addressed by having a fixed (but large) dimensionality for all representations. If **HAPPY** is a particular 500-dimensional vector, and **EMPLOYED** is a different 500-dimensional vector, then **HAPPY⊛EMPLOYED** gives a new 500-dimensional vector (a tensor product would give a 250,000-dimensional vector). Importantly, in high-dimensional spaces, the newly created vector is highly likely to be (almost) orthogonal to the original vectors. This gives a close approximation to all of the required orthogonality requirements mentioned in the target article, but does not lead to an unlimited explosion of dimensions as representations get more complicated. As we have shown elsewhere, adult human vocabularies fit well within 500 dimensional spaces (Eliasmith, in press).

Given cognitive theories expressed in terms of vector symbolic architectures, we have created large-scale neural models that implement those theories. In particular, we use the neural engineering framework (NEF; Eliasmith & Anderson 2003), which gives a principled method for determining how realistic spiking neurons can represent vectors, how connections between neurons can implement computations on those vectors, and how recurrent connections can be used to provide memory and other dynamics. This allows us to turn abstract descriptions of cognitive processing into specific brain mechanisms, connecting a plethora of neural data (functional magnetic resonance imaging [fMRI], electroencephalograms [EEG], single cell recordings) to cognitive function.

In the NEF, distributed representations of vectors are made by generalizing the standard notion of each neuron having a particular preferred direction vector (e.g., Georgopoulos et al. 1986). Whereas Hebbian learning rules can be used to adjust connection weights, we can also directly solve for the desired connection weights, as this kind of distributed representation allows a much larger range of functions to be computed in a single layer of connections than is seen in typical connectionist models. This makes it straightforward to create models that accurately compute linear operations (such as the dot product), and even more complex functions such as a full 500-dimensional circular convolution. These models are robust to neuron death and exhibit realistic variability in spiking behavior, tuning curves, and other neural properties.

Although these techniques have not yet been used on the specific tasks and theories presented in the target article, all of the operations mentioned in the article have been implemented and scaled up to human-sized vocabularies (e.g., Eliasmith 2005; Stewart et al. 2011). Furthermore, we have shown how to organize a neural control structure around these components (based on the cortex–basal ganglia–thalamus loop) so as to control the use of these components (e.g., Eliasmith, in press; Stewart & Eliasmith 2011). This architecture can be used to control the process of first projecting the current state onto one vector (**HAPPY**) and then on to another (**EMPLOYED**), before sending the result to the motor system to produce an output. These neural models generate response timing predictions with no parameter tuning (e.g., Stewart et al. 2010), and show how the neural implementation affects overall behavior. For example, the neural approximation of vector normalization explains human behavior on list memory tasks better than the ideal mathematical normalization (Choo & Eliasmith 2010).

Although the NEF provides a neural mechanism for all of the models discussed in the target article, it should be noted that this approach does not require Gleason's theorem, a core assumption of QP (sect. 4.3). That is, in our neural implementations, the probability of deciding one is **HAPPY** can be dependent not only on the length of the projection of the internal state and the ideal **HAPPY** vector, but also on the lengths of the other competing vectors, the number of neurons involved in the representations, and their neural properties, all while maintaining the core behavioral results. Resolving this ambiguity will be a key test of QP.

Why quantum probability does not explain the conjunction fallacy

doi:10.1017/S0140525X12003123

Katya Tentori[a] and Vincenzo Crupi[b]
[a]*Center for Mind/Brain Sciences (CIMeC), Department of Psychology and Cognitive Sciences (PSC), University of Trento, Palazzo Fedrigotti, Corso Bettini, n. 31, 38068 Rovereto, Italy;* [b]*Dipartimento di Filosofia e Scienze, dell'Educazione, Via S. Ottavio 20, 10124 Torino, Italy.*
katya.tentori@unitn.it
vincenzo.crupi@unito.it
http://www.unitn.it/en/cimec/11750/katya-tentori

Abstract: We agree with Pothos & Busemeyer (P&B) that formal tools can be fruitfully employed to model human judgment under uncertainty, including well-known departures from principles of classical probability. However, existing findings either contradict P&B's quantum probability approach or support it to a limited extent. The conjunction fallacy serves as a key illustration of both kinds of problems.

Pothos & Busemeyer's (P&B's) argument in favor of a quantum probability (QP) approach to cognitive modeling relies on the following premises:

1. A large amount of empirical findings accumulated in the last 40 years shows that human judgment often violates key aspects of classical probability (CP) theory.
2. Heuristic-based approaches, although interesting, are often limited in their applicability and explanatory scope.
3. It is possible to model probability judgment on the basis of formal tools, and use these to re-express and sharpen popular heuristics.

We agree with P&B on all the mentioned assumptions. However, we depart from P&B in our assessment of the potential of their QP approach for achieving a better understanding of human judgments under uncertainty. We will illustrate our perspective with reference to the conjunction fallacy (CF) (Tversky & Kahneman 1982; 1983). The CF plays a key role in P&B's argument because P&B claim that this prominent violation of CP laws has a natural and straightforward explanation in their QP approach. In what follows, we will illustrate two main problems that arise with regard to P&B's interpretation of the CF results.

The first problem is that the QP approach is contradicted by empirical data. To begin with, it is unable to accommodate *double* conjunction fallacies (e.g., the *mile run* scenario in Tversky & Kahneman 1983, p. 306) – that is, all those situations in which $h_1 \wedge h_2$ is ranked over each of h_1 and h_2 appearing in isolation (Busemeyer et al. 2011, p. 202). Several single conjunction fallacy results are also demonstrably inconsistent with the QP approach. For example, suppose that, given some evidence e, the most likely statement has to be chosen among three, namely, a single hypothesis h_1 and two conjunctions, $h_1 \wedge h_2$ and $h_1 \wedge \sim h_2$, as in the following scenario:

K. is a Russian woman.	[e]
Which of the following hypotheses do you think is the most probable?	
– K. lives in New York.	[h_1]
– K. lives in New York and is an interpreter.	[$h_1 \wedge h_2$]
– K. lives in New York and is not an interpreter.	[$h_1 \wedge \sim h_2$]

Tentori et al. (2013) observed that $P(h_2|e \wedge h_1) < P(\sim h_2 \,|e \wedge h_1)$ for the majority of participants. One can just as safely assume that $P(h_2|e) < P(\sim h_2|e)$, as clearly only a tiny fraction of Russian women are interpreters. On these assumptions, the QP account of the conjunction fallacy demonstrably predicts that the judged probability of $h_1 \wedge h_2$ must be lower than that of $h_1 \wedge \sim h_2$ (see Fig. 1), and, therefore, that fallacious choices for $h_1 \wedge h_2$ must be less than those for $h_1 \wedge \sim h_2$.

However, in contrast to this prediction, a significant majority (70%) of the fallacious responses in the Russian woman scenario

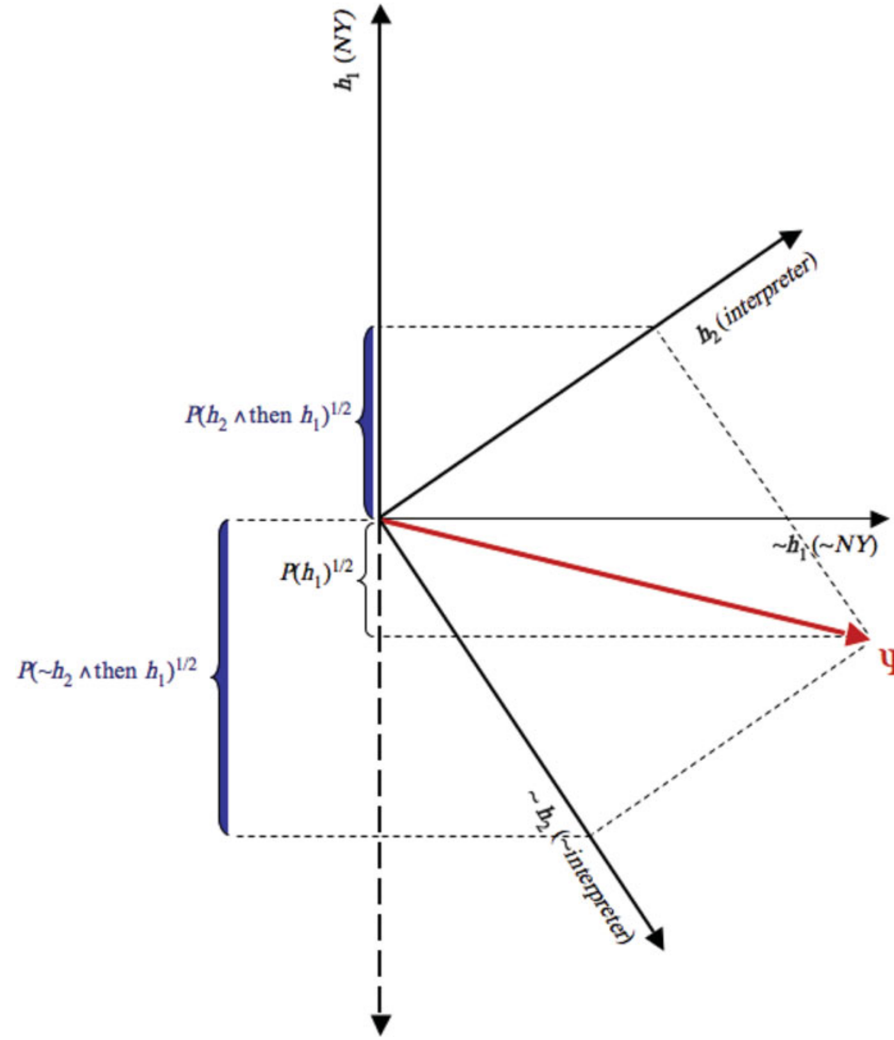

Figure 1 (Tentori & Crupi). A QP representation of the Russian woman scenario. To simplify notation, P is taken to already encode the evidence e (*Russian*). In line with participants' judgment, the basis vectors are displayed as to imply $P(h_2|h_1) < P(\sim h_2|h_1)$, whereas the position of the state vector implies $P(h_2) < P(\sim h_2)$, as clearly only a tiny fraction of Russian women are interpreters. Under these conditions, if we assume – as it seems plausible – that $P(h_1) \geq P(h_2)$, the QP approach does not allow for the conjunctive probability of h_1 and h_2 to rank higher than the single conjunct h_1. However, the QP approach yields the wrong prediction even if we assume $P(h_1) < P(h_2)$. As the figure illustrates, in this case, contrary to the participants' judgment, it is the conjunctive probability of h_1 and $\sim h_2$, not that of h_1 and h_2, which must rank highest.

concerned $h_1 \wedge h_2$ rather than $h_1 \wedge \sim h_2$. More generally, one can prove that, for the QP account, if $P(h_2|e) \leq P(h_3|e)$ and $P(h_2|e \wedge h_1) \leq P(h_3|e \wedge h_1)$, the judged probability of the conjunction $h_1 \wedge h_2$ on the assumption of e cannot be higher than that of $h_1 \wedge h_3$. In a series of four experimental studies (which employed different elicitation procedures, experimental designs, classes of problems, and content), Tentori et al. (2013) documented a robust pattern of results inconsistent with this implication of the QP model, as well as with other recent proposals, such as the averaging (Nilsson et al. 2009) and the random error (Costello 2009) hypotheses, which imply that CF rates should rise as the perceived probability of the added conjunct does. The results supported a different explanation of the CF, based on the notion of inductive confirmation (Crupi et al. 2008). Much like the QP approach, this explanation relies on a well-defined formalism (Bayesian confirmation theory) while avoiding the limitations (e.g., post-hoc parameters) that P&B (sect. 4.1) ascribe to other Bayesian models.

The second problem is that, even when logically consistent with the empirical data, P&B's treatment nonetheless receives limited support. The QP modeler is typically left with a number of choices that are unconstrained by the model itself. Lacking independent and clearly defined empirical input, the modeling exercise does not achieve explanatory relevance. The Linda scenario serves as an illustration. For their QP approach to account for the reported judgment $P(\textit{bank teller} \wedge \textit{feminist}) > P(\textit{bank teller})$, P&B need to make various assumptions (sect. 3.1) on the angle between the basis vectors, as well as on the position of the state vector. Some of these assumptions are uncontroversial. Others are quite subtle, however, and have non-trivial consequences. As the left column of Figure 2 shows, keeping basis vectors equal, a small shift in the position of the state vector is enough to reverse the predicted ranking between $P(\textit{bank teller} \wedge \textit{feminist})$ and $P(\textit{bank teller})$ even if the perceived probability of the *feminist* conjunct remains much higher than that of *bank teller*.

A similar situation arises with the Scandinavia scenario (Tentori et al. 2004):

Suppose we choose at random an individual from the Scandinavian population.	$[e]$
Which do you think is the most probable?	
– The individual has blond hair.	$[h_1]$
– The individual has blond hair and blue eyes.	$[h_1 \wedge h_2]$
– The individual has blond hair and does not have blue eyes.	$[h_1 \wedge \sim h_2]$

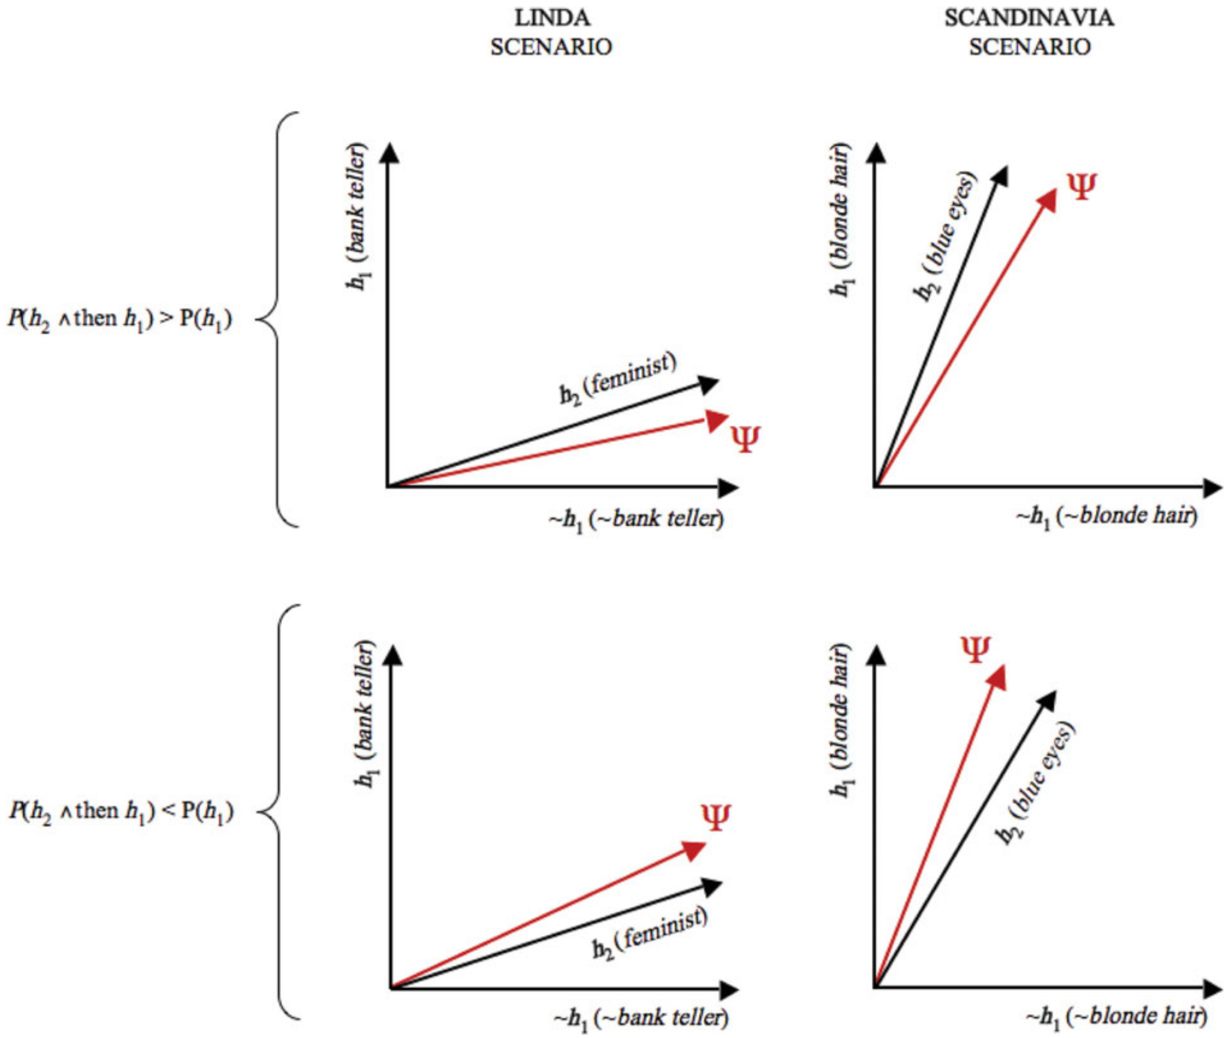

Figure 2 (Tentori & Crupi). Two different plausible QP representations of the Linda and Scandinavia scenarios. The positioning of the vectors in the top half is compatible with the observed conjunction fallacy judgment $P(h_1 \wedge h_2) > P(h_1)$, whereas that in the bottom half is not.

Does the QP approach predict the observed prevalence of fallacious choices for $h_1 \wedge h_2$? This crucially depends on how we determine the vector space.

The right column of Figure 2 shows two different representations of the Scandinavia scenario (the state vector and the h_2 basis vector are simply switched). The QP approach allows for both, and both appear reasonable. However, yet again, in the two cases, opposite orderings of $P(h_1 \wedge h_2)$ and $P(h_1)$ follow from the QP approach. For an observed judgment to be taken as properly supporting the QP explanation of the CF, the corresponding vector space representation needs to be constrained on independently motivated empirical grounds. Otherwise, we can only say that the QP approach can be made compatible with some (and not all) of the CF data. However, for a putatively comprehensive theoretical framework, being able to accommodate some empirical results does not equal predicting and explaining a phenomenon.

ACKNOWLEDGMENTS

Research relevant to this work has been supported by the MIUR grant "Problem solving and decision making: Logical, psychological and neuroscientific aspects within criminal justice" (PRIN, n.2010RP5RNM_006) and by Grant CR 409/1-1 from the Deutsche Forshungsgemeinshaft (DFG) as part of the prority program New Frameworks of Rationality (SPP 1516).

Authors' Response

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

doi:10.1017/S0140525X12003226

Emmanuel M. Pothos[a] and Jerome R. Busemeyer[b]

[a]*Department of Psychology, City University London, London EC1V 0HB, United Kingdom;* [b]*Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405.*

emmanuel.pothos.1@city.ac.uk
jbusemey@indiana.edu
http://www.staff.city.ac.uk/~sbbh932/
http://mypage.iu.edu/~jbusemey/home.html

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,