Fully quantum Bayes' rule from the minimum change principle

Francesco Buscemi, Nagoya University

Shenzhen-Nagoya Workshop on Quantum Science 2025, 25 September 2025

Abstract

Bayes' rule, which is routinely used to update beliefs based on new evidence, can be derived from a principle of minimum change. This principle states that updated beliefs must be consistent with new data, while deviating minimally from the prior belief. Here, we introduce a quantum analog of the minimum change principle and use it to derive a quantum Bayes' rule by minimizing the change between two quantum input-output processes, not just their marginals. This is analogous to the classical case, where Bayes' rule is obtained by minimizing several distances between the joint input-output distributions. When the change maximizes the fidelity, the quantum minimum change principle has a unique solution, and the resulting quantum Bayes' rule recovers the Petz transpose map in many cases. This is work done in collaboration with Ge Bai and Valerio Scarani.

Fully quantum Bayes' rule from the minimum change principle

Francesco Buscemi, Nagoya University

Shenzhen-Nagoya Workshop on Quantum Science 2025
25 September 2025

works this talk is based upon

- F.B., D. Fujiwara, N. Mitsui, and M. Rotondo, *Thermodynamic reverse bounds for general open quantum processes*. Physical Review A, vol. 102, 032210 (2020).
- F.B. and V. Scarani, *Fluctuation theorems from Bayesian retrodiction*. Physical Review E, vol. 103, 052111 (2021).
- C.C. Aw, F.B., and V. Scarani, *Fluctuation theorems with retrodiction rather than reverse processes*. AVS Quantum Science, vol. 3, 045601 (2021).
- F.B., J. Schindler, and D. Šafránek, *Observational entropy, coarse-grained states, and the Petz recovery map: information-theoretic properties and bounds*. New Journal of Physics, vol. 25, 053002 (2023).
- G. Bai, D. Šafránek, J. Schindler, F.B., and V. Scarani, *Observational entropy with general quantum priors*. Quantum, vol. 8, 1524 (2024).
- G. Bai, F.B., and V. Scarani, *Quantum Bayes' rule and Petz transpose map from the minimum change principle*. Physical Review Letters 135, 090203 (2025).
- G. Bai, F.B., and V. Scarani, *Fully quantum stochastic entropy production*. Preprint arXiv:2412.12489 (2024).
- T. Nagasawa, K. Kato, E. Wakakuwa, and F.B., *Macroscopic states and operations: a generalized resource theory of coherence*. Preprint arXiv:2504.12738 (2025).

What is Bayesian inference?

And what is quantum Bayesian inference?

what this talk is not about

philosophical debates (e.g., Bayesianism VS Frequentism, interpretations of QM such as QBism, etc.)

we are postmodern Bayesians!

what this talk is about

$$\underbrace{P(H|D_o)}_{\text{posterior}} = \frac{\overbrace{P(H)}^{\text{prior}}\,\overbrace{P(D_o|H)}^{\text{likelihood}}}{\underbrace{P(D_o)}_{\text{prop. constant}}}$$

This talk is about Bayes' Rule and its "unreasonable pervasiveness" throughout science

The **wrong answer**: it is general because it is a trivial consequence of the law of total probability, detailed balance, etc.

The **correct question**: why Bayes' Rule provides a good update rule?

possible justifications of the pervasiveness of Bayes' Rule

"Consistency" arguments by De Finetti, (Harold) Jeffreys, Savage, and Cox.

(Richard) Jeffrey's "probability kinematics" and Pearl's "virtual evidence method".

parenthesis: the Bayes–Jeffrey–Pearl update

- consider a classical discrete noisy channel $P(i|x)$ and a prior $\gamma(x)$ on the input
- when the receiver observes a definite value i_o , (vanilla) Bayes' Rule says that their posterior should be updated to $R_P^\gamma(x|i_o) := \dfrac{\gamma(x)P(i_o|x)}{[P\gamma](i_o)}$
- but what if the observation is noisy and returns some p.d. $\sigma(i)$ instead?

Theorem (Jeffrey 1965, Pearl 1988)

Given a channel $P(i|x)$ and a prior $\gamma(x)$, the result of a noisy observation $\sigma(i)$ is updated to

$$\widetilde{\sigma}(x) := \sum_i \boxed{R_P^\gamma(x|i)}\ \sigma(i) \ .$$

Note: the usual Bayes' Rule is recovered for $\sigma(i) = \delta_{i,i_o}$.

The problem with these derivations is that they are based on axioms, which may appear compelling to some but less so to others.

Alternative approach: can Bayes' Rule be derived as the (optimal, unique) solution to a concrete task?

variational principles are nice

To avoid unwarranted bias and remain maximally non-committal, the updated belief should be consistent with the new information (the result of the observation), while deviating as little as possible from the initial belief.

formalization: the principle of minimum change

- given channel $P(i|x)$ and prior $\gamma(x)$, construct the forward process $[P \star \gamma](i,x) := P(i|x)\gamma(x)$
- given the new information as $\sigma(i)$, consider the program

$$\min_R \mathbb{D}(P \star \gamma, R \star \sigma) \ ,$$

where $\mathbb{D}(\bullet,\bullet)$ is a suitable information divergence, and the minimum is taken over all channels $R \equiv R(x|i)$

Then, for many reasonable choices of $\mathbb{D}$ (e.g., the KL-divergence), it turns out that

$$\arg\min_R \mathbb{D}(P \star \gamma, R \star \sigma) = R_P^\gamma \ ,$$

where $R_P^\gamma \equiv R_P^\gamma(x|i) = \frac{[P\star\gamma](i,x)}{[P\gamma](i)}$ is Bayes' inverse.

towards a quantum generalization

The minimum change principle is formulated using the *joint* input-output distributions.

Hence, the central idea is that the “change” to be minimized is the change relative to the *whole input-output stochastic process*, not just its marginals.

But this is a **problem** in the quantum case...

quantum joint i/o distributions

Given a channel $\mathcal{E} : A \to B$ define:

- the Choi operator: $C_{\mathcal{E}} := \sum_{i,j} \mathcal{E}(|i\rangle\langle j|)_B \otimes |i\rangle\langle j|_A$
- the joint i/o state: $\mathcal{E} \star \gamma := (\mathbb{1}_B \otimes \sqrt{\gamma_A^T}) C_{\mathcal{E}} (\mathbb{1}_B \otimes \sqrt{\gamma_A^T})$

Note that:

- $\mathrm{Tr}_B[\mathcal{E} \star \gamma] = \gamma_A^T$
- $\mathrm{Tr}_A[\mathcal{E} \star \gamma] = \mathcal{E}(\gamma)_B$
- when all operators are diagonal, we obtain the classical joint i/o probability distribution

the principle of minimum change: the quantum case

Given a channel $\mathcal{E} : A \to B$, a prior state γ_A some "new information" σ_B, consider the program

$$\min_{\mathcal{R}} \mathbb{D}(\mathcal{E} \star \gamma, (\mathcal{R} \star \sigma)^T) ,$$

where the minimum is taken over channels $\mathcal{R} : B \to A$.

Theorem (arXiv:2410.00319, PRL to appear)

For $\mathcal{E} \star \gamma > 0$ and $\sigma > 0$, when the divergence is chosen to be the quantum fidelity $\mathbb{F}(x, y) := \left\|\sqrt{x}\sqrt{y}\right\|_1$,

$$\arg\max_{\mathcal{R}} \mathbb{F}(\mathcal{E} \star \gamma, (\mathcal{R} \star \sigma)^T) = \mathcal{R}_{\mathcal{E},\gamma,\sigma} ,$$

where

$$\mathcal{R}_{\mathcal{E},\gamma,\sigma}(\bullet) := \sqrt{\gamma}\, \mathcal{E}^\dagger \left(\sqrt{\sigma} \frac{1}{\sqrt{\sqrt{\sigma}\mathcal{E}(\gamma)\sqrt{\sigma}}} (\bullet) \frac{1}{\sqrt{\sqrt{\sigma}\mathcal{E}(\gamma)\sqrt{\sigma}}} \sqrt{\sigma} \right) \sqrt{\gamma} .$$

some open questions

In general, when $[\mathcal{E}(\gamma), \sigma] \neq 0$ the dependence of the posterior $\mathcal{R}_{\mathcal{E},\gamma,\sigma}(\sigma)$ on the new data σ is not linear: bug or feature?

What happens when other divergences are used instead of the fidelity?

What about multi-partite situations, locality restrictions, ...?

“Killer app”: quantum state inference (retrodiction) from measurements outcomes

motivation: von Neumann’s “other” entropy

von Neumann recognized that “his entropy” was not a good measure for thermodynamic entropy, which should instead be a quantity *relative to the observer’s knowledge*.

Modern version: observational entropy (OE)

For a density matrix ϱ and a positive operator-valued measure (POVM) $\mathbf{P} = \{P_i\}_i$

$$S_{\mathbf{P}}(\varrho) := -\sum_i p(i) \log \frac{p(i)}{V(i)} \ ,$$

where $p(i) := \mathrm{Tr}[\varrho \, P_i]$ and $V(i) := \mathrm{Tr}[P_i]$.

What is the meaning of OE?

the fundamental bound

Umegaki relative entropy

For density matrices $\varrho \geqslant 0$ and $\gamma > 0$, the Umegaki relative entropy $D(\varrho\|\gamma)$ is defined as $\mathrm{Tr}[\varrho(\log\varrho - \log\gamma)]$. We can thus write

$$S_{\mathbf{P}}(\varrho) = \log d - D(\mathcal{P}(\varrho)\|\mathcal{P}(u)) \ ,$$

where $\mathcal{P}(\bullet) := \sum_i \mathrm{Tr}[\bullet\, P_i]\, |i\rangle\langle i|$, and $u := d^{-1}\mathbb{1}$.

Theorem (NJP, 2023)

For any d-dimensional density matrix ϱ and any POVM $\mathbf{P} = \{P_i\}_i$,

$$S(\tilde{\varrho}_{\mathbf{P}}) - S(\varrho) \geqslant \underbrace{S_{\mathbf{P}}(\varrho) - S(\varrho)}_{D(\varrho\|u) - D(\mathcal{P}(\varrho)\|\mathcal{P}(u))} \geqslant D(\varrho\|\tilde{\varrho}_{\mathbf{P}}) \ ,$$

where $\tilde{\varrho}_{\mathbf{P}} := \sum_i \mathrm{Tr}[\varrho\, P_i]\, \frac{P_i}{V_i}$.

In particular, $\log d \geqslant S(\tilde{\varrho}_{\mathbf{P}}) \geqslant S_{\mathbf{P}}(\varrho) \geqslant S(\varrho)$.

OE tells us something about how much ϱ and $\tilde{\varrho}_{\mathbf{P}}$ "differ" from each other.

Hence, the question: what is the meaning of

$$\tilde{\varrho}_{\mathbf{P}} := \sum_i \mathrm{Tr}[\varrho\, P_i]\, \frac{P_i}{\mathrm{Tr}[P_i]} \qquad \textbf{???}$$

Petz's transpose/recovery map

Definition

Given a channel $\mathcal{E}$ and a prior state γ, the corresponding *Petz's transpose* (or *Petz's recovery*) channel is defined as

$$\mathcal{R}^{\gamma}_{\mathcal{E}}(\bullet) := \sqrt{\gamma}\, \mathcal{E}^{\dagger}\Big[\, \mathcal{E}(\gamma)^{-1/2}\, (\bullet)\, \mathcal{E}(\gamma)^{-1/2} \,\Big] \sqrt{\gamma}\ .$$

Fact: $\tilde{\varrho}_{\mathbf{P}}$ is the state "recovered" from the measurement's outcome

In terms of the measurement channel $\mathcal{P}(\bullet) := \sum_i \mathrm{Tr}[P_i\, \bullet]\, |i\rangle\langle i|$, it turns out that

$$\tilde{\varrho}_{\mathbf{P}} = [\mathcal{R}^{u}_{\mathcal{P}} \circ \mathcal{P}](\varrho) = \frac{1}{d}\, \mathcal{P}^{\dagger}\Big[\, \mathcal{P}(u)^{-1/2}\, \mathcal{P}(\varrho)\, \mathcal{P}(u)^{-1/2} \,\Big]\ .$$

(Note that in this case $\gamma = u = d^{-1}\mathbb{1}$.)

So, the real question is: what is the meaning of Petz's transpose map?

Petz's transpose map and quantum Bayes rule

Recall the form of the minimum change channel:

$$\mathcal{R}_{\mathcal{E},\gamma,\sigma}(\bullet) \coloneqq \sqrt{\gamma}\, \mathcal{E}^{\dagger}\left(\sqrt{\sigma}\frac{1}{\sqrt{\sqrt{\sigma}\mathcal{E}(\gamma)\sqrt{\sigma}}}\,(\bullet)\,\frac{1}{\sqrt{\sqrt{\sigma}\mathcal{E}(\gamma)\sqrt{\sigma}}}\sqrt{\sigma}\right)\sqrt{\gamma}\,.$$

Fact

Whenever $[\mathcal{E}(\gamma), \sigma] = 0$, then $\mathcal{R}_{\mathcal{E},\gamma,\sigma} \equiv \mathcal{R}_{\mathcal{E},\gamma}$, i.e., Petz's transpose map coincides with the minimum change channel.

Since the measurement channel $\mathcal{P}$ is quantum-to-classical, $\tilde{\varrho}_{\mathbf{P}}$ is the unique retrodicted state that satisfies the minimum change principle.

Hence, the entropy difference (or *observational deficit*) $S_{\mathbf{P}}(\varrho) - S(\varrho)$ measures the amount of "irretrodictable information".

macroscopic = fully retrodictable

Definition (macroscopic states)

Recalling the fundamental bound $S_{\mathbf{P}}(\varrho) - S(\varrho) \geqslant D(\varrho \| \tilde{\varrho}_{\mathbf{P}})$ with $\tilde{\varrho}_{\mathbf{P}} = [\mathcal{R}^u_{\mathcal{P}} \circ \mathcal{P}](\varrho)$, we say that a state ϱ is macroscopic w.r.t. measurement **P** and uniform prior u whenever $\varrho = \tilde{\varrho}_{\mathbf{P}}$.

More generally, for non-uniform prior γ, we denote the set of macroscopic states as $\mathfrak{M}^\gamma_{\mathbf{P}} := \{\varrho : \varrho = [\mathcal{R}^\gamma_{\mathcal{P}} \circ \mathcal{P}](\varrho)\}$.

Theorem (arXiv:2504.12738)

A state ϱ is in $\mathfrak{M}^\gamma_{\mathbf{P}}$ if and only if there exists a PVM $\mathbf{\Pi} = \{\Pi_j\}_j$, with $\Pi_j = \sum_i \mu(j|i) P_i$, such that $[\Pi_i, \gamma] = 0$, together with coefficients $c_j \geqslant 0$, such that $\varrho = \sum_j c_j \Pi_j \gamma$.

Note that $\gamma \in \mathfrak{M}^\gamma_{\mathbf{P}}$ by construction.

Which evolutions do not make retrodiction harder?

In other words: which evolutions do not increase “irretrodictability”, i.e., “microscopicity”?

macroscopic operations (idea)

Resource destroying map (RDM)

Recalling the form of macroscopic states $\varrho = \sum_j c_j \Pi_j \gamma$, the map

$$\Delta^\gamma_{\mathbf{P}}(\bullet) := \sum_j \mathrm{Tr}[\Pi_j \, \bullet] \frac{\Pi_j \gamma}{\mathrm{Tr}[\Pi_j \gamma]}$$

is such that $\Delta^\gamma_{\mathbf{P}}(\sigma) \in \mathfrak{M}^\gamma_{\mathbf{P}}$ for all σ, while $\varrho \in \mathfrak{M}^\gamma_{\mathbf{P}} \implies \Delta^\gamma_{\mathbf{P}}(\varrho) = \varrho$.

RDM-covariant operations

A CPTP linear map $\mathcal{N}$ is macroscopic (RDM-covariant) whenever

$$\mathcal{N} \circ \Delta^\gamma_{\mathbf{P}} = \Delta^\gamma_{\mathbf{P}} \circ \mathcal{N} \ .$$

The above framework contains the case of coherence, i.e., $\mathfrak{M}^\gamma_{\mathbf{P}} = \{\text{diagonal states}\}$, or athermality, i.e., $\mathfrak{M}^\gamma_{\mathbf{P}} = \{\gamma\}$.

resolving the paradox of thermodynamic entropy increase in closed (Hamiltonian) systems

- let the initial state of the system at time $t = t_0$ be a macrostate $\mathfrak{M}^u_{\mathbf{P}} \ni \varrho^{t_0} \neq u$
- its initial OE is as small as possible, i.e., it satisfies $S_{\mathbf{P}}(\varrho^{t_0}) = S(\varrho^{t_0})$; let's see how it changes in time
- the system evolves unitarily, so that $S(\varrho^{t_1}) = S(U\varrho^{t_0}U^\dagger) = S(\varrho^{t_0})$; however,

$$\begin{aligned} S_{\mathbf{P}}(\varrho^{t_1}) &= -\sum_i \mathrm{Tr}\Big[P_i\,(U\varrho^{t_0}U^\dagger)\Big] \log \frac{\mathrm{Tr}\big[P_i\,(U\varrho^{t_0}U^\dagger)\big]}{\mathrm{Tr}[P_i]} \\ &= -\sum_i \mathrm{Tr}\Big[(U^\dagger P_i U)\,\varrho^{t_0}\Big] \log \frac{\mathrm{Tr}\big[(U^\dagger P_i U)\,\varrho^{t_0}\big]}{\mathrm{Tr}[U^\dagger P_i U]} \\ &= S_{U^\dagger \mathbf{P} U}(\varrho^{t_0}) \\ &\geqslant S(\varrho^{t_0}) = S_{\mathbf{P}}(\varrho^{t_0}) \end{aligned}$$

- summarizing: in general, $S_{\mathbf{P}}(\varrho^{t_1}) \geqslant S_{\mathbf{P}}(\varrho^{t_0})$ even in closed systems, with equality if and only if $U\varrho^{t_0}U^\dagger \in \mathfrak{M}^u_{\mathbf{P}}$

an "H-theorem" for OE

Theorem (PRR, 2025)

In a d-dimensional system, choose a state ϱ and a POVM $\mathbf{P} = \{P_i\}_i$ with a finite number of outcomes. Choose also a (small) value $\delta > 0$. For a unitary operator U sampled at random according to the Haar distribution, it holds:

$$\mathbb{P}_H \left\{ \frac{S_{\mathbf{P}}(U \varrho U^\dagger)}{\log d} \leqslant (1-\delta) \right\} \leqslant \frac{4}{\kappa(\mathbf{P})} e^{-C\delta\kappa(\mathbf{P})^2 d \log d} ,$$

where $\kappa(\mathbf{P}) = \min_i \mathrm{Tr}[P_i\, u]$ and $C \approx 0.0018$.

Remark. A similar statement holds for unitaries sampled from an approximate 2-design.

$\Longrightarrow$ in the eyes of the observer, the state of a randomly evolving system **quickly** becomes **indistinguishable** from the maximally uniform one, regardless of the system's initial state.

entropy increase = lack of retrodictability (Watanabe's thesis)

"The phenomenological onewayness of temporal developments in physics is due to irretrodictability, and not due to irreversibility."

Satosi Watanabe (1965)

Conclusions

take-home messages

1. we have derived a quantum analogue of Bayes rule, based on the principle of minimum change
2. the quantum Bayes rule coincides with Petz's transpose map for channels with commutative output
3. this gives an operational meaning to the retrodicted quantum state inferred from a measurement's outcome
4. the difference between observational entropy and von Neumann entropy, i.e., the observational deficit, quantifies the amount of irretrodictable information
5. the second law is a statement about the generic loss of retrodictability in time
6. irretrodictability (i.e., "microscopicity") can be framed as a resource theory, generalizing those of coherence and athermality

The End: Thank You!

References

1. F.B., D. Fujiwara, N. Mitsui, and M. Rotondo, *Thermodynamic reverse bounds for general open quantum processes.* Physical Review A, vol. 102, 032210 (2020).

2. F.B. and V. Scarani, *Fluctuation theorems from Bayesian retrodiction.* Physical Review E, vol. 103, 052111 (2021).

3. C.C. Aw, F.B., and V. Scarani, *Fluctuation theorems with retrodiction rather than reverse processes.* AVS Quantum Science, vol. 3, 045601 (2021).

4. F.B., J. Schindler, and D. Šafránek, *Observational entropy, coarse-grained states, and the Petz recovery map: information-theoretic properties and bounds.* New Journal of Physics, vol. 25, 053002 (2023).

5. G. Bai, D. Šafránek, J. Schindler, F.B., and V. Scarani, *Observational entropy with general quantum priors.* Quantum, vol. 8, 1524 (2024).

6. G. Bai, F.B., and V. Scarani, *Quantum Bayes' rule and Petz transpose map from the minimum change principle.* Physical Review Letters 135, 090203 (2025).

7. G. Bai, F.B., and V. Scarani, *Fully quantum stochastic entropy production.* Preprint arXiv:2412.12489 (2024).

8. T. Nagasawa, K. Kato, E. Wakakuwa, and F.B., *Macroscopic states and operations: a generalized resource theory of coherence.* Preprint arXiv:2504.12738 (2025).