	type-g	type-b
Certain (*xd*)	0.28	0.40	0.58	0.69
Uncertain (*cd*)	0.33	0.37	0.53	0.61

The marginal probability of attacking is found by combining the equations for category judgment and decision processes. The marginal probability of attacking a type-b face in the *cd* condition is given by:

$$\begin{aligned}\Pr_{\text{cd}}(A = a \mid F = \text{tb}) = (1 - j) \cdot [c_k \cdot a_{\text{tb},g} + (1 - c_k) \cdot a_{\text{tb},b}] \\ + j \cdot [c \cdot a_{\text{tb},b} + (1 - c) \cdot a_{\text{tb},g}]\end{aligned}$$

The marginal probability of attacking a type-g face in the *cd* condition is given by:

$$\begin{aligned}\Pr_{\text{cd}}(A = a \mid F = \text{tg}) = j \cdot [c \cdot a_{\text{tg},g} + (1 - c) \cdot a_{\text{tg},b}] \\ + (1 - j) \cdot [c \cdot a_{\text{tg},b} + (1 - c) \cdot a_{\text{tg},g}]\end{aligned}$$

Decision Only Condition In the d condition, subjects simply make the decision to attack or withdraw from each face. The JRM assumes that an implicit categorization precedes the decision to attack. The marginal probability of attacking a type-b face in the d condition is given by:

$$\begin{aligned}\Pr_d(A = a \mid F = \text{tb}) = (1 - j) \cdot [c \cdot a_{\text{tb},g} + (1 - c) \cdot a_{\text{tb},b}] \\ + j \cdot [c \cdot a_{\text{tb},b} + (1 - c) \cdot a_{\text{tb},g}].\end{aligned}$$

The equation above provides four paths leading to a decision to attack. The first two paths begin with categorizing a type-b face as good with probability $1 - j$. In the first path, a person is certain in the category judgment with probability c and continues without revision. From there, a person attacks with probability $a_{\text{tb},g}$. In the second path, a person is uncertain in the initial category judgment with probability $1 - c$ and revises it from good to bad. From there, a person attacks with probability $a_{\text{tb},b}$.

The other two paths begin with categorizing a type-b face as bad with probability j. In the third path, a person is certain in the category judgment with probability c and continues without revision. From there, a person attacks with probability $a_{\text{tb},b}$. In the fourth path, a person is uncertain in the initial category judgment with probability $1 - c$ and revises it from bad to good. From there, a person attacks with probability $a_{\text{tb},g}$. The marginal probability of attacking a type-g face in the d condition is given by:

$$\begin{aligned}\Pr_d(A = a \mid F = \text{tg}) = j \cdot [c \cdot a_{\text{tg},g} + (1 - c) \cdot a_{\text{tg},b}] \\ + (1 - j) \cdot [c \cdot a_{\text{tg},b} + (1 - c) \cdot a_{\text{tg},g}].\end{aligned}$$

Parameter Space Partitioning

We found that the JRM and BAE provide similar quantitative fits to the data, so we focus instead on comparing their prediction spaces. A model that predicts any pattern provides little evidence for a theory, no matter how well it fits a particular data set (Roberts & Pashler, 2000). Thus, it is important to know the range of patterns a model can and cannot produce. For this reason, we compare the prediction space of both models using a qualitative model comparison method called parameter space partitioning (PSP; Pitt et al., 2006). PSP explores the parameter space of a model to identify regions associated with different qualitative data patterns. In contrast to model fitting which assess the quantitative fit of a model to a specific data set, the goal of PSP is to understand the behavior of the model across its entire parameter space. In addition, PSP uses volume estimation to determine the prevalence of various patterns in the parameter space.

In total, the paradigm can produce a maximum of 81 possible interference effect patterns. Specifically, the interference effect is computed as the difference between the left hand and right hand side for the definition of the LOTP in Equation 1. The resulting difference yields three types of interference effects: positive, negative and absent (i.e. a approximate difference of zero). An interference effect is computed in four conditions by crossing face type (type-g,type-b) and condition

(xd, cd). Thus, in total, there are $3^4 = 81$ possible patterns in the present paradigm. Our criteria for classifying an effect as absent was a small effect: $|\Pr_d(A = a|F = x) - \Pr_z(A = a|F = x)| \leq 0.01$, where $x \in \{\text{tg}, \text{tb}\}$ and $z \in \{\text{xd, cd}\}$.

We analyzed two versions of the BAE and the JRM:

(1) a relatively constrained version denoted by subscript c, and (2) a relatively unconstrained version denoted by subscript u. In the JRM$_c$, we constrained the judgment certainty parameters to be equal: $c_k = c$. In the JRM$_u$, we allowed $c_k > c$. In the BAE$_c$ model, we constrained $\mu_{\text{tg},b} = -\mu_{\text{tg},g}$ as described the original paper (Wang & Busemeyer, 2016). In the BAE$_u$ model, no such constraint was imposed. Except where constraints apply, the allowable parameter ranges were $j \in [0,1]$ and $\mu_{\text{tb},b}, \mu_{\text{tb},g}, \mu_{\text{tg},g}, \mu_{\text{tg},b}, \gamma \in [-2,2]$ in the BAE, and $[0,1]$ for all parameters in the JRM.

Results

Flexibility

One way to assess flexibility is to count the number of patterns a model can produce. As expected, Table 3 shows that the constrained BAE$_c$ model produced $3^3 = 27$ patterns because it cannot produce interference effects for type-g faces in cd. By contrast, the BAE$_u$ can produce all 81 possible patterns. As expected, the JRM$_c$ only produced the 9 interference effects in *xd* condition. However, the JRM$_u$ can produce the same 27 patterns as the BAE$_c$ model.

One limitation with using pattern counts to assess flexibility is that it does not take into account the volume of regions associated with a data pattern. Although two models may produce the same number of data patterns, one model may concentrate most of its volume on a small subset of patterns whereas a highly flexible model might produce a uniform distribution of volume across patterns. We used the Gini coefficient (Gini, 1921)—an economic measure of income inequality—to better quantify the flexibility of the models. A value of 0 corresponds to maximal flexibility (i.e., a uniform distribution) whereas a value of 1 indicates minimal flexibility (i.e., all volume assigned to one pattern). As shown in Table 3, the Gini coefficient varies markedly across models, but all models are far from maximal flexibility. Although the JRM$_c$ is the least flexible model, it cannot account all empirical patterns (e.g, it cannot produce an interference effect in *cd* for type-b faces). In agreement with the pattern count, the BAE$_u$ model is the most flexible model. Although the BAE$_c$ model and the JRM$_u$ model produce the same patterns, the BAE$_c$ model is less flexible.

Volume

Next, we analyze the volume of regions associated with different patterns, which are normalized as a percentage of the volume for the entire parameter space. One challenge with comparing the volume of patterns between the models is the large number of patterns (81). Our solution to this problem is to analyze volume according to three factors: the type of interference effect (positive, negative, or absent), the number of interference effects, and the condition.

Table 3: A summary of the qualitative pattern of interference effects produced by the BAE and JRM models. n is the number of possible patterns for the model. Gini is a coefficient of inequality. Volume % for patterns with at least one positive interference effect, at least one negative interference effect, and at least one absent effect.

model	n	Gini	positive	negative	absent
BAE$_c$	27	.868	80.1%	76.6%	100.0%
BAE$_u$	81	.656	81.4%	83.8%	72.3%
JRM$_c$	9	.910	58.1%	56.1%	100.0%
JRM$_u$	27	.791	73.0%	73.2%	100.0%

Table 3 shows the volume associated with positive, negative and absent interference effects. For example, a pattern was considered positive if at least one interference effect in the four conditions was positive. Volume for positive and negative interference effects were similar within each model. Volume for positive and negative interference effects was higher for BAE models compared the JRM models. The volume for at least one absent interference effect was high across all models.

Across all models, the volume estimates in Table 4 indicate that volume for interference effects in the *xd* condition was larger than for the *cd* condition. The volume in the *xd* condition was greater for the BAE models compared the the JRM models. As expected, the JRM$_c$ did not produce any interference effects in the *cd* condition. Only the BAE$_u$ model had sufficient flexibility to produce interference effects in the *cd* condition for type-g faces.

Table 4: Volume % as a function of condition and face type.

model	xd type-b	xd type-g	cd type-b	cd type-g
BAE$_c$	94.6%	97.1%	44.7%	0.0 %
BAE$_u$	94.1%	95.4%	46.3%	46.3%
JRM$_c$	72.6%	71.3%	0.0%	0.0%
JRM$_u$	70.2%	70.4%	63.3%	0.0%

Table 5 shows the estimated volume as a function of number of interference effects (positive or negative) for each model. As expected, the JRM$_c$ produced a maximum of two interference effects; the JRM$_u$ and the BAE$_c$ produced a maximum of three interference effects, and the BAE$_u$ produced a maximum of four interference effects. Generally speaking, the JRM models tend to predict a smaller number of interference effects than the BAE models.

Discussion

Our goal was to develop a MPT model of the interference effect and compare its qualitative predictions to those of the

Table 5: Volume % as a function of number of interference effects for each model.

interference effects	BAE_c	BAE_u	JRM_c	JRM_u
0	0.6%	0.8%	14.2%	9.4%
1	5.1%	5.0%	27.7%	18.5%
2	51.8%	33.1%	58.1%	30.8%
3	42.6%	33.4%	0.0%	41.3%
4	0.0%	27.7%	0.0%	0.0%

BAE quantum probability model. Our MPT model, termed the JRM, is based on three cognitive processes: a categorization process, a category revision process, and a decision process. Although the JRM is based on CPT, it can produce interference effects if the judgment certainty can differ across conditions.

We used PSP to compare the models in terms of the data patterns they can and cannot produce. This is important because a model's ability to account for an observed data pattern is less impressive if it can predict many rather than few patterns (Roberts & Pashler, 2000). Our PSP analysis produced three noteworthy findings. First, an unconstrained version of the BAE can produce all qualitative interference effect patterns, and the JRM with constraints fails to produce the observed pattern of interference effects in the *cd* condition. Second, although the unconstrained JRM and the constrained BAE produce the same patterns of the interference effect, the BAE is less flexible because the volume across patterns is less diffuse compared to the JRM. Third, the volume analysis indicates that the JRM tends to generate fewer interference effects compared to the BAE. In summary, the JRM shows promise as an alternative to the BAE, as it can also produce the empirical pattern of interference effects. However, the BAE has the advantage of being less flexible according to the PSP analysis.

Limitations

We note a few limitations. One limitation is that PSP implicitly assumes the prior distribution across parameters is uniform. An extension of PSP incorporating information about the prior probability of parameters may yield different conclusions. The JRM has at least one limitation. In contrast to the BAE, the JRM does not generalize to experiments with different reward rates or associations among features and categories because it uses a parameter for each decision probability. One possible solution to this problem would be replacing the attack probability parameter a with a utility function mapping stimulus inputs to decision probabilities.

Conclusion

One advantage of quantum cognition is its ability to account for a wide range of phenomena, such as order effects and interference effects, with similar mechanisms (Busemeyer et al., 2011). A unified account of these phenomena based on CPT has yet to emerge. Instead, modeling efforts, including this one, have focused on demonstrating that models based on CPT can produce effects that are relatively easy for models of quantum cognition to produce. Recently, for example, several CPT-based models of order effects (which violate the commutative law of CPT) have been proposed, including a MPT model (Kellen et al., 2018), an ACT-R model (Fisher et al., 2021), and a Bayesian network model (Moreira & de Barros, 2021). The wide variety of models in these demonstrations indicates that the current challenge is not one of feasibility. Indeed, models based on different assumptions can produce the effects. Instead, this lack of consensus points to a deeper theoretical challenge in providing an alternative unified account of order effects, interference effects, and other phenomena. A viable alternative to quantum cognition must ultimately seek to provide a unified account. Nonetheless, developing an alternative model of interference effects is a necessary first step in this direction.

Acknowledgments

The opinions expressed herein are solely those of the authors and do not necessarily represent the opinions of the United States Government, the U.S. Department of Defense, the Department of the Air Force, or any of their subsidiaries or employees. This research was supported by Air Force Office of Scientific Research grant 21RHCOR080. Distribution A Approved for public release. Case number: AFRL-2022-2015.

References

Busemeyer, J. R., Pothos, E. M., Franco, R., & Trueblood, J. S. (2011). A quantum theoretical explanation for probability judgment errors. *Psychological Review*, *118*(2), 193–218.

Fisher, C. R., Borghetti, L., Houpt, J. W., & Stevens, C. (2021). An ACT-R model of order effects. In *Proceedings of the 19th international conference on cognitive modeling.*

Gini, C. (1921). Measurement of inequality of incomes. *The Economic Journal*, *31*(121), 124–126.

Kellen, D., Singmann, H., & Batchelder, W. H. (2018). Classic-probability accounts of mirrored (quantum-like) order effects in human judgments. *Decision*, *5*(4), 323–338.

Moreira, C., & de Barros, J. A. (2021). Order effects in Bayesian updates. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 46).

Pitt, M. A., Kim, W., Navarro, D. J., & Myung, J. I. (2006). Global model analysis by parameter space partitioning. *Psychological Review*, *113*(1), 57–83.

Riefer, D. M., & Batchelder, W. H. (1988). Multinomial modeling and the measurement of cognitive processes. *Psychological Review*, *95*(3), 318–339.

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, *107*(2), 358–367.

Wang, Z., & Busemeyer, J. R. (2016). Interference effects of categorization on decision making. *Cognition*, *150*, 133–149.