

Fisher information in quantum statistics

O E Barndorff-Nielsen† and R D Gill‡§

† MaPhySto, Aarhus University, 8000 Aarhus-C, Denmark

‡ Mathematical Institute, University of Utrecht, Box 80010, 3508 TA Utrecht, Netherlands

§ EURANDOM, Box 513, 5600 MB Eindhoven, Netherlands

Received 12 November 1999, in final form 22 February 2000

Abstract. Braunstein and Caves (Braunstein S L and Caves C M 1994 *Phys. Rev. Lett.* **72** 3439–43) proposed to use Helstrom’s *quantum information* number to define, meaningfully, a metric on the set of all possible states of a given quantum system. They showed that the quantum information is nothing other than the maximal Fisher information in a measurement of the quantum system, maximized over all possible measurements. Combining this fact with classical statistical results, they argued that the quantum information determines the asymptotically optimal rate at which neighbouring states on some smooth curve can be distinguished, based on arbitrary measurements on n identical copies of the given quantum system.

We show that the measurement which maximizes the Fisher information typically depends on the true, unknown, state of the quantum system. We close the resulting loophole in the argument by showing that one can still achieve the same, optimal, rate of distinguishability, by a two-stage adaptive measurement procedure.

When we consider states lying not on a smooth curve, but on a manifold of higher dimension, the situation becomes much more complex. We show that the notion of ‘distinguishability of close-by states’ depends strongly on the measurement resources one allows oneself, and on a further specification of the task at hand. The quantum information matrix no longer seems to play a central role.

1. Introduction

Braunstein and Caves (1994) have clarified the relation between the classical Fisher expected information number $i(\theta)$, for the unknown parameter θ of a probability distribution $p(x; \theta)$, and the analogous concept of expected quantum information $I(\theta)$ for a quantum system in state $\rho = \rho(\theta)$ on some Hilbert space. They showed that $I(\theta)$ is the maximal Fisher information $i(\theta; M)$ in the distribution of the outcome of a measurement M , over all measurements of the state. Thereby they supplied a new proof of Helstrom’s (1967) quantum Cramér–Rao bound: no unbiased estimator of θ , based on any measurement, has variance smaller than $I(\theta)^{-1}$. Recall that the classical bound states that no unbiased estimator of θ based on the outcome of the measurement M has variance smaller than $i(\theta; M)^{-1}$.

For n identical copies of a quantum system, and for n independent and identically distributed observations from a probability distribution, quantum and Fisher information are both n times the corresponding quantities for $n = 1$. By classical statistical theory, the quantum bound is therefore *asymptotically* achieved, as $n \rightarrow \infty$, by the maximum likelihood estimator of θ based on the outcomes of the measurement maximizing the Fisher information for $n = 1$, applied to each of n copies of the quantum system separately.

In this paper, we analyse the conditions for equality of the quantum and Fisher information. We show that in general there does not exist a measurement M such that $i(\theta; M) = I(\theta)$ for

Processing

A/109552/PAP

Focal Image

Printed: 10/05/2000

To author

.....

CRC data

| | | | |
|-----------|---|------|-------------|
| File name | A | .TEX | First page |
| Date req. | | | Last page |
| Issue no. | | | Total pages |

all θ simultaneously, studying the pure-state, spin-half case in detail. In that case the model describes a curve on the surface of the unit sphere, specifying the direction of the spin as a function of θ . We show that one has uniform attainability if and only if the curve is a segment of a great circle. We show how (in general) adaptive measurements still allow one to asymptotically achieve the quantum information bound for a scalar parameter—though not in the vector case, where the picture is rather complicated and the quantum information matrix inadequate to describe what is possible.

In section 2 we recapitulate some of the theory of classical and quantum information. Next, in section 3, we specialize the conditions for attainability of the information bound, first to pure states, then further to spin-half models. Unless the model specifies a great circle, no measurement achieves the bound uniformly in the parameter θ . In section 4 we explore the consequences of this result. We show that one can in effect achieve $i(\theta; M_n) \approx nI(\theta)$ for all θ simultaneously, when we measure n identical copies of the quantum system in one joint measurement M_n . This result gives support to Braunstein and Cave's interpretation of the quantum information number $I(\theta)$ as a measure of statistical distinguishability between neighbouring quantum states. Finally, we turn to the case when the parameter is a vector. Both quantum and classical information numbers have matrix generalizations, and inequality between them still holds, in the sense of positive semi-definite matrices. The inequality is sharp but however no longer attainable. For a completely unknown spin-half pure state we show that the optimal rate at which one can distinguish between different states does not follow from the quantum information matrix in the way one would expect from analogy with classical Fisher information. Moreover, it depends on some weighting of the different aspects of the states which one wants to distinguish. Major open problems remain, and the role of the quantum information does not appear to be primary.

A preliminary version of this paper appeared as Barndorff-Nielsen and Gill (1998).

2. Expected classical and quantum information

On a given Hilbert space, consider a quantum state (density operator) $\rho = \rho(\theta)$, which depends on an unknown scalar parameter θ . Consider also a generalized measurement (operator-valued probability measure, POVM) M with outcomes in a measurable space $(\mathcal{X}, \mathcal{A})$. Thus the outcome of a measurement of M on ρ is a random variable X taking values in \mathcal{X} , such that for each measurable subset A of \mathcal{X} , i.e. for each $A \in \mathcal{A}$, we have $\Pr_\theta\{X \in A\} = \text{tr } \rho(\theta)M(A)$. Suppose that M is dominated by a sigma-finite measure μ on $(\mathcal{X}, \mathcal{A})$, i.e. for each $A \in \mathcal{A}$,

$$M(A) = \int_A m(x)\mu(dx)$$

where the operator $m(x)$ is, for each x , non-negative and selfadjoint, and $\int_{\mathcal{X}} m(x)\mu(dx) = \mathbf{1}$. (This is no restriction for finite-dimensional quantum systems, for which one can always take μ to be the measure defined by $\mu(A) = \text{tr } M(A)$.) Under the domination assumption, the outcome X of a measurement of M on ρ has probability density, with respect to μ , given by

$$p(x; \theta) = \text{tr } \rho(\theta)m(x).$$

Under sufficient smoothness, the expected Fisher information number, for θ , from this measurement, is defined by

$$i(\theta; M) = E_\theta \dot{l}(\theta)^2 = \int_{\mathcal{X}} (\dot{l}(x; \theta))^2 p(x; \theta)\mu(dx)$$

where

$$l(\theta) = l(X; \theta) = \log p(X; \theta)$$

is the *log likelihood* and

$$\dot{l}(\theta) = \frac{\partial}{\partial \theta} l(X; \theta)$$

is the *score function* for θ .

Now, let $\lambda = \lambda(\theta)$ denote the symmetric logarithmic derivative of ρ with respect to θ : that is, the self-adjoint operator given implicitly by

$$\dot{\rho} = \frac{1}{2}(\rho\lambda + \lambda\rho). \quad (1)$$

We call λ the *quantum score* for θ . From the relation $\text{tr } \rho = 1$ one finds, by differentiating, $\text{tr } \rho\lambda = 0$. The expected *quantum information* number for θ is defined by

$$I(\theta) = \text{tr } \rho\lambda^2.$$

Note that this quantity is defined without reference to any particular measurement M .

For future reference, we mention that when θ is a vector parameter, the Fisher information matrix is defined in the obvious way, while the quantum information matrix has ij -component $\frac{1}{2} \text{tr } \rho(\lambda_i\lambda_j + \lambda_j\lambda_i)$ where $\lambda_i(\theta)$ is the quantum score for θ_i keeping the other components of θ fixed. For completeness, we mention that there exist other generalizations of quantum information; see Yuen and Lax (1973), Belavkin (1976) and the books by Helstrom (1976) and Holevo (1982).

Define

$$\mathcal{X}_0 = \mathcal{X}_0(M, \theta) = \{x : p(x; \theta) = 0\} \quad (2)$$

and let \mathcal{X}_+ be its complement. One can express the Fisher information $i(\theta; M)$ in terms of the quantum score for ρ :

$$i(\theta; M) = \int_{\mathcal{X}_+} p(x; \theta)^{-1} (\text{Re } \text{tr}(\rho\lambda m(x)))^2 \mu(dx).$$

This follows on noting that

$$\begin{aligned} \dot{l}(\theta) &= p(x; \theta)^{-1} \text{tr } \dot{\rho} m(x) \\ &= p(x; \theta)^{-1} \frac{1}{2} \text{tr}((\rho\lambda + \lambda\rho)m(x)) \\ &= p(x; \theta)^{-1} \text{Re } \text{tr}(\rho\lambda m(x)). \end{aligned}$$

The usual proof of the quantum Cramér–Rao inequality, see Helstrom (1976) or, for a more abstract and precise version, Holevo (1982), follows closely the lines of the usual proof of the classical bound: write down the unbiasedness relation, differentiate under the integral sign, and apply the Cauchy–Schwarz inequality; Holevo (1982) gave a more rigorous proof on the same lines. Braunstein and Caves (1994) noted how the quantum bound could be obtained from the classical bound together with their new inequality $i(\theta; M) \leq I(\theta)$ for all measurements M . The derivation is a chain of three inequalities and therefore leads to a set of three necessary and sufficient conditions for equality (though they did not notice the third). Before presenting the derivation we list the three ingredients. For a given $x \in \mathcal{X}_+$, let $A = m(x)^{1/2} \rho^{1/2}$, $B = m(x)^{1/2} \lambda \rho^{1/2}$, and $z = \text{tr}(A^*B)$. The first inequality step uses the trivial $(\text{Re } (z))^2 \leq |z|^2$ with equality if and only if $\text{Im } (z) = 0$. The second uses the Cauchy–Schwarz inequality $|\text{tr}(A^*B)|^2 \leq \text{tr}(A^*A) \text{tr}(B^*B)$ with equality if and only if $\text{tr}(A^*A)B = \text{tr}(A^*B)A$. The third inequality step, $\text{tr}(M(\mathcal{X}_+)\lambda\rho\lambda) \leq \text{tr}(\rho\lambda^2)$, follows from the fact that $M(\mathcal{X}_+) = \mathbf{1} - M(\mathcal{X}_0)$ where $M(\mathcal{X}_0) \geq \mathbf{0}$. The three ingredients are put together as follows:

$$i(\theta; M) = \int_{\mathcal{X}_+} p(x; \theta)^{-1} (\text{Re } \text{tr}(\rho\lambda m(x)))^2 \mu(dx)$$

$$\begin{aligned}
&\leq \int_{\mathcal{X}_+} p(x; \theta)^{-1} |\operatorname{tr}(\rho \lambda m(x))|^2 \mu(dx) \\
&= \int_{\mathcal{X}_+} |\operatorname{tr}((m(x)^{\frac{1}{2}} \rho^{\frac{1}{2}})^* (m(x)^{\frac{1}{2}} \lambda \rho^{\frac{1}{2}}))|^2 (\operatorname{tr}(\rho m(x)))^{-1} \mu(dx) \\
&\leq \int_{\mathcal{X}_+} \operatorname{tr}(m(x) \lambda \rho \lambda) \mu(dx) \\
&= \operatorname{tr}(M(\mathcal{X}_+) \lambda \rho \lambda) \\
&\leq \operatorname{tr}(\rho \lambda^2) = I(\theta).
\end{aligned} \tag{3}$$

With A , B and $z = \operatorname{tr}(A^* B)$ as above (depending on x), necessary and sufficient conditions for equality at the first two inequality steps in (3) together are equivalent to: for $\mu(dx)$ almost all x in \mathcal{X}_+ , $\operatorname{tr}(A^* B)$ is real and $A \propto_{\mathbb{R}} B$, by which we mean $A = rB$ or $B = rA$ for some real number r . But if $A \propto_{\mathbb{R}} B$ then automatically $\operatorname{tr}(A^* B)$ is real. Thus we have equality in (3) if and only if the following two conditions are satisfied: firstly, for $\mu(dx)$ almost all x in \mathcal{X}_+

$$m(x)^{\frac{1}{2}} \lambda \rho^{\frac{1}{2}} \propto_{\mathbb{R}} m(x)^{\frac{1}{2}} \rho^{\frac{1}{2}} \tag{4}$$

and secondly,

$$\operatorname{tr}(M(\mathcal{X}_0) \lambda \rho \lambda) = 0. \tag{5}$$

Obviously a sufficient condition for (5) is that $M(\mathcal{X}_0) = \mathbf{0}$, and a sufficient condition for that is $p(x; \theta) > 0$ for all μ almost all x . Braunstein and Caves remark that a sufficient condition for (4) is that each $m(x)$ is proportional to a projector onto an eigenspace of λ . In particular, if the measurement M is a simple (von Neumann) measurement of the observable λ then (4) is satisfied. However, this is not a necessary condition for attainability. Thus the obvious fact that, in general, $\lambda(\theta)$ varies with θ , does not show that there are no measurements attaining the bound (3) for all θ simultaneously. We will do this by a further study of condition (4) in a special case.

3. Attainability of the quantum information bound

In this section we concentrate on models for pure states, $\rho = |\psi\rangle\langle\psi|$ where $|\psi\rangle = |\psi(\theta)\rangle$. In this case, the quantum score can be computed explicitly and condition (4) simplifies. Define the (unnormalized) state $|a\rangle = 2|\dot{\psi}\rangle$. Since $\rho^2 = \rho$, we have $\dot{\rho} = \rho\dot{\rho} + \dot{\rho}\rho$. The defining equation (1) for the quantum score therefore tells us that $\lambda = 2\dot{\rho} = |a\rangle\langle\psi| + |\psi\rangle\langle a|$. Now let $|1\rangle = |\psi\rangle$ and let $|2\rangle$ be a normalized orthogonal state such that $|a\rangle$ is in the subspace spanned by $|1\rangle$ and $|2\rangle$; write $|a\rangle = a_1|1\rangle + a_2|2\rangle$ where $a_1 = \langle 1|a\rangle$ and $a_2 = \langle 2|a\rangle$. (Note that all these definitions are relative to a given value of the parameter θ .) We find that

$$\rho^{\frac{1}{2}} \lambda = \rho \lambda = 2(\operatorname{Re} a_1)|1\rangle\langle 1| + \bar{a}_2|1\rangle\langle 2|$$

and condition (4) reduces to

$$m(x)^{\frac{1}{2}} (2(\operatorname{Re} a_1)|1\rangle\langle 1| + a_2|2\rangle\langle 1|) \propto_{\mathbb{R}} m(x)^{\frac{1}{2}} |1\rangle\langle 1|.$$

This can be again simplified, resulting in the condition

$$a_2 m(x)^{\frac{1}{2}} |2\rangle\langle 1| \propto_{\mathbb{R}} m(x)^{\frac{1}{2}} |1\rangle\langle 1|$$

or equivalently

$$a_2 m(x)^{\frac{1}{2}} |2\rangle \propto_{\mathbb{R}} m(x)^{\frac{1}{2}} |1\rangle. \tag{6}$$

For spin-half models, thus a Hilbert space of dimension two, a further simplification occurs. One can take $|2\rangle = |\psi\rangle^\perp$, forming an orthonormal basis (depending on θ) with $|1\rangle = |\psi\rangle$.

From (6) it follows that if $m(x)$ satisfies (4), it must have less than full rank, and hence in the two-dimensional case both it and its square root must be proportional (with real constants of proportionality) to $|\xi\rangle\langle\xi|$ for some normalized state $|\xi\rangle = |\xi(x)\rangle$. A minor rewriting yields that (4) is equivalent to the statement: for $p(x; \theta)\mu(dx)$ almost all x , $m(x)$ is proportional to a one-dimensional projector $|\xi(x)\rangle\langle\xi(x)|$ satisfying

$$\langle\xi|2\rangle\langle 2|a\rangle \propto_{\mathbb{R}} \langle\xi|1\rangle. \quad (7)$$

We show that this algebraic condition has a simple geometric interpretation. First note that from the definition of $|a\rangle$ and the fact that $\langle\psi|\psi\rangle = 1$ for all θ , it follows that $2\operatorname{Re}\langle a|1\rangle = \langle a|1\rangle + \langle 1|a\rangle = 0$, hence $\langle a|1\rangle$ is purely imaginary. By multiplying $|2\rangle$ by a suitable phase factor, one can arrange that $\langle a|2\rangle$ is real and (7) becomes $\langle\xi|2\rangle \propto_{\mathbb{R}} \langle\xi|1\rangle$. Note that $2\hat{\rho} = |1\rangle\langle a| + |a\rangle\langle 1|$. It follows that $\langle 1|\hat{\rho}|1\rangle = 0 = \langle 2|\hat{\rho}|2\rangle$, while $\langle 1|\hat{\rho}|2\rangle$ is real. Hence $2\hat{\rho} = r(|1\rangle\langle 2| + |2\rangle\langle 1|)$ for some real number r . Let $\sigma_x = |1\rangle\langle 2| + |2\rangle\langle 1|$, $\sigma_y = -i|1\rangle\langle 2| + i|2\rangle\langle 1|$, $\sigma_z = |1\rangle\langle 1| - |2\rangle\langle 2|$ be the Pauli spin matrices with respect to the basis $|1\rangle, |2\rangle$. In this basis $\rho = \frac{1}{2}(\mathbf{1} + \sigma_z)$ and $2\hat{\rho} = r\sigma_x$. We can write $|\xi\rangle\langle\xi| = \frac{1}{2}(\mathbf{1} + \vec{\alpha} \cdot \vec{\sigma})$ where $\vec{\alpha} = (\alpha_x, \alpha_y, \alpha_z)$ is a unit vector in \mathbb{R}^3 . Note that $\langle 2|\xi\rangle\langle\xi|1\rangle = \alpha_x + i\alpha_y$. Therefore (7) holds if and only if $\alpha_y = 0$.

With respect to an arbitrary *fixed* basis we can write $\rho(\theta) = \frac{1}{2}(\mathbf{1} + \vec{u}(\theta) \cdot \vec{\sigma})$ and $\lambda(\theta) = 2\hat{\rho}(\theta) = \vec{u}(\theta) \cdot \vec{\sigma} = r(\theta)\vec{v}(\theta) \cdot \vec{\sigma}$ where $\vec{u}(\theta)$ and $\vec{v}(\theta)$ are orthogonal unit vectors, $r(\theta) = \|\vec{u}(\theta)\|$ is real and non-negative, and $\vec{\sigma}$ is the vector of the Pauli spin matrices with respect to the fixed basis. Using the familiar relations $\sigma_x^2 = \mathbf{1}$, $\sigma_x\sigma_y = -\sigma_y\sigma_x = i\sigma_z$, and their cyclic permutations, and the fact that the spin matrices are traceless, one finds that the quantum information $I(\theta) = r(\theta)^2$. The direction $\vec{v}(\theta)$ is uniquely defined if $r(\theta) > 0$ and from now on we assume this is true for all θ . Then (4) is satisfied if and only if for $p(x; \theta)\mu(dx)$ almost all x , $m(x)$ is proportional to a projector for a spin direction in the plane spanned by $\vec{u}(\theta)$ and $\vec{v}(\theta)$. In particular, *any* simple (von Neumann) measurement of spin in a direction in this plane attains equality in (3) if both outcomes have positive probability.

Let $\mathcal{C}(\theta)$ denote the great circle on the unit sphere formed by the intersection of the sphere with the plane $\mathcal{P}(\theta)$ spanned by $\vec{u}(\theta)$ and $\vec{v}(\theta)$, and let $\vec{n}(\theta)$ be the normal unit vector to this plane. We suppose that these objects vary smoothly with θ . Recall that $\vec{u}(\theta)$ moves in the direction $\vec{v}(\theta)$. Now as θ varies, either $\mathcal{P}(\theta)$, $\mathcal{C}(\theta)$, and $\vec{n}(\theta)$ are all fixed or they all vary. In particular, the intersection over all θ of the planes $\mathcal{P}(\theta)$ is either a fixed plane \mathcal{P} (equal to $\mathcal{P}(\theta)$ for all θ), or a fixed straight line \mathcal{L} , or the origin $\vec{0}$. In the second case the normal $\vec{n}(\theta)$ and the circle $\mathcal{C}(\theta)$ must be rotating about the fixed line \mathcal{L} . The rotation must have nonzero speed for values of θ in a set of positive measure. If we assume that the model is identified, so that $\rho(\theta)$ is a one-to-one function of θ , then at most for two values of θ can the true spin $\vec{u}(\theta)$ lie in \mathcal{L} . So there exists a θ for which the rotation has nonzero speed and $\vec{u}(\theta)$ does not lie in the axis of rotation. But then at this point the derivative of $\vec{u}(\theta)$ must have a nonzero component in the direction orthogonal to the plane $\mathcal{P}(\theta)$, which is a contradiction.

So we only have two possibilities: either the plane $\mathcal{P}(\theta)$ is fixed and $\vec{u}(\theta)$ is moving on the great circle \mathcal{C} in the plane, or $\vec{u}(\theta)$ moves on some other curve and the intersection of all $\mathcal{P}(\theta)$ contains only the origin. In the first case any measurement with all components proportional to projectors of directions in this plane, and with $p(x; \theta) > 0$ for all x and θ , achieves the inequality (3) uniformly in θ . Conversely, under the positivity of $p(x; \theta)$, only such measurements uniformly achieve the bound. In the second case a measurement which uniformly achieves the bound would have to have all $m(x)$ equal to zero, which is impossible. Thus, there is no uniformly attaining measurement in this case.

For example, consider a spin-half particle in the pure state $|\psi\rangle = |\psi(\eta, \theta)\rangle$ given by

$$|\psi\rangle = \begin{pmatrix} e^{-i\theta/2} \cos(\eta/2) \\ e^{i\theta/2} \sin(\eta/2) \end{pmatrix}. \quad (8)$$

This pure state has density matrix $\rho = \frac{1}{2}(\mathbf{1} + \vec{u} \cdot \vec{\sigma})$ where $\vec{u} = \vec{u}(\eta, \theta)$ is the point on the unit sphere in \mathbb{R}^3 with polar coordinates (η, θ) . Suppose the colatitude $\eta \in [0, \pi]$ is known and exclude the degenerate cases $\eta = 0$ or $\eta = \pi$; the longitude $\theta \in [0, 2\pi)$ is the unknown parameter.

We have a pure state so $\lambda = 2\dot{\rho} = 2\dot{\vec{u}} \cdot \vec{\sigma} = \sin(\eta)\dot{\vec{u}}(\pi/2, \theta + \pi/2) \cdot \vec{\sigma} = r(\theta)\dot{\vec{v}}(\theta)$. The quantum information is $r^2 = \sin^2 \eta$. As θ varies, $\vec{u}(\theta)$ traces out a great circle if and only if $\eta = \pi/2$. Consequently, for $\eta \neq \pi/2$, no measurement M exists with Fisher information $i(\theta; M)$ equal to the quantum information $I(\theta)$ whatever the value of the unknown parameter θ . If $\eta = \pi/2$ it is possible to achieve the bound uniformly in θ . Any measurement with everywhere positive density and all components proportional to projector matrices for spin directions in the plane $\eta = \pi/2$ will do the job. A simple measurement of spin in one particular direction in that plane attains the information bound at all θ except for θ equal to that direction or opposite to it. At these points the distribution of the outcome is degenerate and the Fisher information not defined. However, since the Fisher information is continuous (indeed, constant) in θ this is a non-essential singularity.

4. Asymptotic attainability and vector parameters

We have shown, for the case of a one-dimensional parameter, that only for rather special models will a measurement M exist such that $i(\theta; M) = I(\theta)$ for all parameter values θ simultaneously. It is on the other hand possible to find a measurement M such that at a given parameter value, $i(\theta; M) = I(\theta)$, as Braunstein and Caves indicate: take each $m(x)$ proportional to a projector onto an eigenspace of the quantum score $\lambda(\theta)$. They do not remark on the possible dependence of M on θ . However, if all we know is that $\rho = \rho(\theta)$ for some θ , we do not know which measurement to use. The eigenspace decomposition of λ generally depends on θ so this does not define a measurement M which achieves the bound uniformly in θ . This is not the only way to achieve the bound, but the previous section shows that one cannot in general expect there to be a uniformly attaining measurement.

Note that the classical information based on n independent and identically distributed realizations from a given density $p(x, \theta)$ is equal to n times the information for one realization. Similarly, the quantum information in the state $\rho(\theta)^{\otimes n}$ corresponding to n identical particles each in state $\rho(\theta)$ is n times the quantum information for one particle.

Braunstein and Caves' aim was to define a statistical distinguishability metric between quantum states. Suppose the measurement M on a single particle satisfies $i(\theta; M) = I(\theta)$. Then the maximum likelihood estimator of θ based on n separate measurements of M on identical copies of the given quantum system, by classical results in mathematical statistics, is generally an asymptotically unbiased estimator with asymptotic variance $(ni(\theta; M))^{-1} = (nI(\theta))^{-1}$. By the quantum Cramér–Rao bound applied to the joint system of n particles, no estimator based on any measurement whatsoever on $\rho^{\otimes n}(\theta)$ can do better. Thus $I(\theta)$ appears to exactly characterize the rate at which one can determine θ .

However, this argument is flawed since the measurement M involved will be a different measurement for each θ , and the whole point is that θ is not known in advance. The question therefore remains: does there exist a measurement procedure *not depending on* θ on the state $\rho^{\otimes n}$, on the basis of which an estimator of θ can be constructed having asymptotic variance $(nI(\theta))^{-1}$? If the answer is ‘yes’, then Braunstein and Caves' proposed role for the quantum information $I(\theta)$ in defining a statistical distinguishability metric is well motivated.

It seems rather natural to try a two-stage procedure: first estimate the parameter using a perhaps inefficient procedure on a vanishing proportion of the particles, say $n_0 = n^\alpha$ ($0 < \alpha < 1$) out of the total of n ; now carry out the ‘estimated optimal measurement’ on the

remaining ones. In both stages only simple or von Neumann measurements (measurements of classical observables) on separate particles are needed.

In our example (8) this would reduce to the following. Measure the spin σ_x on $k = \frac{1}{2}n_0$ of the copies. The number of +1s observed is binomially distributed with parameters k and $p = \frac{1}{2}(\sin \eta \cos \theta + 1)$. Similarly, for another k measurements of the spin σ_y we get a binomial number of '+1' with parameters k and $p = \frac{1}{2}(\sin \eta \sin \theta + 1)$. This allows us consistent estimation of both $\sin \theta$ and $\cos \theta$ and hence of $\theta \in [0, 2\pi)$. Denote such an estimator by $\tilde{\theta}$. We saw that λ in this example was proportional to the spin in the direction $(\pi/2, \theta + \pi/2)$. Let us use the remaining $n' = n - n_0$ particles to measure this spin with θ replaced by $\tilde{\theta}$. Given $\tilde{\theta}$, this results in a binomial number X of '+1' with parameters n' and $p = \frac{1}{2}(1 - \sin \eta \sin(\theta - \tilde{\theta}))$. Let

$$\hat{\theta} = \tilde{\theta} + \arcsin((n' - 2X)/(n' \sin \eta)).$$

Analysis of this 'final' estimator shows that $\hat{\theta}$ has asymptotically the $\mathcal{N}(\theta, (n \sin^2(\eta))^{-1})$ distribution (the normal distribution with indicated mean and variance), whatever θ , so that the quantum information bound is asymptotically achievable by our two-stage procedure.

This approach will work in wide generality in problems with a one-dimensional parameter θ . Suppose, as typically will be the case, that one can construct a consistent estimator $\tilde{\theta}$ based on certain measurements on a vanishing proportion of the particles. Compute the quantum score at $\theta = \tilde{\theta}$, and measure it on each of the remaining particles. Compute the maximum likelihood estimator $\hat{\theta}$ of θ based on the new data, whose probability distribution depends on the unknown θ (as well as on $\tilde{\theta}$, which is at this stage fixed). We argue as follows that $\hat{\theta}$ has approximately the $\mathcal{N}(\theta, (nI(\theta))^{-1})$ distribution, thus this estimator asymptotically achieves the quantum information bound. Let $i(\theta; \tilde{\theta})$ denote the Fisher information for θ in a measurement, on one particle, of the quantum score at $\tilde{\theta}$; thus $i(\tilde{\theta}; \tilde{\theta}) = I(\tilde{\theta})$ for all values of $\tilde{\theta}$, but generally $i(\theta; \tilde{\theta}) < I(\theta)$. Now for n large, $\tilde{\theta}$ is close to θ . By the classical results for maximum likelihood estimators, given $\tilde{\theta}$, $\hat{\theta}$ has approximately the $\mathcal{N}(\theta, (ni(\theta; \tilde{\theta}))^{-1})$ distribution. So if ρ depends on θ smoothly enough that $i(\theta; \tilde{\theta})$ is close to $i(\theta; \theta) = I(\theta)$ for $\tilde{\theta}$ close to θ , we have that unconditionally $\hat{\theta}$ has approximately the $\mathcal{N}(\theta, (nI(\theta))^{-1})$ distribution, hence asymptotically achieves the bound.

Consider now the case of vector parameters. Both quantum and Fisher information numbers are naturally generalized to information matrices; see Helstrom (1976) and Holevo (1982). The Braunstein and Caves result generalizes to the following result: the quantum information matrix is larger (in the sense that the difference is positive semi-definite) than the Fisher information matrix based on the outcome of any measurement M . However, the bound is no longer attainable. As we saw above, a best measurement for each parameter separately is a measurement of the quantum score operator. Typically these do not commute and hence cannot be measured simultaneously. On the other hand, consideration of all smooth one-dimensional sub-models of a given model $\rho = \rho(\theta)$ shows that the quantum information matrix is the smallest matrix larger than the Fisher information matrix of any measurement on a single particle.

For instance, suppose we want to simultaneously estimate both parameters η, θ of the pure-state, spin-half system (8); in other words, we have a completely unknown pure state. Rename θ as ϕ , and let θ from now denote the vector parameter with elements η, ϕ . Suppose we may dispose of a large number of identical copies of this system. Let $I(\theta)$ denote the 2×2 quantum information matrix, and $i(\theta; M)$ denote the Fisher information matrix based on the outcome of a measurement M , both for a single copy of the quantum system. The quantum scores, for a single particle, for the two parameters η and ϕ are $\sigma_{\eta+\pi/2, \phi}$ and $\sin \eta \sigma_{\pi/2, \phi+\pi/2}$ respectively. After a small proportion of measurements we know roughly the location of

the parameter, and it is sufficient to investigate optimal measurement at a ‘known’ parameter value.

Without loss of generality let this be the special point θ_0 with $\eta_0 = \pi/2$, $\phi_0 = 0$; $\rho = \frac{1}{2}(\mathbf{1} + \sigma_x)$. At this point the quantum scores are σ_y and $-\sigma_z$, and the quantum information matrix is the identity $\mathbf{1}$. The intersection of the xy and the xz planes is the x -axis, so it seems that in order to maximize Fisher information both for η and for ϕ we should essentially measure spin in the x direction. However, the Fisher information matrix $i(\theta; M)$ based on the outcome of this measurement is not defined at $\theta = (\pi/2, 0) = \theta_0$ since the distribution of the outcome is degenerate. The singularity is essential since $i(\theta; M)$ is not continuous at $\theta = \theta_0$. In particular, as one moves towards θ_0 along either of the two great circles formed by varying one of the two components of θ , $i(\theta; M)$ converges to each of the diagonal matrices with diagonal elements 1, 0 and 0, 1. These are the two Fisher information matrices corresponding to von Neumann measurements of the two quantum scores, each giving maximal information about the corresponding component of θ and zero information about the other.

Since different components of the parameter vector have incompatible quantum scores, it is clear that for different loss functions, different measurements will be optimal. No single procedure will (asymptotically) dominate all others. Moreover, since we cannot achieve the quantum information bound by measurements on single particles, it is possible that joint measurements on several particles simultaneously could give larger Fisher information (per particle) than measurements on separate particles.

In some very special cases, an optimal procedure is known. An appealing loss function in the completely unknown pure spin-half model is one minus the squared inner product between the true state vector and its estimate. This equals one minus the squared cosine of half the angle between the points on the Poincaré sphere representing the two states. At the special point under consideration therefore, the loss function is asymptotically equivalent to one-quarter times the sum of the squares of the errors in η and ϕ . Massar and Popescu (1995), in response to a problem posed by Peres and Wootters (1991), exhibited a measurement, optimal in the Bayes sense, with respect to this loss function and a uniform prior distribution. It had an asymptotic mean square error $4/n$. This was a genuine generalized measurement of the composite system $\rho^{\otimes n}$. They showed that for the case of $n = 2$ there were no measurement methods of the two particles separately which were as good as the optimal method, and this is expected to hold for all n .

Instead of this exactly optimal procedure (with respect to the given loss function and under a uniform prior) consider taking, with probability one-half, measurements of σ_y and σ_z , independently on each particle. We find that the Fisher information matrix (based on one observation) for η , ϕ , at $\eta = \pi/2$, $\phi = 0$, is $\frac{1}{2}\mathbf{1}$, or one-half of the quantum information matrix. The inverse of this matrix, $2\mathbf{1}/n$ is an asymptotically achievable lower bound to the covariance matrix of (asymptotically unbiased) estimators of η , ϕ based on n of such measurements. The maximum likelihood method would provide estimators asymptotically achieving this bound. The sum of the variances is $4/n$, the same as what is achieved by the Massar and Popescu procedure.

Thus the following two-stage procedure, similar to what we proposed in the one-parameter case, should have asymptotically equivalent covariance matrix to that of the Massar and Popescu procedure, and will also be optimal with respect to a uniform prior distribution and any smooth loss function, invariant under rotations of the sphere. First carry out measurements of each of σ_x , σ_y and σ_z on a small proportion of separate particles. Compute from the results a consistent estimate of (η, ϕ) . With respect to a rotated coordinate system putting the estimated value at $(\eta, \phi) = (\pi/2, 0)$, measure alternately σ_y and σ_z on the remaining particles. Estimate

(η, ϕ) in the new coordinate system by the method of maximum likelihood using the second-stage observations. Finally rotate back to the original coordinate system. Note that if we modified this scheme by just measuring σ_x in the second stage, then overall the procedure determines the radial distance of θ from $\hat{\theta}$ with precision of the order $n^{-1/2}$, but says nothing about the direction, so that finally θ has not been localized to this precision at all.

This conjecture has been confirmed by recent further work of Gill and Massar (1999). Consider a sequence of measurements on n identical copies of a spin-half state, on which is based a sequence of estimators $\hat{\theta}$. The parameter θ might be one- or two-dimensional for a pure-state model, one-, two- or three-dimensional for a mixed state. Suppose that the estimators are asymptotically unbiased and have covariance matrices asymptotically of the form $V(\theta)/n$. Then it is now known that the collection of attainable V is precisely $\{V : \text{tr } I(\theta)^{-1}V(\theta)^{-1} \leq 1 \text{ for all } \theta\}$ in any of the following cases: the parameter is one-dimensional, or the state is pure, or the measurement can be implemented by separate measurements on separate particles. These optimal limiting covariance matrices can all be achieved using a two-stage adaptive procedure of the type described above. The collection of attainable V corresponds to the collection of attainable inverse Fisher information matrices for measurements on single particles. Using joint measurements on mixed states with more than one unknown parameter, one can attain strictly smaller asymptotic covariance matrices. But a clean description of what is attainable is not known. One would like to describe the collections of scaled information matrices $\{i(\theta; M_m)/m\}$ for $m = 1, 2, \dots$, where M_m is an arbitrary joint measurement on m particles. These sets are all convex, they grow with m ; we know the set for $m = 1$, and each set is included in the set of matrices less than or equal to $I(\theta)$. The inverses of these information matrices will be the achievable (scaled) asymptotic covariance matrices based on measuring a large number n of particles in groups of m at a time.

To conclude, in the multiparameter case, the bound implied by the quantum information matrix is not even asymptotically achievable. The rate at which one can distinguish between more than two neighbouring quantum states does not correspond to the rate at which one can distinguish between just two; it depends on what aspect of the quantum states is important, and it depends on whether one may use joint measurements or only separate measurements. The quantum information matrix only plays a role in special cases.

Acknowledgments

We are grateful for useful discussions with Alessandra Luati, Klaus Mølmer, and Peter Jupp. This research was supported by the Danish National Research Foundation through MaPhySto (Centre for Mathematical Physics and Stochastics), and by the Oberwolfach Mathematical Institute's Research in Pairs programme. Richard Gill is also grateful for the hospitality of the Department of Mathematics and Statistics, University of Western Australia.

References

- Barndorff-Nielsen O E and Gill R D 1998 [An example of non-attainability of expected quantum information](https://arxiv.org/abs/quant-ph/9808009) *Preprint* [xxx.lanl.gov quant-ph/9808009](https://arxiv.org/abs/quant-ph/9808009)
- Belavkin V P 1976 Generalized uncertainty relations and efficient measurements in quantum systems *Teor. Mat. Fiz.* **26** 316–29 (Engl. Transl. 1976 **26** 213–22)
- Braunstein S L and Caves C M 1994 Statistical distance and the geometry of quantum states *Phys. Rev. Lett.* **72** 3439–43
- Gill R D and Massar S 1999 [State estimation for large ensembles](https://arxiv.org/abs/quant-ph/9902063) *Phys. Rev. A* at press (Gill R D and Massar S 1999 *Preprint* [xxx.lanl.gov quant-ph/9902063](https://arxiv.org/abs/quant-ph/9902063))

Note 1

Note 2



- Helstrom C W 1967 Minimum mean-square error of estimates in quantum statistics *Phys. Lett. A* **25** 101–2
———1976 *Quantum Detection Theory* (New York: Academic)
- Holevo A S 1982 *Probabilistic and Statistical Aspects of Quantum Theory* (Amsterdam: North-Holland) Russian original 1980
- Massar S and Popescu S 1995 Optimal extraction of information from finite quantum ensembles *Phys. Rev. Lett.* **74** 1259–63
- Peres A and Wootters W K 1991 Optimal detection of quantum information *Phys. Rev. Lett.* **66** 1119–22
- Yuen H P and Lax M 1973 Multiple-parameter quantum estimation and measurement of nonselfadjoint observables *Trans. IEEE* **19** 740–50

Author queries

Note 1 (Page 9):

Please update?

Note 2 (Page 9):

Please update?