



City Research Online

City, University of London Institutional Repository

Citation: Pothos, E. M. & Busemeyer, J. R. (2013). Quantum principles in psychology: The debate, the evidence, and the future. *Behavioral and Brain Sciences*, 36(3), pp. 310-327. doi: 10.1017/S0140525X12003226

This is the unspecified 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/2429/>

Link to published version: <https://doi.org/10.1017/S0140525X12003226>

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.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Quantum principles in psychology: the debate, the evidence, and the future

Emmanuel M. Pothos¹ & Jerome R. Busemeyer²

1. City University London, Department of Psychology, London EC1V 0HB, UK, UK; tel.: 0044 20 7040 8500; emmanuel.pothos.1@city.ac.uk; <http://www.staff.city.ac.uk/~sbbh932/>

2. Indiana University, Department of Psychological and Brain Sciences, 10th St., Bloomington, IN 47405, USA; jbusemey@indiana.edu; <http://mypage.iu.edu/~jbusemey/home.html>

Running text word count: 10,071 words; number of new references: 8

Abstract

The attempt to employ quantum principles for modeling cognition has enabled the introduction of several new concepts in psychology, such as the uncertainty principle, incompatibility, entanglement, and superposition. For many commentators, this is an exciting opportunity to question existing formal frameworks (notably classical probability theory) and explore what is to be gained by employing these novel conceptual tools. This is not to say that major empirical challenges are not there. For example, can we definitely prove the necessity for quantum, as opposed to classical, models? Can the distinction between compatibility and incompatibility inform our understanding of differences between human and non-human cognition? Are quantum models less constrained than classical ones? Does incompatibility arise as a limitation, to avoid the requirements from the principle of unicity, or is it an inherent (or essential?) characteristic of intelligent thought? For everyday judgments, do quantum principles allow more accurate prediction than classical ones? Some questions can be confidently addressed within existing quantum models. A definitive resolution of others will have to anticipate further work. What is clear is that the consideration of quantum cognitive models has enabled a new focus on a range of debates about fundamental aspects of cognition.

1. Beyond classical probability theory: the potential of quantum theory in psychology

As we mentioned in our main article, quantum probability theory simply refers to the theory for assigning probabilities to events developed in quantum mechanics, without any of the physics (cf. Aerts, Broekaert, Gabora, & Sozzo). We are interested in whether the computational properties of quantum models provide a natural way to describe cognitive processes. Quantum cognitive models do not require or assume that any aspect of brain neurophysiology is quantum mechanical. The uniqueness of quantum approaches relative to classical vector models has more to do with its use of projections on subspaces to compute probabilities rather than the fact that complex vector spaces

are employed (Behme). Some quantum cognitive models can be realized with only real vector spaces but others need to use complex vector spaces. The key features of quantum theory arise from certain mathematical theorems, such as Heisenberg's uncertainty principle for incompatible observables, Gleason's theorem (establishing the necessity for computing probabilities from the squared length of projections) and the Kochen-Specker theorem (establishing the necessity for superposition states within a vector space). Because of such theorems, we have properties like incompatibility and entanglement, which have no analogues in classical probability theory or classical vector space theory. It is possible to have heuristic versions of such ideas. But, a formal implementation with a theory of probability is not possible, unless one employs quantum probability theory. The need for complex vector spaces in particular tends to arise when the empirical result involves violations of the law of total probability, since such violations indicate the presence of interference terms. Relatedly, do quantum models provide a unique insight about cognition? Notions such as incompatibility, superposition, and entanglement have been alien to cognitive psychology (the early relevant insight of William James, that influenced the founding fathers of quantum mechanics, had little impact on psychologists), until quantum cognitive models started to emerge. These are all ideas that have no underpinning in any of the other formal theories.

Lee and Vanpaemel question whether any insight about cognition can be provided from quantum models. They argue that "quantum theory assumes deterministic causes do not exist, and that only incomplete probabilistic expressions of knowledge are possible," implying that classical theory does allow deterministic causes. We think this distinction is wrong. If I know that Clinton is honest, then there is a probability $Probability(Gore|Clinton)$, that Gore is honest too. What determines the resolution of my opinion that Gore is honest, given that Clinton is honest? This is as mysterious in probabilistic models of cognition, as in quantum theory. The behavior we observe is indeed probabilistic and so the models must generate predictions that are probabilities. Moreover, a sense of probabilistic determinism can arise in quantum theory in a way analogous to that of classical theory: in quantum theory, if it is likely that thinking that Gore is honest makes Clinton likely to be honest too, then the subspaces for the corresponding outcomes are near to each other. It will not always be the case that a person thinking that Gore is honest will also think that Clinton is honest, but, on average, this will be the case. The corresponding prediction/ implication from classical theory would seem identical. Also, the Schrödinger equation (which is used for dynamics in quantum theory) is a deterministic differential equation, analogous to the Kolmogorov forward equation (which is used in stochastic models of decision making). In quantum theory, the only source of non-determinism comes from the reduction of the state vector.

Lee and Vanpaemel note that "effects of order, context and the other factors... might be hard to understand but adopting the quantum approach requires us to stop trying." On the contrary, quantum models provide a formal mathematical framework within which to understand context, order effects. Such effects no longer have to be relegated to ad hoc (e.g., order) parameters or conditionalizations, rather they can be modeled in a principled way, through the relation between the subspaces corresponding to different outcomes. A related point is that "discussion of entanglement, ... makes it impossible to construct a complete joint distribution between two variables, and thus impossible to model how they interact with each other." The first part of the sentence is correct, but not the second: entanglement does not mean that it is impossible to model the interaction between two variables, rather that it is impossible to do so in a classical way. This is because, for entangled variables, the interaction exceeds classical boundaries. Contrary to what Lee

and Vanpaemel suggest, the objective of quantum cognitive models is exactly to provide insight into those aspects of cognitive process for which classical explanation breaks down.

Gelman and Betancourt note that, in social science, a theory which allows the complete joint to change after each measurement makes sense and outline some situations which might reveal violations of the law of total probability. These observations support the use of quantum theory in corresponding models, though Gelman and Betancourt also questioned some aspects of current quantum cognitive models.

Pleskac, Kvam, and Yu also wondered whether the notion of superposition might reduce insight into quantum models. This is not the case. Note first that a superposition state would have no specific value with respect to a particular question; with respect to another question, the same state may well have specific values. Second, the fact that a superposition state has no specific values (relative to a question) does not mean that this state contains no information. On the contrary, the relation of the state with the various subspaces of interest is the key determining factor in any probabilistic estimate based on the state. In other words, a superposition state can still have highly principled, well-defined relations, with respect to all questions which we are considering (that is, all the relevant subspaces in the Hilbert space in which the state is defined).

The debate in the target article focused on the relation between quantum and classical probability theory, because these are the most comparable quantities. True, we criticize explanations based on individual heuristics. This is because, if one adopts a heuristics approach in cognitive modeling, then there are few constraints over what heuristics can be included and the relation between heuristics. Inevitably, this endows a heuristics approach with considerable flexibility. By contrast, with formal probability approaches (classical or quantum), before one writes one line of his or her model, there are plenty of relevant constraints. But, we do not wish to dismiss heuristics (Khalil): rather, we prefer to explore whether heuristics can be interpreted within a more formal framework. Also, we certainly do not dismiss bounded rationality. Indeed, this idea is one of the main ways to motivate a perspective of quantum, as opposed to classical, rationality.

Gonzalez and Lebiere rightly point out that cognitive architectures, such as ACT-R, go some way towards addressing our criticisms of approaches based on individual heuristics. This is because a cognitive architecture would have in place several constraints to prevent (some) post hoc modifications. One advantage of cognitive architectures is that their conceptual foundation is usually closer to intuition about psychological process. For example, a cognitive architecture could involve memory and association processes, which are obviously related to human cognition. We agree that cognitive architectures have a high appeal in relation to psychological explanation. One unique advantage of modeling on the basis of formal probability theory is that all the elements of a corresponding model either have to arise from the axioms of the theory (and so are tightly interconnected) or have to correspond to straightforward assumptions in relation to the specification of the modeling problem. Arguably, this is why some of the most famous empirical demonstrations in psychology concern demonstrations that human behavior is inconsistent with the axioms of formal (classical) probability theory (such as Tversky & Kahneman, 1983), while corresponding results in relation to cognitive architectures are less prominent. Overall, while we agree that models based on cognitive architectures are more appealing than ones based on individual heuristics, we still contend that formal probability theory, classical or quantum, provides a more principled approach for psychological explanation. Note also that, as described in Busemeyer and Bruza (2012), it is possible to embed quantum probability theory into a larger quantum information processing system, using concepts from quantum computing (Nielsen & Chuang, 2000).

Converging the merits of formal probability theory with cognitive architectures certainly has a priori explanatory appeal.

2. Misconceptions on limitations

Even given the broad description of the theory in the target article, we were impressed that some commentators were able to develop their own variations of quantum models. This is encouraging in relation to how accessible quantum theory is to a general cognitive psychology audience. Equally, inevitably, there were some misperceptions.

At a general level, we assessed quantum models against results which have been at the heart of the decision making literature for nearly three decades now and have had a profound impact on the development of the debate regarding the application of formal probability theory to cognition. Whether they are simplistic or not, they constitute the first testing ground for any relevant theory (Behme). Also, we do not use “physics as a standard against which to evaluate theories in psychology” (Shanteau & Weiss). We do not use physics, rather a scheme for assigning probabilities to events *from* physics. Also, we do not use quantum theory as a standard. Rather, we employ it as an alternative theory, in situations in which it looks like classical theory fails.

Kaznatcheev and Shultz (cf. Franceschetti & Gire) point out that “the quantum probability (QP) approach is a modeling framework: it does not provide guidance in designing experiments or generating testable predictions”. Of course, the same can be said of classical probability theory, neural networks, or mental models. But, does this mean that QP (or any of the other frameworks) do not provide guidance for predictions? We do not think this is the case. For each modeling framework, a researcher is called to scrutinize the unique features of the framework and so explore how these unique features can lead to novel predictions. In the case of QP theory, there are properties such as superposition, incompatibility, and entanglement, which offer possibilities for cognitive modeling, which are unique in relation to other frameworks, such as classical probability theory and neural networks. Relatedly, one cannot just adopt a general formal framework and expect predictions about psychological performance to emerge. One needs to develop specific models (taking advantage of the unique characteristics of the framework) and such models will incorporate assumptions about psychological process that go beyond the prescriptions of the framework. The value of the psychological model is then a function of the relevance of the mathematical theory, as well as the plausibility and motivation of these additional assumptions. An example is our assumption in the quantum model for the conjunction fallacy (and related effects) that more likely predicates are evaluated first. Is this assumption problematic (Kaznatcheev & Shultz)? All we require is that the decision maker has some mechanism of providing an ordinal ordering of overlap.

A key characteristic of quantum models is that they are not constrained by the unicity principle because they allow for incompatible questions. Perhaps then quantum probability models can be partly reduced to models of bounded cognition (cf. Rakow)? It is true that we motivate properties of quantum theory, notably incompatibility, partly by appeal to the unrealistic demands from the principle of unicity. In other words, we suggest that the principle of unicity is an unrealistic requirement for cognitive representations, because it is unlikely that we can develop complete joint probability distributions for the information in our environment. This recommends representations that are not constrained by the principle of unicity, i.e., incompatible representations. But, there are other motivations completely independent of considerations relating to the principle of unicity. For

example, processing one question plausibly interferes with knowledge about another; the available empirical results strongly indicate this to be the case, at least in some cases (e.g., the conjunction fallacy; Busemeyer et al., 2011). Two independently good reasons for some action appear to interfere with each other (e.g., violations of the sure thing principle in prisoner's dilemma; Pothos & Busemeyer, 2009). The meaning of individual constituents appears to determine that of a conceptual combination in a way that shows evidence for entanglement (Aerts & Sozzo, 2011; Bruza et al., 2008). All these lines of evidence could indicate properties of information processing in the cognitive system, plausibly independent of processing limitations, and which are (still) inconsistent with a classical probabilistic description of cognition.

Ross and Ladyman say that quantum cognitive models "offer no analogous law of time evolution of cognitive states." However, Schroedinger's equation does exactly this. For example, we have assumed that time evolution can correspond to a thought process, whereby the initial state develops to a state, which reflects the outcome of the thought process (e.g., Pothos & Busemeyer, 2009; Trueblood & Busemeyer, 2012; Wang & Busemeyer, in press). In such quantum cognitive models, the Hamiltonian may not represent the energy of the system, but it does represent its dynamical properties, as is appropriate. Even though there is no corresponding Planck's constant in cognitive models, Planck's constant has a unit-scaling role in physical models. Relatedly, Kaznatcheev and Shultz question the nature of time development in quantum cognitive models since "given enough deliberation time a participant will always return to a mental state indistinguishable from the one before deliberation." While this is true, deliberation in quantum models typically reflects some sort of ambivalence between e.g. various options in a decision-making problem. Importantly, resolving this deliberation (e.g., accepting an option as a solution to a decision making problem) typically involves a reduction of the state vector to a basis state. Collapsing a state, which is a superposition state (in relation to possible outcomes for a question or problem), to a corresponding basis state is an irreversible process, which breaks the periodicity in quantum time evolution. This is the assumption we made in all the quantum models just mentioned. Equally, the specific characteristics of quantum evolution in quantum models do not always map well onto cognitive processes and in some cases classical models appear more successful (Busemeyer, Wang, & Townsend, 2006).

More specific issues arise in various suggestions from commentators. Newell, Ravenzwaaij, and Donkin note in their example of multiple projections that "it seems unlikely that the believability of any president should necessarily decrease ...as more questions are asked..." We entirely agree, but what these commentators are computing is

there is still tension with classical probability. Note first that $Probability(\text{Clinton honest}) * Probability(\text{Gore honest} | \text{Clinton honest})$ is the probability of deciding that Clinton is honest and Gore is honest. Consistently with Rakow, the conditional tells us whether Gore is honest, taking into account the information that Clinton is honest. But, $Probability(\text{Clinton honest}) * Probability(\text{Gore honest} | \text{Clinton honest}) = Probability(\text{Clinton \& Gore})$. Since conjunction is commutative, the classical approach cannot model relevant question order effects, without extra post hoc order parameters.

3. Empirical and theoretical extensions

Aerts, Broekaert, Gabora, and Sozzo point out that it is not just quantum probability which is relevant. Rather, there many aspects of quantum theory which are potentially relevant to the modeling of cognition. By quantum probability we do not imply a particular aspect of quantum theory, rather the entire formalism for how to assign probabilities to events. Having said this, in pursuing the application of quantum theory to cognition, our starting point has been the more basic elements of quantum theory. We agree with Aerts et al. and the other commentators, who pointed out aspects of quantum theory, which could suitably extend existing models.

For example, Shanteau and Weiss note that happiness (in the target article example) is a continuum and so should not be modeled in a binary way. We only use binary values to make the explanations as simple as possible, and it is straightforward to develop models for more continuous responses (see, e.g., Busemeyer, Wang, Townsend, 2006, for an application to confidence ratings). For example, our demonstrations could be extended in a way such that the question of interest can have an answer along a continuum, rather than a binary yes – no. Such extensions could then be used to explore differences in performance depending on differences in response format, as Pleskac, Kvam, and Yu suggest.

Another issue concerns the output of quantum models. For example, when we evaluate $Probability(\text{Bank teller})$ for Linda in the conjunction fallacy problem, what exactly does this probability correspond to (Kaznatcheev & Shultz; Pleskac, Kvam, & Yu)? We can use a pure state in a Hilbert space to represent the state of an individual, e.g., after a particular experimental manipulation (such as reading Linda's description in Tversky & Kahneman's, 1983, experiment). From this pure state, we can estimate the probability that the individual will assign in relation to, e.g., the bank teller property for Linda. We make the reasonable assumption (common to probabilistic models of cognition, whether classical or quantum) that individual response probabilities translate to sample proportions for choosing one option, as opposed to others. Relatedly, it is not clear what is to be gained by using mixed states, in situations where the emphasis is on individual responses and sample proportions are considered only as indicators of individual behavior (Kaznatcheev & Shultz).

By contrast, if the emphasis is on ensembles (e.g., how the behavior of a whole group of people changes), then perhaps mixed states are more appropriate (Franceschetti & Gire). Mixed states combine classical and quantum uncertainties. They concern an ensemble of related elements, e.g., particles in a particular physical system or students in an undergraduate class. The representation for each element is a pure (perhaps superposition) state. But, we do not really know how many elements are in specific pure states. Therefore, the overall system is described in terms of all the possible pure states, weighted by their classical probability of occurrence. Such mixed states appear particularly useful in situations where it is impossible to consider a psychological situation in

terms of individual elements. For example, mixed states appear appropriate when modeling convergence of belief in a classroom or classification of new instances in a categorization model. Extending quantum models in this way is an exciting direction for future research.

MacLennan raises some interesting points in relation to how the basis vectors corresponding to different questions are determined/ learned. Currently, we make the following minimal, but reasonable assumptions. Our knowledge space would be populated with several possible subspaces, corresponding to questions that relate to our world knowledge. For example, we would in general have subspaces corresponding to properties like being a feminist or a bank teller. Reasonably, the creation of knowledge would involve the creation of new subspaces in our knowledge space. Is there a sense in which the overall dimensionality of the knowledge space would increase, to perhaps reflect an overall quantitative increase in knowledge? Perhaps, and this would require tools from quantum theory, which go beyond those employed in most current models, such as Fock spaces (Aerts et al.). Another issue is how novel information could alter the relation between subspaces. Quantum theory provides a mechanism for doing so, through unitary evolution, though how to specify unitary matrices which can capture the impact of a particular learning experience is beyond current models (but see Ch. 11 in Busemeyer and Bruza, 2012, for some proposals). Such extensions can be a basis for quantum models of learning as well (Castro), although some corresponding insights are possible from current models too. For example, Castro discusses the idea that learning can involve a process of forgetting. In current quantum models, after a projection, the initial state loses out on some of the original information (which is like forgetting), but new insights can be acquired about the initial state (see also Bruza et al., 2009, for the applicability of quantum probability in the understanding of human memory).

Some commentators (Love; Marewski & Hoffrage) observed that incorporating process assumptions in quantum models could potentially greatly increase the scope and testability of the models. We agree and offer some preliminary responses, which show promise in this direction. First, one distinguishing characteristic between quantum and classical models is that, in the former, for incompatible questions, all operations have to be carried out sequentially. This naturally makes computation sequential. Indeed, in some cases, we have had to make particular assumptions regarding projection order, e.g., in the conjunction fallacy. This assumption does not follow from the axiomatic structure of quantum theory, rather it is required because there is no other way to go from the basic assumptions of quantum theory to a mechanism that can be used to predict behavioral results (cf. Kaznatcheev & Shultz). If one proposes that projection order, as specified at the computational level of a quantum model, maps onto a psychological time parameter for the actual computations at the algorithmic level, then projection order could be tested as a process aspect of quantum models. Second, in classical probability models in decision-making (such as random walk/ diffusion models) it is common to employ a time development process (from the forward Kolmogorov equation), but in this case, the constraints at baseline (e.g., the law of total probability) would also apply at any subsequent time. However, in quantum models, time evolution (from Schrodinger's equation) can be quite different. The states at baseline may obey the law of total probability, but after time evolution the law of total probability may break down. We have fruitfully employed this approach in a variety of situations (Pothos & Busemeyer, 2009; Trueblood & Busemeyer, 2012; Wang & Busemeyer, in press). In such quantum models, the time development of the initial state into another one (for which the law of total probability can break down) is necessarily interpreted as a process of deliberation of the relevant evidence. Thus, such models do have a process component, although a lot more work needs to be done along these lines. Finally,

Busemeyer, Wang, and Townsend (2006) explored a random walk quantum model for a decision making task. Thus, there is some groundwork on how to employ quantum principles to implement process models proper.

Behme and Corr highlighted the issue of how quantum cognitive models scale up to more realistic empirical situations. It is probably fair to say that both the application of classical theory and quantum theory have been focused on idealized, laboratory situations, because of the complexity of taking into account a full set of realistic constraints in a principled way. An example of progress in this respect is the model for behavior in Prisoner's Dilemma games of Pothos and Busemeyer (2009; see also Trueblood & Busemeyer, 2012; Wang & Busemeyer, in press). In both the classical and the quantum versions of this model, it was possible to incorporate the idea of cognitive dissonance (Festinger, 1957), so that beliefs and actions would converge to each other. The weight of this process, in relation to evaluating payoff, could not be predicted a priori and so was set with a free parameter. So, 'more realistic' mechanisms can in principle be implemented in both classical and quantum models. There is still a question of whether these mechanisms allow successful coverage of empirical data. One point is that for empirical data that violate the basic axioms of classical theory, even postulating additional realistic mechanisms (such as cognitive dissonance), is still insufficient.

Specifically in the area of preference formation, Kusev and van Schaik suggest that preferences are often constructed in an ad hoc, content-dependent way and Corr also highlights the potential impact of mood or other emotional variables on preference formation for the same option. At face value, incompatibility in quantum theory can allow evaluation of probabilities or preferences in a perspective-dependent way. Corr also points out that introducing considerations relating to mood, emotion etc. may lead to "nonlinear dynamical effects". Because probability in quantum models is assessed by squaring the length of the state vector, there is more potential for capturing any corresponding non-linearities in the dynamics of a psychological system, though clearly whether this is actually the case or not can only be assessed in the context of specific models. The work of Pothos and Busemeyer (2009), Trueblood and Busemeyer (2012), and Wang and Busemeyer (in press) takes advantage of non-linearities in the development of probabilities for particular events, to capture results that are puzzling from a classical perspective.

A more general point regarding the scalability of quantum vs. classical models concerns the requirements from the principle of unicity for the latter. Classically, there is the demanding requirement that there is a complete joint probability distribution for the outcomes of all questions. This requirement means that every time an additional question is included, the complete joint needs to be augmented to take into account the event and, also, all corresponding marginal probability distributions need be specified as well. Therefore, additional questions greatly increase the complexity of the required representation. By contrast, representations of incompatible questions in a quantum model can co-exist in a low dimensionality space. The introduction of further incompatible questions does not necessarily increase the dimensionality of the space and the subspaces for different questions can be related to each other through a unitary operation. Overall, representation of additional questions (for a hypothetical system) in a quantum model implies lower representational demands than in a classical one.

A promising new domain of application of quantum theory is learning and memory (as discussed, see Castro; Bruza et al., 2009). Baldo also discusses how the main elements of signal detection theory can be represented in a Hilbert space. This is an interesting extension to current work with quantum cognitive models. One advantage is that this provides an intuitive visual representation of some related visual phenomena. An issue for future research would be whether

there are empirical findings in signal detection theory, which necessitate a Hilbert space representation, for example, in relation to violations of the law of total probability or order/ context effects. In the more general area of perception, there has been some work along such lines, by Conte et al. (2009) and Atmanspacher and Filk (2010).

4. Empirical challenges

Whether researchers accept the quantum framework as a viable alternative to classical probability theory is partly an empirical issue. We review some important empirical challenges, as highlighted by the commentators.

Hampton points out the quantum model for the conjunction fallacy cannot be applied to understand the guppy effect – a situation in which the conjunction is more probable than both the individual constituents. Note first that we agree that researchers should be looking to understand similarity and decision-making processes within the same modeling framework (Shafir, Smith, & Osherson, 1990). However, the mechanics of projection in the Linda problem are quite different to those for the guppy effect. In the Linda problem, the initial state reasonably corresponds to the impression participants' form after the Linda story. Then, this initial state is assessed against either the single property of bank teller or the conjunction of bank teller and feminist properties. By contrast, in the guppy problem, the initial state is a guppy, a pet fish. If we were to adopt a representation analogous to the one for the Linda problem, we would need an initial state which exists within a subspace for pet fish, which corresponds to the concept of guppy. Then, a guppy has maximum alignment with the pet fish subspace (probability=1) and less alignment with either the pet or the fish subspaces. In other words, in the guppy case it makes no sense to consider the conjunction of pet and fish properties for guppy (as in the Linda case), since by definition guppy is constructed to be a pet fish; the pet fish characterization for guppy is tautological, rather than a conjunction. By contrast, Linda is not by definition a bank teller and a feminist, rather Linda corresponds to a particular state vector, related to all these properties. So, the guppy effect can be reproduced with the quantum model for the Linda problem, albeit in a trivial way. A problem with Hampton's model is that, like for averaging models generally, it fails to account for violations of independence, that is, findings such that $A_1+B_1>A_2+B_1$, but $A_1+B_2<A_2+B_2$ (Myamoto, Gonzalez, & Tu, 1995). Finally, an alternative, not trivial, quantum model for the guppy effect has been proposed by Aerts (2009; Aerts & Gabora, 2005). This model is different in structure from the Linda model (for example, it relies on Fock spaces, which generalize Hilbert spaces in a certain way), nevertheless both models are based on the same fundamental quantum principles.

Tentori and Crupi discuss an interesting finding according to which, for a Russian woman, *Probability* (NY and I) is higher than *Probability* (NY and ~I), whereby NY means living in New York and I being an interpreter. They show how an application of the quantum model for the conjunction fallacy in this case appears to break down. However, their demonstration is based on a two dimensional representation; but quantum models are not constrained to two dimensions. The dimensionality of the representation will be determined by the complexity of the relation between the corresponding questions. Dimensionality is not a free parameter: we just need to ask whether all the information in a problem can have an adequate representation in (say) two versus three dimensions. Note that if one parameterizes the angles between subspaces, then dimensionality can increase the number of parameters. However, with most current demonstrations, the focus has

been on showing whether a result impossible to predict with (typically) classical probability principles can emerge from a quantum model.

Regarding the problem of the Russian interpreter living in New York, that a two dimensional representation is inadequate can be seen by noting that the state vector, assumed to be within the Russian subspace, has to be very close to both the $\sim I$ and $\sim NY$ rays, as, clearly, both of these are extremely unlikely properties for a Russian woman. Moreover, the I and NY rays cannot be close to each other, since knowing that a person is an I does not really tell us very much about whether he/she lives in New York (but equally it does not preclude that the person is living in New York). One can satisfy these tricky constraints (along with the simpler one that the state vector has to be closer to the $\sim NY$ ray than the $\sim I$ one) in a three dimensional space. Note, finally, that the quantum model for the conjunction fallacy assumes that the more likely predicate is evaluated first. Then, it emerges that $\text{Prob}(I \text{ and then } NY) > \text{Prob}(\sim I \text{ and then } NY)$, as required. Figures 1a, 1b and 1c illustrate the representation, but is also possible to verify the prediction directly. For example, setting $NY=[0.39, 0.92, 0.047]$, $\sim NY=[0.087, -0.087, 0.99]$, $I=[-0.39, 0.92, 0.047]$, $\sim I=[-0.087, -0.087, 0.99]$, $\text{Russian}=[0.029, 0.087, 0.99]$, we have that $\text{Prob}(I \text{ and then } NY)=0.0065$ and $\text{Prob}(\sim I \text{ and then } NY)=0.0043$. We note that this is a toy demonstration, which is meant to neither reproduce specific numerical values nor cover all the relevant constraints of Tentori and Crupi's example (for these, a more general model would be needed). Nevertheless, the demonstration does capture their main result.

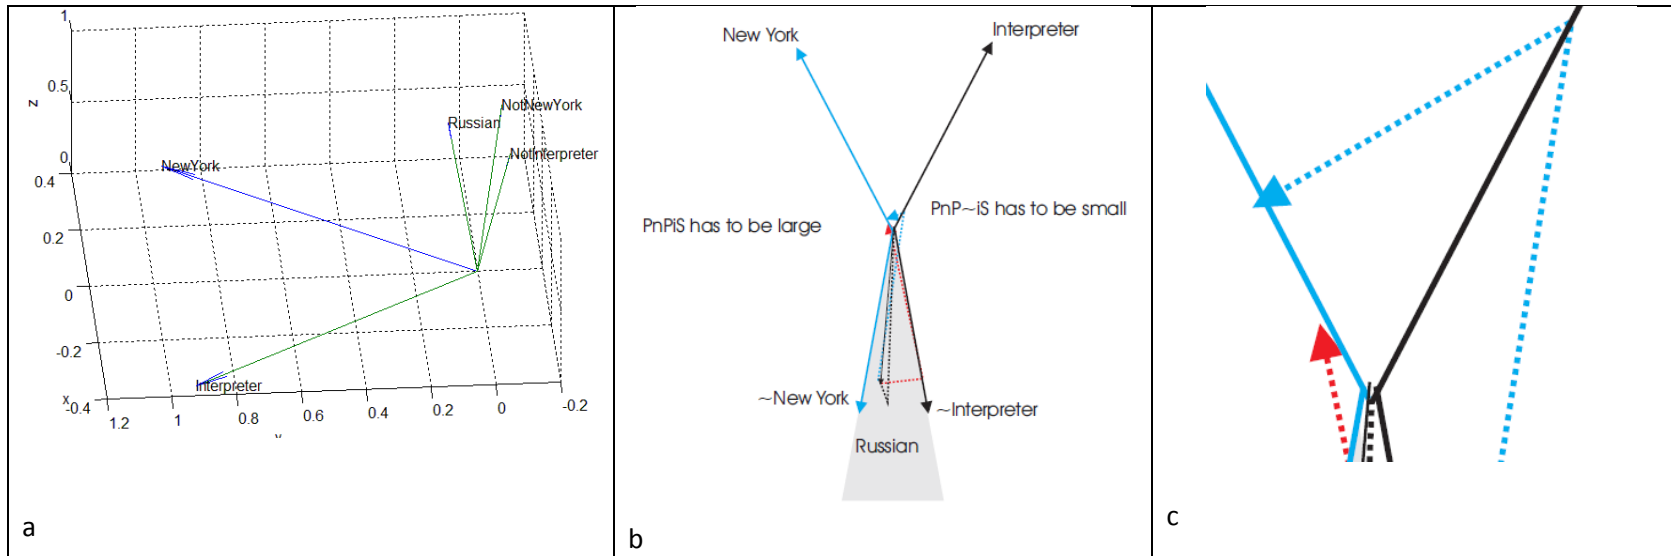


Figure 1. A representation which can account for the main finding discussed by Tentori and Crupi, that a Russian woman is more likely to be living in New York and be an interpreter, than be living in New and not be an interpreter. Figure 1c shows a magnification of the central part of Figure 1b.

A key question is exactly how definitive is the evidence we consider in favor of quantum theory. Kaznatcheev and Shultz rightly point out that there are rarely clear cut answers. They discuss the case of Aerts and Sozzo's (2011) demonstration of entanglement and conclude that one cannot necessarily infer the need for quantum theory, over and above classical theory. Likewise, violations of the law of total probability are amenable to classical explanations, for example, with appropriate conditionalizations on either side of a law of total probability decomposition. We agree, an empirically observed violation of the law of total probability does not definitely prove the inapplicability of classical probability. However, classical explanations would rely on post hoc parameters or conditionalizations, with unclear explanatory value. In quantum models, because of the order and context dependence of probabilistic assessment, it is often the case that corresponding empirical results can emerge naturally, (nearly) just from a specification of the information in the problem (as in the model for the conjunction fallacy; Busemeyer et al., 2011). Relatedly, Shanteau and Weiss questioned the complexity of quantum cognitive models. On the contrary, for the reasons just discussed, it seems clear that quantum models for empirical results such as the conjunction fallacy are actually *simpler* than matched classical ones.

Quantum models can involve both incompatible and compatible questions, classical models only the latter. Are quantum models more flexible than classical ones? Do quantum models subsume classical ones (Houston & Wiesner; Kaznatcheev & Shultz)? There are three responses. First, in all the quantum models that have been employed so far, quantum transition (unitary) matrices need to obey double stochasticity, a constraint not applicable to classical transition matrices. Second, as Atmanspacher notes (cf. Atmanspacher & Roemer, in press), it is possible to test in a very general way whether a quantum approach can apply to a given situation. Consider a particular question and the conjunction of this question with another one, assumed to be incompatible. If the individual question fails to be compatible with the conjunction (i.e., there are order effects between the question and the conjunction), then a Hilbert space representation is precluded. Third, there is the law of reciprocity, according to which the probability of transiting from one pure state to another must equal the probability of transition in the opposite direction. In other words, if we imagine a state vector in a unidimensional subspace A ,

Fuss; Newell, Ravenzwaaj, & Donkin). This point is overstated. First, the default assumption is that questions about distinct properties are incompatible. Noori and Spanagel provide a good expression of the relevant idea: “incompatibility between questions refers to the inability of a cognitive agent in formulating single thoughts for combinations of corresponding outcomes.” Conversely, compatible properties are such that it is possible to combine them in a single corresponding thought. It is usually possible to empirically determine incompatibility between two questions, since conjunction is usually subject to order effects. In the context of Gore, Clinton question order effects, Khalil argues that the questions “involve symmetrical and independent information. Information with regard to each question can be evaluated without the appeal to the other” to conclude that these questions cannot be incompatible. All these reasons are exactly why it is so puzzling that empirically we do observe order effects in the Clinton, Gore question. Crucially, the considerations Khalil provides are not really the ones which are relevant in determining incompatibility, rather these would relate to the presence of order effects.

Second, once incompatibility/ compatibility is resolved, a QP model is limited in terms of both the relation between the various subspaces and the location of the state vector. These relations can be often determined in a principled way, via correlations between answers to relevant questions and considering the participants’ frame of mind in the study, and do not prejudice the accuracy of model predictions. For example, Tentori and Crupi note that variations in the representation for the quantum model for the conjunction fallacy can provide predictions inconsistent with empirical results. Our representation for this model (Figure 2 in the target article) was created to be consistent with three constraints:

probability model can naturally produce a conjunction fallacy, because a fundamental constraint in such models is

or not can be partly replaced with the more testable/ specific question of whether measurement can change a cognitive state.

5. Neural basis

It is possible to utilize quantum theory to build models of neural activity. MacLennan suggests that the distribution of neural activity across a region of cortex could be described by a quantum wave function. While this possibility is intriguing, one has to ask whether there is any evidence that such aspects of brain neurophysiology do involve quantum effects or not. Banerjee and Horwitz offer a preliminary perspective in this respect. They note that to study “auditory-visual integration one would design unimodal control tasks (factors) employing visual and auditory stimuli separately, and examine the change of brain responses during presentation of combined visual-auditory stimuli”. This is effectively a factor based approach, but one which, in a classical analytical framework, does not take into account potential interference effects. Rather, the standard assumption is such that the brain regions engaged by each unimodal task have an additive impact in relation to the multimodal task. Relatedly, Banerjee and Horwitz point out that, when considering functional networks, it is often the case that any “any alternation in even a single link results in alterations throughout the network”. This raises the possibility that the function of nodes in a functional network cannot be described by a joint probability function, so that states of the network may exhibit entanglement-like effects. Clearly, we think that such applications are very exciting (cf. Bruza et al., 2009).

The issue of quantum neural processing, as above, is distinct to that of the neural implementation of quantum cognitive models. According to Mender, if a quantum approach to cognition was tied to neural/ physical theory of implementation, then perhaps pairs of incompatible observables could be extracted from such a theory directly. By contrast, in our current, mostly top-down, approach we have to infer incompatible observables in purely empirical ways. Moreover, in such a cognitive approach it is not really possible to specify an analogue to Planck’s constant. How much does the quantum cognition programme suffer from the lack of a clear theory of neural implementation? We believe not very much. We seek to infer the computational (and possibly process; Love; Marewski & Hoffrage) principles of cognitive processing. Such principles would be consistent with a variety of implementational approaches at the neural level. While ultimately cognitive computation and neural implementation have to constrain each other, these constraints are often fairly loose and in actual practice do not typically impact on the specification of computational models. Mender is right that perhaps neural implementation considerations might help us with the determination of pairs of incompatible properties. But, even just focusing on the fact that certain pairs of observables are incompatible, and information about their relation (which can be inferred from general properties of the observables), seems to provide us with considerable explanatory power.

There is a deeper issue here of whether there is something about the computations involved in quantum cognitive models which *preclude* neural instantiation. Noori and Spanagel argue that brain neurobiology is deterministic and, moreover, it deterministically specifies behavior. So, how can we reconcile neural/ biological determinism with behavioral stochasticity? First, the work reviewed in Noori and Spanagel’s commentary shows a causal relationship between neurochemical changes in the brain and behavior, which is perhaps unsurprising. But, the further step of assuming complete determinism between brain neurochemical reactions and behavior is unwarranted. Such a

step would necessitate assumptions along the lines that human behavior is completely predictable from the current brain state. Given the number of neurons involved, determinism in such a sense, even if 'real', would be impossible to prove or disprove. Second, we do not really know what is the neural reality of stochasticity in behavioral probabilistic models (quantum or classical). For all we know, at the neuronal level, it may be the case that such stochasticity is resolved in a deterministic way. Interestingly, Shanteau and Weiss criticize quantum models in exactly the opposite way, that is by arguing that they are too deterministic for modeling the unpredictable human behavior. We are unconvinced by such arguments. Whichever perspective one adopts in relation to the determinism of brain neurophysiology, at the behavioral level there is an obvious need to represent uncertainty in human life. Therefore, at that level, it makes sense to adopt modeling frameworks which formalize uncertainty (whether classical or quantum).

A related issue is whether quantum cognitive models require quantum processes at the neuronal level (and so a quantum brain). This does not appear to be the case. Two commentaries (Atmanspacher; Blutner and beim Graben) discuss how incompatibility, a key, if not the fundamental, property of quantum cognitive models can arise in a classical way. This can happen if measurement is coarse enough so that epistemically equivalent states cannot be distinguished. Hameroff, however, strongly argues in favor of a quantum brain as the source of quantum behavior (as he says -- if it walks and thinks like a duck then it is a duck). In the past, his work has mainly addressed the problem of consciousness. But in this comment, he advances his ideas by providing a detailed discussion for how his Orch-OR model can be extended to incorporate subspaces and projections for complex thoughts, as would be required in quantum cognitive models. This is a detailed proposal for how quantum computations can be directly implemented in a way that the underlying neuronal processes are quantum mechanical. Such ideas are intriguing, and he mentions a growing literature on quantum biology, but currently this direction remains very controversial.

Finally, Stewart and Eliasmith note how the Neural Engineering Framework can be employed in relation to vector representations and operations. An important point is that the problem of neural implementation of quantum models is not just one of representing vectors and standard vector operations, such as dot products. One also needs to consider how characteristics uniquely relevant to QP theory could be implemented, notably superposition and incompatibility. It is these latter properties which present the greatest challenge. In other words, for example, one would need to consider what in the nature of neural representation gives rise to incompatibility between different properties (as opposed to merely implementing subspaces with basis vectors at oblique angles). Stewart and Eliasmith also note that allowing tensor product operations in a behavioral model raises questions in relation to how an increasing number of dimensions are represented. But, for example, could one not envisage that there are different sets of neurons to represent the information in relation to the Happy and Employed questions? Then, when a joint representation is created, the two sets of neurons are simply brought together. It is also worth noting that Smolensky et al. (in press) have argued against the perception that increases in dimensionality due to tensor products are problematic. It is currently unclear whether such dimensionality increases present a problem in relation to neural implementation.

6. Rationality

The question of whether an account of human rationality (or not) should emerge from quantum cognitive models partly relates to their intended explanatory level. Navarro and Fuss argue that

quantum theory does not make sense as a normative account of everyday inference. Instead, their view is that classical probability theory is the better approach for computational level analysis, since “it is the right tool for the job of defining normative inferences in everyday data analysis.” However, the latter statement is a belief, not an empirically established truth. Many complex systems are sensitive to measurement, and a difficulty arises for classic probability theory when measurement changes the system being measured. Suppose we obtain measurements of a variable B under two different conditions – one in which we measure only B and another in which we measure A before B. If the prior measurement of A affects the system, then we need to construct two different CP models for B, one when A is measured and one when A is not measured. Unfortunately, these two CP distributions are stochastically unrelated, which means that the two distributions have no meaningful relation between each other (cf. Dzhamfarov and Kujala) and a more general theory is required that can construct both distributions from a single set of elements. This is what quantum probability theory is designed to do, although as Gelman and Betancourt point out there may be other ways to achieve this generalization.

A related point is that, if for the moment one does away with normative considerations, both CP and QP theory have analogous structure and objectives with respect to cognitive explanation. Both theories make sense as top-down (axiomatic) theories: in neither case are there assumptions, for example, that the cognitive system represents uncertainty in particular ways. Rather, the claim is that behavior, however it is produced, is consistent with the constraints of CP (or QP) theory. So, we think that a characterization as top down is a more accurate way to label both CP and QP explanations, without reference to normative considerations. As noted in response to other commentators (Love; Marewski & Hoffrage), we aspire to extend quantum models in a way that process assumptions are incorporated, and we think that quantum models offer promise in this respect. But, such extensions start from the ‘explanatory top’, so to say.

Oaksford expresses several intuitions regarding the rational status of classical probability theory. For example, he notes that “If one follows the laws of probability in the Kolmogorov axioms, one will never make a bet one is bound to lose, and therefore it can provide a rational standard.” But, as we argued in the target article, an appeal to the Dutch book theorem cannot be used to motivate the rational status of CP theory, since this is only valid with the psychologically unrealistic assumption of utility as value maximization. He also notes that “If one follows the laws of logic one will never fall into contradiction, and therefore logic can provide a rational standard.” But this argument is undermined by Goedel’s incompleteness theorems, since for many axiomatic systems either not all conclusions can be axiomatically derived or, if they can, they will be inconsistent.

Oaksford also observes that classical normative theory is based on clear intuition. For example, according to standard logic and classical probability theory, it is irrational to believe both p and not-p are true. If two events, A versus B, are mutually exclusive, then this is true for QP theory as well and the two events have to be compatible, in the quantum sense. Also, the same event B cannot be both true and not true simultaneously. The subtle issue arises when events A versus B are incompatible, for in this case, knowing A is true prevents us from concluding that B is true or B is not true. If A is incompatible with B, then we cannot discuss their truth values at the same time. Suppose that we model the truth and falsity of a logical statement as incompatible events, though clearly related in a specific way (one could employ suitably opposite probability distributions). Then, the events ‘statement A is true’ and ‘statement A is false’ no longer correspond to orthogonal subspaces and so it is no longer the case that $Prob(\text{statement A is true and then statement A is false})$ has to be 0. By contrast, Blutner, Pothos, & Bruza (in press) showed that the same approach could

never produce *Probability* (statement A is true & statement A is false)

Lee and Vanpaemel present another objection to quantum theory. They note the extent to which a limited number of physicists have objections to quantum theory. They provide a telling quote from Jaynes, in which he strongly questions the value of quantum theory in physics. (However, we recommend reading Bub, 1999, rather than Jaynes, for a more comprehensive interpretation of quantum theory.) To clarify this issue, physicists do not object to the formal (mathematical) form of quantum theory. They debate its interpretation. Our applications to cognition have used the mathematics, and we have avoided taking any stand on the interpretation of quantum theory. Leaving aside the fact that no other physical theory has had such a profound impact in changing our lives (e.g., through the development of the semiconductor and the laser), few if any physicists think that quantum theory is going into retirement soon. For completeness, it is worth noting that Aspect's work famously and definitively supported quantum theory against Einstein's classical interpretation of Bell's hypothetical experiment. Any introductory quantum mechanics text will outline the main ideas (e.g., see Isham, 1989). Quantum theory is a formal theory of probability: it remains one of the most successful in physics and we wish to explore its possible utility in other areas of human endeavor.

In conclusion, the wide variety of thought provoking comments, ranging across criticisms to empirical challenges to debates about fundamental aspects of cognition, attest to Sloman's view that "quantum theory captures deep insights about the workings of the mind" (this is part of his review for Busemeyer & Bruza's, 2012, book).

Additional References

- Aerts, D. & Sozzo, S. (2011). Quantum structure in cognition: why and how concepts are entangled. In *Proceedings of the Quantum Interaction Conference*, 118-129. Springer.
- Alxatib, S., & Pelletier, J. (2011). On the psychology of truth-gaps. In R. Nouwen, R. van Rooij, U. Sauerland & H.-C. Schmitz (Eds.), *Vagueness in Communication* (pp. 13-36). Berlin, Heidelberg: Springer-Verlag.
- Blutner, R., Pothos, E. M., & Bruza, P. (in press). A quantum probability perspective on borderline vagueness. *Topics in Cognitive Science*.
- Bruza, P., Kitto, K., & McEvoy, D. (2008). Entangling words and meaning. In *Proceedings of the Second Quantum Interaction Conference*, pp 118-124. Springer.
- Bub, J. (1999). *Interpreting the Quantum World*. Cambridge University Press.
- Miyamoto, J. M., Gonzalez, R., & Tu, S. (1995). Compositional anomalies in the semantics of evidence. J. Busemeyer, R. Hastie, & D. Medin (Eds.), *Decision making from a cognitive perspective*. (Volume 32 of the *Psychology of Learning and Motivation*). New York: Academic Press.
- Shafir, E. B., Smith, E. E., & Osherson, D. N. (1990). Typicality and reasoning fallacies. *Memory & Cognition*, 18, 229-239.
- Smolensky, P., Goldrick, M., & Mathis, D. (in press). Optimization and quantization in gradient symbol systems: a framework for integrating the continuous and the discrete in cognition. *Cognitive Science*