
Prequential Quantum Dynamics

Kathryn Blackmond Laskey
George Mason University
4400 University Drive
Fairfax, VA 22032
klaskey@gmu.edu

Abstract

Importation of methods from statistical physics into machine learning has led to rapid advances in methods for efficient learning of good representations for complex problems. This paper examines the potential for cross-fertilization in the other direction. The Stapp ontology for quantum dynamics can be coupled with evolutionary formulations of subjective decision theory to yield a unified theory of the physical and mental aspects of reality. The resulting ontology fills an acknowledged gap in contemporary physics and at the same time provides a natural bridge connecting the physical to the biological and social sciences. The implications of this ontology for artificial intelligence and quantum computing are explored. An explicitly probabilistic quantum logic is proposed as the foundation for a post-classical theory of computing.

1. INTRODUCTION

The relationship between mind and matter has been a persistent puzzle since the dawn of science. The course of Western thought was profoundly affected by Rene Descartes' articulation of the distinction between mind and matter. According to Descartes, the mental and the material were two separate aspects of reality that interacted in an ill-understood manner. Since that time, the prevailing metaphysical stance in the West has been that the material universe evolves autonomously, follows a lawful dynamics independent of our thoughts, and can be described by empirically testable mathematical theories. How it is that we seem to be able to affect this material world by thinking about what we want to do and willing it to happen is poorly understood. That there are limits to our influence is clear, but the nature of and reason for the limits is unclear.

The scientific revolution has been profoundly successful at describing those aspects of the world we label as material. As a result, our ability to manipulate

the material world through technology has exploded. The hypothetico-deductive approach and the commitment to empirical evaluation that mark the scientific attitude have moved beyond the purely material to the biological and social sciences. The scientific approach has led to theoretical innovations and improved practical procedures in fields as diverse as ecology, clinical psychology, and management.¹ The computer revolution has raised the possibility that intelligence itself can be understood scientifically, formalized, and engineered into physical devices.

Despite enormous practical success, science remains unclear about how the mind that formulates and understands scientific theories, and then designs and conducts experiments to test its theories, is related to the material world it studies. The past century has brought broad appreciation that there exist at least statistical regularities underlying the seeming complexity of biological and social phenomena (Gigerenzer et al., 1990). Public policies that shape the context in which private decisions are made are increasingly based on a scientific approach to evaluating their likely effects. The Internet and the personal computer have become ubiquitous features of the office and home, changing the way we do business, interact professionally, educate ourselves, and live our personal lives. As these changes percolate through society, the question of the appropriate role in scientific theories for free will and deliberate choice takes on increased urgency. To deny that it is within the realm of scientific possibility in the relative near term to stumble upon engineered conscious intelligence, is to bury one's head in the sand. The consequences could be disastrous if our understanding of the role of conscious deliberate choice in nature matures less rapidly than our technology for creating engineered intelligence.

This paper presents a unified ontology for science in which conscious experience, learning, and deliberate choice play a central and fundamental role, as distinct from their epiphenomenal role in the classical ontology.

¹ That mathematical models have often been applied inappropriately to yield disastrous results does not negate the considerable successes of the scientific approach properly applied.

The ontology suggested here is fully consistent with our current knowledge of the workings of the physical universe (Stapp, 1999). This ontology fills an acknowledged gap in currently popular ontologies for quantum theory and at the same time fills a complementary gap in current theories of neurobiology, psychology and artificial intelligence. As such, it provides a natural foundation for a unified scientific approach to the physical, biological and social sciences. It also shows promise as the foundation for a post-classical theory of computing founded on explicitly non-deterministic and irreversible quantum systems.

2. UNIFIED ONTOLOGY OF SCIENCE

2.1 QUANTUM THEORY

Classical mechanics is a dynamically complete theory with no role for conscious thought and efficacious deliberate action. Once initial conditions are given, a classical system follows a definite trajectory that, at least in principle, can be deterministically predicted indefinitely into the future. Of course, in practice this predictability is limited by approximation and measurement error in the specification of both the initial conditions and the parameters of the dynamical equations. Nevertheless, *in principle*, the evolution of a classical system is perfectly deterministic.

Early in the 20th century it was discovered that the classical picture of a world of perfectly deterministic physical systems evolving according to local influences was incorrect. The classical picture of the world was replaced by the explicitly statistical quantum theory. Quantum theory deals not with definite trajectories of systems with precisely specified states, but with a *wave function* that describes the probabilities of the outcomes of various measurements a scientist might make on the system. The degree of accord between theoretical prediction and empirical measurements is stunning. Nevertheless, many physicists remain uncomfortable with the theory. There are three major reasons for this discomfort. First, quantum theory makes only statistical predictions, and many scientists are uncomfortable with a picture of Nature that has an intrinsically random component. Second, the theory is nonlocal. That is, there are correlations between spacelike separated events that cannot be explained by a hidden variable theory with strictly local influences. Third, the theory contains a major explanatory gap known as the “measurement problem,” in which deterministic evolution of the wave function is inexplicably punctuated with periodic “collapses” for which physics has no explanation.

2.2 THE MEASUREMENT PROBLEM

The state of a quantum system at any time is described by a mathematical entity called the *quantum state* or *wave function*. According to quantum theory, the quantum state evolves in time according to three distinct processes (Stapp, 1999; Penrose, 1994;

Shankar, 1994). The first process concerns how the state evolves in the absence of interactions with its environment. The second and third processes involve a discontinuous change in state called *state reduction* or *wave function collapse*. The second process determines the time at which reduction occurs and the set of possibilities for the quantum state after reduction. The third process selects one specific possibility to be actualized. Quantum theory describes the first and third processes to a high degree of accuracy. Between measurements, the quantum state evolves deterministically according to a differential equation called the *Schrödinger equation*. Although the quantum state after reduction cannot be predicted with certainty, there is a precise rule for calculating the probabilities of the different possibilities. Quantum theory as currently formulated has nothing at all to say about the second process. It appears that state reduction occurs when a quantum system interacts with its environment. In particular, we know reduction occurs when a scientist takes a measurement that “amplifies” a particular feature of the microscopic quantum world to produce a detectable change in a macroscopic measurement device. For this reason, the lack of a theory for state reduction is called the “measurement problem.”

Thus, the dynamic behavior of a quantum system depends in macroscopically observable ways on a process for which there exists no theory. In particular, a scientist can choose which of several distinct macroscopic effects to actualize by choosing which aspects of the system to observe. Thus, the theory contains a contingent element. It specifies behavior of the system *given* the actions an agent external to the theory takes to observe the system. This dependence of the predictions of the theory on an aspect of reality for which there is no theory has caused consternation among physicists.

2.3 ONTOLOGIES OF QUANTUM THEORY

The orthodox semantics for the equations of quantum theory is associated with Bohr (1934) and is called the Copenhagen interpretation. According to the Copenhagen interpretation, quantum theory replaces a classical theory that refers to an external material universe with a new theory that refers only to the experience of observers. In particular, conditional on a choice of experimental set-up that defines the macroscopically detectable possibilities available to the system, quantum theory predicts the probability that each of these possibilities will occur. Proponents of the Copenhagen interpretation make no ontological commitments regarding the existence of material entities that give rise to the experienced sequence of observations. It is often asserted that it is meaningless to speak of the “actual state” of a quantum system. The wave function is asserted to be nothing but a device for organizing the experiences of observers and enabling the computation of accurate predictions of the outcomes of experiments.

Although the Copenhagen interpretation is the standard view, most physicists prefer, at least informally, to operate with an ontology that connects the terms in the theory to a physical reality that is reflected in the experience of observers. One popular ontology, known as *many worlds*, asserts that the system actually realizes *all* possibilities open to it, but each occurs in a separate reality inaccessible to the other realities. This interpretation appears to be dominant in the field of quantum computing. Another is the pilot wave ontology due to Bohm, which is a nonlocal deterministic theory that includes both classical-like particles and a wave function that guides their evolution. Penrose (e.g., 1994) adopts a realist ontology he attributes to Heisenberg, in which wave function collapse represents the singling out of an actual event to occur, but the process for this is currently unknown. Penrose hypothesizes that a theory for this process will eventually emerge, and will be closely linked with the properties of the nonlinear gravitational force. All of these ontologies are observationally equivalent, and none provides a fully satisfactory resolution to the question of how and why reductions occur.

This paper proposes a new theory we call *prequential quantum dynamics*.² It is grounded in an ontology proposed originally by von Neumann (1934) and further elucidated by Wheeler (1967). It takes an actual event realist view of state vector reduction, but the theory of state vector reduction is different from and, we argue, more physically plausible, than that suggested by Penrose. Prequential quantum dynamics treats the observer and the system being observed as a single quantum system, itself subject to the rules of quantum theory. Complementary explanatory gaps in physics and psychology are filled by allowing an interaction between the informational structure represented by the quantum state and the informational structure of conscious experience. Stapp (1999) argues that such an interaction allows consciousness to become efficacious without disturbing any of the precepts or rules of quantum theory.

Thus, with no extension to existing physics and minimal metaphysical assumptions, it is possible to make the leap from the Copenhagen ontology-free view to a theory that connects physical reality in a plausible way to conscious experience and deliberate choice. The commonly held view that quantum theory is nothing but a recipe for calculating probabilities and says nothing directly about physical reality can be abandoned in favor of a well grounded, unapologetically realist theory. Prequential quantum dynamics is a fully integrated theory of mind and matter in which the equations of quantum theory describe both macroscopic conscious observers and the natural world they inhabit. Because the evolution of a quantum system depends on

the choice and timing of questions, the prequential ontology incorporates deliberate conscious choice into physics at exactly the place where current physics lacks a theory. This ontology thus provides science with a theory of both the physical and informational aspects of nature, and describes how deliberate choices of conscious agents affect both the agents themselves and the world they inhabit. Moreover, as argued below, the physical constraints necessary for deliberate choice to be operative in this way appear to be satisfied by the conditions occurring in live animal brains.

Stapp (1998) gives examples grounded in physics of macroscopically detectable differences in behavior resulting from different choices of what is observed and when. In particular, the quantum Zeno effect (Itano, et al., 1990) predicts that observations taken sufficiently rapidly can keep a quantum system within a constrained region of phase space. He argues that an organism might use the quantum Zeno effect to keep its brain state within a given basin of attraction sufficiently long to trigger behaviors the organism desires to bring about. The quantum Zeno effect has been confirmed experimentally and is thought to occur at time and frequency scales consistent with patterns of electrochemical activity occurring in brains.

Stapp suggests that the quantum states for different responses to a question posed by an observer are associated with different psychological states, or *qualia* of experience. For human observers, qualia are multifaceted, highly complex gestalts that defy simple description. Nevertheless, human observers are able to select among questions to ask in order to bring about brain states associated with qualia they prefer. Stapp proposes that this occurs via a physical mechanism corresponding to what psychologists call *will* or *attention* (Anderson, 1999; James, 1890). That is, an organism complex enough to be labeled conscious can anticipate with some degree of accuracy the qualia associated with the different available question-asking policies, identify those that are most desirable, and focus attention on bringing about one of the most preferred policies. To do this, the brain is hypothesized to encode a "body-world schema" that represents the body, the environment, and the predicted effects of alternative question-asking policies. The organism uses its body-world schema to direct its focus of attention to bring about desired qualia.

There is no requirement that all state reductions be associated with questions asked by conscious observers. Although it is conceivable that some form of the property we call consciousness at the human level exists throughout the natural world, Stapp's theory does not require it. Whatever one's stance on the level of complexity at which consciousness can be said to exist, the question arises of why and how evolution would select for increasingly sophisticated ability to direct actions via conscious intent. Stapp argues that because choice of question affects the evolution of a quantum

² The term *prequential* is taken from Dawid and Vovk's (1999) sequential predictive theory of probability. Prequential theory forms the basis for player learning.

system, the ability to focus on asking questions that will bring about desired outcomes would be expected to have survival value. That is, it seems reasonable that evolution would tend to select for organisms that could form accurate representations of the choices available to them and select those options likely to lead to survival.

3. PREQUENTIAL PLAYERS

3.1 PHYSICS AND INFORMATION

The 20th century saw a classical world of matter with no place for mind displaced by the quantum world of mind with no place for matter. As the new century dawns, advances in computing and information theory have created a climate in which synthesis is possible. The new fields of artificial intelligence, computational psychology, and knowledge representation provide language and tools that can be used to elucidate the connection between the experience of observers as modeled explicitly by quantum theory and the external reality those experiences refer to. A growing consensus is emerging in the field of knowledge representation that at the top level an ontology requires categories for entities in the world, concepts, and relationships between concepts and entities in the world (Sowa, 1998). Science in the nineteenth century focused exclusively on creating categories for entities in the world. The emergence of quantum theory marked an emerging awareness that the edifice science had constructed was a representation, and therefore intrinsically incomplete in its correspondence to the world it described. Theories of representation are becoming more mature and their connections to physics more deeply understood. The time is ripe for consensus to emerge on how to connect the constructs of a combined physical/informational theory to referents in the world. Prequential quantum dynamics is not only a plausible explanation of observed evidence, but provides clear answers to puzzles that appear intractable under alternative ontologies. Moreover, the prequential ontology respects fundamental limits on self-representation first articulated by Plato and formally elucidated by later thinkers such as Russell, Gödel and Turing.

3.2 PLAYERS IN THE PREQUENTIAL GAME

The objective of this paper is to amplify and extend Stapp's ontology by recasting quantum theory as a *prequential game* (Dawid and Vovk, 1999) among interacting systems called *players*.³ Players are the fundamental ontological entities of prequential quantum dynamics. Players are modeled as stochastic processes that make independent decisions. There is no *a priori* commitment regarding whether players are conscious. The fundamental requirement is that a player is an independent causal agent that makes choices as a

³ I am indebted to Henry Stapp for suggesting this term.

unitary whole. For example, a pair of particles with coupled spins cannot be treated as two separate players, even if the particles are spacelike separated, at least while the coupling endures. However, a single electron in a double-slit experiment could be a (non-conscious) player. Human beings are players, but to our best current understanding families, corporations and countries are not. Ants are most likely players, but it is doubtful that an anthill is a player. Bacteria are probably players, but it is unclear whether a neuron or a skin cell is a player. Recent work on the use of directed graphical models to represent and draw inferences about causality might give rise to observational criteria for inferring whether a system is a player or can be treated as a player for modeling purposes.

The standard representation of a quantum state in the mathematical physics literature is a *density operator*.⁴ We consider the global system consisting of a player and the environment it is observing as a quantum state represented by the density operator S_G . The state S_A of Player A would be represented as a projection of S_G , denoted by $S_A = P_A S_G P_A$, where P_A is a projection operator.⁵ The rank of S_A , assumed finite, is the dimension of the subspace associated with Player A's state, or the number of degrees of freedom available to Player A. It seems reasonable to require that any player qualifying as conscious would have a subsystem responsible for representing the effects of its possible actions. The representational subsystem is the hardware that gives rise to the qualia associated with conscious experiences. The state of the representational subsystem is denoted by another projection $R_A = P'_A S_A P'_A$. The rank of R_A corresponds to the degrees of freedom associated with Player A's representation.

The simplest kind of question that can be posed to the system is a binary choice. Because questions with more than two answers can be thought of as a rapid sequence of yes/no questions, we restrict attention to binary choices. A binary choice can be represented as projecting the system onto one of two orthogonal subspaces. A "yes" answer ($Q=1$) corresponds to projection operator P_Q . A "no" answer ($Q=0$) corresponds to projection operator $(I-P_Q)$, where I is the identity operator. Suppose the state of the system just prior to measurement is $S(t-)$. Quantum theory predicts that the probabilities of the two answers are given by:

Answer	New state	Probability
$Q=1$	$S_1(t+) = P_Q S(t) P_Q$	$\text{Tr}(S_1(t+))/\text{Tr}(S(t-))$
$Q=0$	$S_0(t+) = (I-P_Q) S(t) (I-P_Q)$	$\text{Tr}(S_0(t+))/\text{Tr}(S(t-))$

⁴ Most mathematical physicists prefer to represent states as density operators rather than state vectors because the latter depend on an arbitrary phase. A density operator can be thought of as a Hermitian matrix, possibly infinite dimensional. A Hermitian matrix is a symmetric matrix that is equal to its adjoint, or conjugate transpose.

⁵ P is a projection operator if $PP=P$.

In other words, posing question Q to the system results in a “jump” from state $S(t_-)$ to $S_1(t_+)$ with probability $p_1 = \text{Tr}(S_1(t_+)/\text{Tr}(S(t_-)))$ and $S_0(t_+)$ with probability $p_0 = \text{Tr}(S_0(t_+)/\text{Tr}(S(t_-)))$. Between observations the state evolves deterministically according to the Schrödinger equation. If the initial state is $S(t_0+)$ and the next observation occurs at t_1 , then the evolution equation is:

$$S(t_1-) = U(t_1-t_0) S(t_0+) U(t_1-t_0)^* \quad (1)$$

where

$$U(t) = \exp\{-iHt/\hbar\} \text{ and} \quad (2)$$

$$U(t)^* = \exp\{iHt/\hbar\},$$

and H is a Hermetian operator called the *Hamiltonian*.

To summarize, suppose the player/environment system was in state $S(t_0+)$ when last measured, it is now time t_1 , the Hamiltonian is H , and there are k binary questions the player could ask, represented by projections P_{Q_1}, \dots, P_{Q_k} . Then the current state is $S(t_1-)$ as given by (1) and if question P_{Q_i} is asked the system will be in state $P_{Q_i}S(t_1-)P_{Q_i}$ with probability $\text{Tr}(S_1(t_1+)/\text{Tr}(S(t_1-)))$ and $(I-P_{Q_i})S(t_1-)(I-P_{Q_i})$ with probability $\text{Tr}(S_0(t_1+)/\text{Tr}(S(t_1-)))$. To predict the state of the system with as much precision as allowed by Nature would require knowledge of the operators $S(t_0-)$, H , P_{Q_1}, \dots , and P_{Q_k} , together with the mathematical rules of quantum mechanics. Physicists have tested the theory by conducting experiments on systems with sufficiently few degrees of freedom that precise knowledge of the operators is possible. The results of such empirical tests have been accurate to an extremely high degree of precision.

Thus, the state of a quantum system is represented as an operator, and the states of subsystems are represented as operators acting on operators. Schrödinger evolution is also represented as the action of an operator. In other words, the state of the world, the player's representation of the state of the world, how the state of the world changes, and the connection (projection operation) between states of the world and representations are all modeled as entities of the same kind: density operators. This is desirable if the informational component of the theory is to be capable of self-representation.

For a number of reasons, it seems reasonable to assume that the state space of observables is countable. First, the integers form the smallest set sufficient for self-representation. It is desirable not to extend beyond this minimal set unless necessary. Second, observables in quantum theory are discrete events with (at least in practically realizable experiments) finite sample spaces. Third, energy in nature appears to be quantized in discrete packets that come in integral multiples of a fundamental unit called the quantum of action. Finally, contemporary information theory is based on discrete logic. Therefore, it is assumed that the sample space of density operators for player representations is countable. Later it will be proposed that an appropriate state space consists of counts of events in the system's

memory, which form sufficient statistics in the player's current world model.

3.3 PREQUENTIAL QUANTUM LOGIC

Figure 1 shows the Stapp model of quantum dynamics from the point of view of a given player. The diagram represents the player's choice of observation event as an influence diagram. The state evolves according to Schrödinger's equation between observation events. The player chooses when to observe the system and what question to ask. Nature then chooses which subspace the state is projected onto. The player receives an answer to the question that was posed. There is also a change in the player's knowledge state as information is acquired from the observation event and/or lost due to thermodynamic considerations. The degree of satisfaction the player experiences with the outcome is modeled as a value node. This value node represents the psychological criteria a player might use to select among observation policies.

In this model, then, a quantum system consisting of a player and its environment is represented by:

- a Hamiltonian operator governing the system's evolution when the player is not posing questions (i.e., is “operating without conscious control”);
- a projection operator corresponding to the player's body-world schema;
- a finite set of projection operators corresponding to possible questions the player can pose to nature.

Let us simplify this picture down to the barest essentials needed to build up a quantum logic. Consider a player with a single question Q it knows how to ask and to which it can recognize the answer. Denote the corresponding projection operator by P_Q . If the player's current state is $S(t_-)$, then the answer to the question is “yes” if $S(t_+) = P_Q S(t_-) P_Q$ and “no” otherwise. The corresponding probabilities are as given in the above table.

Stapp (1998) states that if the question is asked sufficiently rapidly, the quantum Zeno effect implies that the Hamiltonian H can be replaced (to extremely close approximation) by an “uncoupled” Hamiltonian given by $P_{Q_0} H P_{Q_0} + (I - P_{Q_0}) H (I - P_{Q_0})$. That is, by asking the question frequently, the player keeps the “yes” and “no” states from becoming entangled. Slowing down the rate of question-asking restores entanglement. If the same question is repeated in quick succession there is a high probability that the two answers will be the same (Stapp, 1998). This is because evolution of quantum states is local and the state will not have had sufficient time to spread out. Therefore, this simple one-question player can exert control over both itself and its environment by controlling the rate at which it asks its question.

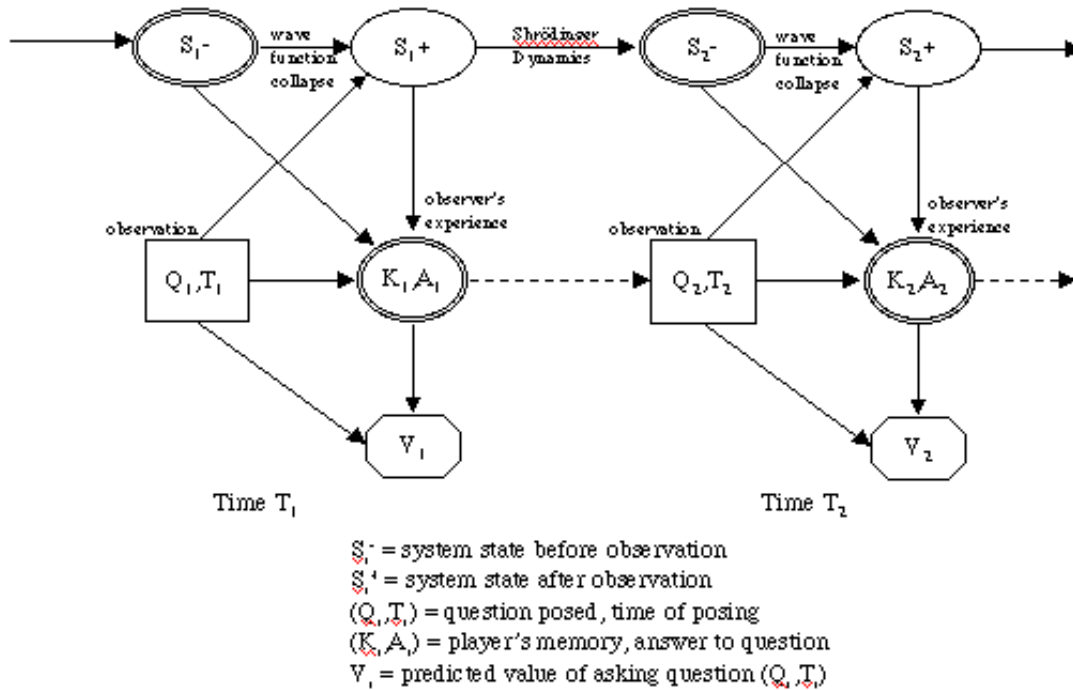


Figure 1: Player's Decision Problem in Prequential Quantum Dynamics

Suppose both the “yes” and “no” answers to Q are moderately stable within the range the player can control. That is, suppose event rates achievable within the range of control of the player’s attention keep the answer highly stable, but a moderate slowing of the event rate produces appreciable spreading. Physically, this would correspond to moderately deep “potential wells,” or quasi-stable electrochemical states in the brain. If this assumption holds, it gives the player a lever of control over evolution of the combined player/environment system. By speeding up the event rate, the answer can be kept fixed at what it was previously; by slowing down the event rate the probability of reversal can be allowed to accumulate. In other words, the player can vary the effort devoted to question asking in order to control the probability with which the answer to Q stays the same or “flips.” We have constructed a simple quantum logic gate that enables the player to choose whether to execute either a “read and copy” or “read and randomize” operation. This simple quantum logic gate corresponds to the influence diagram fragment shown in Figure 2. Note that the effort devoted to asking the question is under the player’s control, but the answer to the question is provided by Nature. By controlling the pattern of effort, the player can choose either the “copy” or the “randomize” operation.


It is this “randomize” operation that differentiates the quantum computer from a digital computer. The two basic operations of a classical Turing computer are “read and copy” and “read and negate.” We propose founding post-classical computing theory on two very simple changes: (1) replacing “read and negate” by “read and randomize;” and (2) permitting a small probability of error in the “read and copy” operation. Not only are these changes required if computing theory is to respect physical constraints, but the randomization operation provides the statistical basis for a genuinely causal learning theory. Moreover, the random elements in the quantum computer gates are free choices of nature, raising the possibility that such computers could at some point evolve to achieve genuine consciousness.

Stapp (1998) argues that animal brains could realistically be expected to exhibit quantum logic gates such as that of Figure 1. That is, at event rates consistent with the time scale of neural pulses, quantum Zeno blocking of transitions between “yes” and “no” answers would not be disrupted by the strong decoherence effects expected to occur in brains. Tegmark (forthcoming) argues that the approach of Penrose and Hemeroff to associate mind with quantum theory cannot succeed because of the strong decoherence effects likely to be present in the brain. However, his argument does not apply to the approach advocated here (Stapp, personal communication).

To summarize, the assumptions needed for a model associating conscious control with a quantum logic gate such as that of Figure 2 are:

1. The player can recognize yes and no answers to its question;
2. The player can control the rate of occurrence of observation events by exerting or relaxing effort;
3. The experimentally verified quantum Zeno effect occurs at achievable event rates for the player's question Q;
4. Appreciable entanglement of answers to the player's question Q occurs at achievable event rates.

These assumptions are fully consistent with our current knowledge of the physics of brains.



$$\begin{aligned}
 P(\text{Answer_Q}(n+1) = \text{Yes} \mid \text{Answer_Q}(n) = \text{Yes}, \text{Effort_Q}(n) = \text{High}) &= \text{LARGE} \\
 P(\text{Answer_Q}(n+1) = \text{Yes} \mid \text{Answer_Q}(n) = \text{No}, \text{Effort_Q}(n) = \text{High}) &= \text{SMALL} \\
 P(\text{Answer_Q}(n+1) = \text{Yes} \mid \text{Answer_Q}(n) = \text{Yes}, \text{Effort_Q}(n) = \text{Low}) &= \text{MODERATE} \\
 P(\text{Answer_Q}(n+1) = \text{Yes} \mid \text{Answer_Q}(n) = \text{No}, \text{Effort_Q}(n) = \text{Low}) &= \text{MODERATE}
 \end{aligned}$$

Figure 2: Elementary Operation in Prequential Quantum Logic

A full quantum logic can be built up out of network fragments like that of Figure 1. For example, the player could execute a “read and negate” operation by relaxing effort until a flip occurs and then focusing on holding the negated answer in place. An “and” operation could be executed by examining the answers to the last two questions that were asked, exerting high effort if both answers were “Yes,” and otherwise executing the “read and negate” operation. Similarly, an “Or” operation can be compiled out of “Not” and “And” operations, and so on. Thus, out of our player’s original question Q we can build up a large array of questions the player can ask, and therefore behaviors the player can exhibit. The player achieves these behaviors by focusing its attention in appropriate patterns of high and low effort.

Prequential logic operations are built up by composing policies in which the player predicts approximately when a desired subgoal state is likely to occur, relaxes attention until the occurrence is predicted to take place, watches for the event, and then focuses on holding it in place until an operator to execute the next subgoal can

be spawned. Complexity of achievable behaviors is limited by the degrees of freedom of the player’s internal state and the structure of the Hamiltonian in the local space-time region the player occupies. Clearly, the ability to construct representations that accurately reflect the local structure of the Hamiltonian would tend to have survival value. Perfect control by players will not be achievable. It is natural to assume that players are subject to fatigue, and therefore that the probability of error increases with the total amount of effort required to produce a given compiled operation. We describe in the next section a physically motivated informational analogue of fatigue.

4. THE PREQUENTIAL GAME

4.1 PREQUENTIAL LEARNING

This section describes an extension to the player model presented above into a decision theoretic learner. To do this, we give the player a memory and specify a belief and decision model for the player. We imagine that the player’s memory consists of traces of past experiences, and we interpret question-asking as an exchange in which information is traded for energy. Specifically, we assume the player stores counts of the number of times it remembers each of the 16 possible current/next effort/answer combinations. When an experience occurs, the count in the appropriate cell is incremented, and then with a small probability chosen by Nature, one of the counts in the player’s memory is randomly decremented. This small probability is interpreted as the energy cost of information. We can model fatigue by assuming the count is decremented more frequently (i.e., more energy is dissipated as entropy) when the effort is high. We assume the player has player-specific value function that rewards “1” answers and low effort, with a tradeoff rate that may vary with the number of observations in the player’s memory. We assume a Dirichlet conjugate prior distribution for which the hyperparameters are also player-specific adjustable parameters. Using this model, the predictive distribution for value and the player’s optimal action can be computed in closed form.

Clearly, more complex informational structures are possible, in which players remember a longer history and are capable of greater look-ahead. The following section describes an evolutionary *prequential game* in which players of increasing complexity can emerge.

4.2 RULES OF THE GAME

In prequential quantum dynamics each player is characterized by a question, an internally controllable “effort lever,” a prior distribution and utility function, and an internal memory where it stores information from past experience. Using its memory and its internal representation, it makes predictions and evaluates the

desirability of different effort allocation policies. Players interact with their environment by exerting effort to ask questions in exchange for information about the environment. Players prefer "Yes" answers and low effort. Therefore, players are better off to the extent they can keep the answer to their question in the "yes" range without expending too much effort.

In this section, we consider interaction among players. Specifically, suppose the "Yes" subspaces for two players are highly correlated. Then these players can benefit from "cooperating" by sharing the effort needed to stay in the "Yes" subspace. However, each player would prefer to let the other player expend more of the effort. This suggests an interaction architecture among players that has the structure of a mixed-motive game. Theoretical, computational, and empirical evidence has demonstrated that mixed-motive games of appropriate structure can lead to the emergence of highly complex behavior patterns.

To fill out a full specification of an evolutionary information exchange game, we need to specify rules for the birth, death and reproduction of players. Although our understanding is rudimentary of how this occurs in natural systems, theory and simulations of complex adaptive system models has provided important insights. Myers et al (1999) describe a physically and biologically inspired learning method called population Markov Chain Monte Carlo, which could be developed into a theoretical framework for evolving populations of players. Such an approach could provide the basis for designing prequential game architectures. A missing piece is a theoretical basis for the distinction between multiple player "ecologies" and "emergent players." The former are self-reinforcing patterns of behavior executed by collections of players acting in concert, whereas the latter represent emergence of a genuinely new type of "executive player" that orchestrates an ecology and causes it to act as a single player. We hypothesize that emergence of a new player is a genuine quantum phenomenon, as yet poorly understood, but perhaps based on a similar physics as the familiar spin-coupling of particles. If such a phenomenon occurs in physics, evolution would tend to select for emergent players when they arose because of the efficiencies possible with executive control. However, to remain stable, emergent players would need to retain the flexibility and adaptability of characteristic of the ecology.

It appears that this process by which ecologies evolve into unitary players, although extremely rare, has happened at a number of occasions in the past. Examples include the emergence of life, the incorporation of prokaryotes as organelles in eukaryotic cells, and the emergence of multicellular organisms. The process of reproduction, or "splitting off" new players from existing players, is by contrast quite common in nature. Reproduction may be analogous to the disentanglement of spin-coupled systems. Our

understanding of the physics of birth, reproduction, death, and emergence of new types of players is quite rudimentary. It is hoped that the viewpoint espoused here may lead to useful empirically testable hypotheses regarding the mechanisms by which they operate.

5. PREQUENTIAL COMPUTING

There is a major philosophical difference between the approach advocated here and standard work in quantum computing (e.g., Gershenfeld, and Chuang, 1998). Standard quantum computing architectures attempt to create isolated, internally coherent many-qbit systems. Such systems appear to be quite difficult to engineer. Prequential theory suggests a design based on the kind of loose coupling that appears to be common in natural complex adaptive systems. In prequential quantum computing, information influences would be engineered by design, and might also be designed to adapt with experience. Parallelism in a prequential quantum computer would be achieved by deliberate randomization of very small numbers of qbits at any given gate. It is the loosely coupled architecture that gives rise to complex behavior.

To implement a quantum logic gate such as Figure 2 in hardware would require identifying a physically achievable system for which the Hamiltonian corresponded to the desired probabilities. The figure specifies a distribution only for the answer bit. This distribution is conditional on effort (i.e., sampling rate of a detector) and the previous answer. Which of its two actions the effort node of Figure 2 selects is, to our best current physical knowledge, a free choice of nature. However, empirically the choice appears to be governed to extremely high accuracy by the probability laws of quantum theory. Therefore, to complete the specification of how such a quantum system would behave, we would need to specify the information available to the quantum gate in making its choice, and engineer the connections appropriately. Physically this would be implemented by allowing the potential field influencing $Effort_Q(n)$ to be affected by the information predecessors to $Effort_Q(n)$.

The feasibility of hardware implementation of a computing architecture such as the one described here is unknown. If such architectures are realistically achievable given current design technology, we hypothesize that appropriately designed quantum randomization could result in major efficiencies over pseudo-random digital simulations for the same problem. Moreover, if quantum computers could be built up out of a logic in which the fundamental "rules" could be interpreted as causally efficacious free choices by autonomous "players," this would make the design of natural human-computer interaction languages much more straightforward.

The approach presented here is a framework for designing a system of interacting players that make free choices, and structuring their information flows in such

a way that their behavior computes something we wish to know. Some may prefer to think of the “atomic players” of prequential logic as proto-conscious agents exercising free will; others may prefer to think of them as natural random number generators. In either case, a well-designed “society” of interacting players might be capable of computing things not possible using today’s digital computers. The possibility exists that genuine engineered intelligence could emerge at some point in a prequential quantum computer.

Acknowledgements

The research reported in this paper was sponsored in part by a career development fellowship from the Krasnow Institute at George Mason University. I am indebted to Henry Stapp for an extensive email dialogue on the ideas presented here. This paper would not have been possible without his contribution. Acknowledgement is also extended to Phil Barry, Menas Kafatos, Paul Lehner, Tod Levitt, Randall Morck, Harold Morowitz, Jim Myers, Fred Alan Wolf, and many others too numerous to list here. These individuals deserve much of the credit. Any blame to be assigned is mine alone.

Bibliography

Anderson, J.R. (1999) *Cognitive Psychology and Its Implications*, Worth Publishing.

Bohr, N. (1934), *Atomic Theory and the Description of Nature*. Cambridge: Cambridge University Press.

Dawid, A.P. and Vovk, V.G. (1999), Prequential Probability: Principles and Properties, *Bernoulli*, 5: 125-162.

Gershenfeld, N. and Chuang, I. (1998) Quantum Computing with Molecules, *Scientific American*, June 1998.

Gigerenzer, G., Swijtink, Z., and Daston, L. (1990) *The Empire of Chance : How Probability Changed Science and Everyday Life*, Cambridge: Cambridge University Press.

W. Itano, D. Heinzen, J. Bollinger, D. Wineland (1990), *Phys. Rev.* 41A, 2295-2300.

James, W. (1890) *Principles of Psychology*, Dover.

Laskey, K.B. and Mahoney, S.M. (1997), Network Fragments: Representing Knowledge for Constructing Probabilistic Models. *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, San Mateo, CA: Morgan Kaufmann.

Myers, J., Laskey, K.B. and Levitt, T. (1999) Population Markov Chain Monte Carlo, *Proceedings of the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence*

Penrose, R. (1994) *Shadows of the Mind*, Oxford University Press.

Searle, John R. (1980), *Minds, Brains, and Programs*, Behavioral and Brain Sciences 3: 417-457. (Lockwood and SEL, QP360 .B425)

Shankar, R. (1994) *Principles of Quantum Mechanics*, Plenum.

Sowa, *Knowledge Representation*.

Stapp, H. (1999), Science of Consciousness and the Hard Problem. *The Journal of Mind and Behavior* 18, 171-193.

Stapp, H. (1998) Whiteheadian Process and Quantum Theory, conference on Whitehead.

Sutherland, K. (1995), Conversations with Zombies. *Journal of Consciousness Studies* 2(4), 312-372.

Tegmark, M. (forthcoming)

Wigner, E. (1967), Remarks on the Mind-Body Problem. in *Symmetries and Reflections*, Indiana University Press, Bloomington, 171-184.

von Neumann, J. (1932), *Mathematical Foundations of Quantum Mechanics* (Ch. VI), Princeton, NJ., Princeton University Press.