

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Towards an Empirical test of Realism in Cognition

Permalink

<https://escholarship.org/uc/item/23m8f70t>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

Authors

Yearsley, James M
Pothos, Emmanuel M

Publication Date

2015

Peer reviewed

Towards an Empirical test of Realism in Cognition

James M. Yearsley (james.yearsley.1@city.ac.uk)

Emmanuel M. Pothos (emmanuel.pothos.1@city.ac.uk)

Department of Psychology, City University London
London, EC1V 0HB UK

Abstract

We discuss recent progress towards an empirical test of ‘realism’ in cognition, ‘realism’ in this context being the property that cognitive variables always have well defined (if unknown) values at all times. Our main result is an inequality obeyed by realist theories, which could be tested by a suitable experiment. We focus our attention in this contribution on two particular issues. The first is the exact notion of realism which is to be tested, as this has received less attention in earlier work. The second is an important technical issue about the inequality we use; in earlier work Atmanspacher and Filk (2010) considered a different expression, and we explain why our inequality is more suitable for use under realistic experimental conditions.

Keywords: cognition; quantum probability; time perception; memory.

Introduction

In this contribution we give an overview of recent work that seeks to address the question of whether models of cognitive processes can be realist, in relation to their time evolution. We will define exactly what we mean by realist below, but the key finding is that given a suitable definition this is an empirical question. The interest in this question arises in part from attempts to model certain aspects of cognition using quantum probability theory (QT), which is famously non-realist (for a review see Pothos and Bussemeyer (2013)). Experimental evidence of a violation of realism would therefore provide strong evidence for the suitability of the QT approach.

This contribution (see also Yearsley and Pothos (2014)), focuses on two particular issues. The first is the exact notion of ‘realism’ being tested, since the notion of realism itself is a novel one in cognitive psychology. This is particularly pressing because the empirical test we propose, which has been discussed before (Atmanspacher and Filk, 2010), is borrowed from the physics literature, and it is not immediately clear how this test is to be interpreted in the context of cognitive models. The second issue we shall address concerns what is at first sight a rather technical difference between our test and an earlier one by Atmanspacher and Filk (2010). However the difference is important because our test is suitable for application in a realistic experimental set up, with noisy or imperfect measurements.

To save space and allow for the maximum of conceptual discussion we omit technical details in this contribution; interested readers may consult Yearsley and Pothos (2014).

The rest of this contribution is structured as follows; in Section 2 we discuss the notion of realism in cognitive models in a general way and in Sections 2 and 3 we introduce the two smaller assumptions that together make up the assumption of

realism proper. In Section 4 we make some very brief comments on the empirical test of realism we propose, and in Section 6 we explain how our test differs from earlier work. In Section 7 we then briefly discuss the options for cognitive modelling should our empirical test rule out realism. We conclude in Section 8.

Realism in Cognitive Models

Ours thoughts, feelings, memories and decisions arise ultimately from the physical matter of our brains. It is generally assumed that if it were possible to know exactly the physical specification of a person’s brain at any moment of time, we should also be able to tell what that person was feeling and predict their judgments. Of course, this is the stuff of science fiction rather than current neuroscience. However the key principle, that the behaviour of cognitive variables such as feelings and judgments can be reduced to statements about the physical specification of the brain does manifest itself in an important way in current cognitive models. In brief, most cognitive models have a property we term ‘realism’ - it is assumed in these models that cognitive variables have definite values at all times (cf Raijmakers and Molenaar, 2004).

This assumption arises in a natural way when we consider the link between cognitive processes at the level of thoughts and feelings, and the underlying neurophysiology of the brain which we assume gives rise to these thoughts and feelings. For the purposes of this contribution we assume the most fundamental processes in the brain relevant for cognition may be described by classical physics. It is a key feature of classical physics that the positions, electric charges, etc of all classical particles are definite at all times, that is, whilst the values of these quantities may be *subjectively* uncertain (since we have only limited knowledge of them) they are nevertheless *objectively* certain. Thus, one can argue, if cognition is ultimately determined by brain neurophysiology, and if the most fundamental variables at the neurophysiological level have definite but unknown values, then presumably all cognitive variables must also have definite if unknown values. We will argue that this assumption is in fact highly questionable.

To make the argument more concrete consider a simple example of a subject’s preference for crisps over chocolate. At any given moment of time our subject will have a preference for either crisps over chocolate, or vice versa. Let us denote this cognitive variable by the function $C(t)$, which takes values between +1 (definitely prefer crisps) and -1 (definitely prefer chocolate.) The key assumption of classical models of cognition is that the variable $C(t)$ has a well defined value

at all times. This may seem reasonable over the course of some short lab experiment, but does it really make sense over longer periods of time? What happens if we distract the subject so that they are no longer thinking about food? Does this cognitive variable still possess a well defined value?

The alternative to such a realist account of cognitive variables is one where these quantities do not possess values until they are measured, that is, measuring such cognitive variables is a *constructive* process. The idea that measuring the value of a cognitive variable can change that value has been considered before (see e.g. Ariely and Norton, 2008), and can be easily incorporated within classical models of cognition. However what we are suggesting here is something different, it is not a question of measurements changing the values of *existing* quantities, rather the process of measurement *creates* those values, where previously there were none. Indeed it is not so hard to imagine that a subject's preference for crisps over chocolate simply isn't *defined* at time when they are not, consciously or unconsciously, thinking about eating. Thus it seems reasonable that there are at least *some* times when this cognitive variable is not defined.

Should we care whether our cognitive theories are realist or not? We argue that we should. There are two main reasons; firstly it turns out there are certain types of behaviour possible in non-realist models which are impossible in realist ones. This means there are limits to the types of cognitive processes realist theories are able to describe. The second reason is potentially more important; our aim in constructing cognitive models is not just to describe or even predict cognitive processes, but at some level to understand them. For this reason it is important to have confidence that the structural features of our models match or map in some sense the way cognition happens in the brain. Although our understanding of the physiology of cognition is currently far too limited to be used to impose detailed constraints on cognitive models, there are nevertheless some basic constraints that we can impose that do limit the classes of cognitive models we should consider acceptable. One of these concerns the idea of 'bounded cognition'; we would argue a second one concerns realism (for some work in this direction see Jones and Love, 2011.) In summary, 'realism' is a property that we can choose to include, or not, in our models of cognition. The way cognition happens in the brain may or may not display this property, and if we try to model non-realist processes using realist models then we may be severely limiting our ability to construct accurate and faithful models of cognition.

Is it possible to prove whether cognition is realist or not? We hinted at the answer earlier; there are some types of behaviour that are impossible to reproduce within a realist cognitive model. In this contribution we will outline a test which allows us to determine whether a given set of judgments can be described by a realist cognitive model¹. Before we do this

¹Of course, many cognitive models are concerned with variables other than judgment outcomes, e.g. reaction times, error rates, neural activity levels etc. It is an open question whether we can be realist about these, our test does not immediately apply to these variables.

however we need to be clearer about exactly what we mean by realism in cognitive models.

The next two sections introduce two reasonable assumptions which together form the joint assumption of 'realism' in cognitive models. We will spend some time discussing these assumptions in depth, because they are really the most important part of this work. Since our test of 'realism' is really a joint test of these assumptions its significance depends entirely on whether one believes these assumptions really capture the correct notion of realism in cognition. In addition, because there are two separate assumptions any purported failure of 'realism' leaves us the option of retaining one of them. If we want to understand which one (if either) we should retain we need first be clear on their meaning. Once we have done this our empirical test of realism follows by some algebra, which we shall skip, interested readers are invited to consult Yearsley and Pothos (2014).

Realism Part 1: Cognitive Realism

Let us set out our first assumption which, together with the assumption discussed in the next section, together define 'realism' in cognitive models.

Cognitive Realism: The reason for any judgement at the cognitive level is ultimately (in principle, if not in practice) reducible to processes at the neurophysiological level.

Cognitive Realism is perhaps what one might think of if one is asked to characterise realism. Indeed it might seem like this completely captures the notion of realism, we will explain why this is not the case below. For now let us instead introduce some notation to help put this assumption on the required mathematical footing needed for our empirical test. Consider again our example cognitive variable $C(t)$. Let us denote the complete neurophysiological state of a given subject as λ . Cognitive Realism means that there is a function which, given that the neurophysiological state of the subject is initially λ , will tell us the value of $C(t)$, let us denote this by $c(\lambda, t)$. This is what we mean when we say that realism means that, in principle, were we to know the physical state of a subject's brain (λ) we would know all their feelings and be able to predict their judgments ($C(t)$). However in practice of course we cannot know a subject's exact neurophysiological state, the best we can do is give some probability distribution based on the limited knowledge we do have. Let us denote the probability distribution representing our knowledge of a subject's λ as $\rho(\lambda)$. Then our best guess about the value of $C(t)$ given our knowledge of the neurophysiological state is,

$$\langle C(t) \rangle = \sum_{\lambda} c(\lambda, t) \rho(\lambda), \quad (1)$$

that is, the expected value of $C(t)$ is just the expectation value of $c(\lambda, t)$, given the probability distribution $\rho(\lambda)$.

Let us make few comments about this assumption, and its mathematical consequence Eq.(1).

- The observant reader may find the time dependence in Eq.(1) rather odd, in that it is contained in the cognitive variable rather than in the distribution over neurophysiological states. This is purely notational, the current notation fits present purposes better.
- There is no expectation that we *can* know a subject's λ , and also no requirement that we know the function $c(\lambda, t)$. Even if Cognitive Realism is true a subject's λ need not be knowable even in principle, but the λ 's should be well defined and $c(\lambda, t)$ must exist.
- Suppose we have unreliable measurements. That is, there is some uncertainty in the measured value of $C(t)$ that comes from the measurement process, and not from $\rho(\lambda)$. This can be modelled by via $c(\lambda, t)$, so that even if the cognitive variable can only take values, e.g. ± 1 , $c(\lambda, t)$ need not equal ± 1 for all λ, t .
- It is very difficult to see how this assumption could fail to be valid at some level. After all, if the values of cognitive variables are not determined by the brain, what are they determined by?

It may seem as though Cognitive Realism totally captures the notion of realism in cognition. However there is something missing from the discussion so far. What one does in practice when constructing a cognitive model is not to consider all possible cognitive variables, but rather to model only a tiny subset of them. Cognitive Realism by itself does not guarantee that we can do this in a self consistent way. We need a guarantee that we can take finite collections of cognitive variables and model them without having to continually reference the neurophysiological state. In other words, Cognitive Realism is the assumption that the cognitive level can be connected to the neurophysiological level; what we also need is an assumption that the cognitive level can be *disconnected* from the neurophysiological level, and modelled on its own (for related ideas see Marr, 1982). That is the content of our second assumption.

Realism Part 2: Cognitive Completeness

Our second assumption is a little harder to state than our first. It concerns the cognitive state of a subject. This is defined to be the object within a given cognitive model that encodes the information needed to make predictions about a particular subject (or group of subjects.) For dynamical models it is equivalent to a set of initial conditions. The cognitive state is therefore equivalent to an exhaustive set of probabilities for future measurement outcomes within the context of a particular model². The exact form the cognitive state takes will depend on the model, and we want to state our assumption without reference to any particular form.

²The idea of defining the state of a system in this way occurs frequently in physics, see e.g. Hardy (2001).

Cognitive Completeness: The probabilities for the outcomes of any judgment within a cognitive model may be expressed in terms of a so-called cognitive state. The cognitive state of a person responding to such a set of judgements can be determined by a finite set of probabilities for the judgement outcomes.

This assumption comes in two parts. The first part is simply the assumption that we can write the probabilities for judgment outcomes in terms of an object we call the cognitive state. This is just an algebraic step and is in fact always true, but it is useful to include it as part of the assumption anyway³. The second part of this assumption is the crucial part, here we demand that not only does the cognitive state encode the probabilities for judgment outcomes, but it can itself be determined from a finite set of them. That is, we assume observing participant behaviour can fully determine the underlying cognitive state, without the need to invoke neurophysiological variables. The 'finite number' caveat ensures that we need only a finite number of judgment outcomes to determine the cognitive state, i.e. we might need to observe participants make 3,4 or 100 judgments, but after some point we can say we have enough information to begin to *predict* future judgments.

The reason this assumption is more vague than the first is that we have not defined exactly what the cognitive state is supposed to be. Generally this will depend on the model. However whatever the form of the cognitive state, if this is the object that allows us to predict judgment outcomes then it is important that it can be determined entirely in terms of them, otherwise it is not possible to establish this state empirically, making prediction impossible.

This assumption has an important consequence. Consider any measurement made on a group of participants that does not change the probabilities for the outcomes of any future judgement in the relevant cognitive model. Let us call such measurements non-disturbing. Whether a given measurement is non-disturbing can always be established empirically.

Cognitive Completeness means that, as long as a measurement is non-disturbing, it can be assumed to have no effect on the neurophysiological state of a participant. This is because Cognitive Completeness tells us that the cognitive state of the participants may be fully determined by knowledge of the outcomes of all judgements in the relevant cognitive model. Thus, at most, a non-disturbing measurement may change the underlying neurophysiological state in a way that gives rise to the same cognitive state. However, any such change is undetectable by any measurement relevant to the cognitive model, and thus for the purposes of our model we may assume no change in the neurophysiological state occurred.

It is useful to express this in a more mathematical way. Cognitive Completeness means that every cognitive model

³Crucially we do not assume that the cognitive state is a probability distribution over judgment outcomes, this is a non-trivial assumption, see below.

defines a set of similarity classes on the set of all probability distributions over the neuropsychological variables, with two distributions $\rho(\lambda)$ and $\rho'(\lambda)$ being similar, $\rho(\lambda) \sim \rho'(\lambda)$, if they lead to the same predictions for all judgements contained in the cognitive model. In general measurement of the cognitive variable $C(t_1)$ at t_1 will change the distribution of neurophysiological variables so that a subsequent measurement of, e.g. $C(t_2)$ with $t_2 > t_1$, will depend on whether or not the first measurement was made. Denote the new distribution over the λ after measurement at t_1 as $\rho(\lambda; t_1)$. Then joint measurement of $C(t_1)$ and $C(t_2)$ yields,

$$\langle C(t_1)C(t_2) \rangle = \sum_{\lambda} c(\lambda, t_1)c(\lambda, t_2)\rho(\lambda; t_1) \quad (2)$$

However if the measurement at t_1 was non-disturbing this is equal to

$$\langle C(t_1)C(t_2) \rangle = \sum_{\lambda} c(\lambda, t_1)c(\lambda, t_2)\rho(\lambda). \quad (3)$$

This is the mathematical result used in the derivation of our empirical test. For details see Yearsley and Pothos (2014).

Given the assumptions of Cognitive Realism and Cognitive Completeness we may show that the cognitive state is in this case equivalent to a probability distribution over the cognitive variables. Details are given in Yearsley and Pothos (2014b).

The Empirical Test

Now we have stated our assumptions we describe how they can be empirically tested. The test takes the form of a set of inequalities satisfied by realist systems but which may be violated by non-realist ones. These inequalities may be derived from the mathematical expressions of Cognitive Realism and Cognitive Completeness (for the derivation see the appendix of Yearsley and Pothos (2014)). We quote the result in terms of our example variable $C(t)$, which recall takes values ± 1 .

$$|\langle C(t_1)C(t_2) \rangle + \langle C(t_2)C(t_3) \rangle + \langle C(t_3)C(t_4) \rangle - \langle C(t_1)C(t_4) \rangle| \leq 2 \quad (4)$$

Eq.(4) is one of a collection of inequalities known as the temporal Bell⁴, or Leggett-Garg inequalities, first derived as constraints on physical systems by Leggett and Garg (1985).

What would a concrete experimental set up to test these inequalities look like? Firstly we require a cognitive variable which takes two distinct values and whose expected value we can manipulate. Depending on the variable we use it may be more accurate to think of the 't' variable in Eq.(4) as a parameter rather than as a physical time. The next ingredient is a reliably non-disturbing measurement process, in the sense outlined above. This might be hard to invent in general, but it is easy to establish whether a given measurement process is non-disturbing, so it presents no problem in principle. We mention in passing that it is not necessary that the measurement process be completely non-disturbing, being able to

⁴This name comes from the similarity between these expressions and the usual Bell inequalities (see e.g. Bell (2004)).

bound the disturbance to some low level is sufficient (Yearsley and Pothos (in preparation)).

The final ingredient is a cognitive variable which we expect to behave in a non-classical way. It may be possible to use variables which have previously been shown to behave in non-classical ways, such as the ones investigated in quantum cognitive models (see e.g. Busemeyer et al (2011), Pothos and Busemeyer (2009), Trueblood and Busemeyer (2011), Wang and Busemeyer (2013)), or variables which have previously been seen to deviate from the prescription of classical probability theory (e.g. Tversky and Kahneman, 1983).

Some Comments on an Earlier Inequality of Atmanspacher & Filk

We wish to comment briefly on the difference between our inequality, Eq.(4) and an earlier one proposed by Atmanspacher and Filk (2010). The inequality they considered was⁵,

$$\langle C(t_1)C(t_2) \rangle + \langle C(t_2)C(t_3) \rangle - \langle C(t_1)C(t_3) \rangle \leq 1 \quad (5)$$

Roughly this expression is analogous to Bell's original formulation of the Bell inequalities (Bell, 2004), while our expression is analogous to the CHSH inequalities (Clauser et al, 1969.) On the face of it Atmanspacher and Filk's expression seems more attractive, since it requires measurements at only three rather than four times. Why then should we prefer our expression? We can see the problem with Eq.(5) by considering how one might derive it from Eq.(4). We can do this in two steps, firstly we choose $t_4 = t_3$. With this choice of times Eq.(4) includes the term,

$$\langle C(t_3)C(t_3) \rangle = \sum_{\lambda} c(\lambda, t_3)^2 \rho(\lambda). \quad (6)$$

In order to arrive at Atmanspacher and Filk's expression we need to assume that Eq.(6) is equal to 1⁶. The justification is that this is the correlation between two measurements performed immediately after one another. Whatever the result of the first experiment (± 1), performing it again immediately should yield the same result, therefore the correlation should be perfect. However this justification fails in general because for a realistic experiment there will be unavoidable noise in the measurement of a judgment. This means that even if we knew a participant's λ exactly we could predict only probabilities for the measurement outcomes. In terms of the mathematics this means we have $-1 < c(\lambda, t) < 1$, so that $c(\lambda, t)^2 < 1$. Thus Eq.(6) will be less than 1 and Eq.(5) will not hold for realistic experiments.

We have seen that for realistic experiments with unavoidable experimental noise, our inequality is a better place to look for violations of realism than the original one proposed by Atmanspacher and Filk (2010). This should serve as a

⁵Atmanspacher and Filk actually gave their expression in terms of a different variable.

⁶Atmanspacher and Filk (2010) did not obtain their inequality this way. However the direct derivation of their expression still requires the right hand side of Eq.(6) equal 1 (see Bell 2004 for details.)

cautionary tale not to oversimplify the analysis in the effort to make the modelling look more appealing. As the quote attributed to Einstein⁷ goes, “Everything should be made as simple as possible, but not simpler.”

What Should we Conclude if ‘Realism’ Fails?

Suppose we were to conduct a test of realism in the way outlined above and find a convincing violation of Eq.(4). What should we conclude? Assuming one accepts the arguments which lead to Eq.(4) then the only conclusion is that one or both of our assumptions, Cognitive Realism and Cognitive Completeness, must be incorrect. But which one?

It might seem as though one could salvage some notion of realism at the price of dropping Cognitive Completeness. The problem with this approach is that it is Cognitive Completeness that means that the cognitive state can be empirically determined, and thus it is Cognitive Completeness that ensures that any model has genuine predictive power.

Nevertheless one might argue that this problem can be circumvented. If we cannot fix the cognitive state in terms of the outcomes of judgments contained in our cognitive model, can we not simply add more judgments, the probabilities for which *would* be enough to fix the cognitive state? The answer is that we cannot. The full argument is given in Yearsley and Pothos (2014), but the essence is that adding cognitive variables which can be measured in a non-disturbing way simply gives an extended cognitive model from which the original one can be recovered by coarse-graining, but since the original model isn’t realist the extended one cannot be either. Adding in cognitive variables which cannot be measured in a non-disturbing way solves this problem, but having cognitive variables which cannot *in principle* be measured in a non-disturbing way means the new model still lacks predictive power. In summary, Cognitive Completeness is possibly even more central to cognitive modelling than realism.

So if we cannot drop Cognitive Completeness, can we drop Cognitive Realism? With certain caveats, discussed below, the answer is yes. Instead of using classical theories, we can model cognition with non-realist theories like quantum probability theory (QT) that include a constructive role for judgment. QT is often described as quantum theory without the physics (see e.g. Aerts and Aerts, 1995, Atmanspacher et al, 2006), and is potentially applicable in any situation where there is a need to quantify uncertainty (see e.g. beim Graben and Atmanspacher, 2009). Indeed there has been no small measure of success modelling some aspects of cognition in this way (e.g. Busemeyer et al, 2011, Pothos and Busemeyer, 2009, Trueblood and Busemeyer, 2011, Wang and Busemeyer, 2013, Bruza et al, 2009. For an overview see Busemeyer and Bruza, 2011, Pothos and Busemeyer, 2013.) We note in passing that QT does satisfy Cognitive Completeness; the cognitive state is the quantum state $|\psi\rangle$ and can be determined via a finite number of probabilities for judgments.

However we need to be cautious. There are theories other

⁷Alas probably apocryphal.

than QT which could account for a violation of our test of realism. Indeed our experimental test of realism could be used to rule out not just realist theories of cognition but also quantum ones!⁸ Even if the results of our test are in agreement with QT our test may rule out realist models of cognition, but it cannot ‘rule in’ quantum models. We need to search elsewhere for convincing evidence for the correctness of quantum approaches to cognition (one promising possibility is the q-test developed by Wang and Busemeyer (2013)).

It is also important to note that, unlike the corresponding case in physics, rejecting the idea that cognitive variables can be modelled in a realist way does not force us to reject the claim that these judgments or feelings arise from neurophysiology. This is a subtle point, but the essence is that the values of cognitive variables, i.e. thoughts, feelings, judgments, are brought into being by the interplay between both the neurophysiological state and the process of measurement. This is what we mean when we say that judgment is a constructive process in these non-realist theories.

Finally we should mention that a failure of realism in cognition could have great significance for models of memory. If judgment is a constructive process then it is easy to imagine that memory retrieval may also be modelled constructively in a similar way (this has been suggested before, e.g. Howe and Courage (1997)). This could open up exciting new possibilities for modelling memory processes.

Conclusion

We have discussed some of the conceptual issues involved in recent work on the question of realism in cognitive models.

What can we conclude from this discussion? We have argued that the standard notion of realism in cognition might be well motivated, but it is open to empirical challenge. The success of the QT programme to date suggests, although it does not prove, that realism may have to be abandoned as an assumption in models of cognition. The proposed empirical test will hopefully bring us closer to resolving the issue. This test is tricky to implement, but should be possible with the right choice of cognitive variable and measurement.

If our tests do rule out realism, this is not by itself reason to adopt QT models of cognition. However such models can give valuable insight into more general non-realist approaches. In particular contextually and constructive judgments are central parts of QT (Kitto, 2008, Busemeyer and Bruza, 2011, White et al, 2014) and these will also be key features of any non-realist theory, quantum or otherwise.

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 Materiel Command, USAF, grants no. FA 8655-13-1-3044. The US Government is authorized to reproduce and distribute

⁸An analogue of Eq.(4) holds in quantum theory, but with the bound on the right hand side equal to $2\sqrt{2}$, see Tsirelson, 1980.

reprints for Governmental purpose notwithstanding any copyright notation thereon.

References

- Aerts, D. & Aerts, S. (1995). Applications of quantum statistics in psychological studies of decision processes. *Foundations of Science*, 1, pp. 85-97.
- Ariely, D & Norton, M.I., (2007). How actions create - not just reveal - preferences. *Trends in Cognitive Sciences*, 12, 13.
- Atmanspacher, H., Romer, H., & Wallach, H. (2006). Weak quantum theory: formal framework and selected applications. *Weak quantum theory: complementarity and entanglement in physics and beyond. Foundations of Physics*, 32, 379-406.
- Atmanspacher, H., & Filk, T. (2010). A proposed test of temporal nonlocality in bistable perception. *Journal of Mathematical Psychology*, 54, 314-321.
- Bell, J.S. (2004). *Speakable and Unspeakable in Quantum Mechanics*. Cambridge University Press: Cambridge, UK.
- Bruza, P. D., Kitto, K., Nelson, D. & McEvoy, C. L. (2009). Is there something quantum-like about the human mental lexicon? *Journal of Mathematical Psychology*, vol. 53, pp. 362-377.
- Bruza, P.D., Kitto, K., Ramm, J.R. & Sitbon, L. (2013) A probabilistic framework for analysing the compositionality of conceptual combinations, arXiv preprint arXiv:1305.5753 available at arxiv.org.
- beim Graben, P. & Atmanspacher, H. (2009). Extending the philosophical significance of the idea of complementarity. In H. Atmanspacher & H. Primas (Eds.). *Recasting Reality. Wolfgang Pauli's Philosophical Ideas and Contemporary Science* (pp. 99-113). Berlin: Springer.
- Busemeyer, J. R. & Bruza, P. (2011). *Quantum models of cognition and decision*. Cambridge University Press: Cambridge, UK.
- Busemeyer, J. R., Pothos, E. M., Franco, R., & Trueblood, J. (2011). A quantum theoretical explanation for probability judgment errors. *Psychological Review*, 118, 193-218.
- Clauser, J.F., Horne, M.A., Shimony, A. & Holt, R.A. (1969) Proposed Experiment to test Local Hidden-Variable Theories. *Phys.Rev.Lett.* 23, 880.
- George, R, Robledo, L, Maroney, O, Blok, M, Bernien, H, Markham, M, Twitchen, D, Morton, J, Briggs, A & Hanson, R (2013). *Proceedings of the National Academy of Sciences*, 110, 3777-3781.
- Hardy, L. (2001). Quantum theory from five reasonable axioms. arXiv preprint arXiv:quant-ph/0101012 available at arxiv.org.
- Howe, M. L. & Courage, M. L. (1997). The emergence and early development of autobiographical memory. *Psychological Review*, 104, 499-523.
- Jones, M. & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34, 169, 231.
- Kitto, K., (2008). *Why Quantum Theory?*, *Proceedings of the Second Quantum Interaction Symposium*, pp. 11-18, College Publications.
- Leggett, A. J. & Garg, A. (1985). Quantum mechanics versus macroscopic realism: is the flux there when nobody looks? *Physical Review Letters*, 54, 857-860.
- Marr, D. (1982). *Vision: a computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.
- Pothos, E. M. & Busemeyer, J. R. (2009). A quantum probability explanation for violations of 'rational' decision theory. *Proc. R. Soc. B*, 276, 2171-2178.
- Pothos, E. M. & Busemeyer, J. R. (2013). Can quantum probability provide a new direction for cognitive modelling? *Behavioral & Brain Sciences*, 36, 255-327.
- Raijmakers, M. E. J. & Molenaar, P. C. M. (2004). Modelling developmental transitions in adaptive resonance theory. *Developmental Science*, 7, 149-157.
- Shafir, E. & Tversky, A. (1992). Thinking through uncertainty: nonconsequential reasoning and choice. *Cognitive Psychology*, 24, 449-474.
- Shiffrin, R. M. (1970). Forgetting, trace erosion or retrieval failure? *Science*, 168, 1601-1603.
- Tenenbaum, J. B, Kemp, C., Griffiths, T. L., & Goodman, N. (2011). How to grow a mind: statistics, structure, and abstraction. *Science*, 331, 1279-1285.
- Trueblood, J. S. & Busemeyer, J. R. (2011). A comparison of the belief-adjustment model and the quantum inference model as explanations of order effects in human inference. *Cognitive Science*, 35, 1518-1552.
- Trueblood, J. S., Brown, S. D., Heathcote, A., & Busemeyer, J. R. (in press). Not just for consumers: context effects are fundamental to decision-making. *Psych. Science*.
- Tsirelson, B.S. (1980). Quantum Generalizations of Bell's Inequality. *Letters in Mathematical Physics*, 4, 93.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunctive fallacy in probability judgment. *Psychological Review*, 90, 293-315.
- Wang, Z., & Busemeyer, J. R. (2013). A quantum question order model supported by empirical tests of an a priori and precise prediction. *TICS*, 5, 689-710.
- White, L. C., Pothos, E. M., & Busemeyer, J. R. (2014). Sometimes it does hurt to ask: the constructive role of articulating impressions. *Cognition* 133, 48.
- Yearsley, J. M. (2013). The Leggett-Garg inequalities and non-invasive measurability. Pre-print available at arXiv.org/abs/1310.2149
- Yearsley, J.M. and Pothos E.M. (2014). Challenging the classical notion of time in cognition: a quantum perspective. *Proc. R. Soc. B* 281, 1781, 20133056.
- Yearsley J.M. and Pothos E.M. (in preparation). An Empirical Test of Temporal Realism in Cognition.