

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A Quantum Probability Explanation for Violations of Symmetry in Similarity Judgments

Permalink

<https://escholarship.org/uc/item/6vc4f2q2>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 33(33)

ISSN

1069-7977

Authors

Pothos, Emmanuel
Busemeyer, Jerome

Publication Date

2011

Peer reviewed

A Quantum Probability Explanation for Violations of Symmetry in Similarity Judgments

Emmanuel M. Pothos (e.m.pothos@swansea.ac.uk)

Department of Psychology, Swansea SA2 8PP, UK

Jerome R. Busemeyer (jbusemey@indiana.edu)

Department of Psychological and Brain Sciences, Bloomington 47468 Indiana, USA

Abstract

A model of similarity is presented which is based on Quantum Probability (QP) theory. The model is applied to the case of violations of symmetry in similarity judgments, as demonstrated by Tversky (1977). The QP similarity model can predict such violations, on the basis of the same underlying intuitions as Tversky (1977). Moreover, we discuss how the model can be extended to account for violations of the triangle inequality and also the empirical findings in relation to Tversky's diagnosticity principle.

Keywords: Similarity; symmetry; quantum probability; representation.

Similarity and Violations of Symmetry

Similarity is a key theoretical construct in many areas of cognitive psychology (principally categorization, as nearly all formal accounts of categorization involve similarity, but also memory, decision making, and attention). One of the most intriguing empirical findings in relation to similarity is Tversky's (1977) demonstration of violations of symmetry in similarity judgments. Tversky asked participants to indicate which of two phrases 'they preferred to use', country a is similar to country b, or country b is similar to country a. For example, 66 out of 69 participants judged the similarity between Korea and China (denoted as *Similarity* (Korea, China) or just *sim*(Korea, China)) as higher than that of China and Korea (denoted as *Similarity* (China, Korea); note that Tversky employed several other pairs of countries and stimuli from other domains). This has been a hugely influential finding in the development of similarity research (his 1977 paper has been cited more than 2,200 times) and, as we shall shortly see, presents a challenge for the dominant approaches to similarity.

One of the main ways in which similarity has been understood is as a function of distance in a coordinate space. Such an approach is embodied in most formal models of categorization, such as exemplar and prototype theory. It is also the basis for Shepard's (1987) celebrated derivation of a similarity law in psychological spaces. Unfortunately, if psychological similarity is a function of distance in some coordinate, representation space, then it must be symmetric, since distance is symmetric. Nosofsky (1991) suggested the use of a 'directionality' parameter, p_{AB} , so that the distance between A and B would be written as $p_{AB} \cdot d_{AB}$. This parameter might take different values, depending on whether we consider the distance from A to B or B to A.

This approach can account for an asymmetry in similarity, though is not satisfactory in the absence of an independent way to predict the value of the directionality parameter.

Tversky's (1977) own proposal is also dependent on the appropriate setting of parameters. Tversky suggested that $similarity(A, B) = \theta f(A \cap B) - af(A - B) - \beta f(B - A)$, where θ, a, β are parameters, $A \cap B$ denotes the common features between A and B, $A - B$ the features of A which B does not have and $B - A$ the features of B which A does not have. Let's say that $\theta = 1$, $a = 1$, and $\beta = 0$. Then, $similarity(China, Korea) = f(A \cap B) - f(A - B)$, which is low, since China has many features which Korea does not have. By contrast, $similarity(Korea, China)$ would be high, since Korea has very few features which China does not have. So, such a setting of parameters in Tversky's similarity model predicts an asymmetry in similarity judgments in the observed direction. However, setting $\theta = 1$, $a = 0$, and $\beta = 0$, would predict no asymmetry. Thus, the ability of Tversky's similarity model to account for his key empirical finding is dependent on particular parameter choices in his similarity model.

Specifying a formal approach to similarity which can predict asymmetries in similarity judgments in a parameter-free way, has been the focus of intense effort (Ashby & Perrin, 1988; Bowdle & Gentner, 1997; Hahn et al., 2009; Krumhansl, 1978). We build on this effort and describe a formal similarity model which can predict violations of symmetry in similarity judgments, without parameters. The model is based on quantum probability theory (QP). QP theory is a framework for assigning probabilities to observables, much like classic probability theory (Isham, 1989). It has been favored by physicists for over 100 years over classic probability theory, because of certain fundamental properties of QP theory, such as its order and context dependence. It is exactly these properties that we believe make QP theory a suitable framework for understanding many psychological processes as well (see also Aerts & Gabora, 2005; Atmanspacher, Filk, & Romer, 2004; Busemeyer, Wang, & Townsend, 2006; Busemeyer et al., in press; Bruza, 2010; Khrennikov, 2004; Pothos & Busemeyer, 2009; Trueblood & Busemeyer, in press).

QP Theory and Similarity

Perhaps contrary to intuition, the basics of QP theory are extremely straightforward. The current knowledge state, ψ , is a unit length vector in a multidimensional space, which

corresponds to, broadly speaking, whatever a person is thinking at a particular time (we will also refer to this as the initial state vector). If we employ Dirac notation, then $|\psi\rangle$ is a column vector and $\langle\psi|$ is the adjoint (conjugate transpose) of this vector (we will often drop the bracket for convenience and refer to $|\psi\rangle$ as just ψ). Then, $|\psi\rangle\langle\psi|$ indicates an outer product and is the projector onto the one-dimensional subspace defined by $|\psi\rangle$. A projection operator is a linear operator, typically expressed as a matrix, which identifies the part of a vector which is restricted/ contained in a particular subspace. Also, $\langle\psi|\psi\rangle$ indicates a standard dot product. In this model, different elements of our knowledge (such as 'Korea' or 'China') correspond to different subspaces. This is a key departure from traditional geometric models of representation and similarity, in which different elements are individual points. An important construct in QP theory is that of a projector (or projection operator), which is a linear operator taking a vector and projecting it onto a particular subspace. For example, suppose that P_{korea} is the projector to the Korea subspace. Then, $P_{korea} \cdot \psi$ corresponds to the part of the vector ψ which is contained in the Korea subspace and $|P_{korea} \cdot \psi|^2$ (the squared magnitude of the projection of vector ψ onto the Korea subspace) corresponds to the probability that ψ is about Korea (this is one of the fundamental axioms of QP theory and a key result differentiating QP theory from linear algebra). This probability reflects the extent to which the vector and the subspace are consistent with each other and so is a measure of similarity (cf. Tenenbaum & Griffiths, 2001).

Evaluating a conjunction of probabilities in QP theory is not as straightforward as in classic probability theory, because it is typically the case that in QP theory two observables cannot be evaluated concurrently (such observables are called incompatible ones). Thus, following Busemeyer et al. (in press), we suggest that $|P_{china} \cdot P_{korea} \cdot \psi|^2$ is the joint probability that vector ψ is consistent with the Korea subspace and that the projection of ψ to the Korea subspace is consistent with the China subspace. In fact, $|P_{china} \cdot P_{korea} \cdot \psi|^2 = |P_{china} \cdot \psi_{korea}|^2 |P_{korea} \cdot \psi|^2$, where $\psi_{korea} = \frac{P_{korea} \cdot \psi}{|P_{korea} \cdot \psi|}$.

The above concerns basic assumptions of QP theory, not specific to psychology. The link with psychological process is made if we assume that the conjunction of probabilities corresponds to similarity, so that, for example, $|P_{china} \cdot P_{korea} \cdot \psi|^2$, would correspond to the similarity between Korea (the projection which is evaluated first) and China. Note that this proposal can, in fact, be seen as a generalization of Sloman's (1993) proposal that the similarity between two categories, A and B , can be computed as $sim(A, B) = \frac{F(A) \cdot F(B)}{|F(A)| |F(B)|}$, where $F(A)$ and $F(B)$ are the vectors representing the categories, the numerator is a dot product, and $|F(A)| = \langle A|A \rangle^{1/2}$. If one employs normalized vectors and in the special case where the considered subspaces are unidimensional, Sloman's similarity measure and ours are identical. However, if we

only use unidimensional subspaces then the similarity measure is symmetric, and so multidimensional subspaces (e.g., planes or hyper-planes) are necessary. This is one key advance made by using quantum theory. It is impressive that Sloman, using mostly intuitive arguments, was basically led to measures very similar to those in QP theory.

In examining how to compute $|P_{china} \cdot P_{korea} \cdot \psi|^2$, we make the assumption that when asked to evaluate the similarity between two entities, A and B , the initial vector is set so that $|P_A \cdot \psi|^2 = |P_B \cdot \psi|^2$. The intuition for this assumption is that prior to assessing the similarity between A and B , the initial vector is set in a way that is biased neither towards A nor B . The implication of this assumption is that the assessment of the similarity between two elements A and B depends only on the geometric relation between the two, corresponding subspaces, and not on whatever it is that the person may be thinking prior to the similarity assessment. Note that in this case it is possible to derive closed-form expressions for ψ so as to satisfy $|P_A \cdot \psi|^2 = |P_B \cdot \psi|^2$, whereby the A and B subspaces have arbitrary dimensionality, but it would be too much of a diversion to do this here. Finally, the fact that $|P_{china} \cdot P_{korea} \cdot \psi|^2$ depends only on the geometric relation between the two subspaces reveals that this is indeed a reasonable way to define similarity in QP theory.

The most important implication of the definition $Sim(Korea, China) = |P_{china} \cdot P_{korea} \cdot \psi|^2$ is that the outcome of the similarity process is order dependent, so that $Sim(Korea, China)$ may be different from $Sim(China, Korea)$ (as long as $P_{china} \cdot P_{korea} \neq P_{korea} \cdot P_{china}$, which will be generally the case, unless the two subspaces can be expressed with the same basis vectors, or the basis vectors of one subspace form a proper subset of the basis vectors of the other). Thus, the QP formalization of similarity judgments allows for the possibility that similarity judgments will not be symmetrical, as required to account for Tversky's (1977) corresponding empirical finding. Specifically, the QP model would be consistent with empirical results if it predicts that $Sim(korea, china) > Sim(china, korea)$ or, equivalently, $|P_{china} \cdot P_{korea} \cdot \psi|^2 > |P_{korea} \cdot P_{china} \cdot \psi|^2$. But, recall, that we have postulated that $|P_{china} \cdot \psi|^2 = |P_{korea} \cdot \psi|^2$, i.e., the initial state vector is not biased towards Korea or China, so that, without loss of generality, the condition which satisfies empirical observation is $|P_{china} \cdot \psi_{korea}|^2 > |P_{korea} \cdot \psi_{china}|^2$, where ψ_{korea} , ψ_{china} are normalized vectors in the corresponding subspaces.

The next challenge we face is to show how a violation of symmetry can be predicted in a lawful way, from the specification of Tversky's similarity task. Our starting point is the same as Tversky's, namely we assume that his participants had a more extensive knowledge of China than Korea. The way to formalize this in the QP similarity model is by assuming that the subspace corresponding to China has a higher dimensionality than the one corresponding to Korea (a subtlety arises in relation to the meaning of the dimensions in the subspace and the relation of different

dimensions to each other, however, it is not necessary to provide a full consideration of these issues for the development of the model). Of course, there is an infinite number of ways in which the dimensionality of one subspace can be greater than the dimensionality of another. In this work, we consider two particular examples of one subspace having a greater dimensionality than another. More importantly, we also discuss why the model is generally expected to be consistent with violations of symmetry in similarity judgments, under circumstances consistent with those in Tversky's (1977) demonstration.

Application of the QP similarity model

In our first example, the dimensionality of the Korea subspace is just one and the dimensionality of the China subspace is two. In order to compute P_{korea} and P_{china} we need to identify the basis vectors for each subspace (that is, the vectors which span all other vectors in the subspace). Note that all the vectors we consider are normalized. We assumed that both the Korea and the China subspace would be subspaces of the same three-dimensional space (this three-dimensional space is, in turn, assumed to be a subspace of our overall knowledge space). Given this, the basis for the Korea subspace was just a random three-dimensional vector. Two basis vectors are required to span the China subspace, since this is a two-dimensional subspace. The first basis vector for China was another random three-dimensional vector, call it China1. Then, we created another random vector, call it Random. Computing $(I - |China1\rangle\langle China1|) \cdot Random$ (where I is the three-dimensional identity matrix) and normalizing gives us a vector which is orthogonal to China1 (in general, the projector to the orthogonal complement of a subspace W is given by $P_{W^\perp} = I - P_W$). It was verified that the two basis vectors for the China subspace in each iteration of the model were orthogonal to each other (very occasionally, this is was not the case due to rounding error).

Overall, each iteration of the computation involved the specification of projectors for a random one-dimensional subspace (corresponding to Korea) and a random two-dimensional one (corresponding to China). The initial state vector was computed so that $|P_{china} \cdot \psi|^2 = |P_{korea} \cdot \psi|^2$. Then, in each iteration we compared $|P_{china} \cdot P_{korea} \cdot \psi|^2$ and $|P_{korea} \cdot P_{china} \cdot \psi|^2$. It turned out that in 100,000 iterations of this scheme it was always the case that $|P_{china} \cdot P_{korea} \cdot \psi|^2$ was always greater than $|P_{korea} \cdot P_{china} \cdot \psi|^2$, meaning that $sim(Korea, China)$ was always predicted to be larger than similarity $sim(China, Korea)$, as required for a demonstration of Tversky's (1977) empirical observation regarding violations of symmetry in similarity judgments.

In an alternative demonstration, we employed an overall five-dimensional subspace, with China corresponding to a four dimensional subspace and Korea to a random place. If

we let $x_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}$, $x_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}$, etc., then the projector to the

China subspace was defined to be $P_{china} = |x_1\rangle\langle x_1| + |x_2\rangle\langle x_2| + |x_3\rangle\langle x_3| + |x_4\rangle\langle x_4|$. The projector to the Korea two-dimensional subspace was computed as before. In this larger dimensionality case, it is a little more involved to compute an initial state vector which is neutral, but, as noted above, it is still possible to do so analytically. In 100,000 iterations of this scheme it was, again, the case that the $sim(Korea, China)$ was always predicted to be larger than similarity $sim(China, Korea)$.

Note that empirical results for such a task may deviate from the 100% prediction because, e.g., it would not be the case that for all participants the knowledge of China would be greater than the knowledge of Korea. Also, we assume that the requirement of making a similarity judgment sets the initial state vector to be neutral between the two subspaces, but in practice this would not be entirely true.

As a final check of the model, we examined a situation in which both China and Korea corresponded to one-dimensional subspaces (the corresponding basis vectors were computed as random vectors in a three-dimensional subspace), to find $sim(china, korea) < sim(korea, china)$ in 35.8% of all times in 100,000 repetitions of the scheme, with 28.2% of all cases being to exact equalities. Thus, in a case where there is no reason to expect a violation of symmetry, the model correctly predicts symmetrical similarity judgments.

We can explore in more abstract terms why the QP similarity model works. Consider a vector $|k\rangle$ and a projector $P = |x\rangle\langle x| + |y\rangle\langle y|$ and suppose that we are interested in examining how much of $|k\rangle$ is reflected in the subspace corresponding to P . In other words, we need to compute the projection $P|k\rangle = |x\rangle\langle x|k\rangle + |y\rangle\langle y|k\rangle$ (recall that $\langle x|k\rangle$, $\langle y|k\rangle$ indicate the dot products between vector $|k\rangle$ and each of the basis vectors of the P subspace; these basis vectors are $|x\rangle$ and $|y\rangle$). Clearly the amplitude of the projection depends on the absolute magnitude of both $\langle x|k\rangle$ and $\langle y|k\rangle$. By contrast, the projection to the one-dimensional subspace defined by $|x\rangle$ would be $|x\rangle\langle x|k\rangle$ and its amplitude would depend on just the absolute magnitude of $\langle x|k\rangle$. In other words, the larger the subspace, the more likely it is that the resulting projection will be large; at the extreme, if the subspace considered is the entire knowledge space, then in projecting a vector to this subspace we obtain the original vector. It is exactly in this way that the QP similarity model can account for violations of symmetry in similarity judgments, that is, in situation where the entities compared correspond to subspaces of different dimensionality. This prediction closely resonates with Tversky's (1977) intuition of when violations of symmetry in similarity judgments are expected, which is when we have more knowledge about one of the compared entities, relative to the other. But, the QP similarity model could

reach the right prediction without manipulating any parameters. This contrasts with Tversky's (1977) proposal, which requires a specific parameter setting, before it can predict violations of symmetry in the right direction.

Extensions

Tversky's (1977) paper has had a profound influence in the development of similarity research, because it presented a series of (seemingly) puzzling empirical phenomena, which set boundary conditions for any aspiring model of similarity. In this work we have considered violations of symmetry. Other key empirical demonstrations in Tversky's paper concern the violation of minimality, the violation of triangle inequality, and his so-called diagnosticity principle. We consider each of these findings in turn and discuss how the QP model could be extended to account for them.

Minimality, the triangle inequality, and symmetry are together known as the metric axioms, that is, a set of properties which any distance measure in a metric space must obey. According to minimality, the distance between a point and itself is zero and, therefore, the similarity between an entity and itself should be maximal. Tversky (1977) showed that, in some cases, naïve observers would not assign the maximum similarity rating for an identical pair of stimuli, thus violating minimality. However, from a theoretical point of view, the violation of minimality is perhaps less interesting. This is because minimality could be violated by, e.g., noise in the system (so that the same stimulus presented twice would lead to slightly different representations). Therefore, violations of minimality do not lead to strong constraints on a similarity model.

According to the triangle inequality, the distance between two points A and B will always be shorter than the distance between A and C plus the distance between C and B. In other words, the triangle inequality is a statement that the shortest distance between two points is a straight line. In terms of similarities, the triangle inequality states that the *Dissimilarity* (A,B) would always be less than *Dissimilarity* (A,C) plus the *Dissimilarity* (C,B) or the *Similarity* (A, B) would always be greater than the *Similarity* (A, C) plus the *Similarity* (C, B). Tversky (1977) reported an example where the triangle inequality is violated. Consider A=Russia and B=Jamaica, so that *Similarity* (A, B) is very low. Consider also C=Cuba. But, *Similarity* (A, C) = *Similarity* (Russia, Cuba) is high (because of political affiliation) and *Similarity* (C, B) = *Similarity* (Cuba, Jamaica) is also high (in this case because of geographical proximity). Thus, Tversky's example suggests a violation of the triangle inequality. Such a finding goes against any measure of similarity according to which similarity is a linear transformation of distances. But, if one employs a non-linear function of distance as a similarity measure, then violations of the triangle inequality can occur. For example, consider similarity as an exponentially decaying function of distance in a metric space, as is commonly assumed in models of categorization (Nosofsky, 1984; Shepard, 1987). Such a model of similarity can violate the triangle

inequality. For example, consider *Distance* (A,B)=5 units, *Distance* (A,C)=4 units, and *Distance* (C,B)=4 units; these distances clearly obey the triangle inequality. For the similarities to still obey the triangle inequality we would need that $Similarity(A,B) > Similarity(A,C) + Similarity(C,B)$. However, it follows immediately that $e^{-5} < e^{-4} + e^{-4} \Leftrightarrow 0.0067 < 0.018 + 0.018$, thus violating the triangle inequality. Thus, a violation of the triangle inequality does not present a challenge for standard approaches to similarity, even those based on a coordinate representation. However, it is still important to confirm that the QP similarity model is consistent with violations of the triangle inequality. In this paper we provide an outline for how this comes about.

Tversky (1977) explained the violation of the triangle inequality in terms of different similarity judgments eliciting a different context of comparison, so to say, for the compared quantities. For example, when comparing Russia and Cuba, the context of the comparison is one of political alignment. The basis for predicting violations of the triangle inequality with the QP similarity model is analogous. Imagine a geometrical space where different countries and their properties are represented. In one region of the space, we would have the property 'communism' and both Russia and Cuba would be placed in that region. In another region of that space, the property 'in the Caribbean' would be present, as well as Cuba and Jamaica. In fact, Cuba, would have to be in-between the regions corresponding to 'communism' and 'in the Caribbean'. Figure 1 shows a two-dimensional example for how to specify vectors consistent with these intuitions (all three countries are assumed to correspond to one-dimensional subspaces, there is no basis either in Tversky's original work or in terms of general intuition for assuming otherwise). In such a case, specifying directly a neutral initial state vector introduces considerable unnecessary complexity to the model. Thus, we simply assumed that, for example, $sim(Russia, Cuba) = |P_{Cuba}P_{Russia}\psi|^2 = |P_{Cuba}\psi_{Russia}|^2$, whereby $\psi_{Russia} = |Russia\rangle$ and likewise for the other similarity terms. Based on the representation in Figure 1, one readily obtains that $|P_{Cuba}\psi_{Russia}|^2 + |P_{Jamaica}\psi_{Cuba}|^2 > |P_{Jamaica}\psi_{Russia}|^2$, with $|P_{Cuba}\psi_{Russia}|^2 = 0.79$, $|P_{Jamaica}\psi_{Cuba}|^2 = 0.79$, and $|P_{Jamaica}\psi_{Russia}|^2 = 0.33$. In other words, this computation reveals that the similarity between Jamaica and Russia is less than the sum of the similarities for Cuba, Russia and Jamaica, Cuba, as required for demonstrating a violation of the triangle inequality in similarity judgment. This provides an existence proof that the QP similarity model can accommodate violations of the triangle inequality, when there is an intuition that this can happen empirically.

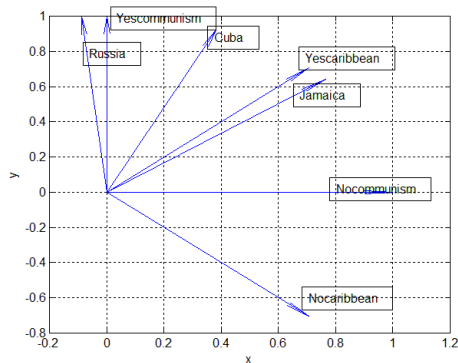


Figure 1: A representation of the countries in Tversky’s (1977) demonstration regarding the triangle inequality.

Perhaps the most significant finding in Tversky’s (1977) paper concerns his so-called demonstration of the diagnosticity principle. Tversky asked participants to pick a country most similar to Austria amongst a set of countries including Sweden, Hungary, and Poland. In such a case, participants tended to prefer Sweden. In another condition, participants were asked to decide which country was most similar to Austria amongst the set of countries Sweden, Norway, and Hungary. In such a case, participants favored Hungary. This is an intriguing phenomenon: how is it possible that the presence of irrelevant (unselected) options affects the similarity between the target item and the preferred item (cf. Roe, Bussemeyer, & Townsend, 2001)? Tversky (1977) suggested that the range of available options establish a context for the similarity judgment and this context, in turn, determines the features along which the similarity judgment takes place (see also Goldstone, Medin, & Halberstadt, 1997). For example, in the case when Austria is compared to Sweden, Hungary, and Poland, ‘Eastern Europe’ emerges as a diagnostic feature, which then makes Austria and Sweden very similar. Tversky’s finding is significant for the study of similarity because it shows that pairwise similarity judgments cannot be modeled in isolation, rather the context of the similarity judgment can have a profound influence on the outcome of the judgment.

The QP similarity model can be extended to cover the empirical findings in relation to the diagnosticity principle, though in this paper we only provide an outline of how this can be done. In brief, a key aspect of the QP similarity model is that in a series of projection operations the penultimate projection effectively establishes a context for the final projection. In the case of assessing the similarity between an isolated pair of items, A and B , we measured similarity as $Sim(A, B) = |P_B \cdot P_A \cdot \psi|^2$. An alternative interpretation of this computation is that it reflects how much of B can be understood in the context of A (Sloman, 1993). Such a scheme could be extended so that where the similarity between A and B is assessed in the context of other elements, these other elements correspond to projection operations *prior* to those for A and B . That such a scheme introduces context dependence is evident in that the

projection from one subspace to another depends on the angle between the two subspaces. Specifically, we suggest that such a scheme is appropriate for predicting the outcome of forced-choice similarity tasks, whereby all the entities involved are fairly similar to each other—this is the structure of Tversky’s (1977) experiments in relation to the diagnosticity principle. Our preliminary computations indicate the QP similarity model, if extended in this way, is consistent with the diagnosticity principle.

Conclusions

We have presented the QP similarity model and some promising analyses in support. One key conclusion is that if we associate different entities with subspaces in a multidimensional space, instead of individual points, then a suitably defined similarity measure becomes naturally (in a parameter-free way) asymmetric. Also, we have seen how a notion of similarity as projection between subspaces makes similarity judgments context dependent. This is most evident in considering diagnosticity. More generally, our work shows that similarity judgments can be understood in a formal geometric framework, a conclusion contrasting with both Tversky’s (1977) arguments and more heuristic approaches to understanding similarity.

Is the QP similarity approach falsifiable? No general framework is directly falsifiable, as particular models can always be augmented with post hoc parameters to accommodate data. The strength of the QP approach lies in the reasonableness of the assumptions which guide the specification of the model and corresponding testable qualitative properties (such as order dependence). No doubt, much additional work will be required before the QP similarity model can be established as a model of human similarity judgments. We are optimistic for a number of reasons.

First, the idea of using dot products and projections in modeling similarity judgments has already been a research focus by psychologists (e.g., Sloman, 1993). The advantage of the QP similarity model is that it draws from QP theory, a theory for assigning probabilities to observables which has been at the forefront of scientific discovery for over 100 years and has been key to some of the most impressive achievements of human science (for example, the transistor, and so the microchip, and the laser). Note that the distance between two vectors, X, Y , is a function of their dot product. The distance between two vectors X, Y (both unit length, in a real space) is given by $|X - Y|^2 = |X|^2 + |Y|^2 - 2\langle X|Y \rangle = 2 - 2\langle X|Y \rangle$. Thus, if X and Y are one-dimensional subspaces, a computation like $|P_Y \psi_X|^2$ depends on the distance between the corresponding points in the knowledge space, so that our proposal can be seen as a generalization of older approaches equating dissimilarity with distance. A key difference between such older approaches and the present proposal for similarity is that the latter is not constrained to equate concepts (or exemplars) with single points in psychological space. Rather, concepts can be subspaces of any dimensionality and, as we have

seen, this allows the prediction of important results (such as the violation of symmetry in similarity judgments).

Second, probabilistic approaches to cognition appear to work. Cognitive models based on QP theory are closely related to models based on Bayesian, classical, probability theory. In the last couple of years, the scientific community has welcomed the emergence of several sophisticated cognitive models based on classical probability theory (e.g., Tenenbaum, Griffiths, & Kemp, 2006). The success of these models attests to the promise of formal probabilistic approaches to cognition in general. Indeed, the predictions from QP theory and classical probability theory often converge. However, there is a difference between the two theories: probability assessment in QP theory is order-dependent, so that, for example, sometimes $P(A \wedge B) \neq P(B \wedge A)$. By contrast, in classic probability theory it has to be that $P(A \wedge B) = P(B \wedge A)$. Some kinds of cognitive processing (such as similarity judgments) display strong order effects. Classical probability theory could be augmented to produce order-dependent predictions if, for example, one postulates that $P(A \wedge B | O_1) \neq P(B \wedge A | O_2)$, where, basically, O_1 and O_2 are two different orders. However, we contend that where order effects do exist in cognitive processes, then QP theory provides a more natural framework for modeling.

Third, the QP theory is a linear theory. In QP models, it is often possible to derive closed-form expressions for major components. Moreover, the key elements of QP theory (in this paper we have seen projection; also, rotation, which has a more natural application in decision making problems and can capture dynamical aspects of such problems; e.g., Pothos & Busemeyer, 2009) can be expressed in basic and intuitive terms. This, we hope, endows QP theory with a transparency and explanatory penetrability which ultimately make corresponding models easier to apply and test.

Overall, it is true that QP theory sometimes looks counterintuitive (and, indeed, physicists applying QP theory for the measurement of physical observables are still puzzled by certain aspects of QP theory models or predictions). Nonetheless, QP theory has been widely adopted in physics because it does provide a very powerful coverage of physical phenomena. Likewise, we hope to have demonstrated in this paper that the QP similarity model (and QP theory more generally) has many promising elements in relation to the description of relevant psychological processes.

Acknowledgements

EMP was partly supported by the British Academy.

References

- Aerts, D., & Gabora, L. (2005). A theory of concepts and their combinations II: A Hilbert space representation. *Kybernetes*, 34, 192-221.
- Ashby, G. F. & Perrin, N. A. (1988). Towards a Unified Theory of Similarity and Recognition. *Psychological Review*, 95, 124-150.
- Atmanspacher, H., Filk, T., & Romer, H. (2004). Quantum zero features of bistable perception. *Biological Cybernetics*, 90, 33-40.
- Bowdle, B. F. & Gentner, D. (1997). Informativity and asymmetry in comparisons. *Cognitive Psychology*, 34, 244-286.
- Bruza, P. D. (2010) Quantum Memory. *Australasian Science*, 31, 34-35.
- Busemeyer, J. R., Wang, Z., & Townsend, J. T. (2006). Quantum dynamics of human decision-making. *Journal of Mathematical Psychology*, 50, 220-241.
- Busemeyer, J. R., Pothos, E. M., Franco, R., & Trueblood, J. (in press). A quantum theoretical explanation for probability judgment errors. *Psychological Review*.
- Goldstone, R. L., Medin, D. L., & Halberstadt, J. (1997). Similarity in context. *Memory & Cognition*, 25, 237-255.
- Hahn, U., Close, J., & Graf, M. (2009). Transformation direction influences shape-similarity judgments. *Psychological Science*, 20, 447-454.
- Isham, C. J. (1989). *Lectures on quantum theory*. Singapore: World Scientific.
- Khrennikov, A. Y. (2004). *Information dynamics in cognitive, psychological, social and anomalous phenomena*. Kluwer Academic.
- Krumhansl, C. L. (1978). Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. *Psychological Review*, 85, 445-463.
- Nosofsky, R. M. (1991). Stimulus bias, asymmetric similarity, and classification. *Cognitive Psychology*, 23, 94-140.
- Pothos, E. M. & Busemeyer, J. R. (2009). A quantum probability explanation for violations of 'rational' decision theory. *Proceedings of the Royal Society B*, 276, 2171-2178.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108, 370-392.
- Shepard, R. N. (1987). Toward a Universal Law of Generalization for Psychological Science. *Science*, 237, 1317-1323.
- Slovan, S. A. (1993). Feature-based induction. *Cognitive Psychology*, 25, 231-280.
- Tenenbaum, J. & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, 24, 629-641.
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, 10, 309-318.
- Trueblood, J. S. & Busemeyer, J. R. (in press). A comparison of the belief-adjustment model and the quantum inference model as explanations of order effects in human inference. *Cognitive Science*.
- Tversky, A. (1977). Features of Similarity. *Psychological Review*, 84, 327-352.