1
Why use quantum theory for cognition and decision? Some compelling reasons

Why should you be interested in quantum theory applied to cognition and decision? Perhaps you are a physicist who is curious whether or not quantum principles can be applied outside of physics. In fact, that is one purpose of this book. Perhaps you are a cognitive scientist who is interested in representing concepts by vectors in a multidimensional feature space. This is essentially the way quantum theory works too. Perhaps you are a decision scientist who is trying to understand how people make decisions under uncertainty. Quantum theory could provide some interesting new answers. Generally speaking, *quantum theory is a new theory for constructing probabilistic and dynamic systems*, and in this book we apply this new theory to topics in cognition and decision. Later in this chapter we will give some specific examples, but let us step back at this point and try to understand the more general principles that support a quantum approach to cognition and decision.

1.1 Six reasons for a quantum approach to cognition and decision

Quantum physics is arguably the most successful scientific theoretical achievement that humans have ever created. It was created to explain puzzling findings that were impossible to understand using the older classical physical theory, and it achieved this by introducing an entirely new set of revolutionary principles. The older classical physical theory is now seen as a special case of the more general quantum theory. In the process of creating quantum mechanics, physicists also created a new theory of probabilistic and dynamic systems that is more general than the previous classic theory (Pitowski, 1989). This book is not about quantum physics per se, but instead it explores the application of the probabilistic dynamic system created by quantum theory to a new domain – the field of cognition and decision making. Almost all previous modelling in cognitive and decision sciences has relied on principles derived from classical probabilistic dynamic systems. But these fields have also encountered puzzling findings that also seem impossible to understand within this limited framework.

Quantum principles may provide some solutions. Let us examine these principles to see why they may be applicable to the fields of cognition and decision.

1.1.1 Judgments are based on indefinite states

According to many formal models (computational or mathematical) commonly used in cognitive and decision sciences (such as Bayesian networks, or production rules, or connectionist networks), the cognitive system changes from moment to moment, but at any specific moment it is in a definite state with respect to some judgment to be made. To make this clearer, let us take a simple example. Suppose you are a member of a jury and you have just heard conflicting evidence from the prosecutor and defense. Your job is to weigh this evidence and come up with a verdict of guilty or not. Suppose your subjective probability of guilt is expressed on a $p \in [0,1]$ probability scale. Formal cognitive models assume that at each moment you are in a definite state with respect to guilt – say a state that selects a value p such that $p > 0.50$ or a state that produces p such that $p \leq 0.50$ (in other words, p is a function of the current state of the system). Of course, the model does not know what your true state is at each moment, and so the model can only assign a probability to you responding with $p > 0.50$ at that moment. But the model is stochastic only because it does not know exactly what trajectory (definite state at each time point) you are following. A stochastic model postulates a sample space of trajectories, along with a measure that assigns probabilities to sets of trajectories. But according to a stochastic model, once a trajectory is sampled (e.g., once a seed is selected to start a computer simulation), then the system deterministically jumps from one definite state (e.g., respond with $p > 0.50$) to another (e.g., respond with $p \leq 0.50$) or stays put across time. The states are pointwise and dispersion free and probabilities only arise from sampling different trajectories across new replications (e.g., starting the computer simulation over again with a new seed). In this sense, cognitive and decision sciences currently model the cognitive system as if it was a *particle* producing a definite sample path through a state space.

Quantum theory works differently by allowing you to be in an *indefinite* state (formally called a *superposition* state) at each moment in time before a decision is made. Strictly speaking, being in an indefinite or superposition state means that the model *cannot* assume either (a) you are definitely in a guilty state (e.g., a state that responds with $p > 0.50$) or (b) you are definitely in a not guilty state (e.g., respond with $p \leq 0.50$) at some moment. You may be in an indefinite state that allows both of these definite states to have *potential* (technically called state amplitudes) for being expressed at *each* moment (Heisenberg, 1958). (This does *not* mean you are definitely in both states simultaneously at each moment.) Intuitively, if you are in an indefinite state, then you do not necessarily think the person is guilty and at the same time you do not necessarily think the person is not guilty. Instead, you are in a superposition state that leaves you *conflicted, or ambiguous, or confused, or uncertain* about the guilty status. The potential for guilt may be greater than the potential for not guilty at one moment, and these potentials (amplitudes) may change from one moment to the next moment, but both answers are potentially available at *each* moment. In quantum theory, there is *no* single trajectory or sample path across time before making a decision, but instead there is a smearing of potentials across states that flows across time. In this sense, quantum theory allows one to model the cognitive system as if it was a *wave* moving across time over the state space until a decision is made. However, once a decision is reached, and uncertainty is resolved, the state becomes definite as if the wave collapses to a point like a particle. Thus, quantum systems require *both* wave (indefinite) and particle (definite) views of a cognitive system.

We argue that the wave nature of an indefinite state captures the psychological experience of conflict, ambiguity, confusion, and uncertainty; the particle nature of a definite state captures the psychological experience of conflict resolution, decision, and certainty.

1.1.2 Judgments create rather than record

According to many formal models, the cognitive system may be changing from moment to moment, but what we record at a particular moment reflects the state of the system as it existed immediately before we inquired about it. So, for example, formal cognitive models assume that if a person watches a disturbing scene and we ask the person a question such as “Are you afraid?”, then the answer reflects the state of the person regarding that question just before we asked it. If instead we asked the person “Are you excited?” then the answer again reflects the state regarding this other question just before we asked it.

One of the more provocative lessons learned from quantum theory is that taking a measurement of a system creates rather than records a property of the system (Peres, 1998). Immediately before asking a question, a quantum system can be in an indefinite state. For example, the person may be ambiguous about his or her feelings after watching a disturbing scene. The answer we obtain from a quantum system is constructed from the interaction of the indefinite state and the question that we ask (Bohr, 1958). This interaction creates a definite state out of an indefinite state. For example, the person may have been ambiguous about their feelings after the disturbing scene, but this state becomes more definite after answering the question about being afraid. If the answer is “Yes, I feel afraid,” then the person acts accordingly. This is, in fact, the basis for modern psychological theories of emotion (Schachter & Singer, 1962). Decision scientists also argue that beliefs and preferences are constructed on line rather than simply being read straight out of memory (Payne *et al.*, 1992). For example, a person may initially be in an indefinite state about a set of paintings on display, but if the person is asked to choose one as a gift, then a preference order is constructed on line for the purpose.

We do not wish to argue that every answer to every question involves the construction of an opinion. For many questions you do have a stored answer that is simply retrieved on demand (e.g., Have you ever read a certain book?). But other questions are new and more complex and you have to construct an answer from your current state and context (e.g., Did you like the moral theme of that book?). So we argue that the quantum principle of constructing a reality from an interaction between the person's indefinite state and the question being asked actually matches psychological intuition better for complex judgments than the assumption that the answer simply reflects a preexisting state.

1.1.3 Judgments disturb each other, introducing uncertainty

According to quantum theory, if one starts out in an indefinite state, and is asked a question, then the answer to this question will change the state from an indefinite state to one that is more definite with respect to the question that was asked. But this change in state after the first question then causes one to respond differently to subsequent questions so that the order of questioning becomes important. Consider the following popular example from social psychology. Suppose a teenage boy is directly asked "How happy are you?" the typical answer is "Everything is great." However, if this teenager is first asked "When was the last time you had a date?" then the answer tends to be "Seems like a long time ago." Following this sobering answer, a later question about happiness tends to produce a second answer that is not so sunny and rosy. Thus, the first question sets up a context that changes the answer to the next question. Consequently, we cannot define a joint probability of answers to question A and question B, and instead we can only assign a probability to the sequence of answers to question A followed by question B. In quantum physics, if A and B are two measurements and the probabilities of the outcomes depend on the order measurement, then the two measurements are non-commutative. In physics, for example, measurements of position and momentum along the same direction are non-commutative, but measurements of positions along the horizontal and vertical coordinates are commutative. Many of the mathematical properties of quantum theory arise from developing a probabilistic model for non-commutative measurements, including Heisenberg's (1927) famous uncertainty principle (Heisenberg, 1958).

Order effects are also responsible for introducing uncertainty into a person's judgments. If the first question A produces an answer that creates a definite state with respect to that question, the state created by A may be indefinite with respect to a different question B. Consider the following consumer choice example. Suppose a man is considering the purchase of a new car and two different brands are in contention: a BMW versus a Cadillac. If he directly asks himself what he prefers, he definitely answers with the BMW. But if he first asks himself what his wife prefers (she definitely wants the Cadillac) and subsequently asks himself what he prefers (after taking on his wife's perspective), then he becomes uncertain about his own preference. In this example, the question about his wife's preference disturbs and creates uncertainty about his own preference. Thus, it may be *impossible* to be in a definite state with respect to two different questions, because a definite state (technically speaking an eigenstate) for one is an indefinite state (superposition) for another. In this case, the questions are said to be *incompatible* and the incompatibility of questions is mathematically implemented by the non-commutativity of quantum measurements. Question order effects are a major concern for attitude researchers, who seek a theoretical understanding of these effects similar to that achieved in quantum theory (Feldman & Lynch, 1988).

1.1.4 Judgments do not always obey classic logic

The classic probability theory used in current cognitive and decision models is derived from the Kolmogorov axioms (Kolmogorov, 1933/1950). These axioms assign probabilities to events defined as sets. Consequently, the family of sets in the Kolmogorov theory obeys the Boolean axioms of logic. Thus, Boolean logic lies at the foundation of current probabilistic models of cognition and decision making. One important axiom of Boolean logic is the distributive axiom: if $\{G, T, F\}$ are events then $G \cap (T \cup F) = (G \cap T) \cup (G \cap F)$. Consider, for example, the concept that a boy is good (G) and the pair of concepts the boy told the truth (T) and the boy did not tell truth (falsehood, F). Suppose you are trying to decide if the boy is good but you do not know if he is truthful. According to Boolean logic, the event G can only occur in one of two ways: either $(G \cap T)$ occurs or $(G \cap F)$ exclusively. This means there are only two mutually exclusive and exhaustive ways for you to think the boy is good: he is good and truthful or he is good and he is not truthful.

From this distributive axiom, one can derive the law of total probability. Define $p(G)$ as the probability of event G, $p(T)$ is the probability of event T, $p(F)$ is the probability of event F, $p(G|T)$ is the probability of event G conditioned on knowing event T, and $p(G|F)$ is the probability of event G conditioned on knowing event F. Then the law of total probability follows from

$$
\begin{aligned}
p(G) &= p\left((G \cap T) \cup (G \cap F)\right) = p\left(G \cap T\right) + p(G \cap F)) \\
&= p(G)p(G|T) + p(F)p(G|F).
\end{aligned}
$$

This law provides the foundation for inferences with Bayes nets. The law of total probability is violated by the results of the disjunction experiment and the category – decision-making experiment in psychology and the two-slit type of experiments in physics, all of which we describe later in this chapter.

Quantum probability theory is derived from the von Neumann axioms (von Neumann, 1932/1955). These axioms assign probabilities to events defined as subspaces of a vector space (more on this in Chapter 2). The definite states form the basis for the vector space, and an indefinite or superposition state can be any point within this vector space. An important consequence of using subspaces is that the logic of subspaces does not obey the distributive axiom of Boolean logic (Hughes, 1989). For example, according to quantum logic, when you try to decide whether a boy is good without knowing if he is truthful or not, you are *not* forced to have only two thoughts: he is good and he is truthful or he is good and he is not truthful. You can have other ambiguous thoughts represented by a superposition over the truthful or not truthful attributes.

The fact that quantum logic does not always obey the distributive axiom implies that the quantum model does not always obey the law of total probability (Khrennikov, 2010). This is why the quantum model can explain the results of the disjunction experiment in psychology and the two-slit experiment in physics. Thus, quantum logic is a generalization of classic logic and quantum probability is a generalized probability theory. We argue that classic logic and classic probability theory are too restrictive to explain human judgments and decisions.

1.1.5 Judgments do not obey the principle of unicity

The classic (Kolmogorov) probability theory, which is used in current cognitive and decision models, is based on the principle of *unicity* (Griffiths, 2003). A single sample space is proposed which provides a complete and exhaustive description of all events that can happen in an experiment.[1] This follows from the Boolean algebra used in classic theory: if A is an event and B is another event from an experiment, then $A \cap B$ must be an event too, and repeated application of this principle leads to intersections that cannot be broken down any further (the atoms or elements or points of the sample space). All events can be described by unions of the atoms or elements or points of the sample space. If you think about this for a while, this is a tremendous constraint on a theory. We argue that it is oversimplifying the extremely complex nature of our world.

Let us examine the consequence of assuming unicity for experiments on human probability judgments. Suppose we do an experiment in which we ask a person to describe the likelihood of various future events with respect to future political history. Perhaps a person has the knowledge to do this within a single sample space. But then we can also ask the same person to describe the likelihood of future events with respect to progress in science. Now it becomes quite a stretch to imagine that the person is able to assign joint probabilities to all historical and scientific events. Instead, the person might need to fall back on one description of events (one sample space) for political futures, but use a different description of events (another sample space) for future scientific progress. To go even further, we could ask about the likelihood of events concerning the romantic and marital relations of Hollywood movie stars. Surely we have passed the capacity of the person who would have little or no idea about how to combine all three of these topics into a unified sample space that assigns joint probabilities to all three kinds of events.[2]

Quantum probability does not assume the principle of unicity (Griffiths, 2003). This assumption is broken as soon as we allow incompatible questions into the theory which cause measurements to be non-commutative (Primas, 2007).

[1] Kolmogorov realized that different sample spaces are needed for different experiments, but his theory does not provide a coherent principle for relating these separate experiments. This is exactly what quantum probability theory is designed to do.

[2] One could try to assume independence between questions about history, science, and Hollywood movie stars. But independence is also an overly severe restriction to impose on human judgments.

Incompatible questions cannot be evaluated on the same basis, so that they require setting up separate sample spaces. Quantum theory allows one to use a partial Boolean algebra: one sample space can be used to answer a first set of questions in a Boolean way and another sample space can be used to answer a different set of questions in a Boolean way, but both Boolean subalgebras are pasted together in a coherent but non-Boolean way. This provides more flexibility for assigning probabilities to events, and it does not require forming all possible joint probabilities, which is a property we believe is needed to understand the full complexity of human cognition and decision.

1.1.6 Cognitive phenomena may not be decomposable

In cognitive science, researchers often take the approach to model a cognitive process for a task by proposing a large collection of random variables, but what is actually observed only corresponds to a small subset of these variables for any single experimental condition. A single *complete* joint distribution, which is a joint probability distribution across all the random variables, is then assumed to exist that can be used to determine the observed marginalized distributions for any subset of variables. This seemingly straightforward and innocuous theoretical approach is what the quantum physicists questioned when trying to experimentally test the existence of quantum entanglement. Entanglement is a surprising phenomenon in which two seemingly distinct and separated systems behave as one – these systems are sometimes referred to as "quantum correlated." It turned out that when systems are entangled it is not possible to construct such a *complete* joint distribution. Intuitively, this result suggests there is an extreme form of dependencies between the systems which goes beyond the dependencies derived from traditional probability theory. Since the initial probabilistic models were developed for the entanglement experiments, there has been a large body of literature explaining the general probabilistic foundations of quantum correlations. This has resulted in analysis methods that are general enough for application to cognitive science. A note of caution, however: assuming the existence of quantum correlations in cognitive phenomena is certainly highly speculative. There is, however, a steadily growing body of pioneering literature putting forward this view. But what do these quantum correlations entail for our cognitive models?

In cognitive science, and science in general, reductionism has been a powerful background philosophy underlying model development. By this we mean the assumption that phenomena can be analyzed by considering their different components separately, and then synthesizing the results obtained. Phenomena, or systems, that can be understood this way are deemed "decomposable." Non-decomposable systems cannot be so straightforwardly understood. The majority of models in cognitive science are decomposable, as they are understood in terms of their constituent parts and how these are related. For example, consider the spreading activation model of words in human memory. Even though the network structure may feature a high degree of interrelationship, the nodes (i.e., the words) are still considered as discrete components in the model. Many models in cognitive science are like this. However, is reductionism always appropriate when modelling cognition? Quantum correlations which appear in cognitive phenomena suggest that this assumption can at least be questioned in specific situations. For example, when a word is studied in a memory experiment, there is a view that a word's associative network arises in synchrony with the word being studied. The intuition behind this view is very similar to that of quantum entanglement – the study word and its associates are behaving as one. The existence of quantum correlations suggests the cognitive model in question is not decomposable in the way we initially assumed and formalized via a particular set of random variables. It forces us to rethink very differently the nature of the phenomena being modelled.

Now that we have identified some general reasons for considering a quantum approach to cognition and decision, let us take a quick look at some simple examples of paradoxical findings from cognition and decision and give a brief idea about how quantum theory can be applied.

1.2 Four examples from cognition and decision

1.2.1 The disjunction effect

The first example is a phenomenon discovered by Amos Tversky and Eldar Shafir called the disjunction effect (Tversky & Shafir, 1992). It was discovered in the process of testing a rational axiom of decision theory called the sure thing principle (Savage, 1954). According to the sure thing principle, if under state of the world X you prefer action A over B, and if under the complementary state of the world ˜X you also prefer action A over B, then you should prefer action A over B even when you do not know the state of the world. For example, suppose you are trying to decide whether or not to make a risky investment right before a presidential election with only two parties, Democrat versus Republican. Assume that if the Democrats win, then you prefer to invest; also, if the Republicans win, then you prefer to invest. Therefore, you should invest without knowing the election outcome. Tversky and Shafir experimentally tested this principle by presenting students with a two-stage gamble; that is, a gamble which can be played twice. At each stage the decision was whether or not to play a gamble that has an equal chance of winning \$200 or losing \$100 (the real amount won or lost was actually \$2.00 and \$1.00 respectively). The key result is based on the decision for the second play, after finishing the first play. The experiment included three conditions: one in which the students were informed that they already won the first gamble, a second condition in which they were informed that they lost the first gamble, and a third in which they did not know the outcome of the first gamble. If they knew they won the first gamble, the majority (69%) chose to play again; if they knew they lost the first gamble, then again the majority (59%) chose to play again; but if they did not know whether they won or lost, then the majority chose not to play (only 36% wanted to play again). What went wrong here? Only one of two

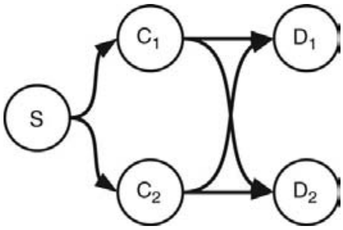

Figure 1.1 Diagram of double-slit experiment showing that a photon, starting from state S, can travel two channels (C_1 or C_2) before hitting one of two detectors (D_1 or D_2).

possible events can occur during the first play: win or lose. These students generally preferred to play if they won and they also preferred to play if they lost. So why do they prefer not to play when they do not know whether they won or lost?

Tversky and Shafir explained the finding in terms of choice based on reasons as follows. If the person knew they won, then they had extra house money with which to play and for this reason they chose to play again; if the person knew they had lost, then they needed to recover their losses and for this other reason they chose to play again; but if they did not know the outcome of the game, then these two reasons did not emerge into their minds. Why not? If the first play is unknown, it must definitely be either a win or a loss, and it cannot be anything else. We try to explain this in more detail in Chapter 9, but let us give a hint here.

Researchers working with quantum models see this finding as an example of an interference effect similar to that found in the double-slit type of experiments conducted in particle physics. Although we do not wish to present a whole lot of physics here, it is worthwhile to briefly discuss this one simple example to compare the experimental design and the statistical results produced by the disjunction experiment in psychology and the double-slit experiment in physics (technically, this description corresponds to a Mach–Zehnder interferometer experiment). Referring to Figure 1.1, a *single* photon is dispersed from a light source (S in the figure) and it is split by a beam splitter off into one of two channels (C_1 and C_2 in the figure) from which it eventually can reach one of two detectors (D_1 or D_2 in the figure). Two conditions are examined: for one condition, the channel through which the photon passes is observed, and for the other condition it is not. The results are the following. When the channels are observed, a single photon either definitely goes through the upper channel or definitely goes through the lower channel, and you never find parts of photons going through both channels. If the photon is observed to pass through the upper channel C_1, then there is a 50% chance of reaching the upper detector D_1; likewise, if the photon is observed to pass through the lower channel C_2, then again there is a 50% chance of reaching the upper detector D_1.

Now consider what happens when the channels are unobserved. Only a single particle ever enters and exits this system, and according to classic probability there is a 0.5 chance of it reflecting off in either direction; so when the channel is not observed, we still should expect 50% of the photons to reach detector D_1. However, in fact, the probability of reaching the upper detector D_1 drops to zero and the probability of reaching the lower channel rises to one. Something is wrong again here! But experimental results are never wrong, so it is our classic thinking that is wrong. Instead, what happens when the channel is unobserved is that the photon enters a superposition state – it is superposed between the two channels. From this superposition state, it cannot make the transition to detector D_1 and must always go to detector D_2. This is an example of an extreme case of interference, and less extreme cases can occur depending on the phases assigned to two channels. We will not go further into the physics to show how this is computed here. Instead, we will wait to see how the probability for the superposition state is computed in a later decision-making example.

By observing the channel through which the photon passes, we break down the superposition state into one of the definite states. The difference between the probability obtained under conditions of no observation versus the probability obtained under conditions of observation is called the *interference* effect. The main point is that interference lowers the probability of being detected at D_1 for the unobserved case far below each of the probabilities for each of the observed conditions. Quantum probability theory was designed to explain these bizarre statistics, which it succeeds in doing extremely well.

At this point, the analogy between the disjunction experiment and the double-slit type of experiment should be clear. Both cases involve two possible paths: in the disjunction experiment, the two paths are inferring the outcome of either a win or a loss with the first gamble; for the double-slit experiment, the two paths are splitting the photon off into the upper or lower channel by a beam splitter. In both experiments, the path taken can be known (observed) or unknown (unobserved). Finally, in both cases, under the unknown (unobserved) condition, the probability (of gambling for the disjunction experiment, of detection at D_1 for the double-slit experiment) falls far below each of the probabilities for the known (observed) cases. So can we speculate that for the disjunction experiment, under the unknown condition, instead of definitely being in the win or loss state, the student enters a superposition state that prevents finding a reason for choosing the gamble? One cannot help but wonder whether the mathematical model that succeeds so well to explain interference statistics in physics could also explain interference statistics in psychology. Let us look at another example to broaden the range of application a little and go into a bit more detail about how the quantum model mathematically works.

1.2.2 Interference of categorization on decision making

James Townsend (Townsend *et al.*, 2000) introduced a new paradigm to study the interactions between categorization and decision making, which we