

Optimal Information Transfer and Real-Vector-Space Quantum Theory

William K. Wootters

1 Introduction

In 1936 Birkhoff and von Neumann initiated an axiomatic approach to the foundations of quantum mechanics, taking as their starting point postulates inspired by classical logic but adapted to the peculiar features of quantum theory [1]. Though they showed that many characteristics of quantum theory could be captured in this way, they could also see that their logical approach would not lead uniquely to standard quantum theory. In particular they noted that along with standard complex-vector-space quantum theory, the postulates could just as well be satisfied by a theory based on a real or quaternionic Hilbert space.¹

Over the years other authors have taken other approaches to axiomatization and have found reasonable assumptions that favor the complex theory over the real and quaternionic models. One successful strategy along these lines has been to insist on the existence of an uncertainty principle of a specific form [5–7]. Another approach put forward by several authors relies on the fact that in standard quantum theory, it is possible to carry out a complete tomographic reconstruction of the state of a multipartite system entirely by means of local measurements on the individual components (taking into account correlations), with no need for global measurements on pairs of subsystems [8–14]. The real-vector-space theory does not have this property; so by adopting local tomography as an axiom, one rules out the real case. Surely, though, much of the appeal of these arguments comes from the fact that they succeed in leading us to what we believe to be the correct answer. If we had found ourselves living

¹In fact, from this logical starting point, it has turned out to be difficult even to narrow the set of possibilities to just these three theories—real, complex, and quaternionic. Major theorems along these lines can be found in Refs. [2, 3]; see also Ref. [4]. Still, nothing in these papers favors the complex theory over the real or quaternionic theory.

W.K. Wootters (✉)

Department of Physics, Williams College, Williamstown, MA 01267, USA
e-mail: William.K.Wootters@williams.edu

© Springer Science+Business Media Dordrecht 2016

G. Chiribella and R.W. Spekkens (eds.), *Quantum Theory:*

Informational Foundations and Foils, Fundamental Theories of Physics 181,
DOI 10.1007/978-94-017-7303-4_2

in a world that seemed to be well described by real-vector-space quantum theory, we would not have regarded it as a logical problem that tomography requires global measurements. It would simply be another peculiar feature of quantum theory, like entanglement.² (I admit, though, that the local tomographic property of the complex theory does feel as if it could be a clue to something deeper.)

In this paper I would like to point out a particular property of real-vector-space quantum theory that I find especially intriguing: the transfer of information from a preparation to a measurement is *optimal* (in a sense to be explained below). Standard quantum theory does not have this property. So if we were trying to find a simple set of axioms that would generate real-vector-space quantum theory, we might well find ourselves adopting optimal information transfer as one of our axioms. This property of the real theory has been known for years—it appears in my 1980 doctoral dissertation [16]—but I would like to give a somewhat simpler and more intuitive presentation of it here.

One motivation for studying real-vector-space quantum theory is simply to shed light on the standard theory by comparison. But I would also like to keep open the possibility that the real-vector-space theory might turn out to be of value in its own right for describing our world. Several authors have given us reasons for not discounting this possibility. In a series of papers published around 1960, Stueckelberg and his collaborators developed an alternative formulation of quantum field theory based on a real Hilbert space [6, 17, 18]. In order to allow the existence of an uncertainty principle, Stueckelberg imposes a specific restriction on all the observables of the real-vector-space theory: every observable is required to commute with a certain operator that we can write as $I \otimes J$, where J is the 2×2 matrix $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ and I is the identity operator. (In the context of Stueckelberg's papers I is the identity on an infinite-dimensional real Hilbert space.) In effect, this restriction forces the matrix representing any observable to be composed of 2×2 blocks of the form $\begin{pmatrix} a & -b \\ b & a \end{pmatrix}$. Such 2×2 blocks add and multiply like complex numbers; so the theory becomes equivalent to the usual complex theory. One of the points Stueckelberg and his collaborators make in these papers is that in this formulation the time-reversal operator becomes linear, rather than antilinear as in the complex formulation. Around the same time, Dyson made the same point and argued that by bringing the time-reversal operator into our formalism, we are in effect basing our quantum theory on the field of real numbers [19].

Recently Gibbons has argued that the complex structure in quantum theory is intimately related to the existence of an arrow of time, and that in spacetimes in which such an arrow cannot be defined, one must fall back on a real-vector-space version of quantum theory [20]. In other work, Myrheim has pointed out that if one wants a version of the canonical commutation relation $[x, p] = i\hbar$ in a discrete system with finitely many values of position and momentum, one cannot use standard

²At least in real-vector-space quantum theory, one never needs to make global measurements involving more than *two* subsystems [15].

complex quantum theory: the trace of any commutator is zero in a finite-dimensional space, but the trace of $i\hbar$ is not zero. On the other hand, if we replace $i\hbar$ with $J\hbar$ (the same J as above), both sides of the equation have zero trace and there is no problem [21]. In the present paper I do not particularly build on any of these observations except insofar as they suggest that a real-vector-space version of quantum theory might be used to describe our actual world, and that the theory is worth studying for this reason as well as for whatever insights it might provide about standard quantum mechanics.

I begin in Sect. 2 by saying what I mean by “real-vector-space quantum theory.” Then in Sects. 3 and 4 I present the property of optimal information transfer, first for a two-dimensional state space and then in d dimensions. As I have said, standard complex quantum theory does not have this property, and it is interesting to ask whether a revised statement of the problem might yield a positive answer even in the complex case. This is the subject of Sect. 5. Section 6 then summarizes our findings.

2 Real-Vector-Space Quantum Theory

One can summarize the basic structure of standard quantum theory in the following four statements:

1. A pure state is represented by a unit vector in a Hilbert space over the complex numbers.
2. An ideal repeatable measurement is represented by a set of orthogonal projection operators whose supports span the vector space. When a state $|s\rangle$ is subjected to the measurement $\{P_1, \dots, P_m\}$, the probability of the i th outcome is $\langle s|P_i|s\rangle$. When the i th outcome occurs, the system is left in a state proportional to $P_i|s\rangle$.
3. A reversible transformation is represented by a unitary operator U . That is, for any initial state $|s\rangle$, the operation takes $|s\rangle$ to $U|s\rangle$.
4. A composite system has as its state space the tensor product of the state spaces of its components.

Of course other states, measurements and transformations are possible. Mixed states are averages of projection operators on pure states, and there also exist non-orthogonal measurements and irreversible transformations. But all such generalizations can be obtained from the cases listed above by applying them to a larger system and possibly discarding part of the system. I have chosen the above formulation partly to keep the discussion simple, but also because I do tend to think of orthogonal measurements and pure states as being more fundamental than their generalizations.

The real-vector-space theory has essentially the same structure, except that all vectors and matrices are limited to real components. The only changes in the above list are that “complex” is to be replaced by “real” in item 1, and “unitary” is to be replaced by “orthogonal” in item 3.

One might wonder what the analogue of the Schrödinger equation is in the real-vector-space theory. The Schrödinger equation generates a unitary transformation through a Hermitian operator, the Hamiltonian:

$$i\hbar \frac{d}{dt}|s\rangle = H|s\rangle. \quad (1)$$

If H is time independent, the unitary operator it generates over a time t is $U(t) = e^{-iHt/\hbar}$, since $|s(t)\rangle = U(t)|s(0)\rangle$ solves the differential equation. The analogous equation in the real-vector-space case should have an antisymmetric real matrix in place of $-iH$, since such a matrix generates orthogonal transformations. We can write the differential equation as

$$\frac{d}{dt}|s\rangle = S|s\rangle, \quad (2)$$

where S is an antisymmetric real operator. I like to call S the ‘‘Stueckelbergian’’ in honor of Ernst Stueckelberg (who of course did not use this term). If S is time independent, then the general solution of Eq. (2) is $|s(t)\rangle = e^{St}|s(0)\rangle$.

Another reasonable question is whether, for example, in a two-dimensional real space the operator

$$R = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (3)$$

should be allowed to count as a possible transformation [22]. It is an orthogonal matrix, so according to the above rules it does count. But there is no 2×2 Stueckelbergian that can generate this operator. This is because the operator R represents a reflection, not a rotation, and there is no continuous set of orthogonal transformations on a two-dimensional real space that takes us from the identity operator to a reflection operator.

Nevertheless, in a real-vector-space world it would still be possible to realize the operation R continuously by bringing in an ancillary two-dimensional system (that is, an ancillary ‘‘rebit’’). To effect the transformation

$$\begin{pmatrix} s_1 \\ s_2 \end{pmatrix} \rightarrow \begin{pmatrix} s_1 \\ -s_2 \end{pmatrix}, \quad (4)$$

we can perform a controlled rotation on our ancillary rebit, conditioned on the state of the original rebit. By rotating the ancillary rebit by half a complete cycle, we can pick up the desired factor of -1 . So it seems reasonable to allow orthogonal matrices with negative determinant to count as possible transformations.

3 Optimal Transfer of Information: The Two-Dimensional Case

Consider the following simple scenario. A stream of photons emerges from a linearly polarizing filter with its preferred axis oriented at an angle θ from the horizontal. Somewhere further along the photons’ path there is a polarizing beam splitter and

a pair of single-photon detectors, which together force each photon to yield either the horizontal outcome or the vertical outcome. The probability of “horizontal” is $p_0(\theta) = \cos^2 \theta$. (The subscript “0” distinguishes this function from other hypothetical functions to be considered shortly.) This function allows someone observing the measurement results to gain information about the angle θ .

This scenario illustrates a typical feature of a quantum measurement: a measurement on a single instance of a system (in this case a single photon) cannot convey complete information about the system’s preparation. But a large statistical sample of measurements on identically prepared copies can eventually home in on the values of the preparation parameters (in this case the single parameter θ). One does not encounter this limitation in classical physics, at least not for pure states: if a particle is placed at position x with momentum p , a measurement can directly reveal those values. This difference between classical and quantum physics reflects the fact that quantum theory is inherently probabilistic.

In our specific example, one can ask how well the information about θ is conveyed through the observed results. Specifically, one can quantify the mutual information between the measurement results and the value of θ . As we will see shortly, given the limitation imposed by the probabilistic nature of the polarization measurement, the transfer of information is optimal in the limit of a large number of trials. That is, in this example anyway, quantum mechanics orchestrates the optimal conveyance of information from the preparation to the measurement outcome.

Before justifying this statement, I want to note that it works only because we are in effect working in the real-vector-space theory. By limiting the possibilities to linear polarizations, we are ruling out all the polarization states one would normally represent with vectors having a nonzero imaginary part (circular and elliptical polarizations). We will return to this point toward the end of this section.

Now let us make the statement precise. We do so by comparing our actual world, in which the probability of “horizontal” is $p_0(\theta) = \cos^2 \theta$, to a fictitious world in which the probability is given by some arbitrary function $p(\theta)$. In such a world, let N photons, each prepared with linear polarization angle θ , be subjected to a horizontal-vs-vertical polarization measurement. Let n be the number of these photons that yield the outcome “horizontal.” The mutual information between the measurement results and the value of θ is based on the Shannon entropy H and is defined to be

$$I(n : \theta) = H(n) - H(n|\theta) = - \sum_{n=0}^N P(n) \ln P(n) + \frac{1}{2\pi} \int_0^{2\pi} \left(\sum_{n=0}^N P(n|\theta) \ln P(n|\theta) \right) d\theta. \quad (5)$$

Here we have assumed a uniform *a priori* distribution of θ over the interval $[0, 2\pi]$. (This is a crucial assumption that we discuss further below.) $P(n|\theta)$ is the probability of getting the horizontal outcome exactly n times if the photons are prepared in the state θ , and $P(n)$ is the probability of getting the horizontal outcome exactly n times in the absence of any information about θ (that is, when θ is uniformly distributed). Both $P(n|\theta)$ and $P(n)$ depend on the function $p(\theta)$. Note that in Eq. (5) we have

written the mutual information as the average amount of information gained about the integer n upon learning the value of θ . It does have this interpretation, but it can alternatively be interpreted as the average amount of information one gains about the value of θ upon learning the value of n . (Mutual information is symmetric in its two arguments.) This latter interpretation is more descriptive of the scenario we are imagining, in which an observer at the polarizing beam splitter is trying to learn about the value of θ .

It turns out that for large N , $I(n : \theta)$ grows as $(1/2) \ln N$. We therefore consider the following limit, which has a finite upper bound:

$$\tilde{I} = \lim_{N \rightarrow \infty} \left[I(n : \theta) - \frac{1}{2} \ln N \right]. \quad (6)$$

We want to show that of all conceivable probability functions $p(\theta)$, the quantum mechanical function $p_0(\theta) = \cos^2 \theta$ gives \tilde{I} its largest possible value.

At this point we could proceed to compute \tilde{I} starting from Eq. (5), but the calculation will be simpler, and I hope clearer, if we abstract the problem away from its quantum mechanical setting. The important point to notice is that $I(n : \theta)$ depends on the probability function $p(\theta)$ only through the *measure* it induces on the binary probability space. That is, before we have any knowledge of θ , we can use $p(\theta)$ and the assumed uniform distribution of θ to figure out how likely it is that the probability of “horizontal” lies in any given interval, and it is this weighting function that figures into $I(n : \theta)$.

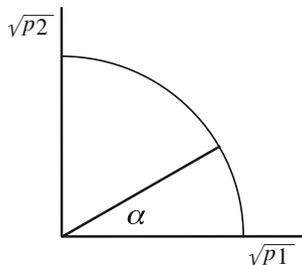
The more abstract problem, then, can be stated as follows. Consider a two-outcome probabilistic experiment, and let (p_1, p_2) denote a point in the binary probability space with p_1 corresponding to outcome #1. The experiment is run N times, and outcome #1 is observed to occur n times. This observation gives the experimenter information about (p_1, p_2) .³ The mutual information I between the value of n and the value of (p_1, p_2) depends on the experimenter’s *a priori* measure on probability space. Our problem is to find the *a priori* measure that maximizes the limit

$$\tilde{I} = \lim_{N \rightarrow \infty} \left[I - \frac{1}{2} \ln N \right]. \quad (7)$$

(The optimal measure will turn out to be unique.) We want to show that this optimal measure is the one induced by the quantum probability function $p_0(\theta) = \cos^2 \theta$ when θ is uniformly distributed.

³The experimenter is trying to determine the value of an unknown probability. It may seem that this problem cannot be framed except in the context of an objective interpretation of the concept of probability, but this is not the case. The representation theorem of de Finetti shows how to express this kind of question within a subjective interpretation [23]. We note that the *quantum* de Finetti theorem does not hold in real-vector-space quantum theory [24], but this fact does not preclude a subjective interpretation of probability in our problem. In our problem the experimenter is trying to refine a distribution over ordinary probability space, to which the classical de Finetti theorem applies.

Fig. 1 The relation between (p_1, p_2) and α



In order to tackle this problem we need to choose a parameterization of the binary probability space. We could use p_1 or p_2 as our parameter, but it turns out to be more convenient to use a different parameter α defined by $(p_1, p_2) = (\cos^2 \alpha, \sin^2 \alpha)$, where $0 \leq \alpha \leq \pi/2$. The relation between α and (p_1, p_2) is illustrated in Fig. 1. (One might object that we seem to be smuggling some quantum mechanics into the calculation here, but we are not. The results will be entirely independent of our choice of parameter. Our choice merely simplifies the calculation.) Let $K(\alpha)d\alpha$ be the *a priori* measure on the set of values of α , normalized so that $\int_0^{\pi/2} K(\alpha)d\alpha = 1$. The mutual information between α and n can be written as

$$I(\alpha : n) = h(\alpha) - h(\alpha|n) = - \int_0^{\pi/2} K(\alpha) \ln K(\alpha) d\alpha + \sum_{n=0}^N P(n) \int_0^{\pi/2} P(\alpha|n) \ln P(\alpha|n) d\alpha, \tag{8}$$

where $h(\alpha)$ and $h(\alpha|n)$ are differential entropies.⁴ Here $P(\alpha|n)$ is the probability distribution the experimenter assigns to α after seeing the value n , and $P(n)$ is the *a priori* probability of the value n as computed from the distribution $K(\alpha)$. (Note that if $K(\alpha)$ is derived from the probability function $p(\theta)$ under the assumption that θ is uniformly distributed, then $I(\alpha : n)$ is exactly equal to the quantity $I(n : \theta)$ given in Eq. (5)). The point of the next paragraph is to show that under modest assumptions about the function $K(\alpha)$, in the limit of very large N the second term on the right-hand side of Eq. (8) becomes *independent* of $K(\alpha)$. So we will only have to think about maximizing the first term.

To evaluate this second term, we need to write down expressions for $P(n)$ and $P(\alpha|n)$. We have

$$P(n) = \int_0^{\pi/2} P(n|\alpha)K(\alpha)d\alpha \tag{9}$$

and

$$P(\alpha|n) = \frac{P(n|\alpha)K(\alpha)}{P(n)}, \tag{10}$$

⁴The differential entropy is not the limit of the entropy of a discretized version of the continuous variable. However, a *mutual information* involving a continuous variable, being the *difference* between two differential entropies, is indeed the limit of the discretized mutual information [25].

where $P(n|\alpha)$ is given by the binomial distribution:

$$P(n|\alpha) = \frac{N!}{n!(N-n)!} p_1^n p_2^{N-n} \quad (11)$$

with $p_1 = \cos^2 \alpha$ and $p_2 = \sin^2 \alpha$. For any value of p_1 strictly between 0 and 1, it is possible to choose N large enough that the binomial distribution is well approximated by a Gaussian:

$$P(n|\alpha) \approx \frac{1}{\sqrt{2\pi N p_1 p_2}} \exp \left[-\frac{N}{2 p_1 p_2} (n/N - p_1)^2 \right]. \quad (12)$$

As N gets very large this distribution, regarded as a function of n/N , becomes arbitrarily highly peaked around $n/N = p_1$. Let $\alpha^{(n)}$ be defined so that $n/N = \cos^2 \alpha^{(n)}$. That is, $\alpha^{(n)}$ is the value of α corresponding to the observed outcome n . Then in the above exponent, we can approximate the quantity $(n/N - p_1)$ as

$$(n/N - p_1) = \cos^2 \alpha^{(n)} - \cos^2 \alpha \approx \frac{d(\cos^2 \alpha)}{d\alpha} \Delta\alpha = (-2 \cos \alpha \sin \alpha) \Delta\alpha = -2\sqrt{p_1 p_2} \Delta\alpha, \quad (13)$$

where $\Delta\alpha = \alpha^{(n)} - \alpha$. This gives us

$$P(n|\alpha) \approx \frac{1}{\sqrt{2\pi N p_1 p_2}} \exp \left[-2N(\Delta\alpha)^2 \right]. \quad (14)$$

Inserting this expression into Eq. (9), we again use the fact that the Gaussian is very highly peaked so that we can (i) extend the integral from $-\infty$ to ∞ without changing its value appreciably and (ii) evaluate everything outside the exponential at $\alpha = \alpha^{(n)}$. Then we get

$$P(n) \approx \frac{K(\alpha^{(n)})}{2N \cos \alpha^{(n)} \sin \alpha^{(n)}}. \quad (15)$$

We now use Eqs. (10), (14) and (15) to approximate $P(\alpha|n)$:

$$P(\alpha|n) \approx \sqrt{\frac{2N}{\pi}} \exp \left[-2N(\Delta\alpha)^2 \right]. \quad (16)$$

Using this expression and again relying on the narrowness of the Gaussian, we get

$$\int_0^{\pi/2} P(\alpha|n) \ln P(\alpha|n) d\alpha \approx \frac{1}{2} \ln \left(\frac{2N}{\pi e} \right). \quad (17)$$

Since this expression does not depend at all on n , it factors out of the sum in Eq. (8), so that the only sum we have to do is $\sum_n P(n)$, which is unity by definition. Putting the pieces together, we arrive at

$$I(\alpha : n) \approx - \int_0^{\pi/2} K(\alpha) \ln K(\alpha) d\alpha + \frac{1}{2} \ln \left(\frac{2N}{\pi e} \right). \quad (18)$$

And then subtracting $(1/2) \ln N$ as in Eq. (7) gives us

$$\tilde{I} = - \int_0^{\pi/2} K(\alpha) \ln K(\alpha) d\alpha + \frac{1}{2} \ln \left(\frac{2}{\pi e} \right). \quad (19)$$

The equality holds as long as our approximations become arbitrarily good as N gets larger. This will indeed be the case if the function $K(\alpha)$ is reasonably well behaved. A sufficient set of conditions on $K(\alpha)$ is that it be positive and differentiable on the interval $[0, \pi/2]$. Then when many trials are run, the range of likely values of α narrows to such a degree that the final distribution $P(\alpha|n)$ does not depend appreciably on the *a priori* distribution $K(\alpha)$.

The problem has now been reduced to finding out what distribution or distributions $K(\alpha)$ maximize the quantity $-\int_0^{\pi/2} K(\alpha) \ln K(\alpha) d\alpha$. The answer to this question is well known: the unique maximizing distribution is the *uniform* distribution $K(\alpha) = 2/\pi$. This result follows from the fact that the function $\phi(x) = -x \ln x$ is a strictly concave function of x for all positive values of x . Jensen's inequality then tells us that

$$(2/\pi) \int_0^{\pi/2} \phi[K(\alpha)] d\alpha \leq \phi \left[(2/\pi) \int_0^{\pi/2} K(\alpha) d\alpha \right] = \phi(2/\pi), \quad (20)$$

with equality holding only for the constant function $K(\alpha) = 2/\pi$.⁵

Now we compare our result to quantum mechanics. Is this uniform distribution over α the one induced by the quantum probability law $p_0(\theta) = \cos^2 \theta$, when θ is uniformly distributed? First consider the values of θ from 0 to $\pi/2$. In that range the law $p_0(\theta) = \cos^2 \theta$ mirrors the definition of α and we have $\alpha = \theta$ (I am taking p_1 to correspond to the horizontal outcome.). So a uniform distribution of θ over this range would induce the uniform distribution of α . In the other three quadrants of the circle, that is, in the rest of the range of θ , the parameter α is not equal to θ but we still have $|d\alpha/d\theta| = 1$ (except at a finite number of points where α "bounces" off one of the endpoints of its range). Thus when θ is uniformly distributed, so is α . This completes our demonstration that the quantum probability function $p_0(\theta) = \cos^2 \theta$ is optimal.

Is the function $p_0(\theta) = \cos^2 \theta$ unique in this respect? The answer is no. Any function $p(\theta)$ that yields the same *a priori* measure on the binary probability space will be equally good. For example, any function of the form $p(\theta) = \cos^2(m\theta/2)$ where m is a positive integer yields the same distribution $K(\alpha) = 2/\pi$. And there

⁵Alternatively, instead of using differential entropies as in Eq. (8), we could have expressed the mutual information $I(\alpha : n)$ in terms of the Kullback-Leibler distances of both $K(\alpha)$ and $P(\alpha|n)$ from the uniform distribution over α . The calculation in Sect. 3 then tells us that \tilde{I} is maximized when the Kullback-Leibler distance between $K(\alpha)$ and the uniform distribution is minimized, that is, when $K(\alpha)$ is itself the uniform distribution.

are many other, less physically interesting examples. Still, a typical function $p(\theta)$ will not have this optimization property.

Looking back over the above argument, one can see that the crucial feature is the exponent in Eq. (14): the coefficient of $(\Delta\alpha)^2$ depends only on N and not on α itself. In other words, the spread in the value of $\alpha^{(n)}$ depends only on the number of trials (when this number is large), and not on the probabilities (p_1, p_2) . This is what is special about parameterizing probability space with the parameter α : it makes the statistical spread uniform. Once we have this fact, it is guaranteed that the final differential entropy $h(\alpha|n)$ will not depend on $K(\alpha)$. Therefore to maximize the mutual information, we want to maximize the initial differential entropy $h(\alpha)$ and we are thereby led to the uniform distribution.

There is perhaps a more direct way of seeing what is special about the function $p_0(\theta) = \cos^2 \theta$. First note that the spread in n itself is *not* uniform over probability space. If one performs the binary experiment N times, the standard deviation in n/N is given by

$$\Delta(n/N) = \sqrt{\frac{p_1 p_2}{N}}, \quad (21)$$

which is smaller near the ends of probability space than near the middle. One can see this dependence in the exponent in Eq. (12). In our polarization experiment, an observer recording the frequency of occurrence of “horizontal” will therefore be more certain of the *probability* of “horizontal” when that probability is close to zero or one. (Again I am assuming that the experimenter’s *a priori* distribution over probability space is reasonably smooth and that the number of trials is large). On the other hand, upon translating the uncertainty in probability to an uncertainty in θ , the observer must use the function $p(\theta)$. For the special case of $p_0(\theta) = \cos^2 \theta$, the slope of this function exactly compensates for the varying size of $\Delta(n/N)$, so that the size of the resulting “region of uncertainty” of θ is independent of the value of n/N . Specifically,

$$\left| \frac{d}{d\theta} \cos^2 \theta \right| = 2 |\cos \theta \sin \theta| = 2\sqrt{p_0(\theta)[1 - p_0(\theta)]}, \quad (22)$$

which perfectly matches the dependence seen in Eq. (21). This compensation is illustrated in Fig. 2. Equalizing the uncertainties in θ also equalizes the uncertainties in α , which, as we have just seen, maximizes the mutual information in the limit of a large number of trials.

We now consider the case in which *all* pure polarization states are possible. The full set of pure states is the Bloch sphere—it includes the circular and elliptical polarizations—and the natural *a priori* measure is the uniform measure on the sphere, since this is the only probability measure invariant under all unitary transformations. We imagine a device that prepares a beam of photons in one of these polarization states, and further along the photons’ path we imagine a person making the horizontal-vs-vertical measurement on each photon. The polarization is now determined by two parameters; for definiteness let us take them to be the polar angle β and the

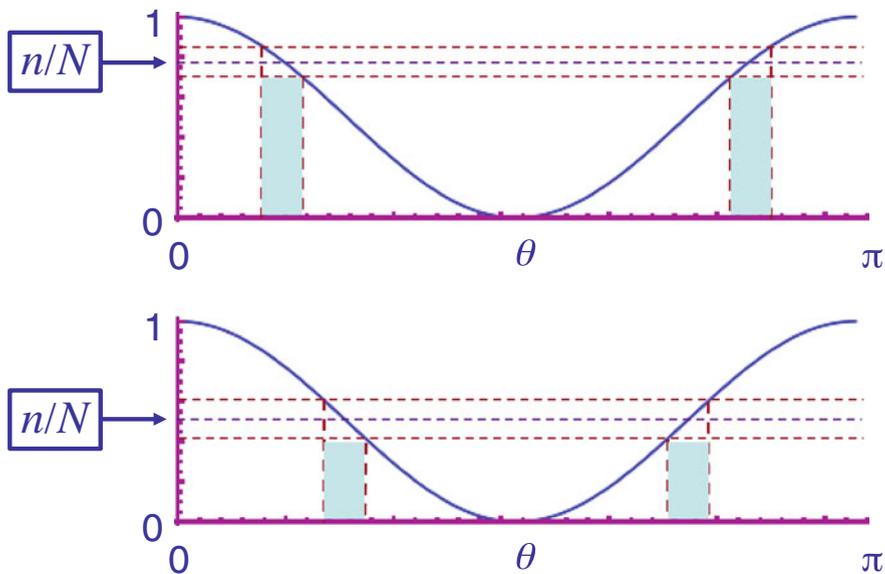


Fig. 2 The uncertainty in θ for two different values of n/N . Notice that the slope of the cosine-squared curve exactly compensates for the varying size of $\Delta(n/N)$, so that the “region of uncertainty” in θ has the same size for all values of n/N . (Here θ is plotted only up to π to make the diagram simpler)

azimuthal angle ϕ , and let the north and south poles of the sphere correspond to horizontal and vertical polarization. It is still possible to define the mutual information between the photons’ preparation and the measurement outcomes; it could be written as $I(n : \beta, \phi)$. Again this mutual information is the same as the quantity $I(\alpha : n)$ given in Eq.(8) and it is maximized only if the *a priori* distribution of α is the uniform distribution $K(\alpha) = 2/\pi$. But now the quantum mechanical law does *not* yield this distribution over the values of α . With the uniform distribution over the Bloch sphere, the parameter $\cos \beta$ is uniformly distributed over the interval $[-1, 1]$, and the quantum mechanical probability of “horizontal,” $p(\theta) = (1/2)(1 + \cos \beta)$, is therefore uniformly distributed over the interval $[0, 1]$. To get the corresponding distribution of α , we use the relation $p_1 = \cos^2 \alpha$ and the assumption that p_1 is uniformly distributed:

$$K(\alpha) = \left| \frac{dp_1}{d\alpha} \right| = 2 \cos \alpha \sin \alpha. \quad (23)$$

Thus, rather than giving us the uniform distribution of α , the full Bloch sphere gives us a distribution that has a maximum in the middle of α ’s range.

We can see directly that this distribution does not allow as much information transfer as the optimal distribution. The relevant quantity is the integral in Eq.(19):

$$- \int_0^{\pi/2} K(\alpha) \ln K(\alpha) d\alpha. \quad (24)$$

For the uniform distribution over α , this quantity has the value $\ln(\pi/2) = 0.452$, whereas for the distribution $K(\alpha) = 2 \cos \alpha \sin \alpha$, we get $1 - \ln 2 = \ln(e/2) = 0.307$.

Just as the uniform measure over the surface of the Bloch sphere is natural because it is invariant under all unitaries, in the real-vector-space theory where the set of pure states in two dimensions traces out a circle rather than a sphere, the uniform distribution over the circle is natural because it is invariant under all orthogonal transformations (rotations and reflections). That is, in the real-vector-space theory, we can use this invariance to justify our original assumption that the angle θ is uniformly distributed over the interval $[0, 2\pi]$.

4 Optimal Transfer of Information: The d -Dimensional Case

The above argument extends to a d -dimensional real vector space. Let a “redit” be a hypothetical quantum object whose pure states are vectors in a d -dimensional vector space over the real numbers. We now imagine an experiment in which a beam of N redits is prepared in a specific pure state $|s\rangle$. At some point further along the beam, an observer makes a fixed complete orthogonal measurement on each redit. The observer records the integers n_1, \dots, n_d , where n_i is the number of times the i th measurement outcome occurs. We ask how much information the observer learns on average about the preparation $|s\rangle$, assuming (crucially) that the vector $|s\rangle$ is initially distributed uniformly over the unit sphere in the d -dimensional space. Again this average information gain is given by the mutual information, which we will write down shortly. The mutual information depends on the law that specifies the probability of the i th outcome given the preparation $|s\rangle$. In real-vector-space quantum theory, this law can be expressed as

$$p_i(|s\rangle) = s_i^2, \quad i = 1, \dots, d, \quad (25)$$

where s_1, \dots, s_d are the components of $|s\rangle$ in the basis defined by the measurement.

As before, what really matters in computing the mutual information is the *a priori* measure on probability space. The uniform measure over the unit sphere in d dimensions, together with Eq.(25), defines some specific *a priori* measure on probability space. We also want to consider other *a priori* measures, in order to show that the one induced by Eq.(25) is optimal. The probability space is now a $(d - 1)$ -dimensional set, since the probabilities must add to one. We could parameterize this set by the probabilities p_1, \dots, p_{d-1} of the first $d - 1$ outcomes, but we instead choose to label the points of probability space by a unit vector $\vec{\gamma} = (\sqrt{p_1}, \dots, \sqrt{p_d})$.

(Note that each $\sqrt{p_i}$ is non-negative; so $\vec{\gamma}$ is confined to the positive part of the unit sphere). We could go further and choose $d - 1$ specific angular coordinates to locate this vector on the sphere (like the α of the preceding section), but we will not need to do so. Let $K(\vec{\gamma})d\vec{\gamma}$ be a generic *a priori* probability measure on the set of vectors $\vec{\gamma}$, where $d\vec{\gamma}$ is an infinitesimal $(d - 1)$ -dimensional surface element on the positive section of the unit sphere. Our goal is to find the distribution $K(\vec{\gamma})$ that maximizes the mutual information.

That mutual information can be written as follows:

$$I(\vec{\gamma} : \vec{n}) = h(\vec{\gamma}) - h(\vec{\gamma}|\vec{n}) = - \int K(\vec{\gamma}) \ln K(\vec{\gamma}) d\vec{\gamma} + \sum P(\vec{n}) \int P(\vec{\gamma}|\vec{n}) \ln P(\vec{\gamma}|\vec{n}) d\vec{\gamma}, \quad (26)$$

where $\vec{n} = (n_1, \dots, n_d)$ specifies the number of times each outcome occurs. The sum is over all vectors \vec{n} for which each n_i is a non-negative integer and $n_1 + \dots + n_d = N$. The mutual information is now based on the multinomial distribution:

$$P(\vec{n}|\vec{\gamma}) = \frac{N!}{n_1! \dots n_d!} p_1^{n_1} \dots p_d^{n_d}. \quad (27)$$

(Here $p_j = \gamma_j^2$.) The functions appearing in Eq. (26) can be obtained from $P(\vec{n}|\vec{\gamma})$ as follows:

$$P(\vec{n}) = \int P(\vec{n}|\vec{\gamma}) K(\vec{\gamma}) d\vec{\gamma} \quad (28)$$

and

$$P(\vec{\gamma}|\vec{n}) = \frac{P(\vec{n}|\vec{\gamma}) K(\vec{\gamma})}{P(\vec{n})}. \quad (29)$$

As N gets large, it will turn out that $I(\vec{\gamma} : \vec{n})$ grows as $[(d - 1)/2] \ln N$. So we will compute the limiting value

$$\tilde{I} = \lim_{n \rightarrow \infty} \left[I(\vec{\gamma} : \vec{n}) - \left(\frac{d-1}{2} \right) \ln N \right]. \quad (30)$$

At this point the calculation is very similar to the one in the preceding section. As we did for the analagous equation in that case, we now show that the second term on the right-hand side of Eq. (26) becomes independent of $K(\vec{\gamma})$ as N approaches infinity.

For any fixed positive values of p_1, \dots, p_d and for large enough N , Eq. (27) can be approximated arbitrarily well by a Gaussian function:

$$P(\vec{n}|\vec{\gamma}) \approx [(2\pi N)^{d-1} p_1 p_2 \dots p_d]^{-1/2} \exp \left[-\frac{N}{2} \sum_{i=1}^d \frac{(n_i/N - p_i)^2}{p_i} \right]. \quad (31)$$

It will be helpful to define the vectors

$$\vec{\gamma}^{(n)} = \left(\sqrt{\frac{n_1}{N}}, \dots, \sqrt{\frac{n_d}{N}} \right) \quad \text{and} \quad \Delta\vec{\gamma} = \vec{\gamma}^{(n)} - \vec{\gamma}. \quad (32)$$

That is, $\vec{\gamma}^{(n)}$ is the vector of square roots of the observed frequencies of occurrence, whereas $\vec{\gamma}$ is the vector of square roots of the probabilities. The difference $\Delta\vec{\gamma}$ between these two vectors is likely to be small when N is large; so we will keep just the lowest-order term in this quantity. We can then rewrite the sum in the exponent of Eq. (31) as follows:

$$\sum_{i=1}^d \frac{(n_i/N - p_i)^2}{p_i} = \sum_{i=1}^d \frac{(\Delta p_i)^2}{p_i} \approx 4 \left| \Delta\vec{\gamma} \right|^2, \quad (33)$$

where Δp_i is defined to be $(n_i/N) - p_i$ and the last step comes from

$$\Delta\gamma_i = \Delta \left(p_i^{1/2} \right) \approx \frac{1}{2} p_i^{-1/2} \Delta p_i. \quad (34)$$

We can therefore approximate Eq. (31) as

$$P(\vec{n}|\vec{\gamma}) \approx (2\pi N)^{-\binom{d-1}{2}} \frac{1}{\gamma_1 \gamma_2 \cdots \gamma_d} \exp \left(-2N \left| \Delta\vec{\gamma} \right|^2 \right). \quad (35)$$

That is, $P(\vec{n}|\vec{\gamma})$, regarded as function of $\vec{\gamma}$, falls off as a Gaussian around the point $\vec{\gamma}^{(n)}$, with a spread that is isotropic and independent of the value of $\vec{\gamma}^{(n)}$. (This function is not, however, a *normalized* distribution of $\vec{\gamma}$; rather, it is normalized with respect to a sum over \vec{n} .)

In approximating the integral in Eq. (28), we rely on the narrowness of the Gaussian: the integral is over a section of a sphere, but we can treat it as being over an infinite flat space having $d - 1$ dimensions—the “plane” tangent to the sphere at the point $\vec{\gamma}^{(n)}$. We also evaluate everything outside the exponential at the point $\vec{\gamma}^{(n)}$. These approximations give us

$$P(\vec{n}) \approx (2N)^{-(d-1)} \frac{K(\vec{\gamma}^{(n)})}{\gamma_1^{(n)} \gamma_2^{(n)} \cdots \gamma_d^{(n)}}. \quad (36)$$

Inserting this expression and Eq. (35) into Eq. (29), we get

$$P(\vec{\gamma}|\vec{n}) \approx \left(\frac{2N}{\pi} \right)^{\frac{d-1}{2}} \exp \left(-2N \left| \Delta\vec{\gamma} \right|^2 \right). \quad (37)$$

We can now do the second integral in Eq. (26), again treating the integral as if it were over an infinite $(d - 1)$ -dimensional flat space. The result is

$$I(\vec{\gamma} : \vec{n}) \approx - \int K(\vec{\gamma}) \ln K(\vec{\gamma}) d\vec{\gamma} + \left(\frac{d-1}{2} \right) \ln \left(\frac{2N}{\pi e} \right). \quad (38)$$

Finally we subtract $[(d - 1)/2] \ln N$ as in Eq. (30) to get

$$\tilde{I} = - \int K(\vec{\gamma}) \ln K(\vec{\gamma}) d\vec{\gamma} + \left(\frac{d-1}{2} \right) \ln \left(\frac{2}{\pi e} \right). \quad (39)$$

Note that Eq. (19) is a special case of this equation, with $d = 2$. As in that case, the expression is maximized by choosing $K(\vec{\gamma})$ to correspond to the uniform distribution:

$$K_{opt}(\vec{\gamma}) = \frac{2^d \Gamma(\frac{d}{2} + 1)}{d \pi^{d/2}}. \quad (40)$$

The constant on the right-hand side of Eq. (40) is the reciprocal of the “surface area” of the positive section of the unit sphere.

The question now is whether the probability rule in real-vector-space quantum theory, $p_i = s_i^2$, induces the measure on probability space given by $K_{opt}(\vec{\gamma})$. The answer is yes, as is easily seen. The state vector $|s\rangle$ ranges over the full unit sphere in \mathbb{R}^d , but consider for now just the section of the sphere in which each s_i is positive. In that section the vector $\vec{\gamma}$ is equal to the vector $|s\rangle$, since $p_i = \gamma_i^2 = s_i^2$. So the uniform distribution of $|s\rangle$ over this section of the sphere induces the uniform distribution of $\vec{\gamma}$. The whole unit sphere in \mathbb{R}^d consists, in effect, of 2^d copies of the positive section. So indeed, the uniform distribution of $|s\rangle$ over the whole sphere does correspond to the uniform distribution of $\vec{\gamma}$ expressed in Eq. (40). That is, the transfer of information from preparation to measurement is optimized in this d -dimensional case, just as it was in the two-dimensional case.

The complications involved in the information-theoretic calculation may obscure what is really a simple underlying fact. Imagine probability space as a $(d - 1)$ -dimensional flat “surface” in a d -dimensional space with orthogonal axes labeled p_1, \dots, p_d . The surface consists of all points (p_1, \dots, p_d) such that each p_i is non-negative and the sum $p_1 + \dots + p_d$ is equal to 1. Around each point on this surface one can imagine a small “region of uncertainty,” representing the spread in the actual frequencies of occurrence of the d possible outcomes in N trials. These regions of uncertainty can be derived from the multinomial distribution, Eq. (27), and for large N their sizes and shapes can be read off the exponent in Eq. (31). (We could, for example, define “region of uncertainty” to be the range of values of $(n_1/N, \dots, n_d/N)$ for which the exponent has magnitude less than 1.) One can see that these regions of uncertainty will have different sizes and shapes, depending on the location in probability space. For example, the largest is at the exact center, as shown in Fig. 3. But if we change the axes of probability space from p_i to $\gamma_i = \sqrt{p_i}$, probability space then looks like a section of a sphere. Again one can speak of a region of uncertainty around each point on this spherical surface, but now it happens that all the regions of uncertainty have the same size and shape—in fact they are all spherical—as we can see in the exponent of Eq. (35) and as is illustrated in Fig. 4. (The issue gets tricky near the edges. The closer one gets to the edge, the higher the value of N must be in order to see this uniformity. But no matter how close one is to the edge—as long as one is not *at* the edge—there is always such a value of N .)

Fig. 3 Regions of uncertainty at different locations in the flat probability space. As one approaches an edge, the uncertainty shrinks along the direction perpendicular to the edge

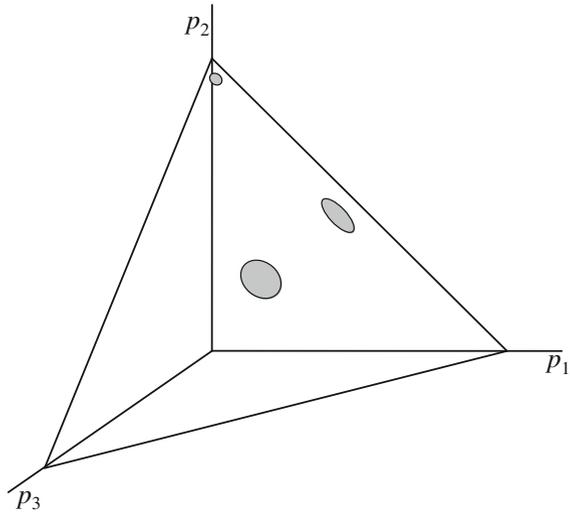
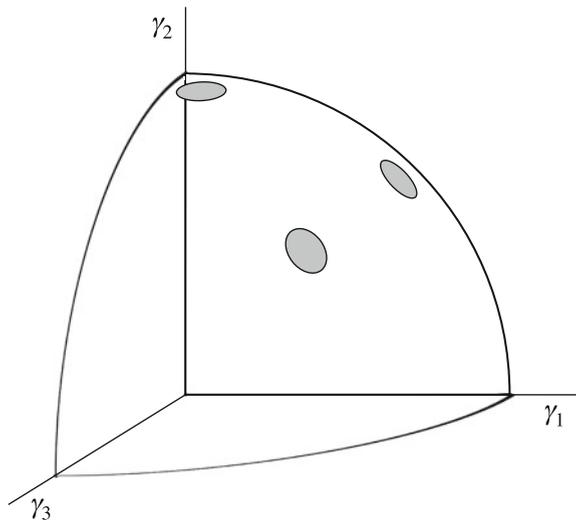


Fig. 4 Regions of uncertainty at different locations in the spherical probability space, with axes corresponding to the square roots of probability. Now all regions of uncertainty are isotropic and of the same size



In this sense, there is something special about representing probability space as a section of a sphere: it captures geometrically the statistical fluctuations in a large sample. What is special about real-vector-space quantum theory is that its set of pure states mirrors this representation of probability space.

As one would expect, the mathematical fact illustrated in Fig. 4 has been well noted in the statistics literature. Bhattacharyya in the 1940s proposed a distance measure between two probability distributions based on the *angle* between their $\vec{\gamma}$ vectors [26]. The square-root construction has been particularly explicit in the genetics literature. One can see diagrams of the positive section of the unit sphere in papers by Cavalli-

Sforza and collaborators from the 1960s, and these authors give credit for the idea to R.A. Fisher [27, 28] (as do Mosteller and Tukey [29]). In the present paper, I have used the square-root construction only to identify a special *measure* on probability space—the uniform measure on the spherical surface traced out by $\vec{\gamma}$. But one can also use it to define a special *metric* on the space, and this is what Bhattacharyya, Cavalli-Sforza and others have done. (One can find in Ref. [30] a review a various “genetic distances,” some of which are based on the square-root construction.) Such a metric has also been used in work on quantum foundations [16, 31–33]. However, I want to emphasize that this special feature of the representation of probability space in terms of square roots of probability arises without any reference to quantum theory. It is simply a matter of statistics.

What about ordinary complex-vector-space quantum theory? In that theory each pure state is represented by a vector $|s\rangle$ in \mathbb{C}^d . The natural *a priori* distribution over pure states is the uniform distribution over the unit sphere in \mathbb{C}^d , that is, the unique distribution invariant under all unitary transformations. (We could just as well speak of a distribution over projection operators $|s\rangle\langle s|$ so as not to have to worry about the irrelevant overall phase factor in the vector $|s\rangle$, but for our purposes either picture leads to the same result.). For a complete orthogonal measurement, the probabilities of the outcomes are given by $p_i = |s_i|^2$, where the s_i 's are the components of $|s\rangle$ in the basis defined by the measurement. We can ask what measure this probability rule, together with the *a priori* distribution of state vectors $|s\rangle$, induces on probability space. That question was answered by Sykora in 1974: it induces the *uniform* distribution, not on the spherical surface defined by $\vec{\gamma}$, but on the flat surface defined by (p_1, \dots, p_d) [34]. This is a remarkably simple and intriguing result, but this distribution is not the one that optimizes the transfer of information from preparation to measurement.

5 Optimal Information Transfer in Standard Quantum Mechanics?

The real-vector-space theory thus has a certain elegance to it, in that there is an optimal correspondence between the set of pure states and the set of probability distributions over the outcomes of a complete orthogonal measurement. The complex theory does not have this property, but one might wonder whether this is because we are not asking the question in the right way. That is, by somehow reframing the problem, might it be possible to see that the usual complex theory does exhibit the property of optimal information transfer in some altered sense?

For example, perhaps we are making a mistake to consider a complete orthogonal measurement. Such a measurement will never reveal the relative phases between the components of the state vector when it is written in the measurement basis. Instead we could consider a special case of a non-orthogonal measurement, namely, a symmetric informationally-complete measurement (a SIC) [35–37]. In d complex dimensions,

such a measurement has d^2 possible outcomes. Each outcome corresponds to a pure state $|m_i\rangle\langle m_i|$, and the inner product between any two of these pure states has the same magnitude: $|\langle m_i|m_j\rangle|^2 = 1/(d+1)$. Numerical evidence strongly indicates that such symmetric measurements exist for all values of d up to 67 [38], and it would be reasonable to guess that they exist for all d . Such symmetric measurements figure prominently in the quantum Bayesian approach to understanding quantum mechanics [39, 40]. Is there a kind of optimal transfer of information from preparation to measurement that occurs when the measurement is a SIC?

With d^2 possible outcomes, one can estimate $d^2 - 1$ independent parameters by repeating the measurement on many identically prepared copies. This is exactly the number of parameters needed to specify a $d \times d$ density matrix, and indeed, any density matrix can be reconstructed with arbitrary precision from a fixed SIC applied to many copies. (This is the meaning of “informationally complete”). To state the question of optimal information transfer, we would need to specify an *a priori* measure on the set of all $d \times d$ density matrices. The measure should be unitarily invariant, but there are many unitarily invariant measures on this set. Is there at least one such measure for which the mutual information between the preparation (of a general mixed state) and the measurement outcomes is optimal?

One can see quickly that there is no such measure. The optimal *a priori* measure on probability space has already been determined in the preceding section. For d^2 possible outcomes, the optimal measure is the uniform measure over the $(d^2 - 1)$ -dimensional spherical surface of probability space, when the axes correspond to the square roots of the probabilities. This measure clearly assigns nonzero weight to every nonzero volume of probability space. But if one performs a SIC on any state, the largest possible value of any probability is $1/d$. Thus the SIC does not make use of the whole probability space; so it is not providing information optimally in our sense, no matter what weighting function we place on the set of density matrices.

Let us try another version of the problem. Suppose we are given a specific entangled state of a pair of qubits, namely, the state

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle). \quad (41)$$

We imagine the first qubit is held by Alice and the second by Bob. Now Alice applies a unit-determinant unitary transformation U to her qubit—an element of $SU(2)$. She then sends her qubit to Bob, who performs a Bell measurement on the two qubits. That is, he distinguishes the four orthogonal states

$$\begin{aligned} |\Phi^+\rangle &= \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle) & |\Phi^-\rangle &= \frac{1}{\sqrt{2}} (|00\rangle - |11\rangle) \\ |\Psi^+\rangle &= \frac{1}{\sqrt{2}} (|01\rangle + |10\rangle) & |\Psi^-\rangle &= \frac{1}{\sqrt{2}} (|01\rangle - |10\rangle) \end{aligned} \quad (42)$$

We imagine this whole procedure is repeated over and over—always with the same initial state, the same U , and the same Bell measurement—so that Bob can try to

gain information about U from the outcomes of his measurements. We assume he already knows the initial state $|\Phi^+\rangle$. (This scenario is like superdense coding [41], except that we are not restricting U to a discrete set. Really what Bob is doing here is a restricted kind of process tomography [42, 43]—trying to infer the process U from the outcomes of measurements.). We can ask whether the transfer of information is optimal between Alice’s choice of unitary transformation and the outcomes of Bob’s measurements.

A general element of $SU(2)$ can be represented as

$$U = \exp [i(\theta/2)\hat{n} \cdot \vec{\sigma}], \quad (43)$$

where \hat{n} is the unit vector defining the axis of the Bloch sphere around which Alice is rotating her qubit, θ is the angle of rotation—it runs from zero to 2π —and $\vec{\sigma}$ is the vector of Pauli matrices. The transformations U are in one-to-one correspondence with the points of a three-dimensional spherical surface, which we can imagine embedded in four dimensions. Specifically, we can label the point corresponding to the above U by the unit vector

$$v_U = (\cos(\theta/2), n_x \sin(\theta/2), n_y \sin(\theta/2), n_z \sin(\theta/2)). \quad (44)$$

The natural measure on $SU(2)$ is the uniform measure over this sphere—it is the unique measure that is invariant under left-multiplication (or right-multiplication) by any group element.

To determine whether the information is transferred optimally from Alice to Bob, we need to compute the probabilities of Bob’s outcomes. It is straightforward to do so, and one finds that the probabilities are, in an arrangement parallel to that given in Eq. (42),

$$\begin{array}{ll} \cos^2(\theta/2) & n_z^2 \sin^2(\theta/2) \\ n_x^2 \sin^2(\theta/2) & n_y^2 \sin^2(\theta/2) \end{array} \quad (45)$$

These probabilities are the squared components of the unit vector v_U given in Eq. (44). Thus the problem is equivalent to the case of real-vector-space quantum mechanics in four dimensions. So indeed, the information is transmitted optimally from Alice to Bob!

Does this example generalize to higher dimensions? The answer is no, at least not in any way that I can see. For example, in three dimensions, we would probably want Alice and Bob to start with the state $|\Phi\rangle = (|00\rangle + |11\rangle + |22\rangle)/\sqrt{3}$. Alice will perform a general unit-determinant unitary transformation U , and then Bob will measure both particles in the generalized Bell basis, which consists of the nine states

$$\begin{array}{lll} |00\rangle + |11\rangle + |22\rangle & |00\rangle + \omega|11\rangle + \omega^2|22\rangle & |00\rangle + \omega^2|11\rangle + \omega|22\rangle \\ |01\rangle + |12\rangle + |20\rangle & |01\rangle + \omega|12\rangle + \omega^2|20\rangle & |01\rangle + \omega^2|12\rangle + \omega|20\rangle \\ |02\rangle + |10\rangle + |21\rangle & |02\rangle + \omega|10\rangle + \omega^2|21\rangle & |02\rangle + \omega^2|10\rangle + \omega|21\rangle. \end{array} \quad (46)$$

Here $\omega = e^{2\pi i/3}$ and I have suppressed the normalization factor $1/\sqrt{3}$. A counting of parameters is initially encouraging: it takes eight real numbers to specify an element of $SU(3)$, and Bob's measurement yields eight independent probabilities. However, one quickly discovers that, as in the case of the SIC, the measurement does not make use of the whole probability space.

Consider specifically the probabilities of the second and third outcomes listed on the first row of Eq. (46); let us call these probabilities p_2 and p_3 (we imagine a list of nine probabilities p_1, \dots, p_9 , of which these are the second and third). In terms of the components u_{ij} of Alice's unitary matrix U , we have

$$p_2 = \frac{1}{9} |u_{00} + \omega^2 u_{11} + \omega u_{22}|^2 \quad \text{and} \quad p_3 = \frac{1}{9} |u_{00} + \omega u_{11} + \omega^2 u_{22}|^2, \quad (47)$$

so that the product has the value

$$p_2 p_3 = \frac{1}{81} |u_{00}^2 + u_{11}^2 + u_{22}^2 - u_{00}u_{11} - u_{00}u_{22} - u_{11}u_{22}|^2. \quad (48)$$

Now, in the whole probability space the maximum value of $p_2 p_3$ is $1/4$, attained when $p_2 = p_3 = 1/2$. But given that each u_{ij} can have a magnitude no larger than 1, one can show that the expression in Eq. (48) cannot exceed the value $16/81 < 1/4$.⁶ Thus a certain region of probability space is inaccessible in the scenario we are imagining. It follows that the information about U is not conveyed optimally to Bob through his measurement outcomes.

Thus, as far as I can tell, the property of optimal information transfer does not easily carry over from the real-vector-space theory to ordinary quantum mechanics.

⁶In proving this inequality, we are free to set u_{00} equal to 1. Then let $u_{11} = -a$ and $u_{22} = -b$ and the desired inequality becomes

$$|1 + a + b + a^2 + b^2 - ab| \leq 4$$

under the assumption that $|a| \leq 1$ and $|b| \leq 1$. (One can see that equality is achieved when $a = b = 1$.) This inequality is equivalent to

$$|A^2 + B^2 + (A - B)^2| \leq 8,$$

where $A = 1 + a$ and $B = 1 + b$. This last inequality can be proved by first noting that

$$|A^2 + B^2 + (A - B)^2| \leq |A^2 + B^2| + |A - B|^2$$

and then writing out the absolute values explicitly. One has to use the fact that A and B are both confined to a circle of unit radius in the complex plane, centered at the value 1.

6 Conclusions

In real-vector-space quantum theory, the number of parameters needed to specify a pure state is equal to the number of independent probabilities in a complete orthogonal measurement: both are equal to $d - 1$ for a state in d dimensions. So by measuring many copies prepared in an unknown pure state, one can hope to pin the state down to a finite number of small regions in state space. In this paper we have seen that this pinning down is in fact optimal, in the sense that the observer gains as much information about the state as could possibly be gained in any probabilistic theory, at least when the number of trials is very large.

Standard quantum theory, based on a complex vector space, does not have this property, and we have not been able to find a restatement of the problem for which the complex theory does achieve such an optimization (except for the case of a unitary transformation applied to a qubit). For our original statement of the problem, one can say that this lack of optimization comes from the fact that for any specific orthogonal measurement, a complete specification of a pure state includes phase factors that have no effect on the probabilities of the outcomes. The presence of these phase factors changes the natural *a priori* measure on probability space, and the mutual information is no longer maximized.

Note that in the complex theory the number of real parameters needed to specify a pure state in d dimensions is $2d - 2$ if we do not count an irrelevant overall phase factor. This number is exactly *twice* the number of independent probabilities an orthogonal measurement can access, and it seems that this doubling of the number of parameters is what spoils the optimization. It is reasonable to ask whether there can be some deeper understanding of this factor of two, but at this point it is hard to have confidence in any particular answer.

In a sense, any axiomatization of quantum mechanics offers a potential answer to this question: whatever assumptions give rise to the structure of quantum theory also give rise to the factor of two. In his axiomatization, Goyal addresses the factor of two directly, formalizing it in his principle of complementarity: for a measurement that at some level generates $2d$ possible events, only d distinguishable outcomes can be observed, each corresponding to a *pair* of fundamental events. This principle, together with a principle of global gauge invariance, leads him to the basic structure of quantum mechanics [33]. One achieves a similar result by assuming that the underlying theory is real-vector-space quantum theory, but that because our knowledge is limited in some fundamental way we do not see the real-vector-space structure; we have access only to those observables that satisfy Stueckelberg's rule. (Goyal in fact relates his work to Stueckelberg's.) Imposing Stueckelberg's rule on a real vector space of $2d$ dimensions reduces the maximum number of orthogonal states from $2d$ to d , and it cuts in half the number of parameters required to specify a maximally pure state.⁷

⁷When standard quantum mechanics is expressed in real-vector-space terms, what we normally call a pure state is represented by a density matrix of rank 2.

While such an interpretation would give an important role to the real-vector-space theory, it raises a difficult question about the status of the main result in this paper. If the limitation on our knowledge is fundamental, then who are the observers for whom the transfer of information from preparation to measurement is optimal? Evidently it is not optimal for us, because whatever the underlying theory may be, the *effective* theory within which we live seems to be complex-vector-space quantum theory.

References

1. G. Birkhoff, J. von Neumann, *Ann. Math.*, second series **37**, 823–843 (1936)
2. C. Piron, *Helv. Phys. Acta* **37**, 439–468 (1964)
3. M.P. Solèr, *Commun. Algebra* **23**, 219–243 (1995)
4. S.S. Holland, *Bull. Am. Math. Soc.* **32**, 205–234 (1995)
5. E.C.G. Stueckelberg, *Helv. Phys. Acta* **32**, 254 (1959)
6. E.C.G. Stueckelberg, *Helv. Phys. Acta* **33**, 727 (1960)
7. P. Lahti, M.J. Maczynski, *J. Math. Phys.* **28**, 1764–1769 (1987)
8. H. Araki, *Commun. Math. Phys.* **75**, 1 (1980)
9. S. Bergia, F. Cannata, A. Cornia, R. Livi, *Found. Phys.* **10**, 723 (1980)
10. W.K. Wootters, in *Complexity, Entropy and the Physics of Information*, ed. by W.H. Zurek (Addison-Wesley, Redwood City, 1990)
11. N.D. Mermin, *Am. J. Phys.* **66**, 753 (1998)
12. L. Hardy, [arxiv:quant-ph/0101012](https://arxiv.org/abs/quant-ph/0101012) (2001)
13. J. Barrett, *Phys. Rev. A* **75**, 032304 (2007)
14. G. Chiribella, G.M. D’Ariano, P. Perinotti, *Phys. Rev. A* **84**, 012311 (2011)
15. L. Hardy, W.K. Wootters, *Found. Phys.* **42**, 454–473 (2012)
16. W.K. Wootters, The acquisition of information from quantum measurements. Ph.D. dissertation, University of Texas at Austin (1980)
17. E.C.G. Stueckelberg, M. Guenin, *Helv. Phys. Acta* **34**, 621 (1961)
18. E.C.G. Stueckelberg, M. Guenin, C. Piron, *Helv. Phys. Acta* **34**, 675 (1961)
19. F.J. Dyson, *J. Math. Phys.* **3**, 1199 (1962)
20. G.W. Gibbons, [arxiv:1111.0457](https://arxiv.org/abs/1111.0457) (2011)
21. J. Myrheim, [arxiv:quant-ph/9905037](https://arxiv.org/abs/quant-ph/9905037) (1999)
22. S. Aaronson, [arxiv:quant-ph/0401062](https://arxiv.org/abs/quant-ph/0401062) (2004)
23. B. de Finetti, *Theory of Probability* (Wiley, New York, 1990)
24. C.M. Caves, C.A. Fuchs, R. Schack, *J. Math. Phys.* **43**, 4537 (2002)
25. T.M. Cover, J.A. Thomas, *Elements of Information Theory* (Wiley, New York, 1991), p. 231
26. A. Bhattacharyya, *Sankhya* **7**, 401–407 (1946)
27. L.L. Cavalli-Sforza, F. Conterio, *Atti. Assoc. Genet. Ital.* **5**, 333 (1960)
28. L.L. Cavalli-Sforza, A.W.F. Edwards, *Am. J. Hum. Genet.* **19**, 233 (1967)
29. F. Mosteller, J.W. Tukey, *J. Am. Stat. Assoc.* **44**, 174 (1949)
30. N. Takezaki, M. Nei, *Genetics* **144**, 389 (1996)
31. W.K. Wootters, *Phys. Rev. D* **23**, 357 (1981)
32. S.L. Braunstein, C.M. Caves, *Phys. Rev. Lett.* **72**, 3439 (1994)
33. P. Goyal, *New J. Phys.* **12**, 023012 (2010)
34. S. Sykora, *J. Stat. Phys.* **11**, 17–27 (1974)
35. G. Zauner, Quantum designs—foundations of a non-commutative theory of designs. Ph.D. dissertation, University of Vienna (1999)
36. C.M. Caves, Symmetric Informationally Complete POVMs, report 9 September 1999, <http://info.phys.unm.edu/~caves/reports/infopovm.pdf>
37. J.M. Renes, R. Blume-Kohout, A.J. Scott, C.M. Caves, *J. Math. Phys.* **45**, 2171 (2004)

38. A.J. Scott, M. Grassl, *J. Math. Phys.* **51**, 042203 (2010)
39. D.M. Appleby, A. Ericsson, C.A. Fuchs, *Found. Phys.* **41**, 564 (2011)
40. C.M. Fuchs, [arxiv:1003.5209](https://arxiv.org/abs/1003.5209) (2010)
41. C. Bennett, S.J. Wiesner, *Phys. Rev. Lett.* **69**, 2881 (1992)
42. I.L. Chuang, M.A. Nielsen, *J. Mod. Opt.* **44**, 2455 (1997)
43. J.F. Poyatos, J.I. Cirac, P. Zoller, *Phys. Rev. Lett.* **78**, 390 (1997)



<http://www.springer.com/978-94-017-7302-7>

Quantum Theory: Informational Foundations and Foils
Chiribella, G.; Spekkens, R.W. (Eds.)
2016, VI, 528 p. 55 illus., 28 illus. in color., Hardcover
ISBN: 978-94-017-7302-7