Quantum Quantum Quantum

by Edward Farhi

Cloning, Money, and Monogamy

and Aram Harrow

In recent years, quantum mechanics has come to be seen not only as a theory of what happens at very short length scales, but also as a theory of information. This is analogous to the way that relativity is not merely a theory of objects traveling close to the speed of light, but also changes the way we define space and time. What does this new perspective offer?

By thinking about quantum mechanics as a theory of information, we gain the ability to understand it in ways that are independent of the specific physical realization. For example, classical computer science is able to describe the same information-processing techniques whether they are implemented on a big abacus or a modern computer cluster. Likewise, quantum information science uses tools that are equally applicable to quantum information processing in superconducting systems, trapped ions, or an array of spin-½ particles. Once we consider quantum information abstractly, we can explore how to use it

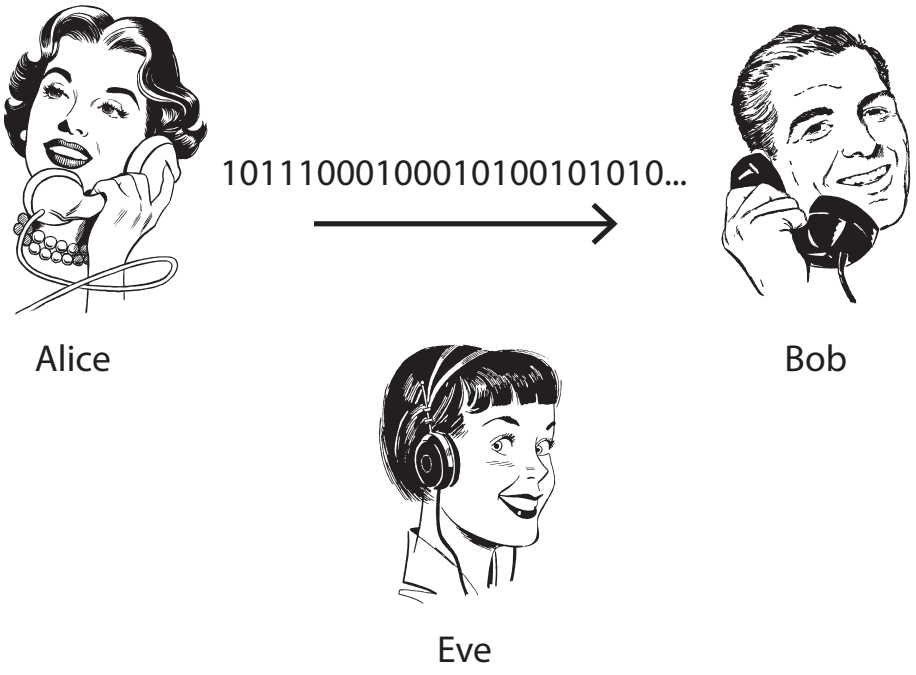

FIGURE 1

Alice wants to send a message to Bob *while preventing Eve from learning its contents. Quantum mechanics can help make this possible.*

for communication or computation tasks, some of which would be unimaginable without quantum mechanics.

One example of this shift in perspective is the venerable Heisenberg uncertainty principle. Originally, this was viewed as a limitation on our ability to simultaneously measure the position and momentum of a particle. The uncertainty principle applies more generally to other pairs of measurable quantities, that we call COMPLEMENTARY OBSERVABLES. For example, consider a spin-½ particle such as an electron. What it means to have spin ½ is that when the Z component of the spin (the electron's intrinsic angular momentum) is measured, the only possible outcomes are $+\hbar/2$ or $-\hbar/2$, instead of the continuous range of possibilities we would expect for a component of a vector. There is nothing special about the Z direction: we can measure the component of the spin along any axis, and only ever obtain outcomes $+\hbar/2$ or $-\hbar/2$. Measurements along different axes, such as the X and Z directions, correspond to complementary observables. A particle known to have spin $+\hbar/2$ in the Z direction cannot have a definite value of the X component of its spin. This means that measuring such a particle in the X direction would yield each of the two possible outcomes $\pm\hbar/2$ with 50% probability. This limitation in our ability to simultaneously know two complementary properties is not a technological failure, but is a fundamental property of quantum mechanics.

Surprisingly, this apparent limitation can be made useful if we understand it in the language of information. Suppose you want to send a message and it is important to you for military, financial or romantic reasons that this message not be read by anyone other than the intended receiver. If it is read, you need to know it. Quantum mechanics can make this possible.

Think of a message as a string of zeros and ones (bits), sent by Alice to another party, Bob *(Figure 1)*. Classically, nothing would stop an eavesdropper (call her Eve), who intercepts the message by copying it and sending it on to Bob without Alice or Bob realizing that the message had been read. However, instead of sending classical bits, Alice could send quantum mechanical bits, which we call QUBITS. These could equally well be spin-½ particles, photon POLARIZATIONS or any other two-level quantum system. A bit value can be encoded into a qubit by representing 0 with spin $+\hbar/2$ in the Z direction, and 1 with $-\hbar/2$ in the Z direction. Another choice is to use the X direction. Since these are complementary observables, a qubit that is encoded in the X direction and measured in the Z direction will give a random answer that reveals nothing about the original bit value. This quantum information theoretic primitive can be used to devise schemes in which any attempt at eavesdropping can be detected. For example, Alice could send her message with some X- and Z-encoded qubits and not tell Bob what her choices were until after he receives the message. Eve's eavesdropping would necessarily corrupt some of the message, and thus can be detected by Alice and Bob. Using photons, messages have been sent over hundreds of kilometers with security ensured by the laws of quantum mechanics.

Suppose you want to send a message and it is important to you for military, financial or romantic reasons that this message not be read by anyone other than the intended receiver. If it is read, you need to know it. Quantum mechanics can make this possible.

Quantum physics has implications for computation as well as for communication. In 1981, Richard Feynman observed that simulating a quantum system on a classical computer requires effort exponential in the size of the quantum system, which means that even small systems are far beyond the reach of supercomputers. However, he suggested that a quantum computer, built to take advantage of quantum law, could simulate quantum systems more efficiently. This was the first clue that quantum computing might be qualitatively different from even very advanced classical computers. Since then, there have been many discoveries of quantum algorithms that promise significant speedup over classical algorithms, even for tasks that on their face have nothing to do with quantum physics. These include algorithms for factoring integers; searching; evaluating the winner of two-player games; and solving large systems of linear equations.

Quantum cloning and quantum money

One simple but deep difference between classical and quantum information is that classical information can in principle always be copied and quantum information cannot. This principle is called the QUANTUM NO-CLONING THEOREM.

Suppose someone gives you an electron but doesn't tell you its spin direction. If we could perfectly copy the state then we could measure the spin in the X direction of one copy and the spin in the Z direction of the other, thereby learning more about the original state than the uncertainty principle permits. This example shows how

the uncertainty principle implies a form of the no-cloning theorem, but in fact it is a more general result: there is no physical process that can be used to copy an unknown quantum state without damaging the original.

The no-cloning theorem is really at the heart of the security of the quantum cryptography scheme discussed above, since when an eavesdropper reads a message and stores it (even in her brain), she has made a copy. The quantum no-cloning theorem forces a tradeoff between how much an eavesdropper can learn about a message and how much damage must be caused in the process. It is this damage, like a resealed envelope, that alerts the receiver to the security breach.

You can carefully examine a paper bill offered to you to convince yourself that it is not a counterfeit and your checking will hopefully not degrade the bill. In the quantum case, any attempt to measure the qubits without knowing the spin axes necessarily damages the qubits.

An intriguing application of quantum no-cloning is the possibility of using it to prevent the counterfeiting of money. The first scheme for quantum money was proposed by Stephen Wiesner, son of MIT president Jerome Wiesner, in the same 1970 paper where he proposed quantum cryptography. This was long before quantum information became an area of study, and several journals rejected Wiesner's manuscript. Weisner's money idea was to have each bill include a bunch of qubits, say spin-½ particles, as part of the bill. The person holding the bill would not know the spin axis of each of the qubits. The quantum-no-cloning theorem guarantees that if the bill holder tries to copy it, he damages it and cannot end up with two good bills. So far, so good, because it is counterfeit-proof. The problem is that the bill holder has no way of verifying that the quantum bill is a valid bill.

You can carefully examine a paper bill offered to you to convince yourself that it is not a counterfeit and your checking will hopefully not degrade the bill. In the quantum case, any attempt to measure the qubits without knowing the spin axes necessarily damages the qubits. Why not then imprint the spin axes on the bill? If the bearer knows the spin axes then he can measure the spins along those directions, telling him whether each one is a + or – $\hbar/2$. However, if he now knows the spin axes and the spin values he can make as many copies as he wants. (The quantum-no-cloning theorem only applies to unknown states.) We see that there is a tension between copy prevention and having the ability to verify.

To get around this, an MIT group [1] proposed a quantum money scheme using highly entangled qubits. Here we sketch their approach. They borrowed some ideas from topology having to do with the properties of knots. A knot is a closed loop of string in a three-dimensional space. The quantum bill contains a bunch of qubits in an entangled superposition of states. Each state encodes a knot and the superposition is over knots that can be deformed into each other, but do not look alike. (Imagine a simple closed loop of string that has been tangled up. Although topologically it is a simple loop, this can be hard to see when it is very tangled.) Such states can be verified by a merchant who possesses a quantum lab by making a particular kind of measurement that will not damage a valid quantum money state. On the other

hand, the no-cloning theorem rules out the possibility of copying the state without making use in some way of its structure. Could a counterfeiter exploit the structure of the knot state in order to copy it? This cannot be ruled out, but it appears to require solving a hard problem in topology, and no scheme for doing so has been proposed to date.

Entanglement and information theory

Entanglement is a feature of quantum mechanics that clearly separates quantum theory from classical probability theory. Measuring subsystems of quantum systems can yield correlations that would be impossible in the purely classical world. This was experimentally demonstrated in the 1970s and 1980s, following a proposal of John Bell's from 1964.

Here we will give an information-theoretic view of entanglement by considering a simple game *(Figure 2)*. In this game, Alice and Bob are working together to maximize the probability of a certain winning outcome, but are isolated from each other (space-like separated). The game show host sends a random bit a to Alice and a random bit b to Bob. Since they cannot communicate, Alice knows a but not b, and similarly Bob knows only b. Next, Alice sends the host a bit x and Bob sends the host a bit y. Alice and Bob win if $(x+y) \bmod 2 = ab$; in other words, if a and b are both 1, then Alice and Bob should output different values, otherwise they should output the same value.

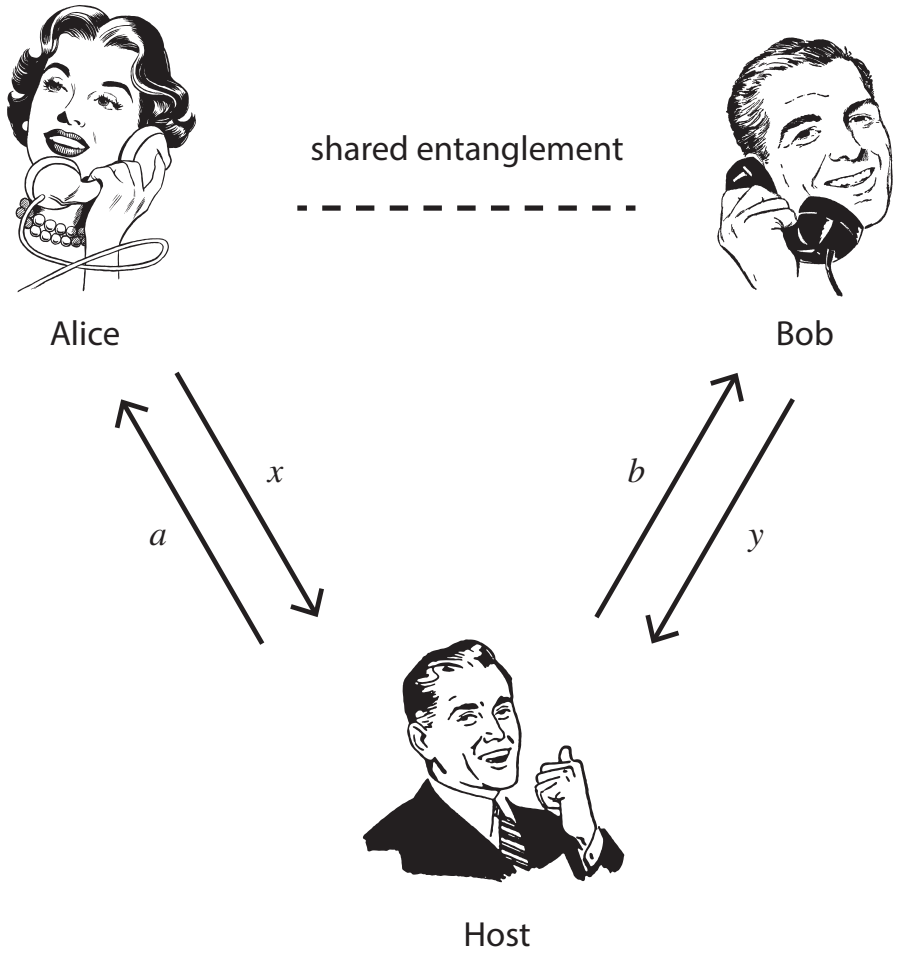

FIGURE 2

Alice and Bob together are cooperating *to try to win at a game. In this game, the host sends a question* a *to Alice and a question* b *to Bob. Alice replies with answer* x *and Bob with answer* y. *It turns out that if Alice and Bob share an entangled quantum state it can improve their probability of winning.*

One simple strategy is for Alice and Bob to always output 0. This will win whenever $ab=0$, which occurs with probability ¾. You can check—by looking at all possible inputs and outputs—that no matter what strategy Alice and Bob have agreed upon before separating, they cannot win with probability greater than ¾. Surprisingly, if Alice and Bob each have a part of an entangled quantum state, they can win this game with probability ½ + 1/√2, which is larger than ¾.

The entangled state used here is the singlet, or spin-0, state of two spin-½ particles. Alice takes one of the spin-½ particles and Bob the other, and then they separate and hear from the host. The fact that the total spin is 0 means that if Alice and Bob measure their own particles along the same axis they find that the outcomes are perfectly anti-correlated. To win this game, Alice chooses a measurement axis that

depends on a, and Bob one that depends on b. Their outputs x and y will come from their measurement outcomes. A careful choice of measurement axes (depending on a and b), gives awinning probability better than anything achievable classically.

Entanglement is more broadly useful than in helping to win hypothetical games. It is at the heart of the speedup achieved by quantum algorithms designed to run on quantum computers. This is because n unentangled systems can be described using n times the amount of information needed to describe a single system, whereas n entangled systems can require an amount of information that is exponential in n. A large quantum computer cannot be simulated by a classical device and can do things a classical computer cannot. Entanglement can also be used for quantum teleportation, cryptography, clock synchronization and many other tasks.

Monogamy of entanglement

While entanglement is more useful than classical correlations for many tasks, it also has limitations that do not apply classically. Correlations can be shared among many systems without restriction. However, entanglement cannot be. A strong correlation between the weather in Cambridge and Boston does not prevent a strong correlation existing between the weather in Cambridge and Somerville.

However, if spins A and B are in a singlet state, then spins A and C cannot also be in a singlet state. This can be seen as a consequence of the uncertainty principle. Suppose that A and B are in a singlet state and that A and C are also in a singlet state. Then we could measure B along the X axis and C along the Z axis, thereby learning the value of A along both of these axes, a contradiction.

The general phenomenon of non-sharability is called the MONOGAMY OF ENTANGLEMENT. This principle can explain why so many quantum systems exhibit behavior that is effectively classical. If a particle is equally entangled with many other particles, then it has very little entanglement with any particular particle. In a lab, if you want to create entanglement between two systems, you need to limit their entanglement with other systems.

So far we have discussed entanglement as a property of pairs of particles, but in many condensed-matter systems there can be entanglement between subsystems, each consisting of many particles. This kind of entangled state cannot be represented as a product of n one-body wavefunctions, or even as a product of $n/2$ two-body wavefunctions. Modeling these wavefunctions can be a difficult task for classical computers, and is one of the hoped-for applications of a quantum computer. Still, we would like to use insights from quantum information theory to learn when manageable representations of the wavefunction can be found.

Many-body quantum Hamiltonians are typically sums of local terms, each involving a pair of particles with a short-range interaction between them. Finding the ground state of a single term is straightforward, but when the local ground states are incompatible, finding the global ground state is a difficult optimization problem. Usually this compromise of competing interactions is seen in terms of energies, but sometimes it can be understood in terms of entanglement. One might expect that short-range interactions give rise to short-range entanglement so that the amount

of entanglement between two regions scales like the surface area of their interface. However, in some cases entanglement can be proportional to *volume*, meaning that large amounts of long-range entanglement exist. These different cases can lead to observably different thermodynamic properties, and it is a major open question in condensed-matter physics to determine which type of fundamental interactions cause which behavior.

While entanglement in many-body quantum systems often leads to unusual physical phenomena, it is also worth understanding when it is actually rather limited. If each particle interacts with sufficiently many other particles, then monogamy of entanglement implies that most of these interactions should be between nearly unentangled pairs. This limits the possible contribution of entanglement to the physics of a many-body system, and is the justification behind the common "mean-field" approximation in which the ground state is approximated by an unentangled state. Using monogamy of entanglement, the mean-field approximation can be rigorously proven to be valid in settings where it was previously only conjectured, such as when no assumption of symmetry is made. And in some cases, intuition from the monogamy of entanglement can be used to even help efficiently find that state. [2] This is an example of how ideas from quantum information can be used to show when quantum systems can be efficiently modeled on classical computers.

While entanglement in many-body quantum systems often leads to unusual physical phenomena, it is also worth understanding when it is actually rather limited.

Another, more speculative, application of the monogamy of entanglement is to the black hole information problem. The problem comes from the fact that (1) general relativity seems to imply that information is lost when a black hole is formed and then evaporates; and (2) quantum mechanics says that information can never be destroyed, only rearranged. An apparent resolution is that information that falls into a black hole is encoded in the outgoing Hawking radiation. However, difficulties arise when we consider throwing one half of an entangled pair into a black hole. Consider the half that remains outside the black hole. General relativity predicts that this should be entangled with degrees of freedom inside the black hole, while quantum mechanics predicts that it will be entangled with the outgoing Hawking radiation. These predictions would be compatible if we were talking about classical correlation, but for entanglement, the monogamy property makes them incompatible. Resolving this problem is an active area of research and has brought together string theorists, cosmologists and quantum information researchers.

REFERENCES

[1] "Quantum Money from Knots," Edward Farhi, David Gosset, Avinatan Hassidim, Andrew Lutomirski, Peter Shor; Proceeding ITCS '12: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, pp. 276-289; ACM New York, NY, USA.

[2] "Product-state Approximations to Quantum Ground States," Fernando G.S.L. Brandão and Aram W. Harrow; Proceedings of the 45th Annual ACM Symposium on Theory of Computing, pp. 871-880, 2013.

EDWARD (EDDIE) FARHI *was trained as a theoretical particle physicist but has also worked on astrophysics, general relativity, and the foundations of quantum mechanics. His present interest is the theory of quantum computation. Farhi has been studying how to use quantum mechanics to gain algorithmic speedup in solving problems that are difficult for conventional computers. He and Sam Gutmann proposed the idea of designing algorithms based on quantum walks, which has been used to demonstrate the power of quantum computation over classical. In 2000, they along with Jeffrey Goldstone and Michael Sipser, introduced the idea of quantum computation by adiabatic evolution, which is the basis of the design of a, somewhat controversial, commercial quantum computer. In 2007, Farhi, Goldstone and Gutmann showed that a quantum computer can determine who wins a game faster than a classical computer. In 2010, he along with Peter Shor and others at MIT introduced a scheme for Quantum Money which so far has resisted attack.*

Eddie Farhi went to New York City's Bronx High School of Science, and Brandeis University, before getting his PhD from Harvard University in 1978. He was then on the staff at the Stanford Linear Accelerator Center and at CERN in Geneva before coming to MIT, where he joined the faculty in 1982. At MIT, he has taught undergraduate courses in quantum mechanics and special relativity. At the graduate level he has taught quantum mechanics, quantum field theory, particle physics and general relativity. Farhi won three teaching awards at MIT and in 2000, 2001, and 2002, he lectured the big freshman physics course, "8.01." In July 2005, he was appointed the Director of MIT's Center for Theoretical Physics, a position he still holds.

ARAM HARROW, *Assistant Professor of Physics, works on the theory of quantum information and quantum computing. He has invented several new tools for quantum communication, including the idea of "coherent classical communication" which is a way of building protocols for quantum communication out of much simpler protocols for classical communication. Harrow has also done foundational work on the role of representation theory in quantum algorithms and quantum information theory. Together with Avinatan Hassidim and Seth Lloyd, he has developed a quantum algorithm for solving large linear systems of equations. More recently he has worked on the theory of quantum entanglement, with applications to testing entangled states and understanding condensed matter systems.*

Harrow grew up in Michigan before attending MIT for his undergraduate (math and physics, 2001) and graduate (physics, 2005) degrees. He then spent seven years away before returning to MIT, during which time he was a lecturer in the math and computer science departments of the University of Bristol for five years and a research assistant professor at the University of Washington for two years.