

## When quantum tomography goes wrong: drift of quantum sources and other errors

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2013 New J. Phys. 15 025024

(<http://iopscience.iop.org/1367-2630/15/2/025024>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 131.215.71.79

The article was downloaded on 14/03/2013 at 14:51

Please note that [terms and conditions apply](#).

## When quantum tomography goes wrong: drift of quantum sources and other errors

S J van Enk<sup>1,2,4</sup> and Robin Blume-Kohout<sup>3</sup>

<sup>1</sup> Department of Physics and Oregon Center for Optics, University of Oregon, Eugene, OR 97403, USA

<sup>2</sup> Institute for Quantum Information, California Institute of Technology, Pasadena, CA 91125, USA

<sup>3</sup> Sandia National Laboratories, Albuquerque, NM 87123, USA  
E-mail: [svanenk@uoregon.edu](mailto:svanenk@uoregon.edu)

*New Journal of Physics* **15** (2013) 025024 (15pp)

Received 31 May 2012

Published 18 February 2013

Online at <http://www.njp.org/>

doi:10.1088/1367-2630/15/2/025024

**Abstract.** The principle behind quantum tomography is that a large set of observations—many samples from a ‘quorum’ of distinct observables—can all be explained satisfactorily as measurements on a *single* underlying quantum state or process. Unfortunately, this principle may not hold. When it fails, any standard tomographic estimate should be viewed skeptically. Here we propose a simple way to test for this kind of failure using the Akaike information criterion. We point out that the application of this criterion in a quantum context, while still powerful, is not as straightforward as it is in classical physics. This is especially the case when future observables differ from those constituting the quorum.

<sup>4</sup> Author to whom any correspondence should be addressed.



Content from this work may be used under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike 3.0 licence](https://creativecommons.org/licenses/by-nc-sa/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

**Contents**

<b>1. Introduction</b>	<b>2</b>
1.1. General remarks . . . . .	2
1.2. Possible errors in standard quantum state tomography . . . . .	3
1.3. Akaike to the rescue . . . . .	5
<b>2. Examples</b>	<b>6</b>
2.1. One qubit, part 1 . . . . .	6
2.2. One qubit, part 2 . . . . .	8
2.3. Two or more qubits . . . . .	9
<b>3. Model selection, the Akaike information criterion and quantum quirks</b>	<b>11</b>
<b>4. Summary and discussion</b>	<b>13</b>
<b>Acknowledgments</b>	<b>14</b>
<b>References</b>	<b>14</b>

**1. Introduction***1.1. General remarks*

The goal of quantum-state tomography [1] is to give a statistically reliable estimate of a quantum state  $\rho$ . Two further questions may come to mind: (i) what is the purpose of that estimate  $\rho$ ? and (ii) why or when are we correct in giving an estimate of *just one* quantum state?

There are at least two answers to the first question: our experiment may be aimed at producing a particular state, say, a cluster state, and we may just want to verify how close  $\rho$  is to the desired state. But that answer provides really just an intermediate goal. The ultimate goal is always to use the desired state for some particular quantum information processing task. So we could say that the goal of producing an estimate  $\rho$  is to be able to predict the future performance in a particular protocol of one or more unmeasured quantum system(s) produced by the same source.

Now there is a nice statistical method for ranking different models according to their ability to predict future measurement results (*not* on how well they fit the past data!), based on the Akaike information criterion (AIC) [2]. That criterion was developed entirely within a classical context, but it ought to apply to quantum-state estimation, too. We show this is true, even though we will point out some interesting differences between classical and quantum statistics.

The motivation behind the second question is as follows. Since we do not have full control over all physical quantities relevant to the quantum-state generation process (for example, even the best laser suffers from phase diffusion; and there are always spatially and temporally fluctuating magnetic and electric fields), the quantum states produced by a quantum source are not all identical. A possible description of the individual states of  $M$  systems  $k = 1, \dots, M$  would be a sequence  $\{\rho_k, k = 1, \dots, M\}$  where each  $\rho_{k+1}$  is a little different from the previous one (even with entanglement or correlation between the different systems, we can define  $\rho_k$  by tracing out all the other systems). So, why would we use just a single estimate  $\rho$  in this case? One aspect of the answer is, of course, that we have no way of estimating each individual  $\rho_k$ . A more positive answer is that multiple measurements of a given observable  $\hat{O}$  only yield estimates of *average* quantities such as  $\langle O \rangle = \overline{\text{Tr } \rho_k \hat{O}}$  or  $p_n = \overline{\text{Tr } \rho_k |O_n\rangle\langle O_n|}$ , where the average is over

those  $k$  on which  $\hat{O}$  was measured, and where  $|O_n\rangle$  denotes an eigenstate of  $\hat{O}$ . These averages being linear in  $\rho_k$  are determined by a *single* density matrix, namely the average density matrix  $\rho = \overline{\rho_k}$ . This simple picture has been made much more rigorous by Renner in [3]. He showed that the crucial ingredient (missing in the simple picture) is permutation invariance. That is, if we randomly permute the sequence of quantum systems, and then trace out some subset, the joint state of the remaining systems is to a good approximation independently and identically distributed (i.i.d.). In our context this means that as long as the quorum of observables is measured in a random order, then to a good approximation any one of the remaining unmeasured systems can be described by a single density matrix  $\rho$ . We now discuss what may go wrong if we measure the observables constituting a quorum in a nonrandom order.

### 1.2. Possible errors in standard quantum state tomography

It is much easier to measure a given observable from the quorum many times in a row, before switching to measurement of the next observable. Such a procedure is standard practice, but it voids Renner's proof, and so it may be that there is not a single density matrix that can be validly assigned to the remaining unmeasured quantum systems.

Let us introduce this problem with a simple example. Given an ensemble of  $3N \gg 1$  qubits that—we assume!—are identically and independently prepared, we want to estimate their density matrix. So we divide them into three equal *and sequential* groups, and measure  $\sigma_x$  on samples  $1, \dots, N$ ,  $\sigma_y$  on samples  $N+1, \dots, 2N$ , and  $\sigma_z$  on the last  $N$ . Now, if the samples are indeed identically prepared in some state  $\rho$ , then we can safely perform the measurements in this order—the state  $\rho^{\otimes 3N}$  is invariant under permutations, so all orderings are equivalent. But if the source is drifting over time, the first  $N$  copies are best described by a mean density matrix  $\bar{\rho}_1$ , while the second and third sets of  $N$  qubits are best described by (possibly different) average states  $\bar{\rho}_2$  and  $\bar{\rho}_3$ , respectively.

For an amusing (albeit extreme) example, consider a situation where the first  $N$  copies are best described by  $\bar{\rho}_1 = |+\rangle\langle+|$ , the second group by  $\bar{\rho}_2 = |+i\rangle\langle+i|$ , and the third by  $\bar{\rho}_3 = |0\rangle\langle 0|$ . The measurement outcomes in this case are not random at all: every single measurement (of  $\sigma_x$ ,  $\sigma_y$ ,  $\sigma_z$ ) will yield eigenvalue  $+1$ . Linear inversion tomography will yield a radically nonpositive state

$$\hat{\rho}_{\text{tomo}} = \begin{pmatrix} 1 & \frac{1+i}{2} \\ \frac{1-i}{2} & 0 \end{pmatrix}, \quad (1)$$

and maximum likelihood estimation (MLE) yields the projector onto  $\hat{\rho}_{\text{tomo}}$ 's positive eigenspace. Although both estimates are plausible answers to 'what single matrix best fits the observed data?', neither one of them is of any predictive use at all! The source is drifting so rapidly and drastically that *this* set of  $3N$  samples really tells us almost nothing about future observations. This is the simplest and best conclusion at which our data analysis should arrive.

This is a rather extreme and contrived example of experimental drift (below we will discuss a more common type of nonrandom experiment where the above cycle of measurements is repeated once: so we measure  $\sigma_x$  on the first  $N/2$  copies, then  $\sigma_y$ , then  $\sigma_z$ , and then  $\sigma_x$ ,  $\sigma_y$ ,  $\sigma_z$  again, each on  $N/2$  sequential copies). More realistic examples show similar behavior, though. The statistics given above are actually more consistent with a different (and still plausible) mechanism: when the measurement apparatus is 'rotated' to perform a different measurement,

the experimenter inadvertently ‘rotates’ the samples as well. A particularly naïve version of this could occur with photon polarization, where one way to physically rotate a polarizer is for the experimentalist to simply rotate his own frame of reference (e.g., by lying down). Such a passive rotation obviously fails to change the relative orientation of samples and apparatus. More realistic examples occur when similar quantum gate devices are used to (i) prepare states (e.g. EPR states) and (ii) implement measurements. In quantum process tomography, this sort of pitfall is well known; it violates the conditions for complete positivity of processes, and causes negative eigenvalues just as in our example above [4].

All of these failures are examples of a single phenomenon: *sample-apparatus correlation*. In process tomography, this is usually explained by correlation between the system and its environment. In state tomography, there is no environment *per se*, but if the state of the  $k$ th sample is (in any way) correlated with the behavior of the measurement apparatus (e.g. with what measurement it is oriented to perform), then tomography goes wrong. Experimental drift is a simple and easy to understand example: the sample state is correlated with time, and if the apparatus setting is also allowed to vary with time, then there will be sample-apparatus correlation. As noted above, this can be eliminated by explicitly randomizing the order of measurements, so that while the samples are still time-dependent, the apparatus is not. Other kinds of sample-apparatus correlation are not so easy to remedy.

In the example given above, the extremity of the data—and the fact that the linear inversion estimate is radically negative—are a dead giveaway. On the other hand, linear inversion can produce negative estimates even with ideal data [5, 6] because of statistical fluctuations. The *raison d’être* of MLE is to fix this negativity, but by constraining the estimate to positive states, MLE also hides the tell-tale signature of failed tomography. Moreover, negative estimates are not (in general) a reliable symptom even of drastic experimental drift. If the drifting states in the example above were a bit more mixed—e.g.  $\bar{\rho}'_k = \frac{1}{2}\bar{\rho}_k + \frac{1}{4}\mathbb{1}$ —then linear inversion and MLE would yield identical and positive density matrices. But, just as in the original example, those estimates would be useless and not predictive.

Fortunately, there is a general solution to this problem. It elegantly generalizes the observation (made above) that a radically negative  $\hat{\rho}_{\text{tomog}}$  should trigger skepticism. It can also diagnose drift in the absence of negativity if the data are sufficiently rich. It is called *model selection*.

The core principle is that, when tomography fails:

1. The standard model for tomography—i.i.d. samples described by a single density matrix—is bad.
2. Some other model will be better.
3. We can quantify ‘bad’ and ‘better’, and use the results to decide whether our tomography went wrong.

Clearly, putting this into practice requires that we come up with alternative models to describe the data. Model *design* is more of an art than a science. Here, we demonstrate alternative models for some simple and relevant problems, and leave the rich problems of general and optimal alternative-model design to future work. Instead, we focus on model *selection*, which means determining whether (i) the standard tomographic model is pretty good, or (ii) some other model (e.g. a drifting source model) is better.

### 1.3. Akaike to the rescue

To accomplish this, we propose, as we mentioned above, to use the AIC [2]. Widely used outside of physics [7, 8], the AIC is relatively unknown within the physics community. However, it has been applied in astrophysics [9], entanglement verification [10] and quantum state estimation [11–13]. Its function is to quantify (by assigning a real number) how well a given model describes the data from a given experiment. The AIC’s absolute value is not meaningful, but the *relative* AIC values for multiple different models have a deep and useful meaning (see following section for a more detailed discussion of the AIC, its meaning, and its derivation). Their simplest use is to rank all the different models, and thus to identify (a) which is the best, and (b) how significantly ‘worse’ the others are.

The AIC assigns a number  $\Omega_k$  to each model  $k$ , given by<sup>5</sup>

$$\Omega_k := \ln \mathcal{L}_k - K_k, \quad (2)$$

where  $\mathcal{L}_k$  is the likelihood of model  $k$ —or, if model  $k$  has adjustable parameters (as is usually the case), the *maximum* of the likelihood over all those parameters—and  $K_k$  is the number of independent model parameters used in model  $k$  to fit the data<sup>6</sup>. The larger the AIC ( $\Omega_k$ ) is, the higher the model is ranked. While  $\Omega_k$ ’s absolute value is meaningless, the difference  $\Delta = \Omega_k - \Omega_{k'}$  represents (roughly speaking) the weight of evidence in favor of  $k$  over  $k'$ , measured in bits. So, for example, if we want to report a weighted average of the two models, the ratio of the weights assigned to models  $k$  and  $k'$  should be  $w_k/w_{k'} = \exp(\Omega_k - \Omega_{k'})$ .

The AIC’s simple form admits a simple interpretation: fitting the data better (higher likelihood) is good, but extra parameters are bad. Additional parameters must justify their existence by improving the likelihood (a measure of goodness-of-fit) by at least a factor of  $e$ . This helps to prevent overfitting. Adding adjustable parameters will always improve a model’s fit—but a good fit to past data is *not* a guarantee that the model will accurately predicting future measurements.

*Example.* If we measure each of  $3N$  qubits, measuring  $\hat{O}_j$  on qubits  $j$  for  $j = 1, \dots, 3N$ , then the best possible fit to the data is to assume that each qubit  $j$  just happened to be in the appropriate eigenstate of  $\hat{O}_j$  so that the probability of the observed data is  $\mathcal{L} = 1!$

Intuitively, this ‘explanation’ is absurd. The AIC quantifies that intuition; that model requires a huge ( $O(N)$ ) number of parameters, and the resulting penalty will overwhelm its higher likelihood, ensuring that its AIC is far worse than that of simpler models.

To apply the AIC to our example of section 1.2, we need an alternative model (the ‘standard model’ just uses a single density matrix for all  $3N$  qubits). A simple alternative that describes experimental drift (as well as some other forms of sample-apparatus correlation) is to use one density matrix for each of the three groups of samples. This alternative model will always fit the data at least as well, but it may use more parameters<sup>7</sup>. The AIC ranks both models, and quantifies how much better one is than the other. We perform and analyze this calculation for

<sup>5</sup> Usually an overall minus sign and an extra factor of 2 appear on the right-hand side of (2), but for our purposes it is more convenient not to include those.

<sup>6</sup> Actually, when the number of measurements (e.g.  $3N$  in the example given) is small, there is a correction term. The *corrected AIC* ( $\text{AIC}_c$ ) is given by  $\Omega_k^c = \Omega_k - K_k(K_k + 1)/(3N - K_k - 1)$ , with  $\Omega_k$  given by 2. Hereafter, we will simply assume  $3N \gg K_k$ .

<sup>7</sup> Intuitively, the alternative model is more contrived, and this should be reflected in its ranking. However, it is not immediately obvious how to quantify this complexity.

our single-qubit example (where just two models are sufficient) in section 2.1, and address more complicated variations on this theme—with multiple alternative models—in section 2.3.

To conclude this (long) introduction, we note that the appearance of maximum likelihoods in the AIC does *not* imply any privileged role for MLE estimation of states or any other physical quantities. The likelihood is a central concept in statistics, and appears in almost every method. In the AIC, it is used specifically to quantify goodness-of-fit, and (obviously) the AIC balances this quantity against another (model complexity). Moreover, the AIC is used only to rank different models. There is no implicit requirement that the highest-ranked model must be chosen exclusively (in fact, a common strategy is to average over high-ranked models), and even if the ‘best’ model is chosen, we remain free to analyze that model without MLE (e.g. via Bayesian averaging).

## 2. Examples

In this section we first treat the example from the introduction, tomography on single qubits, in more detail (section 2.1). In this example, inconsistencies can arise only when the observed average values of  $\sigma_x$ ,  $\sigma_y$ ,  $\sigma_z$  are inconsistent with each other, which in turn can only happen if the density matrix obtained by linear inversion is unphysical. The next example, discussed in section 2.2, also concerns single qubits, but now measurements of  $\sigma_x$ ,  $\sigma_y$ ,  $\sigma_z$  are each repeated once. In this (experimentally more relevant) case inconsistencies can arise when two estimates of the same quantity are statistically different. Ad hoc methods that just consider this particularly simple type of inconsistencies work just as well as the AIC. In the last section, 2.3, we will consider the case of multiple qubits, in which the validity of ad hoc methods is much harder to verify, but the AIC still works in the same manner, thus showing the universality of that method.

### 2.1. One qubit, part 1

We return to tomography of single qubits, where we measure  $\sigma_x$  on the first  $N$  qubits, then  $\sigma_y$  on the next  $N$ , and  $\sigma_z$  on the last  $N$  qubits. Denote the three thusly observed averages by  $X := \langle \sigma_x \rangle_{\text{obs}}$ ,  $Y := \langle \sigma_y \rangle_{\text{obs}}$ , and  $Z := \langle \sigma_z \rangle_{\text{obs}}$ . In order to calculate likelihoods, we need the frequencies of having observed spin up (+) and down (−), respectively. They are given in terms of these averages by

$$f_{\pm}^{(x)} = \frac{1 \pm X}{2}, \quad (3a)$$

$$f_{\pm}^{(y)} = \frac{1 \pm Y}{2}, \quad (3b)$$

$$f_{\pm}^{(z)} = \frac{1 \pm Z}{2}. \quad (3c)$$

A density matrix describing just the first set of  $N$  measurements really uses or needs only one parameter,  $X$  (the other two parameters are, obviously, not at all determined by those data). And no matter what  $X$  is, there is always a perfect fit to the data. The logarithm of the (maximum) likelihood of such a density matrix is, therefore,

$$\ln \mathcal{L}^{(x)} = N f_+^{(x)} \ln(f_+^{(x)}) + N f_-^{(x)} \ln(f_-^{(x)}) = -NH \left( \frac{1+X}{2} \right), \quad (4)$$

with  $H(\cdot)$  the Shannon entropy. The same story holds for the next two sets of measurements, and so there is always a perfect fit to the data when we use the ‘alternative model’ with three density matrices, and that model needs three independent parameters. We conclude that the AIC assigns the following ranking to the alternative model:

$$\Omega_a = -N \left\{ H\left(\frac{1+X}{2}\right) + H\left(\frac{1+Y}{2}\right) + H\left(\frac{1+Z}{2}\right) \right\} - 3. \quad (5)$$

The performance of the ‘standard model’ depends on the value of just one number. If

$$R^2 := X^2 + Y^2 + Z^2 \leq 1, \quad (6)$$

there is a single maximum likelihood density matrix  $\bar{\rho}$  (with purity  $\text{Tr } \bar{\rho}^2 = (R^2 + 1)/2$ ) that describes the whole measurement perfectly, just as the alternative model does. The standard model also needs three parameters in this case, and the maximum likelihood is also the same as for the alternative model. So, in this case there is no real difference between the two models—we could pick  $\bar{\rho}_1 = \bar{\rho}_2 = \bar{\rho}_3 = \bar{\rho}$ —and we have  $\Omega_s = \Omega_a$ . There is no reason to reject the standard model when  $R \leq 1$ .

Now let us suppose that  $R > 1$ . We have then the choice between two descriptions:

1. *Alternative model.* We describe each of the three measurements by their own density matrix. The maximum likelihood estimates of those three states satisfy

$$\text{Tr } \bar{\rho}_1 \sigma_x = X, \quad (7a)$$

$$\text{Tr } \bar{\rho}_2 \sigma_y = Y, \quad (7b)$$

$$\text{Tr } \bar{\rho}_3 \sigma_z = Z. \quad (7c)$$

*Three* independent parameters are needed for this model. ( $\bar{\rho}_1, \bar{\rho}_2, \bar{\rho}_3$  are underdetermined, of course, but for the purpose of finding the maximum likelihood  $\mathcal{L}_a$  the information suffices.)

2. *Standard model.* We use one density matrix to describe all three measurements together. The maximum likelihood estimate of that state will be *pure*. There is no known method to compute it exactly, but a generally good approximation is given by

$$\text{Tr } \bar{\rho}_s \sigma_x = X/R, \quad (8a)$$

$$\text{Tr } \bar{\rho}_s \sigma_y = Y/R, \quad (8b)$$

$$\text{Tr } \bar{\rho}_s \sigma_z = Z/R, \quad (8c)$$

and this state’s likelihood is a strict (but generally pretty tight) lower bound on the maximum likelihood for the standard model. *Two* independent parameters are needed in this model<sup>8</sup>.

<sup>8</sup> One of the authors remains bothered by the decision to separate the standard model, effectively, into two distinct models that contain (i) all the mixed states ( $K = 3$ ) and (ii) the pure states ( $K = 2$ ), respectively. However, this protocol is specifically discussed and justified by Burnham and Anderson [7, section 6.9.6]. They express a similar concern, but also provide some preliminary justification for it. So, while further thought and research seems warranted, so does this choice.

The reason we end up with a *pure* maximum likelihood state in the standard model is that the single matrix fitting the data perfectly lies outside the set of physical states (it has a negative eigenvalue), and the closest physical state lies *on* the boundary [5]. In the case of qubits, this means a pure state. More precisely, if the unphysical best-fit matrix  $\tilde{\rho}$  is written in its diagonal form,  $\tilde{\rho} = \sum_{k=+,-} \lambda_k |\psi_k\rangle\langle\psi_k|$ , with  $\lambda_+ > 1$  and  $\lambda_- < 0$ , then the maximum likelihood estimate would be  $\bar{\rho}_s = |\psi_+\rangle\langle\psi_+|$ . The latter state has the properties (8), as can be easily verified by explicit calculation.

Thus, when  $R > 1$  the alternative model fits the data better but uses one more parameter than does the standard model. We can calculate the maximum likelihoods analytically in each of the two models, and thus obtain the relative AIC score of the two models:

$$\Omega_s - \Omega_a = 1 + N \sum_{M=X,Y,Z} \frac{1}{2} \ln \frac{1 - M^2/R^2}{1 - M^2} + \frac{M}{2} \ln \frac{(R+M)(1-M)}{(R-M)(1+M)}. \quad (9)$$

We accept the standard model as consistent iff  $\Omega_s \geq \Omega_a$ . This will happen only if  $R$  is sufficiently close to 1. If we expand  $R$  around 1, we can Taylor expand the right-hand side of (9) as

$$\Omega_s - \Omega_a \approx 1 - N \sum_{M=X,Y,Z} \frac{(R-1)^2 M^2}{2(1-M^2)}, \quad (10)$$

provided  $(R-1)^2 \ll (1-M^2)$  for  $M = X, Y, Z$ . That is, with this proviso, the standard model is consistent only when

$$(R-1) \leq \frac{C}{\sqrt{N}}, \quad (11)$$

with the constant  $C$  given by

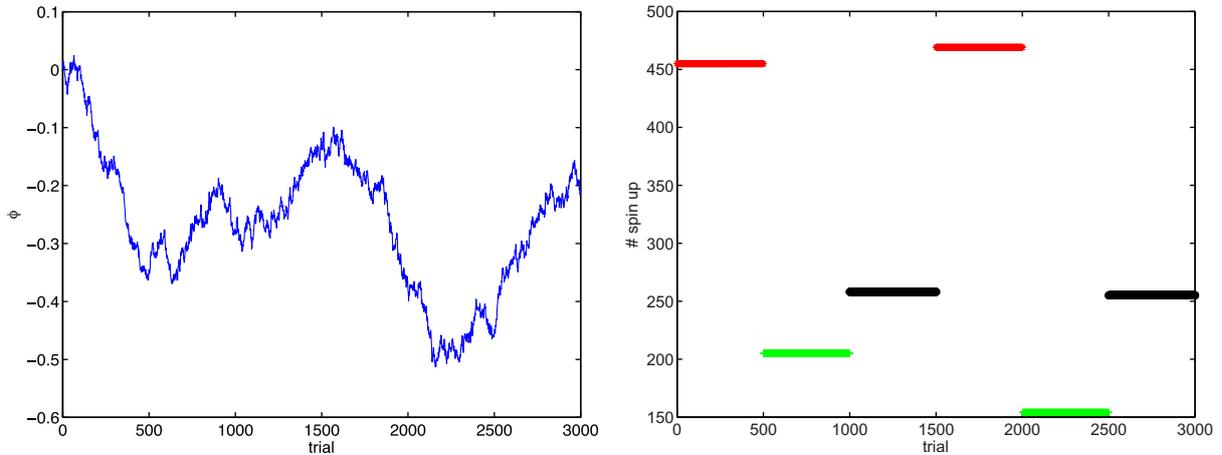
$$C = \frac{1}{\sqrt{\sum_M M^2/2(1-M^2)}}. \quad (12)$$

The dependence of the condition (11) on  $N$  agrees with the simple idea that it is sufficient for  $R$  to be less than about a standard deviation or two above 1 for the standard model to still apply, and that standard deviation, of course, decays like  $1/\sqrt{N}$  for  $N \rightarrow \infty$ .

## 2.2. One qubit, part 2

The implementation of tomography in the previous example is probably too simple and too obviously wrong for it to have been applied in an actual experiment. The straightforward improvement to measure each of  $\sigma_x, \sigma_y, \sigma_z$  in two separate blocks will allow one to detect drift. Let us denote the six observed averages by  $X_{1,2} := \langle\sigma_x\rangle_{\text{obs}1,2}$ ,  $Y_{1,2} := \langle\sigma_y\rangle_{\text{obs}1,2}$ , and  $Z_{1,2} := \langle\sigma_z\rangle_{\text{obs}1,2}$ . Drift can be detected by comparing the pairs of estimates  $X_{1,2}$  with each other,  $Y_{1,2}$  with each other, and  $Z_{1,2}$  with each other. The AIC works as follows: we need again at least two different models for describing the data. One will be the standard model, with one density matrix describing all six measurements. This density matrix will be determined by the three averages  $(X_1 + X_2)/2$ , etc. The alternative model may consist of two independent density matrices (with six parameters in total) or of two density matrices that are not independent with either four or five parameters in total. Let us test the AIC in a simulation of data generated by single-qubit states of the form

$$\rho_{\text{actual}} = p|\psi_\phi\rangle\langle\psi_\phi| + (1-p)\mathbb{1}/2, \quad (13)$$



**Figure 1.** Left: simulation of diffusion of the angle  $\phi$  in the state  $|\psi_\phi\rangle$  over the course of 3000 measurements. Right: the number of ‘spin up’ results for the measurements of  $\sigma_x$  (in red, for measurements 1–500 and 1501–2000), of  $\sigma_y$  (in green, for measurements 501–1000, and 2001–2500), and of  $\sigma_z$  (in black, for the remaining measurements). The numbers for  $\sigma_y$  are statistically different, and this is reflected in the relative ranking the AIC accords to the different models. Here we have  $\Omega_s - \Omega_a = -5.07$ , where the negative sign implies the standard model of a single density matrix is significantly worse than the alternative model (here, two density matrices with five parameters in total, two parameters more than the standard model). Tomography failed in this case.

where the pure state  $|\psi_\phi\rangle$  depends on an angle  $\phi$ , which we assume to undergo a random walk, and with the following meaning:

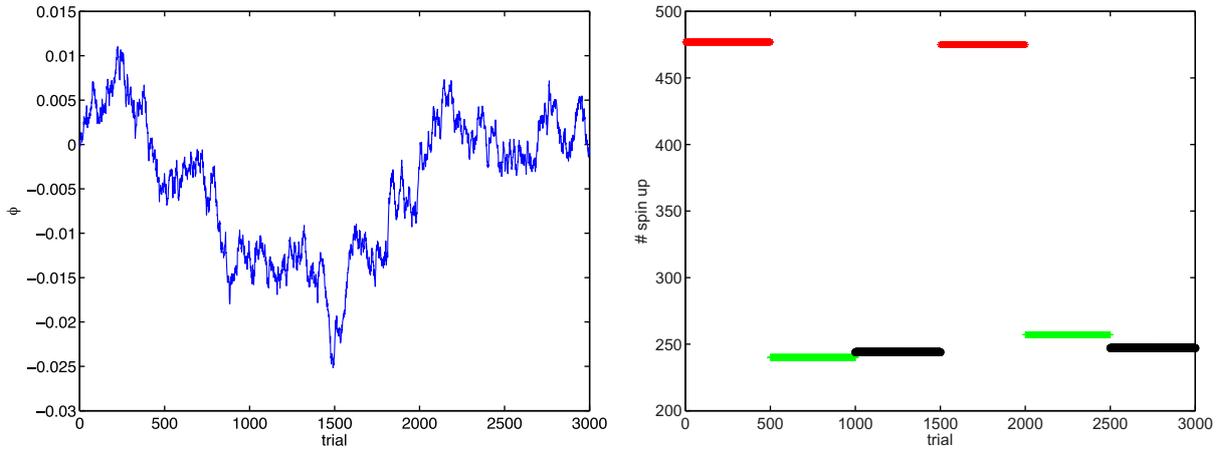
$$\begin{aligned}\langle\psi_\phi|\sigma_x|\psi_\phi\rangle &= \cos\phi, \\ \langle\psi_\phi|\sigma_y|\psi_\phi\rangle &= \sin\phi, \\ \langle\psi_\phi|\sigma_z|\psi_\phi\rangle &= 0.\end{aligned}\tag{14}$$

For  $p$  we take the value  $p = 0.9$ . We perform in total 3000 measurements, divided into six groups of 500, in which we measure  $\sigma_x, \sigma_y, \sigma_z, \sigma_x, \sigma_y, \sigma_z$  in that order.

In figures 1 and 2 we plot two qualitatively different cases. In the first case the diffusion of  $\phi$  is so fast that it leads to noticeably different values of  $X_{1,2}$  and  $Y_{1,2}$ . The AIC in this case gives a very clear preference for the alternative model of using two density matrices with five parameters in total (only the expectation value of  $\sigma_z$  does not change over the course of the experiment). In the second case the drift over the course of the experiment is small enough so that the standard model is still the best, even though there is some drift, and even though the more complicated model does, of course, fit the data slightly better.

### 2.3. Two or more qubits

Consider now a tomographically complete measurement on  $9N$  copies of two qubits, where on the first  $N$  pairs of qubits we measure  $\sigma_x$  on both qubits independently, then on the next  $N$  pairs



**Figure 2.** Same as figure 1, but for a case where the drift is much smaller over the course of 3000 measurements. Here  $\Omega_s - \Omega_a = 1.38$ , so that the standard model is better than the best alternative model (which has two extra parameters). Tomography succeeded.

we measure  $\sigma_x$  on one qubit and  $\sigma_y$  on the other, then on the third set of  $N$  pairs we measure  $\sigma_x$  on the one and  $\sigma_z$  on the other ... until on the last (ninth) set of  $N$  pairs we measure  $\sigma_z$  on both qubits. The first measurement is described by *three* independent averages that are obtained from measuring  $\sigma_x$  on both qubits independently:

$$XX := \langle \sigma_x \otimes \sigma_x \rangle_{\text{obs},1}, \quad (15a)$$

$$IX := \langle \mathbb{1} \otimes \sigma_x \rangle_{\text{obs},1}, \quad (15b)$$

$$XI := \langle \sigma_x \otimes \mathbb{1} \rangle_{\text{obs},1}. \quad (15c)$$

Thus, a two-qubit density matrix perfectly fitting the data of the first measurements needs three parameters. The description of the second measurement of  $\sigma_x$  on one qubit and  $\sigma_y$  on the other is likewise determined by three observed averages

$$XY := \langle \sigma_x \otimes \sigma_y \rangle_{\text{obs},2}, \quad (16a)$$

$$IY := \langle \mathbb{1} \otimes \sigma_y \rangle_{\text{obs},2}, \quad (16b)$$

$$XI' := \langle \sigma_x \otimes \mathbb{1} \rangle_{\text{obs},2}. \quad (16c)$$

The new feature arising here is that we get a second estimate of the same parameter,  $XI$  in this case. That is, if there were only a single two-qubit state in the experiment, the estimates (15c) and (16c) would have to agree (within error bars). Conversely, if they do not agree, we have encountered a new diagnosis of inconsistent tomography.

Writing down all different averages obtained from this particular experiment, we find nine quantities that are measured once, and six other quantities that are measured thrice. It becomes now much harder to judge when all the differences between those different estimates of the same quantities are, in total, statistically significant or not. That is, the generalization of the ad hoc method that worked fine for a single qubit, becomes troublesome. This, of course, becomes exponentially worse for more than two qubits.

On the other hand, the AIC can be applied straightforwardly to various alternative models. It is sufficient to find just one alternative model superior to the standard model in order to have succeeded in diagnosing an inconsistency in our tomographic experiment. Of course there is a large multitude of alternative models, but one can be guided in searching for such models by looking for those estimates of the same quantities that are the least consistent.

### 3. Model selection, the Akaike information criterion and quantum quirks

Data are generally assumed to be generated by some stochastic process<sup>9</sup>—e.g. a probability distribution  $f(x)$  (where  $x$  denotes the sample space containing all possible events). Unfortunately, these ‘true’ probabilities are unknown to us. All we have are some data. So, in order to (i) describe the data; (ii) approximate the underlying process  $f$ ; and (iii) most importantly, predict *future* observations, we use *models*.

A model is just another probability distribution  $g(x)$ . Almost always, the model contains a whole family of *parameterized* distributions  $g_\theta(x)$ , where  $\theta$  comprises the values of  $K$  distinct (real-valued) parameters. One obvious model is the universal one where each of the probabilities  $g(x)$ —for every possible value of  $x$ —is itself a free parameter. This is the richest possible model, with the most parameters. If  $x$  takes on uncountably many values, this model is utterly intractable (and the AIC penalizes it infinitely for its richness). The ubiquity of this problem in statistics motivates the use of restricted parameterized models (e.g. Gaussian distributions) where finitely many parameters can specify  $g(x)$  for every possible  $x$ .

Quantum tomography applications usually involve finitely many parameters, but few-parameter models are still important. This is partly because of the simplification obtained by eliminating many parameters (e.g. when a quantum state in  $2^N$  dimensions is approximated by a matrix product state with  $\text{poly}(N)$  parameters), but even more importantly because it guards against overfitting. This is precisely where well-designed model selection techniques come in, and the AIC is a canonical example. When there is a choice between different candidate models describing one and the same experiment, the AIC provides a numerical ranking of the different models.

The AIC (as given in equation (2)) appears very simple. Moreover, it bears a strong resemblance to quantities that appear in likelihood-ratio (LR) hypothesis testing (see, e.g., [14]). But in fact, the AIC’s theoretical underpinnings are rather different, and remarkably elegant (see [7] for extensive discussion). LRs are a fundamentally frequentist technique: given two competing models, we calculate ahead of time the probability that various values of the LR statistic will be observed *if* one model or the other is ‘correct’, and then we formulate a rule for what to announce upon seeing any given value of the LR statistic. Many canonical results on LR tests require that the models be nested—i.e. that one be a subset of the other. In particular, given this and a few other conditions, it is possible to derive expectation values of the LR statistic that look identical to equation (2) because the loglikelihood ratio is  $\chi_K^2$  distributed, and has mean value  $K$ .

<sup>9</sup> A semi-philosophical note: it is not necessary to invoke intrinsic randomness. The underlying process might be deterministic but so complex that its description is hopeless or just not worth the time. The Bayesian view of probabilities is entirely compatible with this view, and permits us to *describe* data and the processes that generate them without needing to ‘believe in’ randomness.

But despite this similarity, the AIC is derived differently. Akaike began by postulating that ‘goodness’ of a model is quantified by the Kullback–Leibler divergence [15] between the model and the ‘true model’ that actually underlies the data. Then, rather remarkably, he showed that it is possible to estimate this divergence<sup>10</sup> — *even when the true model is unknown!* The AIC is the expected value of the (unknown) Kullback–Leibler divergence between a specified model and the (unknown) true model, conditional upon the data in our possession. So the AIC (i) has a powerful and universal interpretation, and (ii) can be used to compare arbitrary models, without any requirement for nesting.

This is not to say that the AIC is the acme of model selection, nor that it is perfectly adapted to quantum tomography problems. First, there are competing derivations of other model ranking statistics, such as the Bayesian information criterion (again, see [7]). Moreover, the AIC is inherently an asymptotic result—much like, for example, the efficiency of MLE. So, even though there is a finite sample size correction (the AIC<sub>c</sub>), this correction is part of an asymptotic expansion and may be unreliable for any fixed  $N$ .

One significant consequence of this is that, for finite samples, an event  $x$  whose true probability is nonzero may not be observed—in which case a model might assign zero probability to it. (The MLE within the full model, where each probability  $g(x)$  is a parameter, behaves this way.) This results inevitably in an *infinite* Kullback–Leibler divergence. Asymptotically, the probability of such a pathology occurring goes to zero almost certainly. But for any finite sample size it is a concern. So, beware of rank-deficient estimates in tomography!

A related phenomenon is (almost) unique to quantum tomography. Akaike’s derivation assumes that a very good (if not the best) measure of predictive power is the Kullback–Leibler divergence between the true model  $f(x)$  for the observed process  $x$  and the assigned model  $g(x)$ . But in quantum tomography, the observed process ‘ $x$ ’ is some particular (and rather arbitrary) quorum of measurements that the tomographer has performed. We do not necessarily care about predicting *those* measurements! Instead, we care about the underlying quantum state—or, to put it more operationally, we care about a large and unknown set of *other* measurements that might be performed on samples of that state in the future. Quite frequently, we care about measurements of that state’s diagonal basis. This completely undermines Akaike’s assumption (that predicting  $x$  is the goal). This does not mean that the AIC should not be used—but it does strongly suggest that

1. conclusions drawn from the AIC, or any other classical statistical method, should be treated with thoughtful care,
2. better methods may still be derived (e.g. a ‘quantum AIC’),
3. estimates obtained via the AIC should *not* be expected to have good properties with respect to quantum relative entropy (the quantum version of Kullback–Leibler divergence).

Importantly, however, there are cases where our future measurements will be the same as those used for our preliminary quantum tomography experiment. For instance, in the case of quantum computing, where error correction is implemented by CSS codes, all measurements will be Pauli measurements. In such a case, the conclusions of the AIC, applied to a tomography experiment that used Pauli measurements as well, should be trustworthy.

<sup>10</sup> More precisely, this estimate is in fact determined only up to a constant, but that constant is the same for all models, and hence drops out when comparing different models.

#### 4. Summary and discussion

Our central message here is that when the assumptions of tomography fail, it is often due to some sort of sample-apparatus correlation, and that this can be detected with statistical reliability by model selection using the AIC. One particular example, the drifting source, clearly voids the single-density-matrix model, but can be described naturally (and more accurately!) by multiple density matrices associated with different times and/or measurement settings. The AIC is a particularly good and elegant tool for identifying whether the added complexity of this model is justified. Ultimately, the point of model selection (especially using the AIC) is to get better predictions of future measurement outcomes—*not* just better fits to observed data.

While the AIC ranks competing models, by assigning each model  $k$  a number  $\Omega_k$ , through equation (2), we have great flexibility in what to do with that ranking. Small differences in AIC are not significant; if  $|\Omega_k - \Omega_{k'}| \ll 1$ , then both models are equally good. But even when significant differences exist, we may choose to use the ‘best’ model exclusively, or to hedge by mixing it with lower-ranked models (with weights determined by their respective AICs). We could apply Bayesian methods to the highest-ranked model, or use maximum likelihood estimates to choose model parameters. Choosing between these alternatives is beyond the scope of this paper.

If a model-selection (e.g. AIC) analysis finds overwhelming evidence of sample-apparatus correlation (e.g. source drift), it is often possible to go beyond the conclusion ‘tomography has failed!’. What has really failed is the i.i.d. assumption—we have convincing evidence that the samples are not identically distributed. The joint state is therefore not (with high confidence) of De Finetti form (see [16]). But it may be possible to assign states with a relaxed De Finetti form, and thereafter to do tomography with this in mind. For example, if the AIC declares the alternative three-state model much superior to the single-state model, one could assign a state of the form

$$\rho^{(3N)} = \int d\bar{\rho}_1 \int d\bar{\rho}_2 \int d\bar{\rho}_3 P_a(\bar{\rho}_1, \bar{\rho}_2, \bar{\rho}_3) \bar{\rho}_1^{\otimes N} \otimes \bar{\rho}_2^{\otimes N} \otimes \bar{\rho}_3^{\otimes N} \quad (17)$$

to the  $3N$  qubits, where  $P_a(., ., .)$  is a joint probability distribution over three two-dimensional density matrices. This form itself needs to be tested and validated, by comparison to a richer model (e.g. a model with six, nine, or more different states). In general, validating a model requires more sophisticated model design—e.g. to describe more arbitrary forms of source drift—and perhaps different measurements or experiments specifically aimed at detecting those models, as proposed in [17]. But once a given model *is* validated, if it implies a relaxed De Finetti form as in equation (17), then we can in principle perform tomography independently on each of the i.i.d. subsets of the whole sample.

In the simplest case of tomography on single qubits, we discussed two competing models. Either one uses just a single density matrix  $\bar{\rho}$  to describe the experiment (the standard model), or one uses three— $\bar{\rho}_1, \bar{\rho}_2, \bar{\rho}_3$ —one for each set of  $N$  qubits used to measure  $\sigma_x$ ,  $\sigma_y$ , and  $\sigma_z$ , respectively (the alternative model). But what does it mean to use *three* density matrices for predicting future measurement outcomes? The answer is that the predictions refer to measurements on qubits that have not been measured yet (of course). Consider one unmeasured qubit taken from, say, a set of  $N + n$  qubits, from which  $N$  qubits were randomly picked to be measured in the  $\sigma_x$  basis and  $n$  were not measured. In this case, those  $n$  qubits would be assigned

a state of the form

$$\rho^{(n)} = \int d\bar{\rho}_1 \int d\bar{\rho}_2 \int d\bar{\rho}_3 P_a(\bar{\rho}_1, \bar{\rho}_2, \bar{\rho}_3) \bar{\rho}_1^{\otimes n}, \quad (19)$$

valid for any  $n$ , including  $n = 1$ . The mixed model, as mentioned above, would combine the standard and alternative models and assign an even more mixed state. For example, in the case  $n = 1$  it would assign the estimate

$$\rho_{\text{mixed}} = w_a \int d\bar{\rho}_1 \int d\bar{\rho}_2 \int d\bar{\rho}_3 P_a(\bar{\rho}_1, \bar{\rho}_2, \bar{\rho}_3) \bar{\rho}_1 + w_s \int d\bar{\rho} P_s(\bar{\rho}) \bar{\rho}, \quad (20)$$

with  $w_a = \exp(\Omega_a)/(\exp(\Omega_a) + \exp(\Omega_s))$  and  $w_s = 1 - w_a$  the relative weights of the two models, as assigned by the AIC, and with  $P_s(\cdot)$  the standard De Finetti probability distribution over single density matrices.

Although we have avoided discussion of model design here, one simple but powerful technique deserves mention. In the example at the beginning of the paper, we introduced an alternative model wherein each measurement setting is associated with a different density matrix. When the measurements are informationally complete, this alternative model has precisely as many parameters as the standard model. But if they are *overcomplete*, then the alternative model has more parameters. As long as the samples really are i.i.d., we expect the alternative model to fit slightly better, and the AIC to declare them (on average) equally good. However, in the presence of experimental drift, we will find inconsistencies *within* the overcomplete measurement set—i.e. we will *not* be able to fit all the measurements well with a single density matrix! This is a simple test for experimental drift that does not rely on negativity of  $\hat{\rho}_{\text{tomo}}$ .

For the main point of this paper, however, all these complications are unnecessary. All that matters is whether assigning a single density matrix to our tomography experiment constitutes the best model or not. If not, something is amiss, but at least we have diagnosed the problem.

The main issue we left open is the following: is there a sense in which the AIC works reliably if future measurements are different than those used in our tomography experiment? If not, is there a ‘quantum’ version of the AIC that, e.g. takes into account the quorum of observables that have been measured, as well as the set of observables that will be measured?

## Acknowledgments

This work was supported by NSF grant no. PHY-1004219. Sandia National Laboratories is a multi-program laboratory operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the US Department of Energy, National Nuclear Security Administration under contract no. DE-AC04-94AL85000.

*Note added in proof:* Upon completion of this paper [18] appeared, which is similar in spirit to our paper, but which uses  $\chi^2$  tests to detect errors in tomography. It points out, too, the problem with pure-state assignments for those tests. After submission of the proofs we also became aware of a very relevant article [19].

## References

- [1] Paris M and Reháček J 2004 *Quantum State Estimation* (Berlin: Springer)
- [2] Akaike H 1974 *IEEE Trans. Autom. Control* **19** 716

- [3] Renner R 2007 *Nature Phys.* **3** 645
- [4] Weinstein Y S, Havel T F, Emerson J, Boulant N, Saraceno M, Lloyd S and Cory D G 2004 *J. Chem. Phys.* **121** 6117
- [5] Blume-Kohout R 2010 *New J. Phys.* **12** 043034
- [6] James D F V, Kwiat P G, Munro W J and White A G 2001 *Phys. Rev. A* **64** 052312
- [7] Burnham K P and Anderson D R 2002 *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (Berlin: Springer)
- [8] Claeskens G and Hjort N L 2008 *Model Selection and Model Averaging* (Cambridge: Cambridge University Press)
- [9] Liddle A R 2007 *Mon. Not. R. Astron. Soc. Lett.* **377** L74
- [10] Lougovski P and van Enk S J 2009 *Phys. Rev. A* **80** 052324
- [11] Usami K, Nambu Y, Tsuda Y, Matsumoto K and Nakamura K 2003 *Phys. Rev. A* **68** 022314
- [12] Yin J O S and van Enk S J 2011 *Phys. Rev. A* **83** 062110
- [13] Guta M, Kypraios T and Dryden I 2012 arXiv:1206.4032
- [14] Blume-Kohout R, Yin J O S and van Enk S J 2010 *Phys. Rev. Lett.* **105** 170501
- [15] Kullback S and Leibler R A 1951 *Ann. Math. Stat.* **22** 79
- [16] Caves C M, Fuchs C A and Schack R 2002 *J. Math. Phys.* **43** 4537
- [17] Schwarz L and van Enk S J 2011 *Phys. Rev. Lett.* **106** 180501
- [18] Langford N 2012 arXiv:1212.2982
- [19] Rosset D, Ferretti-Schöbitz R, Bancal J-D, Gisin N and Liang Y-C 2012 *Phys. Rev. A* **86** 062325