

Helstrom Theorem

Formal statement

Let ρ and σ be density operators on a finite-dimensional Hilbert space \mathcal{H} . Suppose a source prepares ρ with prior probability

$$p,$$

and prepares σ with prior probability

$$q = 1 - p.$$

The receiver is given one copy of the unknown state and may perform any quantum measurement allowed by the POVM formalism. The task is to guess which state was prepared while minimizing the probability of error, or equivalently maximizing the probability of success.

Define the weighted difference operator

$$\Delta = p\rho - q\sigma.$$

The Helstrom theorem, also called the Holevo-Helstrom theorem in this form, says that the optimal success probability is

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2} (1 + \|p\rho - q\sigma\|_1)$$

and the optimal error probability is

$$P_{\text{err}}^{\text{opt}} = \frac{1}{2} (1 - \|p\rho - q\sigma\|_1).$$

Here

$$\|X\|_1 = \text{Tr} \sqrt{X^\dagger X}$$

is the trace norm. Since Δ is Hermitian, $\|\Delta\|_1$ is simply the sum of the absolute values of the eigenvalues of Δ .

The optimal measurement is obtained from the spectral decomposition of

$$\Delta = p\rho - q\sigma.$$

Let Π_+ be the projector onto the positive eigenspace of Δ , and let Π_- be the projector onto the negative eigenspace. The Helstrom measurement guesses ρ on the positive subspace and guesses σ on the negative subspace. Any zero eigenspace may be assigned arbitrarily, because it contributes equally to both hypotheses.

Thus the theorem is not only a formula for the optimal probability. It also tells us what measurement to perform.

The operational problem

The problem is a binary decision problem. Nature secretly chooses one of two possible quantum states. The receiver is allowed to interrogate the state with a measurement and must output one of two labels:

“it was ρ ” or “it was σ ”.

Because the final decision has only two possible answers, no generality is lost by using a two-outcome POVM. Even if the receiver first performs a measurement with many outcomes, all outcomes that lead to the guess ρ can be grouped into a single effect, and all outcomes that lead to the guess σ can be grouped into the other effect. Therefore the most general strategy can be represented by a POVM

$$\{M, I - M\},$$

where

$$0 \leq M \leq I.$$

The effect M means “guess ρ ,” and the effect $I-M$ means “guess σ .”

For this measurement, the success probability is

$$P_{\text{succ}}(M) = p \text{Tr}(M\rho) + q \text{Tr}((I - M)\sigma).$$

The first term is the probability that ρ was prepared and the measurement led us to guess ρ . The second term is the probability that σ was prepared and the measurement led us to guess σ .

Now expand the expression:

$$\begin{aligned} P_{\text{succ}}(M) &= p \text{Tr}(M\rho) + q \text{Tr}(\sigma) - q \text{Tr}(M\sigma) \\ &= q + \text{Tr}(M(p\rho - q\sigma)) \\ &= q + \text{Tr}(M\Delta). \end{aligned}$$

Thus the entire optimization problem becomes

$$\max_{0 \leq M \leq I} [q + \text{Tr}(M\Delta)].$$

Since q is fixed, we only need to maximize

$$\text{Tr}(M\Delta)$$

over all effects M satisfying $0 \leq M \leq I$.

This is the key reduction. The Helstrom theorem is ultimately a spectral theorem applied to the weighted evidence operator Δ .

Proof

Because Δ is Hermitian, it has a positive-negative decomposition. Write

$$\Delta = \Delta_+ - \Delta_-,$$

where $\Delta_+ \geq 0$, $\Delta_- \geq 0$, and the supports of Δ_+ and Δ_- are orthogonal. Concretely, if

$$\Delta = \sum_j \lambda_j |j\rangle\langle j|$$

is the spectral decomposition of Δ , then

$$\Delta_+ = \sum_{\lambda_j > 0} \lambda_j |j\rangle\langle j|,$$

and

$$\Delta_- = \sum_{\lambda_j < 0} (-\lambda_j) |j\rangle\langle j|.$$

The trace norm is

$$\|\Delta\|_1 = \text{Tr } \Delta_+ + \text{Tr } \Delta_-.$$

Now take any measurement effect M with

$$0 \leq M \leq I.$$

Then

$$\text{Tr}(M\Delta) = \text{Tr}(M\Delta_+) - \text{Tr}(M\Delta_-).$$

Since $M \geq 0$ and $\Delta_- \geq 0$, the term

$$\text{Tr}(M\Delta_-)$$

is nonnegative. Therefore

$$\text{Tr}(M\Delta) \leq \text{Tr}(M\Delta_+).$$

Since $M \leq I$, we also have

$$\text{Tr}(M\Delta_+) \leq \text{Tr } \Delta_+.$$

Thus every possible measurement satisfies

$$\text{Tr}(M\Delta) \leq \text{Tr} \Delta_+.$$

This upper bound is achievable. Choose

$$M = \Pi_+,$$

where Π_+ is the projector onto the positive eigenspace of Δ . Then

$$\text{Tr}(\Pi_+\Delta) = \text{Tr} \Delta_+.$$

So

$$\max_{0 \leq M \leq I} \text{Tr}(M\Delta) = \text{Tr} \Delta_+.$$

Therefore

$$P_{\text{succ}}^{\text{opt}} = q + \text{Tr} \Delta_+.$$

We now rewrite this in trace-norm form. Since

$$\Delta = \Delta_+ - \Delta_-,$$

we have

$$\text{Tr} \Delta = \text{Tr} \Delta_+ - \text{Tr} \Delta_-.$$

Also,

$$\|\Delta\|_1 = \text{Tr} \Delta_+ + \text{Tr} \Delta_-.$$

Adding the two equations gives

$$\|\Delta\|_1 + \text{Tr} \Delta = 2 \text{Tr} \Delta_+.$$

Hence

$$\text{Tr } \Delta_+ = \frac{1}{2} (\|\Delta\|_1 + \text{Tr } \Delta).$$

But

$$\text{Tr } \Delta = \text{Tr}(p\rho - q\sigma) = p - q,$$

because $\text{Tr } \rho = \text{Tr } \sigma = 1$. Therefore

$$\begin{aligned} P_{\text{succ}}^{\text{opt}} &= q + \frac{1}{2} (\|\Delta\|_1 + p - q) \\ &= \frac{p + q}{2} + \frac{1}{2} \|\Delta\|_1 \\ &= \frac{1}{2} (1 + \|p\rho - q\sigma\|_1). \end{aligned}$$

Since

$$P_{\text{err}}^{\text{opt}} = 1 - P_{\text{succ}}^{\text{opt}},$$

we obtain

$$P_{\text{err}}^{\text{opt}} = \frac{1}{2} (1 - \|p\rho - q\sigma\|_1).$$

This proves the theorem.

What the optimal measurement is doing

The operator

$$\Delta = p\rho - q\sigma$$

is the quantum analogue of a weighted likelihood difference. In a classical binary decision problem, for each possible observation x , one compares

$$pP(x|\rho)$$

with

$$qP(x|\sigma).$$

If the first number is larger, one guesses ρ . If the second number is larger, one guesses σ .

In the quantum problem, there may be no single classical sample space in which ρ and σ are simultaneously diagonal. The theorem says to diagonalize the operator

$$\Delta = p\rho - q\sigma$$

instead. Its positive eigenspace is the region of Hilbert space where the weighted evidence favors ρ . Its negative eigenspace is the region where the weighted evidence favors σ .

The Helstrom measurement is therefore the quantum likelihood-ratio test. It does not generally measure ρ or σ directly. It measures the sign of their weighted difference.

This is the operational mental image:

$$\Delta > 0 \Rightarrow \text{guess } \rho,$$

while

$$\Delta < 0 \Rightarrow \text{guess } \sigma.$$

If Δ has a zero eigenspace, then the two hypotheses are exactly balanced there, so any decision rule on that subspace is equally good.

Equal priors and trace distance

A particularly important special case is

$$p = q = \frac{1}{2}.$$

Then

$$\Delta = \frac{1}{2}(\rho - \sigma),$$

so

$$\|\Delta\|_1 = \frac{1}{2}\|\rho - \sigma\|_1.$$

The optimal success probability becomes

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2} + \frac{1}{4}\|\rho - \sigma\|_1.$$

The trace distance between two states is usually defined as

$$D(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1.$$

Therefore

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2}(1 + D(\rho, \sigma)).$$

This gives the trace distance its most important operational meaning. Trace distance is exactly the advantage, above random guessing, that the best possible quantum measurement gives for distinguishing two equally likely states.

If

$$D(\rho, \sigma) = 0,$$

then $\rho = \sigma$, and the best success probability is $1/2$. The states are indistinguishable. If

$$D(\rho, \sigma) = 1,$$

then the states have orthogonal support, and the best success probability is 1. They are perfectly distinguishable.

Example 1: identical states

Suppose

$$\rho = \sigma.$$

Then

$$\Delta = p\rho - q\rho = (p - q)\rho.$$

Since $\rho \geq 0$ and $\text{Tr}\rho = 1$,

$$\|\Delta\|_1 = |p - q|.$$

Thus

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2}(1 + |p - q|) = \max\{p, q\}.$$

This is exactly what should happen. If the two possible quantum states are identical, no measurement can reveal which preparation occurred. The best strategy is to ignore the system and always guess the more likely prior hypothesis.

This example is important because it shows that state distinguishability depends on both the states and the priors. If the states contain no distinguishing information, the prior probabilities still matter.

Example 2: orthogonal pure states

Let

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

These states have orthogonal support. The weighted difference is

$$\Delta = p|0\rangle\langle 0| - q|1\rangle\langle 1|.$$

Its eigenvalues are

$$p \quad \text{and} \quad -q.$$

Therefore

$$\|\Delta\|_1 = p + q = 1.$$

The theorem gives

$$P_{\text{succ}}^{\text{opt}} = 1.$$

The optimal measurement is the projective measurement in the basis

$$\{|0\rangle, |1\rangle\}.$$

If the outcome is $|0\rangle$, guess ρ . If the outcome is $|1\rangle$, guess σ . This example expresses the basic rule that orthogonal quantum states can be perfectly distinguished in a single shot.

Example 3: commuting mixed states reduce to a classical decision rule

Suppose ρ and σ commute. Then they are simultaneously diagonalizable. Write

$$\rho = \sum_i r_i |i\rangle\langle i|, \quad \sigma = \sum_i s_i |i\rangle\langle i|.$$

Then

$$\Delta = \sum_i (pr_i - qs_i) |i\rangle\langle i|.$$

The trace norm is

$$\|\Delta\|_1 = \sum_i |pr_i - qs_i|.$$

Therefore

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2} \left(1 + \sum_i |pr_i - qs_i| \right).$$

This can also be written as

$$P_{\text{succ}}^{\text{opt}} = \sum_i \max\{pr_i, qs_i\}.$$

Indeed, for each classical outcome i , the best decision is to compare pr_i and qs_i . If $pr_i \geq qs_i$, guess ρ . Otherwise, guess σ .

For example, take equal priors

$$p = q = \frac{1}{2},$$

and probability distributions

$$r = (0.8, 0.2), \quad s = (0.3, 0.7).$$

Then

$$\|\Delta\|_1 = \left| \frac{1}{2}(0.8) - \frac{1}{2}(0.3) \right| + \left| \frac{1}{2}(0.2) - \frac{1}{2}(0.7) \right| = 0.25 + 0.25 = 0.5.$$

Thus

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2}(1 + 0.5) = 0.75.$$

Operationally, this is just classical Bayesian hypothesis testing. The quantum theorem becomes classical likelihood comparison when the two density matrices commute.

Example 4: distinguishing $|0\rangle$ from $|+\rangle$

Now consider two nonorthogonal pure qubit states with equal priors:

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

where

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

The overlap is

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

For two pure states with priors p and q , the Helstrom formula simplifies to

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2} \left(1 + \sqrt{1 - 4pq|\langle \psi|\phi \rangle|^2} \right).$$

With equal priors, this gives

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2} \left(1 + \sqrt{1 - |\langle 0|+\rangle|^2} \right) = \frac{1}{2} \left(1 + \frac{1}{\sqrt{2}} \right) \approx 0.8536.$$

Thus a single copy of $|0\rangle$ and $|+\rangle$ can be distinguished better than random guessing, but not perfectly. The failure of perfect discrimination comes from nonorthogonality.

Let us see the measurement. With equal priors,

$$\Delta = \frac{1}{2}|0\rangle\langle 0| - \frac{1}{2}|+\rangle\langle +|.$$

In the computational basis,

$$|0\rangle\langle 0| = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix},$$

and

$$|+\rangle\langle +| = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}.$$

Therefore

$$\Delta = \begin{pmatrix} \frac{1}{4} & -\frac{1}{4} \\ -\frac{1}{4} & -\frac{1}{4} \end{pmatrix}.$$

The eigenvalues are

$$+\frac{1}{2\sqrt{2}} \quad \text{and} \quad -\frac{1}{2\sqrt{2}}.$$

The positive eigenvector is proportional to

$$\cos \frac{\pi}{8} |0\rangle - \sin \frac{\pi}{8} |1\rangle.$$

The negative eigenvector is the orthogonal direction. The optimal measurement is not the computational-basis measurement and not the $|+\rangle, |-\rangle$ measurement. It is the projective measurement halfway, in the correct geometric sense, between the two hypotheses.

This example shows the operational content of the theorem. To distinguish nonorthogonal states optimally, one does not measure in the basis of either state alone. One measures the sign of the weighted difference operator.

Example 5: unequal priors can dominate the measurement

Let

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

but now suppose

$$p = 0.9, \quad q = 0.1.$$

The pure-state formula gives

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2} \left(1 + \sqrt{1 - 4(0.9)(0.1)\frac{1}{2}} \right).$$

Since

$$4(0.9)(0.1)\frac{1}{2} = 0.18,$$

we get

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2}(1 + \sqrt{0.82}) \approx 0.9528.$$

If we ignored the quantum system and always guessed ρ , the success probability would already be 0.9. The measurement improves this to about 0.9528. The theorem tells us exactly how much improvement the quantum evidence can provide beyond the prior bias.

This example is useful because it prevents a common misunderstanding. The optimal measurement is not determined only by the geometric separation between ρ and σ . It is determined by the weighted operator

$$p\rho - q\sigma.$$

Large prior imbalance shifts the decision boundary.

Example 6: mixed states with overlapping support

Consider two qubit states diagonal in the same basis:

$$\rho = \begin{pmatrix} 0.7 & 0 \\ 0 & 0.3 \end{pmatrix}, \quad \sigma = \begin{pmatrix} 0.4 & 0 \\ 0 & 0.6 \end{pmatrix},$$

with equal priors. Then

$$\Delta = \frac{1}{2}(\rho - \sigma) = \begin{pmatrix} 0.15 & 0 \\ 0 & -0.15 \end{pmatrix}.$$

Thus

$$\|\Delta\|_1 = 0.30.$$

The optimal success probability is

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2}(1 + 0.30) = 0.65.$$

The optimal measurement is the computational-basis measurement. If the outcome is 0, guess ρ , because $0.7 > 0.4$. If the outcome is 1, guess σ , because $0.6 > 0.3$.

This example shows how the theorem behaves for mixed states. The states are not perfectly distinguishable because both assign nonzero probability to both outcomes. The trace norm measures the net weighted separation of the two statistical patterns.

How to use the theorem in practice

To apply the Helstrom theorem, one should proceed in a very concrete way. First form the weighted evidence operator

$$\Delta = p\rho - (1 - p)\sigma.$$

Second, diagonalize Δ . Third, build the POVM that projects onto the positive and negative eigenspaces:

$$M_\rho = \Pi_+, \quad M_\sigma = I - \Pi_+.$$

Equivalently, one may use

$$M_\rho = \Pi_+, \quad M_\sigma = \Pi_- + \Pi_0,$$

where Π_0 projects onto the zero eigenspace. The assignment of Π_0 is arbitrary.

Fourth, compute the trace norm

$$\|\Delta\|_1 = \sum_j |\lambda_j|,$$

where λ_j are the eigenvalues of Δ . Then plug it into

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2}(1 + \|\Delta\|_1).$$

In a laboratory interpretation, the theorem says that the optimal detector is a projective measurement of the sign of Δ . In a mathematical proof, the theorem is often used to replace an optimization over all binary POVMs by a trace norm.

Relation to trace distance

For equal priors, the theorem says

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2}(1 + D(\rho, \sigma)),$$

where

$$D(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1.$$

Therefore the trace distance has a direct operational meaning:

$$D(\rho, \sigma) = 2P_{\text{succ}}^{\text{opt}} - 1.$$

It is the distinguishability advantage of the best possible measurement. This is why trace distance is one of the central distance measures in quantum information. It does not merely measure algebraic separation between two matrices. It measures how much better one can do than random guessing in a one-shot discrimination task.

This also explains why trace distance is contractive under quantum channels. If a noisy process \mathcal{N} is applied to both states before the measurement, the receiver can always simulate any later measurement on $\mathcal{N}(\rho)$ and $\mathcal{N}(\sigma)$, but cannot gain information that was not present before the channel. Thus physical processing cannot increase optimal distinguishability.

Relation to POVMs and projective measurements

Although the optimization is over all POVMs, the optimal binary measurement can be chosen to be projective. This is special to the two-state minimum-error problem. The reason is that the optimization reduces to choosing an effect M that maximizes

$$\text{Tr}(M\Delta).$$

The best choice is the sharp projector onto the positive eigenspace of Δ .

For more than two states, there is generally no equally simple closed-form solution. The optimal measurement for discriminating three or more states may be a genuinely nonprojective POVM, and the problem is usually treated using semidefinite programming or special symmetry arguments.

Thus the binary Helstrom theorem is unusually clean. It is the rare case where the full quantum measurement optimization collapses to diagonalizing a single Hermitian operator.

Common mistakes

A common mistake is to use

$$\rho - \sigma$$

instead of

$$p\rho - (1 - p)\sigma.$$

That is correct only for equal priors. With unequal priors, the priors change the decision boundary and must be included inside the Helstrom operator.

Another common mistake is to think the optimal measurement must distinguish the eigenbasis of ρ or the eigenbasis of σ . In general, it does neither. It distinguishes the positive and negative eigenspaces of the weighted difference operator

$$\Delta = p\rho - q\sigma.$$

A third mistake is to think that nonorthogonal pure states can be perfectly distinguished by a clever measurement. They cannot. The Helstrom theorem quantifies the best possible minimum-error strategy. It allows some probability of error. This differs from unambiguous state discrimination, where one allows an inconclusive outcome in exchange for never making a wrong conclusive guess.

A fourth mistake is to forget that the theorem is a one-copy result. If many identical copies are available, one can distinguish

$$\rho^{\otimes n}$$

from

$$\sigma^{\otimes n},$$

and the optimal error probability is governed asymptotically by other results, such as the quantum Chernoff bound. The Helstrom theorem still applies for each fixed n , but the operator becomes

$$p\rho^{\otimes n} - q\sigma^{\otimes n}.$$

Final mental image

The Helstrom theorem says that binary quantum state discrimination is solved by one Hermitian operator:

$$\Delta = p\rho - (1 - p)\sigma.$$

The positive part of Δ is the region where the evidence favors ρ . The negative part is the region where the evidence favors σ . The optimal measurement asks which region the state falls into.

The trace norm

$$\|\Delta\|_1$$

measures the total imbalance between the two weighted hypotheses. If this imbalance is zero, the state gives no useful evidence beyond the priors. If it is one, the hypotheses are perfectly distinguishable. In between, the theorem gives the exact optimal probability:

$$P_{\text{succ}}^{\text{opt}} = \frac{1}{2}(1 + \|\Delta\|_1).$$

Thus the theorem connects geometry, measurement, and decision theory. The trace norm is not just a matrix norm; in this setting, it is the operational distinguishability of quantum states.

References

Helstrom, Carl W. *Quantum Detection and Estimation Theory*. Academic Press, 1976.

Holevo, Alexander S. *Probabilistic and Statistical Aspects of Quantum Theory*. North-Holland, 1982; Springer reprint, 2011.

Holevo, Alexander S. "Statistical Decision Theory for Quantum Systems." *Journal of Multivariate Analysis* 3, no. 4 (1973): 337-394.

Yuen, Horace P., Robert S. Kennedy, and Melvin Lax. "Optimum Testing of Multiple Hypotheses in Quantum Detection Theory." *IEEE Transactions on Information Theory* 21, no. 2 (1975): 125-134.

Nielsen, Michael A., and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 10th anniversary edition, 2010.

Watrous, John. *The Theory of Quantum Information*. Cambridge University Press, 2018. See the section on quantum state discrimination and the Holevo-Helstrom theorem.

Barnett, Stephen M., and Sarah Croke. "Quantum State Discrimination." *Advances in Optics and Photonics* 1, no. 2 (2009): 238-278.

Document information

Helstrom Theorem

Project	[QIT 002] State Distinguishability and Measurement Theorems
Document	Primary document
Author	mujirin
Verifier	Not verified
Downloaded	July 04, 2026 22:36 KST
Status	Working
Document link	https://www.theorytrace.com/projects/state-distinguishability-and-measurement-theorems/documents/helstrom-theorem/