G Quantum probability in cognition

There are interesting connections between the mathematics of quantum mechanics and information processing in the brain. This is not so remarkable when we recall that the foundations of quantum mechanics are matters of information and knowledge.[28]

G.1 Theories of decision making

How do people make decisions under uncertainty? There have been three major phases of models.

(i) Logic. From Aristotle's time, the most common model of human thinking has been formal logic, especially deductive logic. For example, the title of George Boole's book, in which he introduced Boolean algebra, was called *The Laws of Thought*, and that is what he supposed it to be. The first AI program (1956) was the LOGIC THEORIST, and formal deductive logic still dominates many AI systems. Since the 1960s, however, there has been accumulating psychological evidence that classical logic is not a good model of everyday reasoning.

An additional, more technical problem is that classical logic is *monotonic*, that is, the body of derived theorems can only increase. But everyday reasoning is *nonmonotonic*: propositions that were previously taken to be true can become false (either because the facts have changed or an assumption has been invalidated. As a consequence, the body of truths can shrink or change in other ways. Existing truths can be nullified. An additional problem is that much of our reasoning is *inductive* rather than deductive, that is, it moves from more particular premises to more general conclusions, rather than vice versa, as deductive logic does. But after many years of research, there really isn't an adequate inductive logic that accounts for scientific reasoning as well as everyday generalization.

(ii) Classical probability (CP). The most common models of human decision making have been based on classical probability theory (CP) and Bayesian inference. Amos Tversky and Daniel Kahneman were pioneers (from the 1970s) in the study of how people actually make decisions and judgments. (In 2002 Kahneman recieved the Nobel Prize in Economics for

[28]This chapter is based primarily on Emmanuel M. Pothos and Jerome R. Busemeyer, "Can quantum probability provide a new direction for cognitive modeling?" *Behavioral and Brain Sciences* (Pothos & Busemeyer, 2013), including MacLennan (2013).

this work; Tversky had already died.) Since then many other psychologists have confirmed and extended their findings. They concluded that everyday human reasoning follows the laws of neither classical logic nor classical probability theory. "Many of these findings relate to order/context effects, violations of the law of total probability (which is fundamental to Bayesian modeling), and failures of compositionality." (Pothos & Busemeyer, 2013)

(iii) Quantum probability (QP). The mathematics of quantum mechanics provides an alternative system (axiomatization) of probability which has the potential to account for these violations of CP, as we will see. Note that QP is just a probability theory; there is no presumption that physical quantum phenomena are significant in the brain (although they might be). In this sense, our brains appear to be using a kind of quantum computation.

G.2 Framework

G.2.a Questions & outcomes

Just as CP begins by defining a sample space, QP begins by defining a Hilbert space, which defines all possible answers that could be produced for all possible questions (addressed by the model). Corresponding to the *quantum state* is the *cognitive state*, which you can think of as the indeterminate state of the brain before there is a decision or determination to act in some way (such as answering a question). Corresponding to *observables* in quantum mechanics, we have *questions* in QP. More generally, we might refer to *quandries*, that is, unsettled dispositions to act. Corresponding to *projectors* into *subspaces* we have *decisions*. Often the subspaces are one-dimensional, that is, *rays*.

G.2.b Happiness example

Consider asking a person whether they are happy or not. Before asking the question, they might be in an indefinite (superposition) state (Fig. III.43(a)):

$$|\psi\rangle = a|\text{happy}\rangle + b|\text{unhappy}\rangle.$$

It is not just that we do not know whether the person is happy or not; rather the person "is in an indefinite state regarding happiness, simultaneously entertaining both possibilities, but being uncommitted to either" (Pothos & Busemeyer, 2013). More realistically "happy" and "unhappy" are likely to

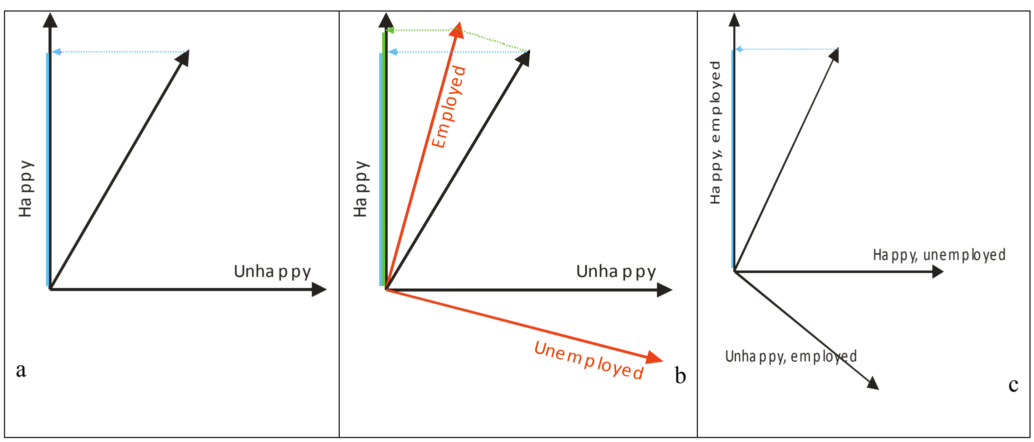

Figure III.43: “An illustration of basic processes in QP theory. In Figure [b], all vectors are coplanar, and the figure is a two-dimensional one. In Figure [c], the three vectors ‘Happy, employed’, ‘Happy, unemployed’, and ‘Unhappy, employed’ are all orthogonal to each other, so that the figure is a three-dimensional one. (The fourth dimension, ‘unhappy, unemployed’ is not shown).” (Pothos & Busemeyer, 2013) be complex subspaces, not rays, but for the sake of the example, we use a 2D outcome space.

Asking the question is equivalent to measuring the state in the "happiness basis," which comprises two projectors P_{happy} and P_{unhappy}:

$$\begin{aligned} P_{\text{happy}} &= |\text{happy}\rangle\langle\text{happy}|, \\ P_{\text{unhappy}} &= |\text{unhappy}\rangle\langle\text{unhappy}|. \end{aligned}$$

The probability that the person responds "happy" is, as expected:

$$\|P_{\text{happy}}|\psi\rangle\|^2 = \||\text{happy}\rangle\langle\text{happy} \mid \psi\rangle\|^2 = |a|^2.$$

Measurement (decision) collapses the indefinite state to a definite basis state, $|\text{happy}\rangle$, with probability $|a|^2$. The judgment or decision is not just a "read out"; it is *constructed* from the state and the question, which actively disambiguates the superposition state.

G.2.c Incompatibility

As in quantum mechanics, questions can be *compatible* or *incompatible.* In fact, Neils Bohr borrowed the notion of *incompatible questions* from the psychologist William James. Compatible questions can be asked in any order; they commute; incompatible questions do not commute.

In CP it is always possible to specify a joint probability distribution over the four possible pairs of answers (the *unicity principle*). In QP you can do this for compatible questions, but not for incompatible ones. Psychologically, in the incompatible case the person cannot form a single thought for all combinations of possible outcomes (because they are linearly dependent). In the incompatible case, asking the first question alters the *context* of the second question, and thus affects its answer. Therefore, in applying QP in psychology, we can ask whether one decision is likely to affect the other.

Suppose we are going to ask a person two questions, whether they are happy or not and whether they are employed or not. It is plausible that happiness and employment are related, so we postulate a single 2D space spanned by both bases (Fig. III.43(b)). The angle between the two bases reflects the fact that happiness is likely to be correlated to employment. Notice that once we get an answer regarding happiness, we will be in an indefinite state regarding employment, and vice versa.

Suppose we ask if the subject is employed and then ask if they are happy. The probability that they answer "yes" to both is given by:[29]

$$\mathcal{P}\{\text{employed \&\& happy}\} = \mathcal{P}\{\text{employed}\} \times \mathcal{P}\{\text{happy} \mid \text{employed}\}.$$

The rules of QP give the first probability: $\mathcal{P}\{\text{employed}\} = \|P_{\text{employed}}|\psi\rangle\|^2$. Asking about employment has collapsed the state, which is now

$$|\psi_{\text{employed}}\rangle = \frac{P_{\text{employed}}|\psi\rangle}{\|P_{\text{employed}}|\psi\rangle\|}.$$

The probability of a happy response is then

$$\mathcal{P}\{\text{happy} \mid \text{employed}\} = \|P_{\text{happy}}|\psi_{\text{employed}}\rangle\|^2.$$

Hence the probability of the two responses is

$$\mathcal{P}\{\text{employed \&\& happy}\} = \|P_{\text{happy}} P_{\text{employed}}|\psi\rangle\|^2.$$

From this example, we can see that the law for conditional probability in QP, called *Lüder's Law*, is:

$$\mathcal{P}\{A \mid B\} = \frac{\|P_A P_B|\psi\rangle\|^2}{\|P_B|\psi\rangle\|^2} = \frac{\mathcal{P}\{B \text{ \&\& } A\}}{\mathcal{P}\{B\}}.$$

Look at Fig. III.43(b). You can see that

$$\mathcal{P}\{\text{happy}\} < \mathcal{P}\{\text{employed \&\& happy}\},$$

which cannot happen in CP (since $\mathcal{P}\{A\} \geq \mathcal{P}\{A \wedge B\}$ always). The psychological interpretation would be that the subject's consciousness of being employed makes her more likely to say she is happy. This is because happiness and employment are correlated, but this correlation does not affect the outcome without the prior question about employment. In general, $\mathcal{P}\{A \text{ \&\& } B\} \neq \mathcal{P}\{B \text{ \&\& } A\}$, which cannot happen in CP. That is, conjunction is not commutative. You can see

$$\mathcal{P}\{\text{happy \&\& employed}\} < \mathcal{P}\{\text{employed \&\& happy}\}.$$

This is because the subject was more uncertain about their happiness than their employment, and therefore the state vector lost a lot of its amplitude

[29] As in C++, " && " should be read "and then" (sequential "and").

via its projection first onto $|\text{happy}\rangle$. "The size of such angles and the relative dimensionality of the subspaces are the cornerstones of QP cognitive models and are determined by the known psychology of the problem. These angles (and the initial state vector) have a role in QP theory analogous to that of prior and conditional distributions in Bayesian modeling." (Pothos & Busemeyer, 2013)

G.2.d Compatible questions

Fig. III.43(c) displays the case where the questions are compatible (only three of the four basis vectors are shown). In this case we have a tensor product between the space spanned by $\{|\text{happy}\rangle, |\text{unhappy}\rangle\}$ and the space spanned by $\{|\text{employed}\rangle, |\text{unemployed}\rangle\}$. For compatible questions the states are composite vectors, e.g.,

$$\begin{aligned}
|\text{H}\rangle &= \eta|\text{happy}\rangle + \eta'|\text{unhappy}\rangle, \\
|\text{E}\rangle &= \epsilon|\text{employed}\rangle + \epsilon'|\text{unemployed}\rangle, \\
|\Psi\rangle &= |\text{H}\rangle \otimes |\text{E}\rangle \\
&= \eta\epsilon|\text{happy}\rangle|\text{employed}\rangle + \eta\epsilon'|\text{happy}\rangle|\text{unemployed}\rangle \\
&\quad +\eta'\epsilon|\text{unhappy}\rangle|\text{employed}\rangle + \eta'\epsilon'|\text{unhappy}\rangle|\text{unemployed}\rangle.
\end{aligned}$$

Then, for example, as in CP the joint probability

$$\mathcal{P}\{\text{happy} \wedge \text{employed}\} = |\eta\epsilon|^2 = \mathcal{P}\{\text{happy}\}\mathcal{P}\{\text{employed}\}.$$

G.2.e Structured representations and entanglement

Many concepts seem to have *structured representations*, that is, components, properties, or attributes,, which are "aligned" when concepts are compared.[30] Structured concepts are naturally represented in QP by tensored spaces representing the concept's components. However QP also permits entangled (non-product) states, such as

$$\alpha|\text{happy}\rangle|\text{employed}\rangle + \beta|\text{unhappy}\rangle|\text{unemployed}\rangle.$$

This represents a state in which happiness and employment are strongly interdependent. It represents a stronger degree of dependency than can be expressed in CP. In CP you can construct a complete joint probability out of pairwise joints, but this is not possible in QP.

[30] Think of the variable components (fields) of a C++ class.

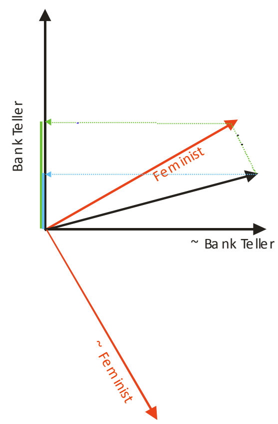

Figure III.44: Hypothetical basis state space of the "Linda experiment." (Pothos & Busemeyer, 2013)

G.2.f Time evolution

Time evolution is CP is defined by "a transition matrix (the solution to Kolmogorov's forward equation)" (Pothos & Busemeyer, 2013). It transforms the probabilities without violating the law of total probability. In QP amplitudes change by a unitary transformation.

G.3 Experimental evidence

The "Linda experiment." In 1983 Tversky and Kahneman reported on experiments in which subjects read a description of a hypothetical person named Linda that suggested she was a feminist. Subjects were asked to compare the probability of two statements: "Linda is a bank teller" (extremely unlikely given Linda's description), and "Linda is a bank teller and a feminist." Most subjects concluded:

$$\mathcal{P}\{\text{bank teller}\} < \mathcal{P}\{\text{bank teller} \wedge \text{feminist}\},$$

which violates CP; it is an example of the *conjunction fallacy*. Many experiments of this sort have shown that everyday reasoning commits this fallacy. Tversky and Kahneman proposed that people use heuristics rather than formal CP, but it can also be explained by QP.

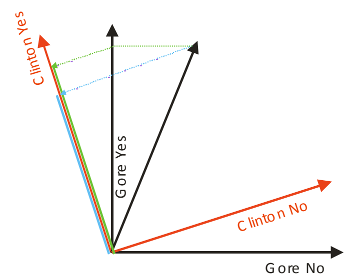

Figure III.45: Example of order effects in Gallop polls. [fig. from PB]

The QP explanation is as follows. We suppose that the written description makes it a priori likely that Linda is a feminist and unlikely that she is a bank teller; these priors are depicted in Fig. III.44. However, notice that being a feminist is largely independent of being a bank teller. In making a judgment like "Linda is a bank teller and a feminist" it is supposed that it is a sequential conjunction, with the most likely judgment evaluated first, in this case, "feminist && bank teller." Look at the figure. The green projection onto $|\text{feminist}\rangle$ and then onto $|\text{bank teller}\rangle$ is longer than the blue projection directly onto $|\text{bank teller}\rangle$. Projection can be thought of as an *abstraction* process, and so the projection of Linda onto $|\text{feminist}\rangle$ throws away details about her (it stereotypes her, we might say), and makes it more likely that she is a bank teller (since there is not a strong correlation between feminists and bank tellers). This may be compared to decoherence and loss of information in a quantum system. "In general, QP theory does not always predict an overestimation of conjunction. However, given the details of the Linda problem, an overestimation of conjunction necessarily follows. Moreover, the same model was able to account for several related empirical findings, such as the disjunction fallacy, event dependencies, order effects, and unpacking effects…" (Pothos & Busemeyer, 2013)

G.3.a Failure of commutativity

The "Clinton-Gore experiment." A Gallup poll asked "Is Clinton honest?" and "Is Gore honest?" Results depended on the order in which they were asked:

order	Clinton	Gore
Clinton — Gore	50%	68%
Gore — Clinton	57%	60%

This is also a common characteristic of everyday judgment; it is also common in the assessment of evidence for a hypothesis. QP explains this as follows. The "Yes" basis vectors have a smaller angle reflecting an expected correlation between the answers (since Clinton and Gore ran together). The initial state vector is a little closer to the $|\text{Gore Yes}\rangle$ vector reflecting the assumption that Gore's honesty is a priori more likely than Clinton's. You can see this by looking at the green projection onto $|\text{Gore Yes}\rangle$, which is longer than its blue projection onto $|\text{Clinton Yes}\rangle$. Note further that the two-step blue projection onto $|\text{Clinton Yes}\rangle$ is longer than the direct projection onto it. That is, judging Gore to be honest increases the probability of also judging Clinton to be honest.

G.3.b Violations of the sure-thing principle

"The sure thing principle is the expectation that human behavior ought to conform to the law of total probability." (Pothos & Busemeyer, 2013) In 1992 Shafir and Tversky reported experiments showing violations of the sure-thing principle in the *one-shot prisoner's dilemma*: The subject has to decide whether to cooperate or defect, as does their opponent. This is a typical payoff matrix; it shows the payoff for you and your opponent for each pair of choices:

	opponent	
↓ you	cooperate	defect
cooperate	3, 3	0, 5
defect	5, 0	1, 1

If you are told what your opponent is going to do, then you should defect. This is what subjects usually do. If you don't know, then the optimal strategy is still to defect. This is the "sure thing": you should defect in either case. However, some subjects decide to cooperate anyway (thus violating the sure-thing principle). One explanation is "wishful thinking." If you have a bias toward cooperation, you might suppose (in the absence of evidence) that your opponent has a similar bias.

The QP explanation is as follows. Suppose $|\psi_C\rangle$ and $|\psi_D\rangle$ are the states of knowing that your opponent will cooperate and defect, respectively. Suppose

P_C and P_D are projections representing your decision to cooperate or defect. Under the condition where you know what your opponent is going to do, the probability of you deciding to defect in the two cases is:

$$\begin{aligned}\mathcal{P}\{\text{you defect}\} &= \|P_D|\psi_C\rangle\|^2,\\ \mathcal{P}\{\text{you defect}\} &= \|P_D|\psi_D\rangle\|^2.\end{aligned}$$

In the unknown condition, we can suppose the state is $|\psi\rangle = \frac{1}{\sqrt{2}}(|\psi_C\rangle + |\psi_D\rangle)$. Hence, in this case the probability of you deciding to defect is:

$$\begin{aligned}\mathcal{P}\{\text{you defect}\} &= \left\|\frac{1}{\sqrt{2}}(P_D|\psi_C\rangle + P_D|\psi_D\rangle)\right\|^2\\ &= \frac{1}{2}(\langle\psi_C| + \langle\psi_D|)P_D^\dagger P_D(|\psi_C\rangle + |\psi_D\rangle)\\ &= \frac{1}{2}\|P_D|\psi_C\rangle\|^2 + \frac{1}{2}\|P_D|\psi_D\rangle\|^2 + \langle\psi_D \mid P_D^\dagger P_D \mid \psi_C\rangle.\end{aligned}$$

The interference term $\langle\psi_D \mid P_D^\dagger P_D \mid \psi_C\rangle$ could be positive or negative, in the latter case decreasing the probability below unity.

G.3.c Asymmetric similarity

In 1977 Tversky showed that similarity judgments violate metric axioms, in particular, symmetry.

China-Korea experiment. For example, North Korea was judged more similar to China, than China was judged to be similar to North Korea:

$$\text{Sim}(\text{North Korea}, \text{China}) > \text{Sim}(\text{China}, \text{North Korea}).$$

The QP explanation is that concepts correspond to subspaces of various dimensions, with the dimension of the subspace roughly corresponding to the number of known properties of the concept (i.e., how much someone knows about it). The judgment of the similarity of A to B is modeled by the projection of the initial state into A and then into B. It's assumed that the initial state is neutral with respect to A and B (i.e., the subject hasn't been thinking about either). If $|\psi\rangle$ is the initial state, then

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

The subjects in this case are assumed to be more familiar with China than with North Korea, so the China subspace is larger (see Fig. III.46). When

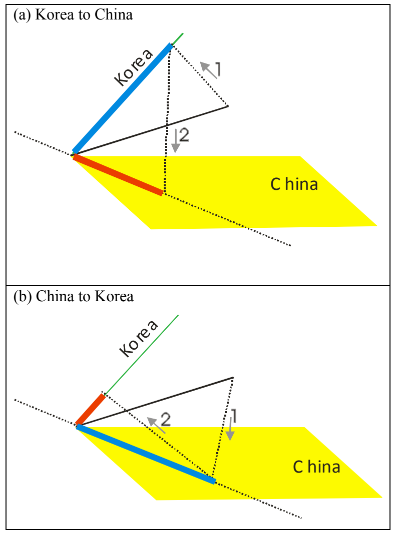

Figure III.46: QP model of China – (North) Korea experiment. [fig. from PB]

North Korea is compared to China, more of its amplitude is retained by the final projection into the higher dimensional subspace corresponding to China: Fig. III.46(a). In the opposite case, the projection into the lower dimensional North Korea subspace loses more amplitude: Fig. III.46(b). This is not universally true.

G.4 Cognition in Hilbert space

Pothos & Busemeyer (2013) defend the application of QP in a function-first or top-down approach to modeling cognition.[31] This is done by postulating vectors in a low-dimensional space. I argue that consideration of the high-dimensional complex-valued wavefunction underlying the state vector will expand the value of QP in cognitive science.

[31]Material in this section is adapted from MacLennan (2013), my commentary on Pothos & Busemeyer (2013).

G.4.a QM premises

To this end, application of QP in cognitive science would be aided by importing two premises from quantum mechanics:

The first premise is that the fundamental reality is the wavefunction. In cognitive science this corresponds to postulating a spatially-distributed pattern of neural activity as the elements of the cognitive state space. Therefore the basis vectors used in QP are in fact basis functions for an infinite (or very high) dimensional Hilbert space.

The second important fact is that the wavefunction is complex-valued and that wavefunctions combine with complex coefficients. This is the main reason for interference and other non-classical properties. The authors acknowledge this, but do not make explicit use of complex numbers in the target article.

G.4.b Possible neural substrates

What is the analog of the complex-valued wavefunction in neurophysiology? There are several possibilities, but perhaps the most obvious is the distribution of neural activity across a region of cortex; even a square millimeter of which can have hundreds of thousands of neurons. The dynamics will be defined by a time-varying Hamiltonian, with each eigenstate being a spatial distribution of neurons firing at a particular rate. The most direct representation of the magnitude and phase (or argument) of a complex quantity is frequency and phase of neural impulses.

G.4.c Projection

Possible neural mechanisms: Pothos & Busemeyer (2013) specify that a judgment or decision corresponds to measurement of a quantum state, which projects it into a corresponding subspace, but it is informative to consider possible mechanisms. For example, the need to act definitely (such as coming to a conclusion in order to answer a question) can lead to mutually competitive mechanisms, such as among the minicolumns in a macrocolumn, which creates dynamical attractors corresponding to measurement eigenspaces. Approach to the attractor amplifies certain patterns of activity at the expense of others. Orthogonal projectors filter the neural activity and win the competition with a probability proportional to the squared amplitude of the patterns to which they are matched. (In the case where the phases of neural impulses encode complex phases, matching occurs when the phases are delayed in such a way that the impulses reinforce.) The winner positively reinforces its matched signal and the loser negatively reinforces the signal to which it is matched. Regardless of mechanism, during collapse the energy of the observed eigenstate of the decision (measurement) operator receives the energy of the orthogonal eigenstates (this is the effect of renormalization). The projection switches a jumble of frequencies and phases into a smaller, more coherent collection, corresponding to the answer (observed) eigenspace.

No inherent bases: The target article suggests that a QP model of a process begins by postulating basis vectors and qualitative angles between alternative decision bases (significantly, only real rotations are discussed). As a consequence, a QP model is treated as a low-dimensional vector space. This is a reasonable, top-down strategy for defining a QP cognitive model, but it can be misleading. There is no reason to suppose that particular decision bases are inherent to a cognitive Hilbert space. There may be a small number of "hard-wired" decisions, such as fight-or-flight, but the vast majority are learned. Certainly this is the case for decisions corresponding to lexical items such as (un-)happy and (un-)employed.

Creation/modification of observables: Investigation of the dynamics of cognitive wavefunction collapse would illuminate the mechanisms of decision making but also of the processes by which observables form. This would allow modeling changes in the decision bases, either temporary through context effects or longer lasting through learning. Many decision bases are ad hoc, as when we ask, "Do you admire Telemachus in the *Odyssey*?" How such ad hoc projectors are organized requires looking beneath a priori state basis vectors to the underlying neural wavefunctions and the processes shaping them.

G.4.d Incompatible decisions

The commutator and anti-commutator: In quantum mechanics the uncertainty principle is a consequence of non-commuting measurement operators, and the degree of non-commutativity can be quantified (see Sec. B.7, p. 93). Two measurement operators P and Q commute if $PQ = QP$, that is, if the operator $PQ - QP$ is identically $\mathbf{0}$. If they do not commute, then $PQ - QP$ measures the degree of non-commutativity. This is expressed in quantum mechanics by the "commutator" $[P, Q] = PQ - QP$. It is relatively easy to show that this implies an uncertainty relation: $\Delta P\,\Delta Q \geq |\langle [P,Q] \rangle|$. That is, the product of the uncertainties on a state is bounded below by the absolute mean value of the commutator on the state. Suppose H is a measurement that returns 1 for $|\text{happy}\rangle$ and 0 for $|\text{unhappy}\rangle$, and E is a measurement that returns 1 for $|\text{employed}\rangle$ and 0 for $|\text{unemployed}\rangle$. If $|\text{employed}\rangle = a|\text{happy}\rangle + b|\text{unhappy}\rangle$, then the commutator is

$$[H, E] = ab \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

The absolute mean value of this commutator (applied to a state) gives a minimum joint uncertainty. If we could measure $[P, Q]$ for various pairs of questions, P and Q, we could make quantitative empirical predictions of the joint uncertainty in decisions.

Might we design experiments to measure the commutators and so quantify incompatibility among decisions? Certainly there are difficulties, such as making independent measurements of both PQ and QP for a single subject, or accounting for intersubject variability in decision operators. But making such measurements would put more quantitative teeth into QP as a cognitive model.

G.4.e Suggestions

Pothos & Busemeyer (2013) do an admirable job of defending QP as a fruitful top-down model of decision making, but I believe it would be more valuable if it paid greater attention to the complex-valued wavefunction that underlies QP in both quantum mechanics and cognition. This would allow a more detailed account of the origin of interference effects and of the structure of both learned and ad hoc decision operators. Finally, the treatment of incompatible decisions can be made more rigorous by treating them quantitatively as noncommuting operators.

G.5 Conclusions

You might wonder why it is so important to understand the less-then-perfect inferential abilities of humans. There are at least two reasons, *scientific* and *technological*. First, it is important to understand human inference as both *pure* and *applied* science. It reveals much about our human nature, and specifically provides hints as to how the brain works. From a more applied perspective, we need to understand how humans determine their actions in order to predict (and even influence) human behavior. In terms of technology, it might seem that the last thing we might want to do would be to emulate in our machine intelligence the "imperfect, fallacious" reasoning of humans. It might be the case, however, that QP-based reasoning is better than CP for real-time purposeful action in natural, complex situations, where the premisses of CP are inaccurate.