To: *Chris Fuchs and Rüdiger Schack*
From: *C. M. Caves*
Subject: **Learning and the de Finetti representation**
2000 October 24

Classical probabilities

The classical de Finetti representation theorem states that *every* ∞-exchangeable probability assignment can be written *uniquely* as a mixture of iid's, i.e., can be generated from a unique probability on single-trial probabilities, or generating function. Formally this means that the map from generating functions to ∞-exchangeable probability assignments is *one-to-one* (uniqueness) and *onto* (every). Another formal way of stating this is that the convex set of ∞-exchangeable probability assignments is a simplex whose extreme points are the product distributions, or iid's. That the product distributions are the entire set of extreme points is the onto property of the map, and that the set is a simplex is the uniqueness of the expansion in terms of extreme points. It should be emphasized that neither of these properties holds for $(N < \infty)$-exchangeable probability assignments: the map from generating functions to N-exchangeable probability assignments is neither one-to-one nor onto.

So why do we care about the one-to-one and onto properties established by the de Finetti representation theorem? The importance of the one-to-one property is, I think, easy to identify. The subjectivist program is aimed at replacing probabilities on probabilities—and their interpretation in terms of a "man in a box"—with primary probability assignments. The one-to-one property is crucial for this program: if there were more than one probability on probabilities for an ∞-exchangeable probability assignment, then probabilities on probabilities would have distinguishing features not captured by the ∞-exchangeable probability assignment; an ∞-exchangeable probability assignment would not characterize fully the behavior of the man in the box.

This reasoning leaves unanswered the question of why it is important that *every* ∞-exchangeable probability assignment correspond to a probability on single-trial probabilities. Why should the subjectivist program care if there are ∞-exchangeable probabilities that have *no* interpretation in terms of a man in a box? The answer, I believe, lies in the need, even for a subjectivist, to maintain a connection between observed frequencies and single-trial probabilities. The subjective view undermines this connection, there being no necessary relation between frequencies and single-trial probabilities, since single-trial probabilities predict long-run frequencies only for an iid. On the other hand, one might still hope that using frequency data to update probability assignments for future trials could be thought of as using the observed frequencies to *learn* about single-trial probabilities, i.e., to learn about a possible objective "mechanism" embodied in the man in the box.

Exchangeability is clearly the condition under which frequency data is a *sufficient statistic*: frequency data, and *only* frequency data, is relevant for updating probability assignments for future trials. The question is whether the updating based on frequency data can be thought of in terms of updating a probability for single-trial probabilities, which can then be used to update the probability assignment for future trials. This kind of updating, which I call updating or learning *through* single-trial probabilities, is the subjectivist connection between frequencies and probabilities, i.e., the connection between observables and possible objective mechanisms. It is clear that when an exchangeable probability assignment, finite or infinite, is derived from a probability on probabilities, all the updating is of this sort. This is why the de Finetti representation theorem says something about learning: the onto property of the representation theorem says that for all ∞-exchangeable probability assignments, updating based on frequency data is equivalent to learning through single-trial probabilities.

More generally, consider a situation where Bayesian updating of probabilities proceeds through updating an underlying parameter (or sufficient statistic). I call this "learning through the parameter." The parameter can be thought of as describing an underlying mechanism—"man in a box"—that determines the probabilities. Whether such a mechanism actually "exists"—i.e., whether you are learning about an objective property—is irrelevant; the point is that this is the method scientists use for finding out about such a mechanism should it exist.

Call the parameter α, and let it be distributed according to a probability density $p(\alpha)$. The updating must proceed through the parameter no matter what $p(\alpha)$ is. Let D_1 and D_2 be two pieces of data, which depend on the parameter according to a joint probability $P(D_1, D_2|\alpha)$. The first piece of data can be used to update the probability for the second piece by using Bayes's theorem:

$$
\begin{aligned}
P(D_2|D_1) &= \frac{P(D_1, D_2)}{P(D_1)} \\
&= \frac{\displaystyle\int P(D_1, D_2|\alpha) p(\alpha)\, d\alpha}{P(D_1)} \\
&= \int P(D_2|D_1, \alpha) \frac{P(D_1|\alpha) p(\alpha)}{P(D_1)}\, d\alpha \\
&= \int P(D_2|D_1, \alpha) p(\alpha|D_1)\, d\alpha\;.
\end{aligned}
\tag{1}
$$

The condition that the Bayesian updating proceed through the parameter α is that

$$
P(D_2|D_1, \alpha) = P(D_2|\alpha) \quad\iff\quad P(D_1, D_2|\alpha) = P(D_1|\alpha) P(D_2|\alpha)\;. \tag{2}
$$

This means that given the parameter, the two pieces of data are statistically independent or, equivalently, that the two pieces of data are correlated only through the parameter. Under this condition, the Bayesian updating of Eq. (1) becomes

$$
P(D_2|D_1) = \int P(D_2|\alpha) p(\alpha|D_1)\, d\alpha\;. \tag{3}
$$

This is what we mean by updating probability assignments through the parameter α.

The learning condition (2) has a nice information-theoretic expression:

$$
H(D_1; D_2|\alpha) = 0\;. \tag{4}
$$

Given α, the mutual information between the two pieces of data is zero; in other words, all the mutual information between D_1 and D_2 flows through the parameter.

The notion of learning through a parameter is easily generalized to any number of pieces of data. What we have shown is that learning through a parameter is equivalent to the statement that the probability of the data is a mixture of product distributions:

$$
P(D_1, \ldots, D_N) = \int P(D_1|\alpha) \cdots P(D_N|\alpha) p(\alpha)\, d\alpha\;. \tag{5}
$$

The application of learning through a parameter to exchangeability is the following: if the parameter is faithful to exchangeability—i.e., if the parameter generates an N-exchangeable distribution for all choices of $p(\alpha)$—then the conditional probability for results $x_1, \ldots, x_L$ in L trials satisfies

$$
P(x_1, \ldots, x_L|\alpha) = P(x_1|\alpha) \cdots P(x_L|\alpha)\;, \tag{6}
$$

where the probabilities on the right are independent of trial, depending only on the alternative. Thus they make up a vector of single-trial probabilities, $\mathbf{p} = (p_1, \ldots, p_L)$. The parameter α is nothing but a label for this vector, so we can write

$$
P(x_1, \ldots, x_N|\mathbf{p}) = p_{x_1} \cdots p_{x_N} = p_1^{n_1} \cdots p_L^{n_L}\;. \tag{7}
$$

The result is that the ability to update probabilities through an exchangeable parameter is equivalent to having a probability on probabilities, i.e.,

$$
P(x_1, \ldots, x_N) = \int P(x_1, \ldots, x_N|\mathbf{p}) p(\mathbf{p})\, d\mathbf{p} = \int p_1^{n_1} \cdots p_L^{n_L} p(\mathbf{p})\, d\mathbf{p}\;. \tag{8}
$$

Notice that this result says nothing about the existence or the uniqueness of an expansion in terms of probabilities on single-trial probabilities.

We can now formulate another way of stating the de Finetti representation theorem: every ∞-exchangeable probability has a unique representation in terms of learning through single-trial probabilities. This establishes the subjective connection between frequencies, which constitute the data for updating, and single-trial probabilities.

It is instructive to spell out what is different for finite exchangeable sequences. The Heath-Sudderth proof shows that any finite exchangeable sequence can be generated by mixing probabilities for drawing from various urns without replacement. The parameter in this case labels the initial properties of the urns. Though this is a case of learning solely from frequency data, it is not a case of learning through the parameter, because updated probabilities after an initial set of trials depend both on the parameter and on the results of the initial trials, which determine how the urns have changed as a consequence of the initial trials.

Quantum mechanics

Now on to the quantum case. The quantum de Finetti representation theorem says that any density operator $\rho^{(N)}$ in an ∞-exchangeable density operator can be written uniquely as a mixture of repetition product density operators (rpdo's), i.e., product density operators of the form

$$\rho^{\otimes N} \equiv \underbrace{\rho \otimes \rho \cdots \rho}_{N \text{ times}} \; . \tag{9}$$

Formally this means that the map from mixtures of rpdo's, or probabilities on single-system density operators, to ∞-exchangeable density operators is one-to-one and onto. In the way we prove the quantum theorem, the one-to-one and onto properties follow directly from applying the classical theorem to the exchangeable probabilities generated by an informationally complete single-system POVM, followed by a relatively simple argument showing that the representation is restricted to positive single-system operators. The crucial property of the POVM is that the product POVM for many systems remains informationally complete. Another formal way of stating the theorem is that the convex set of ∞-exchangeable density operators is a simplex whose extreme points are the infinite rpdo's. We now face the same question we confronted in the classical case: the one-to-one property of the map is sufficient for the subjectivist program of replacing probabilities on single-system density operators by a primary density-operator assignment, so what's important about knowing that *every* ∞-exchangeable density operator corresponds to a probability on single-system density operators?

The answer again lies in considering what is meant by learning through a parameter. The difference is that in quantum mechanics we have a mathematical structure for describing measurement statistics (POVMs) and the updating of quantum states based on the results of measurements (operations). What we want to know is when this structure is consistent with learning through a parameter via Bayes's theorem.

For this purpose consider two systems, labeled 1 and 2, having a density operator

$$\rho = \int p(\alpha)\rho_\alpha \, d\alpha \; , \tag{10}$$

where ρ_α is the density operator conditioned on the parameter. Suppose a measurement on system 1, described by operations $\mathcal{A}_{D_1}$, yields result D_1. The probability for this result is

$$P(D_1) = \mathrm{tr}\big((\mathcal{A}_{D_1} \otimes \mathcal{I}_2)(\rho)\big) \; , \tag{11}$$

and the corresponding probability given α is

$$P(D_1|\alpha) = \mathrm{tr}\big((\mathcal{A}_{D_1} \otimes \mathcal{I}_2)(\rho_\alpha)\big) \; . \tag{12}$$

The state of system 2 after the measurement takes the form

$$\begin{aligned}
\rho'_{D_1} &= \frac{\mathrm{tr}_1\big((\mathcal{A}_{D_1}\otimes\mathcal{I}_2)(\rho)\big)}{P(D_1)} \\
&= \frac{\displaystyle\int \mathrm{tr}_1\big((\mathcal{A}_{D_1}\otimes\mathcal{I}_2)(\rho_\alpha)\big)p(\alpha)\,d\alpha}{P(D_1)} \\
&= \int \frac{\mathrm{tr}_1\big((\mathcal{A}_{D_1}\otimes\mathcal{I}_2)(\rho_\alpha)\big)}{\mathrm{tr}\big((\mathcal{A}_{D_1}\otimes\mathcal{I}_2)(\rho_\alpha)\big)}\frac{P(D_1|\alpha)p(\alpha)}{P(D_1)}\,d\alpha \\
&= \int \rho'_{D_1,\alpha}p(\alpha|D_1)\,d\alpha\ ,
\end{aligned} \tag{13}$$

where

$$\rho'_{D_1,\alpha} = \frac{\mathrm{tr}_1\big((\mathcal{A}_{D_1}\otimes\mathcal{I}_2)(\rho_\alpha)\big)}{\mathrm{tr}\big((\mathcal{A}_{D_1}\otimes\mathcal{I}_2)(\rho_\alpha)\big)} \tag{14}$$

is the state of system 2 after a measurement with result D_1, provided the initial state is ρ_α. The condition for updating through the parameter α is

$$\rho'_{D_1,\alpha} = \rho'_\alpha\ , \tag{15}$$

in which case Eq. (13) becomes

$$\rho'_{D_1} = \int \rho'_\alpha p(\alpha|D_1)\,d\alpha\ . \tag{16}$$

What we mean by updating a density operator by learning through a parameter is that Eq. (15)—and, hence, Eq. (16)—holds for all measurements on system 1.

If the states ρ_α are product states, i.e., $\rho_\alpha = \rho_{1,\alpha}\otimes\rho_{2,\alpha}$, it is obvious that the updating condition (15) is satisfied, with $\rho'_{D_1,\alpha} = \rho_{2,\alpha}$, no matter what measurement operations are applied to system 1. That the states ρ_α must be product states can be seen in the following way. Consider an informationally complete POVM for system 1, with POVM elements E_{μ_1}, and let the operations for system 1 be $\mathcal{A}_{\mu_1} = \sqrt{E_{\mu_1}}\odot\sqrt{E_{\mu_1}}$. With these choices, Eq. (14) becomes

$$\rho'_{\mu_1,\alpha} = \frac{\mathrm{tr}_1(\rho_\alpha E_{\mu_1}\otimes 1_2)}{\mathrm{tr}(\rho_\alpha E_{\mu_1}\otimes 1_2)}\ . \tag{17}$$

Letting F_{μ_2} be the POVM elements of an informationally complete POVM for system 2, we can multiply the numerator of Eq. (17) by $1_1\otimes F_{\mu_2}$ and take a trace over system 2, yielding

$$\mathrm{tr}(\rho_\alpha E_{\mu_1}\otimes F_{\mu_2}) = \mathrm{tr}(\rho_\alpha E_{\mu_1}\otimes 1_2)\mathrm{tr}(\rho'_{\mu_1,\alpha}1_1\otimes F_{\mu_2})\ . \tag{18}$$

When the learning condition (15) is satisfied, the quantity on the left becomes a product of a function of μ_1 and a function of μ_2:

$$\mathrm{tr}(\rho_\alpha E_{\mu_1}\otimes F_{\mu_2}) = \mathrm{tr}(\rho_\alpha E_{\mu_1}\otimes 1_2)\mathrm{tr}(\rho'_\alpha 1_1\otimes F_{\mu_2})\ . \tag{19}$$

Since $E_{\mu_1}\otimes F_{\mu_2}$ is an informationally complete POVM for the joint system, we can conclude that $\rho_\alpha = \rho_{1,\alpha}\otimes\rho_{2,\alpha}$ is a product state. What we have shown is that the quantum learning condition is satisfied if and only if ρ is separable.

The learning condition (15) has a nice interpretation: given the parameter, no measurement on system 1 provides any useful information for updating the state of system 2; in other words, all information relevant for updating runs through the parameter. It is satisfying that the learning condition can only be met by separable states, with the parameter labeling the product states in the mixture.

Applied to N systems, the learning condition is that

$$\rho^{(N)}_\alpha = \rho_{1,\alpha}\otimes\cdots\otimes\rho_{N,\alpha}\ . \tag{20}$$

If the parameter is faithful to exchangeability, the density operators in this product must be identical, and the parameter can be taken to be this single-system operator. We conclude that the only way to learn through a parameter that is faithful to exchangeability is if the parameter is equivalent to a single-system density operator, in which case we have

$$\rho^{(N)} = \int p(\rho)\rho^{\otimes N}\, d\rho\ . \tag{21}$$

This result is independent of the quantum de Finetti theorem, which establishes that any ∞-exchangeable density operator can be updated uniquely through single-system density operators.

Real vector spaces

In real vector spaces something different happens, but it turns out to have only fairly mild effects. In complex quantum mechanics both the de Finetti representation theorem and the quantum learning condition rely on the property that the tensor product of two informationally complete POVMs is an informationally complete POVM for the joint system. This property does not hold in real vector spaces. Though this does not prove that the theorem and the condition fail in real vector spaces, it certainly prejudices one in that direction. Though it might be thought that one is cast completely adrift on the sea of reality, it is easy to confirm the prejudice and to derive all the properties of exchangeable density operators in real vector spaces simply by complexifying the real vector space.

Consider now a D-dimensional real vector space. Density operators and POVM elements are positive, symmetric (i.e., their matrix representations are symmetric) operators in this space. The vector space of symmetric operators, which is $[D(D+1)/2]$-dimensional, is spanned by operators S_j, $j = 1,\ldots,D(D+1)/2$, whereas the vector space of antisymmetric operators, which is $[D(D-1)/2]$-dimensional, is spanned by operators A_j, $j = 1,\ldots,D(D-1)/2$. In a composite system made up of two D-dimensional systems, the space of symmetric operators is spanned by the operators $S_j \otimes S_k$ and $A_j \otimes A_k$, giving a dimension

$$\frac{D^2(D+1)^2}{4} + \frac{D^2(D-1)^2}{4} = \frac{D^2(D^2+1)}{2}\ , \tag{22}$$

which is clearly greater than the dimension, $D^2(D+1)^2/4$, of the product space of symmetric operators.

The $[D^2(D+1)^2/4]$-dimensional subspace spanned by the operators $S_j \otimes S_k$ is called the *symmetric-symmetric (SS) subspace*; SS operators are invariant under partial transposition. The $[D^2(D-1)^2/4]$-dimensional subspace spanned by the operators $A_j \otimes A_k$ is called the *antisymmetric-antisymmetric (AA) subspace*; AA operators change sign under partial transposition and thus have zero partial trace.

The analysis of learning through a parameter proceeds through Eq. (16) just as in standard quantum mechanics, with the condition for updating through a parameter given by Eq. (15). When one proceeds further, however, one finds that Eq. (19) only requires that the SS part of ρ_α be a product. This suggests that any state of the form

$$\rho_\alpha = \rho_{1,\alpha} \otimes \rho_{2,\alpha} + A_\alpha\ , \tag{23}$$

where A_α is an AA operator, ought to satisfy the learning condition. To show this, we need to know that a real quantum operation preserves symmetric and antisymmetric operators. Thus we can write

$$\rho'_{D_1,\alpha} = \frac{\mathrm{tr}_1\big((\mathcal{A}_{D_1} \otimes \mathcal{I}_2)(\rho_\alpha)\big)}{\mathrm{tr}\big((\mathcal{A}_{D_1} \otimes \mathcal{I}_2)(\rho_\alpha)\big)} = \rho_{2,\alpha}\frac{\mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_{1,\alpha})\big)}{\mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_{1,\alpha})\big)} = \rho_{2,\alpha}\ . \tag{24}$$

Since A_α has zero trace, the only constraint on it is that ρ_α be positive. Generalizing the learning condition to N systems is tedious at best. I don't know how to write the general density operator that satisfies the condition, so I don't go into it here. What I do below is to consider learning through a parameter in the case of ∞-exchangeable states.

The analogue of the de Finetti representation theorem for real vector spaces can be gotten directly from the corresponding theorem for the complexification. If $\rho^{(N)}$ is part of an ∞-exchangeable sequence of real density operators, then the complex de Finetti representation theorem says that it can be written uniquely as

$$\rho^{(N)} = \int p(\rho)\rho^{\otimes N}\, d\rho\ , \tag{25}$$

where the integral runs over all *complex* density operators. Since $\rho^{(N)}$ is real, we have that

$$\rho^{(N)} = (\rho^{(N)})^* = \int p(\rho)(\rho^*)^{\otimes N}\, d\rho\ , \tag{26}$$

which, from the uniqueness of the representation, implies that

$$p(\rho^*) = p(\rho)\ . \tag{27}$$

Define

$$\begin{array}{l} \rho_S \equiv \dfrac{1}{2}(\rho+\rho^*) \\ r_A \equiv -\dfrac{i}{2}(\rho-\rho^*) \end{array} \qquad \Longrightarrow \qquad \begin{array}{r} \rho = \rho_S + ir_A \\ \rho^* = \rho_S - ir_A \end{array}\ , \tag{28}$$

where ρ_S is a real (symmetric) density operator and r_A is a real, antisymmetric operator. Equation (27) is equivalent to

$$p(\rho) = p(\rho_S, r_A) = p(\rho_S, -r_A) = p(\rho^*)\ . \tag{29}$$

We can manipulate the de Finetti representation (25) in the following ways:

$$\begin{aligned} \rho^{(N)} &= \int p(\rho_S, r_A)(\rho_S + ir_A)^{\otimes N}\, d\rho_S\, dr_A \\ &= \int' p(\rho_S, r_A)(\rho_S + ir_A)^{\otimes N}\, d\rho_S\, dr_A + \int' p(\rho_S, -r_A)(\rho_S - ir_A)^{\otimes N}\, d\rho_S\, dr_A \\ &= \int' p'(\rho_S, r_A)\frac{1}{2}\Big((\rho_S + ir_A)^{\otimes N} + (\rho_S - ir_A)^{\otimes N}\Big)\, d\rho_S\, dr_A\ . \end{aligned} \tag{30}$$

Here the primed integrals run over half the domain of r_A for each ρ_S and $p'(\rho_S, r_A) = 2p(\rho_S, r_A) = 2p(\rho_S, -r_A)$.

What we have shown is that if $\rho^{(N)}$ is part of an ∞-exchangeable sequence of real density operators, then it has a unique representation of the form (30). A formal way of saying this is the following: whereas the convex set of ∞-exchangeable complex density operators is a simplex whose extreme points are the infinite rpdo's, the convex set of ∞-exchangeable real density operators is a simplex whose extreme points are states of the form

$$\frac{1}{2}\Big((\rho_S + ir_A)^{\otimes\infty} + (\rho_S - ir_A)^{\otimes\infty}\Big)\ . \tag{31}$$

If $r_A \neq 0$, these states are necessarily entangled, since they contain antisymmetric operators that cannot appear in a separable state.

Now let's ask about learning through the parameters ρ_S and r_A in the case of an ∞-exchangeable real density operator. For this purpose consider the N-system state

$$\frac{1}{2}\Big((\rho_S + ir_A)^{\otimes N} + (\rho_S - ir_A)^{\otimes N}\Big)\ . \tag{32}$$

Given result D_1 for quantum operation $\mathcal{A}_{D_1}$ applied to system 1, the state of the remaining systems is

$$\begin{aligned} &\frac{\mathrm{tr}_1\left(\mathcal{A}_{D_1}\otimes\mathcal{I}_{2,\ldots,N}\left(\frac{1}{2}\Big((\rho_S + ir_A)^{\otimes N} + (\rho_S - ir_A)^{\otimes N}\Big)\right)\right)}{\mathrm{tr}\left(\mathcal{A}_{D_1}\otimes\mathcal{I}_{2,\ldots,N}\left(\frac{1}{2}\Big((\rho_S + ir_A)^{\otimes N} + (\rho_S - ir_A)^{\otimes N}\Big)\right)\right)} \\ &\quad= \frac{\frac{1}{2}\Big((\rho_S + ir_A)^{\otimes(N-1)}\mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_S + ir_A)\big) + (\rho_S - ir_A)^{\otimes(N-1)}\mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_S - ir_A)\big)\Big)}{\frac{1}{2}\Big(\mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_S + ir_A)\big) + \mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_S - ir_A)\big)\Big)} \\ &\quad= \frac{1}{2}\Big((\rho_S + ir_A)^{\otimes(N-1)} + (\rho_S - ir_A)^{\otimes(N-1)}\Big)\frac{\mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_S)\big)}{\mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_S)\big)} \\ &\quad= \frac{1}{2}\Big((\rho_S + ir_A)^{\otimes(N-1)} + (\rho_S - ir_A)^{\otimes(N-1)}\Big)\ . \end{aligned} \tag{33}$$

Given the parameters ρ_S and r_A, the updated state does not depend on the data, so the ∞-exchangeable state (30) satisfies the condition for learning through the parameters. Notice, however, that since

$$p(D_1|\rho_S, r_A) = \mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_S + ir_A)\big) = \mathrm{tr}_1\big(\mathcal{A}_{D_1}(\rho_S)\big)\ , \tag{34}$$

the data provides no information about r_A. Two observers starting with different probabilities $p(\rho_S, r_A)$ and making single-system measurements will not to come to agreement on the value of r_A, since they get no information about its value, though they will come to agreement about the probabilities for future single-system measurements, since these probabilities are independent of r_A.

The inability to find out anything about r_A can be remedied by considering two-system measurements. Given result D_{12} for quantum operation $\mathcal{A}_{D_{12}}$ applied to systems 1 and 2, the state of the remaining systems is

$$\begin{aligned}
&\frac{\mathrm{tr}_{12}\left(\mathcal{A}_{D_{12}} \otimes \mathcal{I}_{3,\ldots,N}\left(\frac{1}{2}\Big((\rho_S + ir_A)^{\otimes N} + (\rho_S - ir_A)^{\otimes N}\Big)\right)\right)}{\mathrm{tr}\left(\mathcal{A}_{D_{12}} \otimes \mathcal{I}_{3,\ldots,N}\left(\frac{1}{2}\Big((\rho_S + ir_A)^{\otimes N} + (\rho_S - ir_A)^{\otimes N}\Big)\right)\right)} \\
&\quad = \frac{\frac{1}{2}\left((\rho_S + ir_A)^{\otimes(N-2)}\mathrm{tr}_{12}\Big(\mathcal{A}_{D_{12}}\big((\rho_S + ir_A)^{\otimes 2}\big)\Big) + (\rho_S - ir_A)^{\otimes(N-2)}\mathrm{tr}_{12}\Big(\mathcal{A}_{D_{12}}\big((\rho_S - ir_A)^{\otimes 2}\big)\Big)\right)}{\frac{1}{2}\left(\mathrm{tr}_{12}\Big(\mathcal{A}_{D_{12}}\big((\rho_S + ir_A)^{\otimes 2}\big)\Big) + \mathrm{tr}_{12}\Big(\mathcal{A}_{D_{12}}\big((\rho_S - ir_A)^{\otimes 2}\big)\Big)\right)} \\
&\quad = \frac{1}{2}\left((\rho_S + ir_A)^{\otimes(N-2)} + (\rho_S - ir_A)^{\otimes(N-2)}\right)\frac{\mathrm{tr}_{12}\Big(\mathcal{A}_{D_{12}}\big((\rho_S + ir_A)^{\otimes 2}\big)\Big)}{\mathrm{tr}_{12}\Big(\mathcal{A}_{D_{12}}\big((\rho_S + ir_A)^{\otimes 2}\big)\Big)} \\
&\quad = \frac{1}{2}\left((\rho_S + ir_A)^{\otimes(N-1)} + (\rho_S - ir_A)^{\otimes(N-1)}\right)\ .
\end{aligned} \tag{35}$$

Given the parameters ρ_S and r_A, the updated state does not depend on the data, so the ∞-exchangeable state (30) satisfies the condition for learning through the parameters for two-system measurements. Unlike single-system measurents, two-system measurements do provide information about r_A:

$$p(D_{12}|\rho_S, r_A) = \mathrm{tr}_{12}\Big(\mathcal{A}_{D_{12}}\big((\rho_S + ir_A)^{\otimes 2}\big)\Big) = \mathrm{tr}_{12}\Big(\mathcal{A}_{D_{12}}\big((\rho_S - ir_A)^{\otimes 2}\big)\Big)\ . \tag{36}$$

As a consequence, coming to agreement about predictions for two-system measurements will involve coming to agreement about both ρ_S and r_A.

A real quantum operation $\mathcal{A}$ consists of sums of terms of the form $A \odot A^T$. We have $(AOA^T)^T = AO^TA^T$, so $\big(\mathcal{A}(O)\big)^T = \mathcal{A}(O^T)$, so a symmetric operator stays symmetric, and an antisymmetric operator stays antisymmetric. Moreover, we can also conclude that

$$\mathrm{tr}\big(\mathcal{A}(O)\big) = \mathrm{tr}\Big(\big(\mathcal{A}(O)\big)^T\Big) = \mathrm{tr}\big(\mathcal{A}(O^T)\big)\ . \tag{37}$$