

Quantum Stein's Lemma

Formal statement

Let ρ and σ be density operators on a finite-dimensional Hilbert space \mathcal{H} . We consider the asymmetric hypothesis-testing problem between

$$H_0 : \text{the state is } \rho$$

and

$$H_1 : \text{the state is } \sigma.$$

For n independent copies, the two possible states are

$$\rho_n = \rho^{\otimes n}, \quad \sigma_n = \sigma^{\otimes n}.$$

A binary test is an operator $0 \leq T_n \leq I$. We interpret T_n as the decision "accept H_0 ," meaning "decide that the state was $\rho^{\otimes n}$." The type-I and type-II errors are

$$\alpha_n(T_n) = \text{Tr}[(I - T_n)\rho^{\otimes n}]$$

and

$$\beta_n(T_n) = \text{Tr}[T_n\sigma^{\otimes n}].$$

The type-I error is the probability of rejecting $\rho^{\otimes n}$ when $\rho^{\otimes n}$ was true. The type-II error is the probability of accepting $\rho^{\otimes n}$ when $\sigma^{\otimes n}$ was true.

For $0 < \varepsilon < 1$, define the optimal type-II error under type-I constraint ε by

$$\beta_{n,\varepsilon}(\rho\|\sigma) = \min_{0 \leq T_n \leq I} \{ \text{Tr}[T_n\sigma^{\otimes n}] : \text{Tr}[T_n\rho^{\otimes n}] \geq 1 - \varepsilon \}.$$

The quantum relative entropy is

$$D(\rho\|\sigma) = \text{Tr } \rho(\log \rho - \log \sigma),$$

provided

$$\text{supp}(\rho) \subseteq \text{supp}(\sigma).$$

If this support condition fails, then

$$D(\rho\|\sigma) = +\infty.$$

The quantum Stein's lemma says that for every fixed $0 < \varepsilon < 1$,

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log \beta_{n,\varepsilon}(\rho\|\sigma) = D(\rho\|\sigma).$$

If logarithms are base two, the exponent is measured in bits per copy. If natural logarithms are used, the exponent is measured in nats per copy.

The theorem says that in the many-copy asymmetric testing problem, the best exponential decay rate of the type-II error is exactly the quantum relative entropy, as long as the type-I error is kept below any fixed constant strictly smaller than one.

Operational meaning

Quantum Stein's lemma gives the operational meaning of quantum relative entropy.

The number

$$D(\rho\|\sigma)$$

is not just an algebraic entropy expression. It is the asymptotic rate at which the false-acceptance probability of $\sigma^{\otimes n}$ can be forced to zero while still accepting $\rho^{\otimes n}$ with high probability.

The mental image is this. Suppose ρ is the “real” model and σ is the “wrong” model. You are allowed to make a small but fixed probability of rejecting the real model. Among all tests that keep this type-I error small, quantum Stein's lemma says that the probability of mistakenly accepting the real model when the wrong model σ was true decays like

$$\beta_n \approx 2^{-nD(\rho\|\sigma)}$$

when logs are base two.

So $D(\rho\|\sigma)$ is the number of distinguishability bits per copy that ρ has against σ in the asymmetric many-copy limit.

Relation to the classical Stein lemma

If ρ and σ commute, they are diagonal in the same basis:

$$\rho = \sum_x P(x)|x\rangle\langle x|, \quad \sigma = \sum_x Q(x)|x\rangle\langle x|.$$

Then the problem becomes classical hypothesis testing between two probability distributions P and Q . The quantum relative entropy becomes the classical Kullback-Leibler divergence:

$$D(\rho\|\sigma) = \sum_x P(x) \log \frac{P(x)}{Q(x)} = D(P\|Q).$$

Classical Stein's lemma says that the optimal type-II error exponent under a fixed type-I constraint is $D(P\|Q)$. Quantum Stein's lemma is the noncommutative extension of this fact.

The genuinely quantum difficulty is that ρ and σ may not commute. Then there is no single measurement basis in which the problem is simply classical. The theorem says that even in this noncommutative setting, the correct asymptotic exponent is still the quantum relative entropy.

The key test: positive part of a weighted difference

For a rate R , consider the Neyman-Pearson test

$$T_{n,R} = \{ \rho^{\otimes n} - 2^{nR} \sigma^{\otimes n} \geq 0 \}.$$

This notation means the projector onto the positive spectral subspace of the Hermitian operator

$$\rho^{\otimes n} - 2^{nR}\sigma^{\otimes n}.$$

This is the many-copy version of the quantum Neyman-Pearson test. It accepts $\rho^{\otimes n}$ on the subspace where $\rho^{\otimes n}$ dominates $2^{nR}\sigma^{\otimes n}$, and rejects it otherwise.

On the support of $T_{n,R}$, we have

$$T_{n,R} (\rho^{\otimes n} - 2^{nR}\sigma^{\otimes n}) T_{n,R} \geq 0.$$

Taking traces gives

$$\text{Tr}(T_{n,R}\rho^{\otimes n}) \geq 2^{nR} \text{Tr}(T_{n,R}\sigma^{\otimes n}).$$

Therefore

$$\beta_n(T_{n,R}) = \text{Tr}(T_{n,R}\sigma^{\otimes n}) \leq 2^{-nR}.$$

This proves immediately that the type-II error exponent is at least R , provided the type-I error

$$\alpha_n(T_{n,R}) = 1 - \text{Tr}(T_{n,R}\rho^{\otimes n})$$

goes to zero.

The central asymptotic fact behind Stein's lemma is that this happens exactly when

$$R < D(\rho\|\sigma).$$

In words: under the true state $\rho^{\otimes n}$, the log-likelihood comparison between $\rho^{\otimes n}$ and $\sigma^{\otimes n}$ concentrates around $nD(\rho\|\sigma)$.

Proof of achievability

We first prove that every rate below $D(\rho\|\sigma)$ is achievable. Choose

$$R < D(\rho\|\sigma).$$

Use the Neyman-Pearson test

$$T_{n,R} = \{ \rho^{\otimes n} - 2^{nR} \sigma^{\otimes n} \geq 0 \}.$$

As shown above,

$$\beta_n(T_{n,R}) \leq 2^{-nR}.$$

The nontrivial part is to show that

$$\alpha_n(T_{n,R}) \rightarrow 0.$$

This is the quantum relative typicality statement. It says that, when the true state is $\rho^{\otimes n}$, almost all of the probability mass eventually lies in the subspace where

$$\rho^{\otimes n} \gtrsim 2^{nR} \sigma^{\otimes n}$$

for every $R < D(\rho\|\sigma)$. Equivalently,

$$\text{Tr}(T_{n,R} \rho^{\otimes n}) \rightarrow 1.$$

In the commuting case, this is just the law of large numbers applied to the random variable

$$\log \frac{P(X)}{Q(X)}.$$

In the noncommuting case, this concentration is the hard part of the theorem. One standard method is to pinch $\rho^{(\otimes n)}$ in the spectral decomposition of $\sigma^{(\otimes n)}$. The pinched operator commutes with $\sigma^{(\otimes n)}$, so a classical large-deviation argument can be applied after controlling the error introduced by pinching. The pinching overhead grows only polynomially in n , and therefore contributes only $o(n)$ to the exponent.

Thus, for every $R < D(\rho \parallel \sigma)$, there exists a sequence of tests such that

$$\alpha_n \rightarrow 0$$

and

$$\beta_n \leq 2^{-nR}.$$

Therefore

$$\liminf_{n \rightarrow \infty} -\frac{1}{n} \log \beta_{n,\varepsilon}(\rho \parallel \sigma) \geq D(\rho \parallel \sigma)$$

for every fixed $0 < \varepsilon < 1$.

Proof of the converse

Now we prove that no rate larger than $D(\rho \parallel \sigma)$ is possible under a fixed nontrivial type-I constraint.

Choose

$$R > D(\rho \parallel \sigma).$$

Let

$$P_{n,R} = \{ \rho^{\otimes n} - 2^{nR} \sigma^{\otimes n} \geq 0 \}.$$

The complementary quantum relative typicality statement says that

$$\text{Tr}(P_{n,R} \rho^{\otimes n}) \rightarrow 0.$$

That is, under $\rho^{\otimes n}$, the subspace where $\rho^{\otimes n}$ dominates $2^{nR}\sigma^{\otimes n}$ has asymptotically negligible probability when R is above the relative entropy.

Let T_n be any test. Decompose

$$\text{Tr}(T_n \rho^{\otimes n})$$

as follows:

$$\text{Tr}(T_n \rho^{\otimes n}) = \text{Tr} [T_n (\rho^{\otimes n} - 2^{nR} \sigma^{\otimes n})] + 2^{nR} \text{Tr}(T_n \sigma^{\otimes n}).$$

For any Hermitian operator X and any $0 \leq T \leq I$,

$$\text{Tr}(TX) \leq \text{Tr}(X_+),$$

where X_+ is the positive part of X . Hence

$$\text{Tr}(T_n \rho^{\otimes n}) \leq \text{Tr} [(\rho^{\otimes n} - 2^{nR} \sigma^{\otimes n})_+] + 2^{nR} \beta_n(T_n).$$

The positive-part term is asymptotically negligible for $R > D(\rho \|\sigma)$. Therefore, if a test satisfies

$$\text{Tr}(T_n \rho^{\otimes n}) \geq 1 - \varepsilon,$$

then for large n ,

$$2^{nR} \beta_n(T_n)$$

must remain bounded away from zero. In particular,

$$\beta_n(T_n) \not\ll 2^{-nR}.$$

Since this holds for every $R > D(\rho \|\sigma)$, the type-II error exponent cannot exceed $D(\rho \|\sigma)$. Thus

$$\limsup_{n \rightarrow \infty} -\frac{1}{n} \log \beta_{n,\varepsilon}(\rho \parallel \sigma) \leq D(\rho \parallel \sigma).$$

Combining achievability and converse gives

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log \beta_{n,\varepsilon}(\rho \parallel \sigma) = D(\rho \parallel \sigma).$$

This proves quantum Stein's lemma.

Strong converse interpretation

There is a stronger statement. If one tries to force the type-II error to decay with exponent strictly larger than $D(\rho \parallel \sigma)$, then the type-I error does not merely fail to stay below a chosen small value. It tends to one.

More concretely, if for some $\delta > 0$,

$$\beta_n(T_n) \leq 2^{-n(D(\rho \parallel \sigma) + \delta)}$$

for all large n , then

$$\alpha_n(T_n) \rightarrow 1.$$

This means that a test that suppresses false acceptance of $\sigma^{\otimes n}$ too aggressively will almost always reject $\rho^{\otimes n}$ even when $\rho^{\otimes n}$ is true. Ogawa and Nagaoka proved this strong converse form, strengthening the earlier achievability and converse picture due to Hiai and Petz.

Example: pure state versus maximally mixed state

Let

$$\rho = |0\rangle\langle 0|$$

and

$$\sigma = \frac{I}{2}.$$

Then

$$D(\rho||\sigma) = \text{Tr } \rho(\log_2 \rho - \log_2 \sigma) = 0 - (-1) = 1.$$

For n copies,

$$\rho^{\otimes n} = |0^n\rangle\langle 0^n|$$

and

$$\sigma^{\otimes n} = \frac{I}{2^n}.$$

Use the test

$$T_n = |0^n\rangle\langle 0^n|.$$

Then

$$\alpha_n = 0,$$

because T_n accepts $|0^n\rangle$ with probability one. The type-II error is

$$\beta_n = \text{Tr} \left[|0^n\rangle\langle 0^n| \frac{I}{2^n} \right] = 2^{-n}.$$

Thus

$$-\frac{1}{n} \log_2 \beta_n = 1,$$

exactly matching

$$D(\rho||\sigma) = 1.$$

This is the cleanest example of Stein's lemma. The wrong model σ assigns probability 2^{-n} to the all-zero string, while the true model ρ assigns probability one.

Example: classical binary distributions

Let

$$P = (0.8, 0.2), \quad Q = (0.3, 0.7).$$

The relative entropy in bits is

$$D(P||Q) = 0.8 \log_2 \frac{0.8}{0.3} + 0.2 \log_2 \frac{0.2}{0.7} \approx 0.771.$$

Classical Stein's lemma says that if samples are drawn independently from either P or Q , then the optimal type-II error exponent under a fixed type-I constraint is about 0.771 bits per sample.

The quantum version reduces to this example when

$$\rho = 0.8|0\rangle\langle 0| + 0.2|1\rangle\langle 1|$$

and

$$\sigma = 0.3|0\rangle\langle 0| + 0.7|1\rangle\langle 1|.$$

The optimal test is asymptotically a threshold test on the empirical log-likelihood ratio. Since the states commute, no genuinely quantum measurement is needed.

Example: orthogonal support and infinite exponent

Let

$$\rho = |0\rangle\langle 0|, \quad \sigma = |1\rangle\langle 1|.$$

Then

$$\text{supp}(\rho) \not\subseteq \text{supp}(\sigma),$$

so

$$D(\rho||\sigma) = +\infty.$$

The test

$$T = |0\rangle\langle 0|$$

has

$$\alpha = 0$$

and

$$\beta = 0.$$

The two states can be perfectly distinguished with one copy. Thus the type-II error exponent is infinite. This explains why the support condition appears in the definition of quantum relative entropy. If σ assigns zero support to something that ρ can produce, then observing that component rules out σ completely.

Example: nonorthogonal pure states still give infinite relative entropy

Let

$$\rho = |0\rangle\langle 0|, \quad \sigma = |+\rangle\langle +|,$$

where

$$|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}.$$

These two states are not orthogonal, so they cannot be perfectly distinguished with one copy. Nevertheless,

$$D(\rho\|\sigma) = +\infty,$$

because the support of ρ , namely $\text{span}\{|0\rangle\}$, is not contained in the support of σ , namely $\text{span}\{|+\rangle\}$.

This may feel surprising at first. The point is that Stein's lemma is an asymptotic asymmetric theorem, not a one-copy symmetric theorem. For n copies, one can test whether the state lies outside the one-dimensional support of $\sigma^{\otimes n}$. Under $\sigma^{\otimes n}$, this event has probability zero. Under $\rho^{\otimes n}$, its probability tends to one because

$$|\langle +|0\rangle|^{2n} = 2^{-n}.$$

Thus there are tests with type-II error exactly zero and type-I error tending to zero. Hence the type-II exponent is infinite.

This example is useful because it separates Stein's lemma from the Helstrom theorem. One-copy Helstrom discrimination between $|0\rangle$ and $|+\rangle$ has nonzero error. Asymmetric many-copy Stein testing can achieve zero type-II error eventually while allowing type-I error to vanish.

How to use the theorem

To use quantum Stein's lemma, identify the null state ρ and the alternative state σ . Then compute

$$D(\rho\|\sigma) = \text{Tr } \rho(\log \rho - \log \sigma).$$

If

$$\text{supp}(\rho) \subseteq \text{supp}(\sigma),$$

this number is finite and gives the optimal type-II error exponent:

$$\beta_{n,\epsilon}(\rho\|\sigma) \approx 2^{-nD(\rho\|\sigma)}.$$

If the support condition fails, then $D(\rho\|\sigma)=+\infty$, and the alternative can be rejected with type-II error zero asymptotically while keeping type-I error small.

In proof work, the theorem is usually used to replace an operational hypothesis-testing quantity by relative entropy in the large- n limit. In one-shot information theory, the finite- n quantity

$$D_H^\varepsilon(\rho\|\sigma) = -\log \beta_\varepsilon(\rho\|\sigma)$$

is called the hypothesis-testing relative entropy. Quantum Stein's lemma says that its normalized i.i.d. limit is ordinary quantum relative entropy:

$$\lim_{n \rightarrow \infty} \frac{1}{n} D_H^\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) = D(\rho\|\sigma).$$

Common mistakes

A common mistake is to confuse symmetric and asymmetric hypothesis testing. The quantum Chernoff bound governs the optimal symmetric Bayesian error exponent when both kinds of errors are treated together. Quantum Stein's lemma governs the asymmetric setting where the type-I error is constrained and the type-II error is minimized.

A second mistake is to forget that ε is fixed. The theorem says that for any fixed $0 < \varepsilon < 1$, the exponent is the same. If ε is allowed to depend on n , then refined second-order and moderate-deviation theories are needed.

A third mistake is to ignore the support condition. If $\text{supp}(\rho) \not\subseteq \text{supp}(\sigma)$, the relative entropy is infinite, and the operational exponent is infinite in the Stein sense.

A fourth mistake is to assume the optimal test is a product measurement performed independently on each copy. In the commuting classical case, product measurements are enough. In the noncommuting quantum case, optimal tests may require collective measurements across many copies.

Final mental image

Quantum Stein's lemma says that quantum relative entropy is the asymptotic price paid by the wrong hypothesis.

If the true state is $\rho^{\otimes n}$ and the alternative is $\sigma^{\otimes n}$, then the best tests can make the probability of falsely accepting $\rho^{\otimes n}$ under $\sigma^{\otimes n}$ decay as

$$2^{-nD(\rho\|\sigma)}$$

while keeping the probability of rejecting $\rho^{\otimes n}$ below any fixed $\varepsilon \in (0,1)$.

So the theorem gives quantum relative entropy its most important operational interpretation:

$$D(\rho\|\sigma) = \text{optimal asymmetric distinguishability rate of } \rho \text{ against } \sigma.$$

It is the noncommutative many-copy version of the classical likelihood-ratio law of large numbers.

References

Hiai, Fumio, and Dénes Petz. "The Proper Formula for Relative Entropy and its Asymptotics in Quantum Probability." *Communications in Mathematical Physics* 143, no. 1 (1991): 99–114.

Ogawa, Tomohiro, and Hiroshi Nagaoka. "Strong Converse and Stein's Lemma in Quantum Hypothesis Testing." *IEEE Transactions on Information Theory* 46, no. 7 (2000): 2428–2433.

Hayashi, Masahito. *Quantum Information Theory: Mathematical Foundation*. Springer, 2017.

Watrous, John. *The Theory of Quantum Information*. Cambridge University Press, 2018.

Nussbaum, Michael, and Arleta Szkoła. "The Chernoff Lower Bound for Symmetric Quantum Hypothesis Testing." *Annals of Statistics* 37, no. 2 (2009): 1040–1057.

Li, Ke. "Second-Order Asymptotics for Quantum Hypothesis Testing." *Annals of Statistics* 42, no. 1 (2014): 171–189.

Bjelaković, Igor, and Rainer Siegmund-Schultze. "Quantum Stein's Lemma Revisited, Inequalities for Quantum Entropies, and a Concavity Theorem of Lieb." [arXiv:quant-ph/0307170](https://arxiv.org/abs/quant-ph/0307170).

Document information

Quantum Stein's Lemma

Project	[QIT 002] State Distinguishability and Measurement Theorems
Document	Primary document
Author	mujirin
Verifier	Not verified
Downloaded	July 03, 2026 20:22 KST
Status	Working
Document link	https://theorytrace.com/projects/state-distinguishability-and-measurement-theorems/documents/untitled-document-d0cfbd/