



City Research Online

City, University of London Institutional Repository

Citation: Yearsley, J. M., Pothos, E. M., Hampton, J. A. and Duran, A. B. (2015). Towards a quantum probability theory of similarity judgments. *Lecture Notes in Computer Science*, 8951, pp. 132-145. doi: 10.1007/978-3-319-15931-7_11

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/7151/>

Link to published version: http://dx.doi.org/10.1007/978-3-319-15931-7_11

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Towards a Quantum Probability Theory of Similarity Judgments

James M. Yearsley^(✉), Emmanuel M. Pothos, James A. Hampton,
and Albert Barque Duran

Department of Psychology, City University London,
London EC1V 0HB, UK
james.yearsley.1@city.ac.uk

Abstract. We review recent progress in understanding similarity judgments in cognition by means of quantum probability theory (QP) models. We begin by outlining some features of similarity judgments that have proven difficult to model by traditional approaches. We then briefly present a model of similarity judgments based on QP, and show how it can solve many of the problems faced by traditional approaches. Finally we look at some areas where the quantum model is currently less satisfactory, and discuss some open questions and areas for further work.

1 Introduction

1.1 Background

The study of similarity judgments is central to many branches of psychology (e.g. Goldstone 1994; Pothos 2005), and this is one reason why the various attempts to formalize similarity judgments have received much attention and debate (see e.g. Goodman 1972). Another reason is that similarity is often assumed to correspond to some kind of measure of the ‘distance’ between concepts in psychological space. Any proposed similarity measure based on this concept must obey various restrictions arising from the fact that (dis-)similarity functions as a metric on psychological space. For this reason models of similarity lend themselves particularly well to empirical refutation, and this feature alone may explain some of the popularity of this subject.

The classic demonstration of the failure of similarity judgments to respect the restrictions one would expect of a metric is due to Tversky (1977). Two of the empirical features of similarity judgments that Tversky reported are particularly striking: The first is a lack of symmetry in certain similarity judgments, whilst the second, dubbed the diagnosticity effect, is a particular type of contextuality. We outline both effects below.

1.2 Asymmetry

A similarity judgment is often a directional comparison of one stimulus with another; for example, how similar is A to B? Directionality can arise from the

syntax of the similarity comparison, when it is linguistically framed, but it is often a simple consequence of the fact that the relevant stimuli cannot be simultaneously presented. In the latter case, the temporal ordering of the stimuli imposes directionality structure in the similarity comparison. Whenever there is directionality in a similarity comparison, there is a potential for asymmetry.

Tversky (1977) asked participants to indicate their preference for one of two statements, e.g., '(North) Korea is similar to (Red) China' vs. 'China is similar to Korea'. Most participants preferred the former to the latter statement (this demonstration involved several other pairs of countries and was generalized to other kinds of stimuli). An important insight into why such asymmetries arise relates to an understanding of the similarity process as an interpretative one. Tversky's (1977) participants would know far less about Korea than China. Therefore, asserting that Korea is similar to China is like a process of attempting to understand the more limited representation of Korea in terms of the more extensive representation for China. China is like a cognitive reference point (cf. Rosch 1975) and the statement 'Korea is similar to China' can be considered as more informative or providing more potential for new inferences regarding Korea, on the basis of the more extensive knowledge about China (cf. Bowdle and Medin 2001). An important objective in providing a formal model of similarity asymmetries is exactly to understand how ideas like cognitive reference points or information flow may be modelled.

Most researchers accept Tversky's (1977) claim that the asymmetry in similarity judgments in the Korea, China example arises because of differences in the extent of knowledge between the two stimuli. But, asymmetries in similarity judgments can also arise in other ways. For example, Polk et al. (2002) identified asymmetries based on just differences in the frequency of occurrence of one of the compared stimuli (the highest similarity was observed when comparing the low frequency stimulus with the higher frequency one). Also, Rosch (1975) discussed asymmetries arising when comparing a less prototypical stimulus with a more prototypical one (similarity in this direction higher, than in the reverse direction). It is possible that some such asymmetries can be explained in the same way as asymmetries arising from differences in the amount of knowledge, since we may have more knowledge (in the form of a greater number of associations) for more prototypical stimuli. However, there may be other asymmetries which arise from purely perceptual properties and, in such cases, an approach based on extent of knowledge is inadequate.

It should hardly need mentioned that asymmetries are extremely difficult to reconcile with the idea of similarity-as-distance. Indeed symmetry is one of the basic assumptions of any metric function. It is however possible to modify the similarity measure to explicitly include terms that break the symmetry of similarity judgments, but these modifications have to be included by hand and are thus rather unsatisfactory. What would be preferable is some mechanism that can produce asymmetries in some circumstances in a more natural way. We will see below that the QP approach provides such a mechanism.

1.3 Diagnosticity

Another of Tversky's (1977) seminal proposals is that of the diagnosticity effect. In a typical trial, participants were asked to identify the country most similar to Austria, from a set of alternatives including Hungary, Poland, and Sweden. Participants typically selected Sweden. However, when the alternatives were Hungary, Sweden, Norway, participants typically selected Hungary. Thus, the same similarity relation (e.g., the similarity between Sweden and Austria vs. Hungary and Austria) appears to depend on which other stimuli are immediately relevant, showing that the process of establishing a similarity judgment may depend on the presence of other stimuli, not directly involved in the judgment.

Analogous context effects also appear in decision making. Consider a choice between two options. According to the so-called similarity effect, introducing an option which is equally attractive to one of the existing ones leads to an increase in the probability of the dissimilar option (e.g., Trueblood et al., in press). The diagnosticity effect has been harder to replicate, even though Tversky (1977) did report alternative demonstrations, based on variations of the stimuli. His explanation was that the diagnosticity effect arises from the grouping of some of the options. For example, when Hungary and Poland are both included, their high similarity makes participants spontaneously code them with their obvious common feature (Eastern Europe), which, in turn, increases the similarity of the other two options, through the absence of this common feature (Austria and Sweden would become similar because they are neither in Eastern Europe, rather they are in Western Europe).

As with the case of asymmetries, the diagnosticity effect is difficult to square with the notion of similarity as a distance measure on psychological space. Note however that unlike asymmetries, the diagnosticity effect has proven hard to replicate. This may indicate either that the effect is fragile, or perhaps even that it is not a genuine effect but rather an artefact of the particular set up used by Tversky. We will return to this issue below.

1.4 Discussion

We have discussed two specific empirical challenges to the idea the similarity judgments can be thought of as measuring distance in some psychological space. Of course, there is nothing particularly surprising about this. It is highly improbable that information about concepts is stored and processed in the brain in a way that can be faithfully mapped onto a Euclidean 'concept space.' Thus by the same token it should hardly be surprising if similarity judgments between some concepts resist embedding in such a concept space. Nevertheless such models have proven surprisingly popular, perhaps in part because they provide a lucid account of the cognitive process that leads to a particular similarity judgment. That is, although the work of Tversky (1977) casts doubt on the adequacy of the concept of similarity-as-distance to provide an empirical description of similarity judgments, at least some of the reason for the popularity of the idea is due to the fact it provides a very compelling description of the process of these judgments.

One of the challenges for any alternative theory of similarity judgments is to provide a similarly compelling account of how these judgments arise from simple computations in some appropriate psychological space. We will see that the QP approach, although possessing some attractive features, still has room for improvement.

2 The Quantum Model of Similarity Judgments

In this section we will present an alternative model for similarity judgments based on Quantum Probability theory (QP). The use of QP for modelling these types of judgments follows on from a number of recent attempts to describe various phenomena in psychology, and the social sciences more generally, using non-classical models of probability. In brief, there is some consensus that certain types of probabilistic reasoning, in situations where there is not just uncertainty but also a form of incompatibility between the available options (see e.g. Busemeyer et al. 2011), may be better modelled using QP than by classical probabilities theories such as Bayesian models. For examples and a more detailed justification of the use of QP in this context see e.g. Aerts and Gabora (2005), Atmanspacher et al. (2006), Busemeyer and Bruza (2011), Khrennikov (2010).

Our discussion of the QP model follows closely the account given in Pothos et al. (2013). We will begin with a concise account of the main features of the quantum similarity model. We will then consider some of the details of the model in more depth.

2.1 Outline of the Model

The basic ingredient in our quantum model is a complex vector space H (strictly a Hilbert space), representing the space of possible thoughts, which may be partitioned into (vector) subspaces, H_i , each of which represents a particular concept. The subspace corresponding to concept A may be associated with a projection operator P_A . The set of subspaces relevant to a particular set of similarity judgments need not be disjoint or complete, so that a particular thought may be associated with more than one concept. Although a realistic psychological space may have very high dimensionality, the important features can often be captured by a model with a much smaller effective concept space.

The knowledge state is given by a density operator, ρ on H . It corresponds, broadly speaking, to whatever a person is thinking at a particular time. For example, the knowledge state could be determined by the experimental instructions, or alternatively it could represent the expected degree of knowledge of naïve participants. Note that in some cases it may be more appropriate to model the knowledge state as a pure state $|\psi\rangle$, but this is not the most general possibility and is unlikely to be appropriate for describing an inhomogeneous group of participants.

Finally the similarity between two concepts A and B is computed as

$$\text{Sim}(A, B) = \text{Tr}(P_B P_A \rho P_A), \quad (1)$$

which, if the knowledge state is pure, reduces to

$$\text{Sim}(A, B) = |P_B P_A |\psi\rangle|^2. \quad (2)$$

2.2 Comments

Initial State. We will discuss how this model can reproduce asymmetries in similarity judgments below, but for now note that this effect does not follow by itself from the non-commutation of the operators P_A, P_B etc. Suppose we were to choose as an initial knowledge state the maximally mixed state corresponding to an equal prior probability for any thought, $\rho \sim 1_H$. Whether this a reasonable choice depends of course on the model, but it is easily seen that such a state leads to symmetric judgments of similarity whatever P_A and P_B . We see therefore that the specification of the initial knowledge state is an important part of this model and must be done in a reasonably principled way.

Subspaces. Subspaces of the knowledge space represent different concepts, like China. A subspace could be a ray spanned by a single vector, or a plane spanned by a pair of vectors, or a three dimensional space spanned by three vectors, etc. Suppose that the China subspace is spanned by two orthonormal vectors, $|v_1\rangle$ and $|v_2\rangle$ (that is, the China subspace is two-dimensional; we will shortly consider how meaning may be ascribed to $|v_1\rangle, |v_2\rangle$). That is, $|v_1\rangle$ and $|v_2\rangle$ are basis vectors for the China subspace. Then, the concept of China is basically all the vectors of the form $a|v_1\rangle + b|v_2\rangle$, where $|a|^2 + |b|^2 = 1$ (as is required for a state vector in quantum theory). Note that this statement is different from, though obviously related to, the statement that a category corresponds to a region of psychological space (Ashby and Perrin 1998; Nosofsky 1984). So, to represent China with a subspace is to assume that the concept China is the collection of all thoughts, $a|v_1\rangle + b|v_2\rangle$, which are consistent with this concept. For example, our knowledge of China would include information about culture, food, language etc. The representation of China as a subspace implies that all these properties have to be contained in the China subspace. Therefore, the greater the range of thoughts we can have about a concept (e.g., properties or statements), the greater the dimensionality of the subspace. If we represent China as a two dimensional subspace and Korea as a one dimensional subspace, this means that we can have a greater range of thoughts for China, than for Korea, which is equivalent to assuming that we have greater knowledge for China than for Korea.

Note that a thought of the form $|\psi\rangle = a|v_1\rangle + b|v_2\rangle$ is neither about $|v_1\rangle$ nor $|v_2\rangle$, but rather reflects the potentiality that the person will end up definitely thinking about $|v_1\rangle$ or $|v_2\rangle$ ¹. For example, if $|a| > |b|$, then this means that the person has a greater potential to think of $|v_1\rangle$ than $|v_2\rangle$. In QP theory, states like $a|v_1\rangle + b|v_2\rangle$ are called superposition states and the fact that we cannot

¹ It is often asserted that a superposition state such as $|\psi\rangle$ represents thinking about $|v_1\rangle$ and $|v_2\rangle$ at the same time. This is incorrect. The correct interpretation of such a state is that it represents thinking about neither $|v_1\rangle$ nor $|v_2\rangle$ (Griffiths 2002).

ascribe definite meaning to such states is the result of a famous theorem (the Kochen-Specker theorem).

Since the China concept is represented by a subspace spanned by vectors $|v_1\rangle$ and $|v_2\rangle$, the mathematical expression for China is a projector denoted as $P_{\text{China}} = |v_1\rangle\langle v_1| + |v_2\rangle\langle v_2|$ (although this decomposition is not unique.) Thus, following from the example above, if we think about the Chinese language, then $|\psi\rangle = |\text{Chinese}\rangle$, and $P_{\text{China}}|\text{Chinese}\rangle = |\text{Chinese}\rangle$, showing that this is a thought included in the China concept (but, the China concept would include many other thoughts). More generally the range of thoughts $|\psi\rangle$ such that $P_{\text{China}}|\psi\rangle = |\psi\rangle$ is the range of thoughts consistent with the concept of China or, equivalently, the thoughts which are part of the concept of China.

Finally we consider the meaning of vectors $|v_1\rangle$ and $|v_2\rangle$, in the claim that they span the China subspace. We could consider each such vector as a separate, distinct property of China. However, in general, different subsets of properties of a particular concept are likely to correlate with each other. For example, the properties relating to Chinese food are likely to correlate with properties relating to the general health of the average Chinese person. We so interpret $|v_1\rangle$ and $|v_2\rangle$ linearly independent combinations of all the thoughts that make up the concept of China. How to determine the set of appropriate vectors, properties, or dimensions is an issue common to all geometric approaches to similarity. Recent work, especially by Storms and collaborators (e.g., De Deyne et al. 2008), shows that this challenge can be overcome, for example, through the collection of similarity information across several concepts or feature elicitation. Then, the relatedness of the properties will determine the overall dimensionality of the concept.

Similarity. Given a particular subspace and an appropriate knowledge state vector, we can examine the degree to which the state vector is consistent with the subspace. In quantum theory, this operation is achieved by a projector. A projector can be represented by a matrix, which takes a vector and projects it (lays it down) onto a particular subspace. For example, if P_{China} is the projector onto the China subspace, then the projection $P_{\text{China}}|\psi\rangle$ represents the match between the current knowledge state and China, in other words, it computes the part of the vector $|\psi\rangle$ which is restricted or contained in the China subspace.

It is now easy to measure the consistency between a subspace and a state vector, from the projected vector. The length of the projection squared can be shown to be the probability that the state vector is consistent with the corresponding subspace. For example, the probability that a thought $|\psi\rangle$ is consistent with the China concept equals $|P_{\text{China}}|\psi\rangle|^2 = \langle\psi|P_{\text{China}}|\psi\rangle$. If the state vector is orthogonal to a subspace, then the probability is 0. In the more general language of density matrices this can be written as,

$$p(\text{China}) = \langle P_{\text{China}} \rangle_{\rho} = \text{Tr}(P_{\text{China}}\rho) \quad (3)$$

Thus the probability that the initial knowledge state is consistent with the concept China is given by the expectation value of P_{China} , computed in the

initial knowledge state. We propose that the similarity between two concepts is determined by the sequential projection from the subspace corresponding to the first concept to the one for the second concept. Roughly, this corresponds to the idea that the similarity comparison is a process of thinking about the first of the compared concepts, followed by the second. Similarity in the quantum model is about how easy it is to think about one concept, from the perspective of another. The similarity between, e.g., Korea and China may therefore be written as,

$$\text{Sim}(\text{Korea}, \text{China}) = \text{Tr}(P_{\text{China}}P_{\text{Korea}}\rho P_{\text{Korea}}), \quad (4)$$

or

$$\text{Sim}(\text{Korea}, \text{China}) = |P_{\text{China}}P_{\text{Korea}}|\psi\rangle|^2, \quad (5)$$

in the special case that the initial knowledge state is pure.

2.3 Asymmetry

Suppose we are interested in how similar Korea is to China. When there is no particular directionality implied in the judgment we can either average the result from both directionalities or determine the directionality in another way (Busemeyer et al. 2011). However, similarity judgments are often formulated in a directional way (Tversky 1977). When this is the case, we suggest that the directionality of the similarity judgment determines the directionality of the sequential projection, i.e., the syntax of the similarity judgment matches the syntax of the quantum computation. For example, the similarity of Korea to China would involve a process of thinking about Korea (subject, mentioned first) and then China (object, mentioned second), which corresponds to

$$\text{Sim}(\text{Korea}, \text{China}) = |P_{\text{China}}P_{\text{Korea}}|\psi\rangle|^2 \quad (6)$$

Let us consider the justification for this formula in more detail. Suppose the initial state is $|\psi\rangle$. From this initial state, the probability to think about Korea is $|P_{\text{Korea}}|\psi\rangle|^2$. If the person thinks that the current state matches the Korea subspace, then the new state is revised to become the normalized projection of the previous state onto the Korean subspace, so that $|\psi_{\text{Korea}}\rangle = P_{\text{Korea}}|\psi\rangle/|P_{\text{Korea}}|\psi\rangle|$. Finally, the probability that this conditional state is consistent with China equals $|P_{\text{China}}|\psi_{\text{Korea}}\rangle|^2$. Thus, $|P_{\text{China}}|\psi_{\text{Korea}}\rangle|^2|P_{\text{Korea}}|\psi\rangle|^2$ exactly computes the sequence of probabilities for whether $|\psi\rangle$ is consistent with the Korea subspace and whether the (normalized) projection of $|\psi\rangle$ onto Korea is consistent with the China subspace. The product rule then follows from,

$$\begin{aligned} |P_{\text{China}}|\psi_{\text{Korea}}\rangle|^2|P_{\text{Korea}}|\psi\rangle|^2 &= |P_{\text{China}}(P_{\text{Korea}}|\psi\rangle)/(|P_{\text{Korea}}|\psi\rangle|)^2|P_{\text{Korea}}|\psi\rangle|^2 \\ &= |P_{\text{China}}P_{\text{Korea}}|\psi\rangle|^2 \end{aligned} \quad (7)$$

(Busemeyer et al. 2011).

As we noted above in order to generate asymmetries in similarity judgments we need some principle for fixing the initial state. Usually we will (partly) fix

the initial knowledge state by demanding that it is unbiased, that is, that there is equal prior probability that the initial state is consistent with either, say, Korea or China. Such an assumption is analogous to that of a uniform prior in a Bayesian model. Then, it is straightforward to show that $\text{Sim}(\text{Korea}, \text{China}) \sim |P_{\text{China}}|\psi_{\text{Korea}}\rangle|^2$, whereby the vector $|\psi_{\text{Korea}}\rangle$ is a normalized vector contained in the Korea subspace. Therefore, the quantity $|P_{\text{China}}|\psi_{\text{Korea}}\rangle|^2$ depends on only two factors, the geometric relation between the China and the Korea subspaces and the relative dimensionality of the subspaces.

Although there is not space here for a full discussion, we note briefly that it is possible to argue against Eq.(6) as a viable measure of the similarity between A and B in the following way. Equation(6) is basically the joint probability to think about A and then B . A more natural notion of the ‘distance’ between A and B would rather be the conditional probability to think about B given we are initially thinking about A . In this case it follows that we should divide Eq.(6) by the probability of thinking about Korea, given the initial state. However this gives a result that is symmetric with respect to A and B when these are represented by rays.

The counterargument to this is that these similarities are not best thought of as ‘objective’ distances (even in a psychological space), but rather as subjective or perceived ones. This is already apparent in the fact that the representation of the stimuli depends on the extent of knowledge of these stimuli (high vs low subspace dimensionality in the case of China-Korea), and it is reasonable that the perceived similarity should also depend on the extent to which a subject may be thinking about a stimuli prior to the comparison. That A is similar to B is much less likely to occur to a subject not initially thinking about A . This line of argument is similar to that discussed in Aerts et al. (2011), where asymmetric judgments arise from the existence of a ‘point of view’ vector. Unfortunately space does not allow us to discuss the relationship between these approaches. Likewise, the connections between the emergence of similarity in the QP model and other models which directly include asymmetric similarity metrics (e.g. Jones et al. (2011) and Michelbacher (2011)) must await discussion elsewhere. Both of these issues will be taken up in Yearsley et al. (in preparation).

2.4 Diagnosticity

A modification of the basic similarity calculation to take into account context is motivated by Tversky’s (1977) diagnosticity effect, one of the most compelling demonstrations in the similarity literature. In his experiment, participants had to identify the country most similar to a particular target, from a set of alternatives, and the empirical results showed that pairwise comparisons were influenced by the available alternatives. Such an influence can be accommodated within the quantum similarity model.

Previously, we have assumed that the initial state is unbiased between the stimuli, since we have no reason to assume participants are more likely to be thinking about any particular stimulus. However sometimes what a person is thinking just prior to a comparison cannot be assumed to be irrelevant to the

comparison. Suppose that the similarity of A and B is computed in a way that has to take into account the influence of some contextual information, C, which is represented by a particular subspace. This information C could correspond to the alternatives in Tversky's (1977) diagnosticity task. The similarity between A and B should then be computed as,

$$\text{Sim}(A, B) = |P_B P_A |\psi'\rangle|^2 = |P_B |\psi'_A\rangle|^2 |P_A |\psi'\rangle|^2, \quad (8)$$

where $|\psi'\rangle = |\psi_C\rangle = P_C |\psi\rangle / |P_C |\psi\rangle|$ is no longer a state vector neutral between A and B, but rather one which reflects the influence of information C. If we minimally assume that the nature of this contextual influence is to think of C, prior to comparing A and B, then

$$\begin{aligned} \text{Sim}(A, B) &= |P_B P_A |\psi'\rangle|^2 = |P_B P_A (P_C |\psi\rangle) / (|P_C |\psi\rangle|)|^2 \\ &= |P_B P_A P_C |\psi\rangle|^2 / |P_C |\psi\rangle|^2. \end{aligned} \quad (9)$$

In other words, if the similarity comparison between A and B involves first thinking about A and then about B, then the same similarity comparison, in the context of some other information, C should involve an additional first step of first thinking about C. Additional contextual elements correspond to further prior projections, though note that eventually this process must break down (there must be a limit to how many proximal items can impact on a decision).

As before, the link with probability justifies the choice of $|P_B P_A P_C |\psi\rangle|^2$, since

$$|P_B P_A P_C |\psi\rangle|^2 = |P_B P_A |\psi_C\rangle|^2 |P_C |\psi\rangle|^2 = |P_B |\psi_{AC}\rangle|^2 |P_A |\psi_C\rangle|^2 |P_C |\psi\rangle|^2, \quad (10)$$

where $|\psi_C\rangle = (P_C |\psi\rangle) / (|P_C |\psi\rangle|)$ and $|\psi_{AC}\rangle = (P_A |\psi_C\rangle) / (|P_A |\psi_C\rangle|)$. Therefore, the similarity comparison between A and B is now computed in relation to a vector which is no longer neutral, but contained within the C subspace. Depending on the relation between subspace C and subspaces A and B, contextual information can have a profound impact on a similarity judgment. Also, the term $|P_C |\psi\rangle|^2$ affects the overall magnitude of the similarity comparison, but we assume that a computation like $|P_B P_A P_C |\psi\rangle|^2$ can lead to a sense of similarity in relation to other, matched computations. Such an assumption follows from discussions on the flexibility of similarity response scales, e.g., depending on the range of available stimuli (Parducci 1965).

Compared to the case of asymmetries, the account of the diagnosticity effect in the QP model is more heuristic. One needs to accept a very particular model of the influence of the context items on the similarity judgment, and it is hard to see how this could be convincingly motivated (we take up this challenge in Yearsley et al. (in preparation)). A more reasonable place to include such effects would be in the choice of initial knowledge state. Nevertheless the main empirical findings are reproduced well in this model, and the approach also gives some qualitative predictions about when the effect is likely to be present or absent, based on the geometric relationships between the stimuli in psychological space. Any attempt to go beyond this model will therefore have to meet a stern challenge of both matching or bettering the predictions of this model while also being more convincingly motivated.

3 Open Questions and Areas for Further Work

Below we give a (incomplete) list of problems with the current quantum model and open questions for further research. Some of these are issues which have been raised already, but it is useful to collect them in one place.

3.1 How Do We Fix the Initial State?

One obvious problem with the quantum model as presented is that it relies on a particular choice of initial state in order to reproduce the asymmetry/diagnosticity effects, but it gives little by way of explanation for this choice. Even in set ups where one can partially fix the initial state by demanding it be unbiased, as outlined above, this typically leaves many degrees of freedom unfixed. Furthermore even this partial fixing is somewhat unsatisfactory since it has a very classical feel to it, one is essentially fixing prior probabilities rather than prior amplitudes.

What is needed is firstly a reliable way to determine the knowledge state of a group of participants, and secondly a reliable way to manipulate this knowledge state, i.e. to be able to prime participants to have reasonably arbitrary knowledge states. We noted above that a more convincing way to model the diagnosticity effect would be by direct manipulation of the initial knowledge state to reflect the set of available choices, and this will be one test of any theory that fixes the initial state.

3.2 Can We Model Asymmetries Due to Frequency/Prototypicality?

An important gap in the current quantum model concerns how to deal with asymmetries arising from differences in the frequency of presentation of stimuli, or from their differing prototypicality. This failure is particularly striking when we note that there appears to be an obvious way to include such effects. Presumably what distinguishes a prototypical stimulus from a non-prototypical one, or a stimulus presented many times from one presented only infrequently is the increased potentiality for a participant to think about this stimulus. In other words, suppose $|A\rangle$ corresponds to a more prototypical/frequently presented stimulus and $|B\rangle$ to a less prototypical/frequently presented one. Then one obvious way to encode this difference is to set the initial knowledge state to be $|\psi\rangle = N(a|A\rangle + b|B\rangle)$, with N some suitable normalization factor and $|a| > |b|$. Unfortunately it is easy to show that whilst this approach does lead to asymmetries in predicted similarity judgments, it predicts the opposite effect from that empirically observed, i.e. this model predicts $\text{Sim}(A, B) > \text{Sim}(B, A)$.

We would therefore like an account of how differences in frequency/prototypicality can lead to asymmetries in the quantum model, or at the very least a good explanation of why a simple manipulation of the initial state, as suggested above, is not the right way to proceed.

3.3 Is There Genuine Contextuality in Similarity Judgments?

One of the reasons why the treatment of the diagnosticity effect is currently unsatisfactory in this model is that quantum theory naturally includes a certain amount of contextuality, but this is not what is responsible for diagnosticity in the QP model. As it stands this represents a lost opportunity, since a context effect in similarity judgments that followed from the contextuality of QP would be a very powerful, admittedly a posteriori, prediction. It would be interesting to see if a new explanation for the diagnosticity effect can be devised which makes better use of the properties of QP, or alternatively if the genuine contextuality of QP leads to additional predictions. Of course, it may also turn out that the current approach to context in the diagnosticity model (with its sensitivity to grouping) is the optimal way to proceed.

3.4 Are There Novel Quantum Predictions?

Following on from the previous point, it is important to understand whether the QP model makes any novel predictions about similarity judgments in particular cases. These could either take the form of new qualitative effects, or of quantitatively accurate predictions for similarity judgments between some simple artificial stimuli.

3.5 Can the Quantum Similarity Model Be Extended?

As well as extracting new predictions from the current QP model, it is interesting to ask whether the model may be extended in some way to cope with new types of judgment. Many of the possible extensions are not tied particularly to similarity judgments, but may be incorporated into QP models of many different types of judgments. There is not space here to discuss all the possible extensions of the QP model, but we will instead focus on a single possibility, the role of memory effects in the QP model.

A model of memory could be included in the QP scheme in at least three ways; firstly one might consider a process of forgetting whereby information about the stimuli is gradually lost. This may have the effect of reducing the effective dimension of the knowledge subspace spanned by each stimuli, and so could potentially change the size or even direction of any asymmetry effects. A second possibility is to include memory recall as a constructive process in these models, so that comparing a present stimulus with a past one may depend on whether one is asked to recall the presence or absence of certain features of the stimuli. A final radical possibility is that holding a stimulus in ones short term memory may allow thoughts about that stimulus to interact with other thoughts and memories, potentially resolving ambiguities and collapsing any superpositions of distinct possibilities. Thus it may be that quantum effects are less likely to occur the longer participants have to process the stimuli.

These are just some of the many options for extensions to the QP model. We believe these present exciting possibilities for future research directions.

3.6 Can We Frame Quantum Similarity as a Process Theory?

Perhaps the most serious concern with the quantum model is that it is not currently clear how to extract from the mathematics of the theory a picture of what similarity judgments are really about. Partly this is an inherent difficulty with quantum theory as a model for anything. Indeed, the history of attempts to decipher what quantum theory as applied to physics is really about is long, tortuous and largely unproductive. However there are some difficulties with this model that go beyond the usual problems with the interpretation of QP.

At first glance it seems like an interpretation of similarity judgments in QP in terms of the thought process involved should be obvious, indeed we explicitly motivated the order of the projection operations above by regarding the similarity judgment $\text{Sim}(A,B)$ as a process of thinking about A followed by thinking about B. However in actual fact things are slightly more complicated than this. The first complication is that it is not the order of the projection operators that is important so much as their positions relative to the knowledge state ρ . In the above we jumped the gun somewhat by calling this the initial knowledge state, but really its role is confined to ensuring judgments are not biased. There is nothing in principle to stop us computing similarity by starting from a completely mixed state, thinking about A followed by B and then demanding that our final knowledge state be unbiased in the sense above. This leads an identical expression for $\text{Sim}(A,B)$ but with the opposite ordering of the projection operators. One could also imagine demanding including both an initial and a final knowledge state.

Another difficulty with interpreting the current model is the problem, already mentioned, of establishing the correct initial states and subspaces for particular similarity judgments. However it is possible to argue that this problem is no more severe than that encountered by other approaches to representing stimuli in psychological space.

A final difficulty with QP as a process theory of similarity judgments concerns what happens when we make sequential judgments, of the kind involved in the forced choice tasks of Tversky (1977). The difficulty here is that, according to QP, after judging the similarity between A and B, our knowledge state is no longer ρ , but rather

$$\rho' = (P_B P_A \rho P_A P_B) / \text{Tr}(P_B P_A \rho P_A P_B) \quad (11)$$

That is, performing the similarity judgment between A and B collapses the initial state ρ into the new state ρ' . Such a collapse is not currently included in the quantum model.

Thus we can see that although attractive in many ways, the ‘narrative’ given by the QP theory relating to what happens during a similarity judgment is far from complete. This presents us with a problem but also with an opportunity. It is possible that by focusing on making the QP model a better description of the process of making similarity judgments, we may simultaneously clear up some of the technical problems, such as how to account for other types of asymmetry.

4 Conclusions

So what are we to conclude about the current status of the QP approach to similarity judgments? In this contribution we have been particularly harsh on the approach, and we haven't shirked from pointing out some of the flaws. However it is worth remembering that this approach does deal very well with asymmetries due to differences in the level of knowledge, providing a good qualitative account of the observed similarity judgments as well as the outline of an account of the process by which these judgments are made. In the case of diagnosticity although the details of the model are less well motivated it does provide a good fit to the current data. It is also worth pointing out that the alternatives to the QP model largely involve putting in asymmetry factors by hand. Still, the QP model could not be said to be convincing in its current form. Technical problems aside, the challenge is to convert some of the obvious parallels between similarity judgments and QP (order effects, contextuality etc.), into both a broad range of accurate empirical predictions/explanations and a convincing narrative of the cognitive processes behind similarity judgments. However it would be wrong to be overly pessimistic. The QP approach to similarity judgments is more than just an alternative to a particular classical theory of similarity-as-distance. Instead it is better seen as just one possible application of an entirely new way of thinking about cognition that may also be applied to constructive judgments, belief updating, moral dynamics and many other areas of research in cognition. The QP approach to similarity may be still in its infancy, but one should be prepared to accept such teething troubles when the reward is the possibility of a revolution in our understanding of cognition.

Acknowledgments. E.M.P. and J.M.Y. were supported by Leverhulme Trust grant no. RPG-2013-00. Further, E.M.P. was supported by Air Force Office of Scientific Research (AFOSR), Air Force Material Command, USAF, grants no. FA 8655-13-1-3044. The US Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright notation thereon.

References

- Aerts, D., Gabora, L.: A theory of concepts and their combinations II: a Hilbert space representation. *Kybernetes* **34**, 192–221 (2005)
- Aerts, S., Kitto, K., Sitbon, L.: Similarity metrics within a point of view. In: Song, D., Melucci, M., Frommholz, I., Zhang, P., Wang, L., Arafat, S. (eds.) *QI 2011*. LNCS, vol. 7052, pp. 13–24. Springer, Heidelberg (2011)
- Ashby, G.F., Perrin, N.A.: Towards a unified theory of similarity and recognition. *Psychol. Rev.* **95**, 124–150 (1988)
- Atmanspacher, H., Romer, H., Wallach, H.: Weak quantum theory: formal framework and selected applications. Weak quantum theory: complementarity and entanglement in physics and beyond. *Found. Phys.* **32**, 379–406 (2006)
- Bowdle, B.F., Medin, D.L.: Reference-point reasoning and comparison asymmetries. In: Moore, J.D., Stenning, K. (eds.) *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, pp. 116–121. Psychology Press, New York (2001)

- Busemeyer, J.R., Bruza, P.: *Quantum Models of Cognition and Decision Making*. Cambridge University Press, Cambridge (2011)
- Busemeyer, J.R., Pothos, E.M., Franco, R., Trueblood, J.: A quantum theoretical explanation for probability judgment errors. *Psychol. Rev.* **118**, 193–218 (2011)
- De Deyne, S., Verheyen, S., Ameel, E., Vanpaemel, W., Dry, M.J., Voorspoels, W., Storms, G.: Exemplar by feature applicability matrices and other Dutch normative data for semantic concepts. *Behav. Res. Meth.* **40**, 1030–1048 (2008)
- Goldstone, R.L.: The role of similarity in categorization: providing a groundwork. *Cognition* **52**, 125–157 (1994)
- Goodman, N.: Seven strictures on similarity. In: Goodman, N. (ed.) *Problems and Projects*, pp. 437–447. Bobbs-Merrill, Indianapolis (1972)
- Griffiths, R.B.: *Consistent Quantum Theory*. Cambridge University Press, Cambridge (2002)
- Jones, M.N., Gruenfelder, T.M., Recchia, G.: In defense of spatial models of lexical semantics. In: Carlson, L., Hlscher, C., Shipley, T. (eds.) *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, pp. 3444–3449. Cognitive Science Society, Austin (2011)
- Khrennikov, A.Y.: *Ubiquitous Quantum Structure: From Psychology to Finance*. Springer, Berlin (2010)
- Michelbacher, L., Evert, S., Schtze, H.: Asymmetry in corpus-derived and human word associations. *Corpus Linguist. Linguist. Theor.* **7**(2), 245–276 (2011)
- Nosofsky, R.M.: Choice, similarity, and the context theory of classification. *J. Exp. Psychol. Learn. Mem. Cogn.* **10**, 104–114 (1984)
- Parducci, A.: Category judgment: a range-frequency model. *Psychol. Rev.* **72**, 407–418 (1965)
- Polk, T.A., Behensky, C., Gonzalez, R., Smith, E.E.: Rating the similarity of simple perceptual stimuli: asymmetries induced by manipulating exposure frequency. *Cognition* **82**, B75–B88 (2002)
- Pothos, E.M.: The rules versus similarity distinction. *Behav. Brain Sci.* **28**, 1–49 (2005)
- Pothos, E.M., Busemeyer, J.R., Trueblood, J.S.: A quantum geometric model of similarity. *Psychol. Rev.* **120**, 679–696 (2013)
- Rosch, E.: Cognitive reference points. *Cogn. Psychol.* **7**, 532–547 (1975)
- Trueblood, J.S., Brown, S.D., Heathcote, A., Busemeyer, J.R.: Not just for consumers: context effects are fundamental to decision-making. *Psychological Science* (in press)
- Tversky, A.: Features of similarity. *Psychol. Rev.* **84**, 327–352 (1977)
- Yearsley, J.M., Pothos, E.M., Hampton, J.A., Barque Duran, A.: A quantum approach to context effects in human similarity judgments (in preparation)