

Quantum Neyman-Pearson Lemma

Formal statement

Let ρ and σ be density operators on a finite-dimensional Hilbert space \mathcal{H} . We want to decide between two hypotheses:

$$H_0 : \text{the state is } \rho, \quad H_1 : \text{the state is } \sigma.$$

A binary quantum test is described by an operator T satisfying

$$0 \leq T \leq I.$$

We interpret T as the POVM effect for accepting H_0 , namely for deciding that the state was ρ . The complementary effect $I-T$ corresponds to accepting H_1 , namely deciding that the state was σ .

The type-I error is the probability of rejecting H_0 when H_0 is true:

$$\alpha(T) = \text{Tr}[(I - T)\rho] = 1 - \text{Tr}(T\rho).$$

The type-II error is the probability of accepting H_0 when H_1 is true:

$$\beta(T) = \text{Tr}(T\sigma).$$

For a fixed allowed type-I error ε , the asymmetric hypothesis-testing problem is

$$\beta_\varepsilon(\rho||\sigma) = \min_{0 \leq T \leq I} \{ \text{Tr}(T\sigma) : \text{Tr}(T\rho) \geq 1 - \varepsilon \}.$$

The quantum Neyman-Pearson lemma says that an optimal test is obtained from the positive spectral subspace of a weighted operator difference

$$\rho - \lambda\sigma$$

for a suitable threshold parameter $\lambda \geq 0$. More precisely, for each $\lambda \geq 0$, define

$$\Delta_\lambda = \rho - \lambda\sigma.$$

Let Π_+^λ , Π_-^λ , and Π_0^λ denote the projectors onto the positive, negative, and zero eigenspaces of Δ_λ , respectively. Then every test of the form

$$T_\lambda = \Pi_+^\lambda + Q_0, \quad 0 \leq Q_0 \leq \Pi_0^\lambda,$$

is optimal for the corresponding Lagrange problem

$$\max_{0 \leq T \leq I} \text{Tr}[T(\rho - \lambda\sigma)].$$

If Q_0 is chosen so that

$$\text{Tr}(T_\lambda \rho) = 1 - \varepsilon,$$

then T_λ is optimal for the constrained Neyman-Pearson problem $\beta_\varepsilon(\rho|\sigma)$. If exact equality cannot be reached without boundary randomization, the freedom $0 \leq Q_0 \leq \Pi_0^\lambda$ is precisely the quantum analogue of randomizing on the likelihood-ratio boundary.

This is the asymmetric version of binary quantum state discrimination. The symmetric Bayesian version with priors is the Helstrom-Holevo theorem. The same positive-part idea appears there with the weighted operator $p\rho - (1-p)\sigma$.

What problem the theorem solves

The Helstrom theorem answers the symmetric Bayesian question: if ρ is prepared with prior probability p and σ with prior probability $1-p$, what measurement minimizes the average error probability? The Neyman-Pearson version answers a different question. It says: suppose we care more about one error than the other. Suppose we are willing to tolerate type-I error at most ε , and among all such tests we want the smallest possible type-II error. What measurement should we use?

This asymmetric formulation is extremely important in statistics and information theory. In many scientific or engineering decisions, the two mistakes are not equally costly. In quantum communication, quantum cryptography, and quantum Shannon theory, one often fixes the probability of rejecting the true state ρ and then studies how fast the probability of mistaking σ for ρ can decay over many copies.

The operational mental image is this. The operator

$$\Delta_\lambda = \rho - \lambda\sigma$$

is a quantum likelihood-ratio comparison. On subspaces where Δ_λ is positive, the weighted evidence favors ρ . On subspaces where it is negative, the weighted evidence favors σ . The parameter λ controls how conservative the test is about accepting ρ . A larger λ makes the test demand stronger evidence before accepting ρ , which usually reduces type-II error at the cost of increasing type-I error.

Proof of the positive-part optimization

The core of the theorem is a simple operator inequality. Fix $\lambda \geq 0$, and write the Jordan decomposition of the Hermitian operator

$$\Delta_\lambda = \rho - \lambda\sigma$$

as

$$\Delta_\lambda = (\Delta_\lambda)_+ - (\Delta_\lambda)_-,$$

where

$$(\Delta_\lambda)_+ \geq 0, \quad (\Delta_\lambda)_- \geq 0,$$

and the supports of the two positive operators are orthogonal. The operator $(\Delta_\lambda)_+$ is the positive part of Δ_λ , and $(\Delta_\lambda)_-$ is the negative part with its sign removed.

Let T be any test, so $0 \leq T \leq I$. Then

$$\text{Tr}(T\Delta_\lambda) = \text{Tr}(T(\Delta_\lambda)_+) - \text{Tr}(T(\Delta_\lambda)_-).$$

Because $T \geq 0$ and $(\Delta_\lambda)_- \geq 0$, we have

$$\text{Tr}(T(\Delta_\lambda)_-) \geq 0.$$

Therefore

$$\text{Tr}(T\Delta_\lambda) \leq \text{Tr}(T(\Delta_\lambda)_+).$$

Since $T \leq I$, we also have

$$\text{Tr}(T(\Delta_\lambda)_+) \leq \text{Tr}(\Delta_\lambda)_+.$$

Combining the inequalities gives

$$\text{Tr}(T\Delta_\lambda) \leq \text{Tr}(\Delta_\lambda)_+.$$

This upper bound is achieved by choosing

$$T = \Pi_+^\lambda,$$

the projector onto the positive eigenspace of Δ_λ . More generally, any choice

$$T = \Pi_+^\lambda + Q_0, \quad 0 \leq Q_0 \leq \Pi_0^\lambda,$$

also achieves the same value, because the zero eigenspace contributes nothing to $\text{Tr}(T\Delta_\lambda)$. Thus

$$\max_{0 \leq T \leq I} \text{Tr}(T\Delta_\lambda) = \text{Tr}(\Delta_\lambda)_+.$$

This proves the Lagrange form of the lemma.

From the Lagrange form to the constrained test

Now suppose T_λ is a positive-part test satisfying

$$\text{Tr}(T_\lambda \rho) = 1 - \varepsilon.$$

We show that it minimizes the type-II error among all tests whose type-I error is at most ε .

Let T be any feasible test, so

$$\text{Tr}(T\rho) \geq 1 - \varepsilon = \text{Tr}(T_\lambda\rho).$$

Because T_λ maximizes $\text{Tr}(T\Delta_\lambda)$, we have

$$\text{Tr}(T_\lambda\Delta_\lambda) \geq \text{Tr}(T\Delta_\lambda).$$

Substituting $\Delta_\lambda = \rho - \lambda\sigma$, this becomes

$$\text{Tr}(T_\lambda\rho) - \lambda \text{Tr}(T_\lambda\sigma) \geq \text{Tr}(T\rho) - \lambda \text{Tr}(T\sigma).$$

Rearranging gives

$$\lambda [\text{Tr}(T\sigma) - \text{Tr}(T_\lambda\sigma)] \geq \text{Tr}(T\rho) - \text{Tr}(T_\lambda\rho).$$

The right-hand side is nonnegative because T is feasible. Therefore, when $\lambda > 0$,

$$\text{Tr}(T\sigma) \geq \text{Tr}(T_\lambda\sigma).$$

This says

$$\beta(T) \geq \beta(T_\lambda).$$

Thus T_λ has the smallest possible type-II error among all tests satisfying the type-I constraint. This proves the constrained Neyman-Pearson statement.

The special boundary cases $\lambda=0$, zero eigenspaces, or inactive constraints are handled by the same positive-part principle together with the freedom to randomize on the boundary. In finite dimension the feasible set $0 \leq T \leq I$ is compact and convex, so an optimal test exists; the supporting-threshold argument above identifies its structure.

Connection with the classical Neyman-Pearson lemma

The classical Neyman-Pearson lemma says that the most powerful test at a fixed significance level is a likelihood-ratio test. If the classical distributions are $r(x)$ under H_0 and $s(x)$ under H_1 , then one accepts H_0 when

$$r(x) > \lambda s(x),$$

rejects H_0 when

$$r(x) < \lambda s(x),$$

and possibly randomizes when equality holds.

The quantum theorem is exactly the noncommutative version of this rule. If ρ and σ commute, they can be diagonalized in the same basis:

$$\rho = \sum_x r_x |x\rangle\langle x|, \quad \sigma = \sum_x s_x |x\rangle\langle x|.$$

Then

$$\rho - \lambda\sigma = \sum_x (r_x - \lambda s_x) |x\rangle\langle x|.$$

The positive projector Π_{+}^{λ} is exactly the projector onto those classical outcomes x for which

$$r_x > \lambda s_x.$$

So the quantum test reduces to the classical likelihood-ratio test.

The only new difficulty in quantum theory is that ρ and σ may not commute. Then there is no single classical outcome space on which both states are simultaneously diagonal. The theorem says that the correct replacement for the likelihood ratio is the spectral sign of the Hermitian operator $\rho - \lambda\sigma$.

Example: a classical commuting test

Let

$$\rho = \begin{pmatrix} 0.8 & 0 \\ 0 & 0.2 \end{pmatrix}, \quad \sigma = \begin{pmatrix} 0.3 & 0 \\ 0 & 0.7 \end{pmatrix}.$$

These are diagonal states, so this is a classical problem in quantum notation. Suppose we want type-I error at most

$$\varepsilon = 0.2.$$

The test that accepts H_0 only on the first outcome is

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

Its type-I error is

$$\alpha(T) = 1 - \text{Tr}(T\rho) = 1 - 0.8 = 0.2.$$

Its type-II error is

$$\beta(T) = \text{Tr}(T\sigma) = 0.3.$$

This test is produced by the positive part of $\rho - \lambda\sigma$. For example, take $\lambda=1$. Then

$$\rho - \sigma = \begin{pmatrix} 0.5 & 0 \\ 0 & -0.5 \end{pmatrix}.$$

The positive subspace is spanned by $|0\rangle$, so the Neyman-Pearson test is

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

Operationally, this says: accept H_0 only when the observed outcome is much more likely under ρ than under σ .

Example: the noncommuting two-state test

Let

$$\rho = |0\rangle\langle 0|, \quad \sigma = |+\rangle\langle +|, \quad |+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}.$$

These states do not commute. There is no basis in which both are simultaneously classical. Take $\lambda=1$. Then

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

In the computational basis,

$$\rho = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad \sigma = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix},$$

so

$$\Delta = \begin{pmatrix} 1/2 & -1/2 \\ -1/2 & -1/2 \end{pmatrix}.$$

The eigenvalues of Δ are

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

The positive eigenvector is proportional to

$$|0\rangle + (1 - \sqrt{2})|1\rangle.$$

Therefore the optimal threshold test for $\lambda=1$ is not a measurement in the $|0\rangle, |1\rangle$ basis and not a measurement in the $|+\rangle, |-\rangle$ basis. It is the projective measurement onto the positive and negative eigenspaces of $\rho-\sigma$.

This is the first genuinely quantum lesson of the theorem. The best asymmetric test is not generally obtained by measuring either state in its own eigenbasis. It is obtained by diagonalizing the weighted difference operator.

Example: relation to the Helstrom theorem

If the goal is not to constrain one error but to minimize the average error with priors p and $q=1-p$, then the success probability for a test T that accepts ρ is

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

This can be rewritten as

$$P_{\text{succ}}(T) = q + \text{Tr}[T(p\rho - q\sigma)].$$

Therefore the optimal Bayesian test is the positive projector of

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

This is the Helstrom measurement. Thus the Helstrom theorem is the Bayesian version of the same positive-part principle, while the Neyman-Pearson lemma is the constrained asymmetric version.

For equal priors, $p=q=1/2$, the relevant operator is proportional to

$$\rho - \sigma.$$

The optimal success probability becomes

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

This is the usual trace-distance form of binary quantum state discrimination.

Boundary randomization

In finite classical hypothesis testing, the likelihood-ratio test sometimes needs randomization at the threshold in order to hit a prescribed type-I error exactly. The same thing happens quantum mechanically.

If $\rho - \lambda\sigma$ has a nontrivial zero eigenspace, then vectors in that eigenspace satisfy exact weighted balance between the two hypotheses. Including or excluding that subspace does not change the Lagrange objective

$$\text{Tr}[T(\rho - \lambda\sigma)].$$

Therefore the theorem allows

$$T = \Pi_+^\lambda + Q_0, \quad 0 \leq Q_0 \leq \Pi_0^\lambda.$$

The operator Q_0 is a randomized or partial decision on the boundary subspace. It is chosen to make

$$\text{Tr}(T\rho) = 1 - \varepsilon$$

when exact saturation of the type-I constraint is desired.

Operationally, the zero eigenspace is the region where the test is exactly indifferent between the two weighted hypotheses. Randomization there changes the error probabilities without changing the optimality of the threshold rule.

How this theorem is used in quantum information

The single-shot quantity

$$\beta_\varepsilon(\rho\|\sigma) = \min_{0 \leq T \leq I} \{\text{Tr}(T\sigma) : \text{Tr}(T\rho) \geq 1 - \varepsilon\}$$

is one of the basic objects of one-shot quantum information theory. Its logarithmic version is the hypothesis-testing relative entropy,

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

The Neyman-Pearson lemma tells us the structure of the optimizing test. For many-copy states, one applies the same idea to

$$\rho^{\otimes n} \quad \text{and} \quad \sigma^{\otimes n}.$$

The asymptotic behavior of the optimal type-II error under a fixed type-I constraint leads to quantum Stein's lemma, whose exponent is the quantum relative entropy $D(\rho\|\sigma)$. Quantum Stein's lemma and its strong converse were developed by Hiai-Petz and Ogawa-Nagaoka, and later asymmetric error-exponent refinements include the quantum Hoeffding and Chernoff bounds.

Thus the Neyman-Pearson lemma is not only a one-shot measurement rule. It is the local building block behind asymptotic quantum hypothesis testing, one-shot information theory, converse bounds, coding theorems, and many resource-theoretic distinguishability arguments.

Common mistakes

A common mistake is to confuse the Neyman-Pearson problem with the Helstrom problem. The Helstrom problem minimizes average error with fixed priors. The Neyman-Pearson problem fixes or bounds one kind of error and minimizes the other. Both are solved by positive parts of weighted differences, but the interpretation of the weight is different.

Another common mistake is to forget which POVM element means which decision. In this presentation, T means “accept H_0 ,” so

$$\alpha(T) = 1 - \text{Tr}(T\rho)$$

and

$$\beta(T) = \text{Tr}(T\sigma).$$

Some books use the opposite convention, where the test operator means “reject H_0 .” The formulas then look transposed, but the mathematics is identical.

A third mistake is to assume that the optimal test is a measurement in the eigenbasis of ρ or σ . This is generally false when the states do not commute. The correct operator to diagonalize is

$$\rho - \lambda\sigma.$$

A fourth mistake is to ignore randomization on the zero eigenspace. Boundary randomization is not an artificial trick. It is required when the allowed type-I error lies between two discrete acceptance probabilities.

Final mental image

The quantum Neyman-Pearson lemma is the asymmetric decision rule for two quantum states. It says that the optimal test is obtained by forming a weighted evidence operator

$$\rho - \lambda\sigma,$$

diagonalizing it, accepting ρ on the positive subspace, accepting σ on the negative subspace, and possibly randomizing on the zero subspace.

In one sentence:

the optimal asymmetric quantum test is the sign measurement of a weighted sta

The theorem is useful because it converts an optimization over all binary POVMs into a spectral decomposition of one Hermitian operator. It is the quantum analogue of the classical likelihood-ratio test, and it is the foundation of asymmetric quantum hypothesis testing.

References

Neyman, Jerzy, and Egon S. Pearson. "On the Problem of the Most Efficient Tests of Statistical Hypotheses." *Philosophical Transactions of the Royal Society of London. Series A* 231 (1933): 289-337.

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.

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

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.

Audenaert, K. M. R., M. Nussbaum, A. Szkoła, and F. Verstraete. "Asymptotic Error Rates in Quantum Hypothesis Testing." *Communications in Mathematical Physics* 279 (2008): 251-283.

Document information

Quantum Neyman-Pearson Lemma

Project	[QIT 002] State Distinguishability and Measurement Theorems
Document	Primary document
Author	mujirin
Verifier	Not verified
Downloaded	July 03, 2026 18:20 KST
Status	Working
Document link	https://theorytrace.com/projects/state-distinguishability-and-measurement-theorems/documents/untitled-document-f84944/