# Jensen-Shannon Stability for Sample-Based Testers

A short addendum to the Gaussian halfspace pullback note


Jonathan R. Landers



Draft, June 2026



**Abstract**

The current note [1] proves a simple total-variation stability lemma: if a tester uses $q$ independent samples, then replacing the sampling law $\mu$ by $\nu$ perturbs its acceptance probability by at most $q\|\mu - \nu\|_{\mathrm{TV}}$. This addendum records a parallel information-divergence version. With Jensen-Shannon divergence measured in nats,

$$\left|\mathbb{P}_\mu[A^f = 1] - \mathbb{P}_\nu[A^f = 1]\right| \le \sqrt{2q\,\mathrm{JS}(\mu\|\nu)}.$$

Thus a distributional approximation satisfying $\mathrm{JS}(\mu\|\nu) \le \eta^2$ changes the acceptance probability by at most $O(\sqrt{q}\,\eta)$, improving the naive linear-in-$q$ dependence when Jensen-Shannon is quadratically small in the underlying perturbation. The result is best viewed as a refinement of the approximate-transfer part of the note, not as a replacement for the exact pullback/lognormal theorem.


## 1 Setup

Let $(\mathcal{X}, \mathcal{A})$ be a measurable space and let $f : \mathcal{X} \to \{0, 1\}$ be a fixed Boolean target. A $q$-sample tester $A$ receives

$$(X_1, f(X_1)), \ldots, (X_q, f(X_q)), \qquad X_i \overset{\text{i.i.d.}}{\sim} \mu,$$

has arbitrary internal randomness, and outputs either accept or reject. Write $\mathbb{P}_\mu[A^f = 1]$ for its acceptance probability when the ambient law is $\mu$.

Theorem 1 of [1] gives the bound

$$\left|\mathbb{P}_\mu[A^f = 1] - \mathbb{P}_\nu[A^f = 1]\right| \le q\|\mu - \nu\|_{\mathrm{TV}}.$$

This is sharp for a union-bound argument, but it is not the only useful way to measure distributional closeness.

**Definition 1** (Jensen-Shannon divergence)**.** *For probability measures $\mu, \nu$ on the same measurable space, define*

$$m := \frac{\mu + \nu}{2}, \qquad \mathrm{JS}(\mu\|\nu) := \frac{1}{2}\mathrm{KL}(\mu\|m) + \frac{1}{2}\mathrm{KL}(\nu\|m),$$

*where logarithms are natural. With base-2 logarithms, multiply the right-hand side of the final theorem below by* $\ln 2$ *inside the square root.*

## 2 The Jensen-Shannon stability theorem

**Theorem 1** (Jensen-Shannon stability for $q$-sample testers)**.** *Let $A$ be any $q$-sample tester and let $\mu, \nu$ be probability measures on $\mathcal{X}$. Then for every Boolean target $f : \mathcal{X} \to \{0, 1\}$,*

$$\boxed{\left|\mathbb{P}_\mu[A^f = 1] - \mathbb{P}_\nu[A^f = 1]\right| \le \sqrt{2q\,\mathrm{JS}(\mu\|\nu)}.}$$

*Proof.* Let

$$P_\mu^f := (x \mapsto (x, f(x)))_\# \mu, \qquad P_\nu^f := (x \mapsto (x, f(x)))_\# \nu$$

be the induced laws of one labeled sample. The labeled transcript laws before internal randomness are $(P_\mu^f)^q$ and $(P_\nu^f)^q$. Since the map $x \mapsto (x, f(x))$ retains $x$, it is information-preserving on its image, so

$$\mathrm{JS}(P_\mu^f \| P_\nu^f) = \mathrm{JS}(\mu\|\nu).$$

The tester's accept/reject decision is a randomized measurable function of the transcript. Therefore, by data processing for total variation,

$$\left|\mathbb{P}_\mu[A^f = 1] - \mathbb{P}_\nu[A^f = 1]\right| \le \|(P_\mu^f)^q - (P_\nu^f)^q\|_{\mathrm{TV}}.$$

We now bound the right-hand side by Jensen-Shannon divergence. For arbitrary laws $P, Q$,

$$\|P - Q\|_{\mathrm{TV}} \le \sqrt{2\mathrm{JS}(P\|Q)}.$$

Indeed, if $M = (P + Q)/2$, then $\|P - M\|_{\mathrm{TV}} = \|Q - M\|_{\mathrm{TV}} = \|P - Q\|_{\mathrm{TV}}/2$, and Pinsker's inequality gives

$$\mathrm{KL}(P\|M) \ge 2\|P - M\|_{\mathrm{TV}}^2 = \frac{1}{2}\|P - Q\|_{\mathrm{TV}}^2,$$

with the same bound for $\mathrm{KL}(Q\|M)$. Averaging yields $\mathrm{JS}(P\|Q) \ge \frac{1}{2}\|P - Q\|_{\mathrm{TV}}^2$.

It remains to use the product bound

$$\mathrm{JS}(P^q\|Q^q) \le q\mathrm{JS}(P\|Q).$$

For completeness, here is the short proof. Let $B$ be a fair bit. Conditional on $B = 0$, draw $Z_1, \ldots, Z_q$ i.i.d. from $P$; conditional on $B = 1$, draw them i.i.d. from $Q$. Then

$$\mathrm{JS}(P^q\|Q^q) = I(B; Z_1, \ldots, Z_q).$$

By the chain rule and conditional independence given $B$,

$$I(B; Z_1, \ldots, Z_q) = \sum_{i=1}^q I(B; Z_i \mid Z_{<i}) \le \sum_{i=1}^q I(B, Z_{<i}; Z_i) = \sum_{i=1}^q I(B; Z_i) = q\mathrm{JS}(P\|Q).$$

Applying these two inequalities with $P = P_\mu^f$ and $Q = P_\nu^f$ gives

$$\|(P_\mu^f)^q - (P_\nu^f)^q\|_{\mathrm{TV}} \le \sqrt{2\mathrm{JS}((P_\mu^f)^q\|(P_\nu^f)^q)} \le \sqrt{2q\mathrm{JS}(P_\mu^f\|P_\nu^f)} = \sqrt{2q\mathrm{JS}(\mu\|\nu)}.$$

This proves the theorem. $\square$

## 3 Gap transfer

**Corollary 1** (Jensen-Shannon gap transfer)**.** *Let $\mathcal{Y}$ and $\mathcal{N}$ be two families of Boolean functions on $\mathcal{X}$. Suppose a $q$-sample tester $A$ satisfies, under law $\mu$,*

$$\mathbb{P}_\mu[A^f = 1] \geq c \quad (f \in \mathcal{Y}), \qquad \mathbb{P}_\mu[A^f = 1] \leq s \quad (f \in \mathcal{N}),$$

*with gap $\gamma := c - s > 0$. Under any other law $\nu$, the same tester satisfies*

$$\mathbb{P}_\nu[A^f = 1] \geq c - \varepsilon_{\mathrm{JS}} \quad (f \in \mathcal{Y}), \qquad \mathbb{P}_\nu[A^f = 1] \leq s + \varepsilon_{\mathrm{JS}} \quad (f \in \mathcal{N}),$$

*where*

$$\varepsilon_{\mathrm{JS}} := \sqrt{2q\mathrm{JS}(\mu\|\nu)}.$$

*Hence the retained gap is at least*

$$\gamma - 2\sqrt{2q\mathrm{JS}(\mu\|\nu)}.$$

*Proof.* Apply the theorem separately to the completeness and soundness inequalities. □

**Corollary 2** (Constant-overhead amplification)**.** *Assume in the preceding corollary that $c = 2/3$ and $s = 1/3$. If*

$$\sqrt{2q\mathrm{JS}(\mu\|\nu)} \leq \frac{1}{12},$$

*then one run under $\nu$ has completeness at least $7/12$ and soundness at most $5/12$. A constant number of independent repetitions restores a $2/3$ versus $1/3$ guarantee, so the query complexity remains $O(q)$.*

## 4 Interpretation for the halfspace note

The theorem is an information-divergence analogue of the $q\Delta$ transfer lemma in [1]. The total-variation theorem gives

$$\text{acceptance perturbation} \leq q\|\mu - \nu\|_{\mathrm{TV}},$$

whereas the Jensen-Shannon theorem gives

$$\text{acceptance perturbation} \leq \sqrt{2q\mathrm{JS}(\mu\|\nu)}.$$

Thus, if a perturbative regime gives

$$\mathrm{JS}(\mu\|\nu) = O(\delta^2),$$

then the stability loss is

$$O(\sqrt{q}\,\delta),$$

instead of the $O(q\delta)$ loss one would get by combining Theorem 1 with $\|\mu-\nu\|_{\mathrm{TV}} = O(\delta)$. This is the sense in which the Jensen-Shannon version can capture a quadratic improvement in distributional fidelity, echoing the role of Jensen-Shannon divergence in the beta-estimation paper [2].

**Remark 1** (What this does and does not replace)**.** *This result belongs in the approximate-transfer layer of the current note. It improves the way acceptance probabilities are controlled when $\nu$ is close to $\mu$ in an information-theoretic sense. It does not replace the exact pullback principle. In particular, for the multivariate lognormal application in [1], the stronger statement is still the transport identity $T_\# D = N(0, I_n)$, which preserves the testing problem exactly and incurs no divergence loss at all.*

**Remark 2** (Promise-set caveat). *As in Remark 4 of [1], this theorem controls the algorithmic effect of changing the sampling law for fixed promise sets. If the yes/no promises themselves are defined using the ambient measure, as in relative-error testing, then a full approximate-transfer theorem may also require comparing the induced distance notions under $\mu$ and $\nu$. The exact pullback theorem avoids that additional issue because the relevant probabilities are preserved after conjugation by the transport map.*

## 5 Suggested insertion point

For integration into [1], this material can be placed immediately after Theorem 1 and Corollary 2 as a short subsection titled "A Jensen-Shannon refinement." The narrative would then be:

$$\text{TV stability} \quad \longrightarrow \quad \text{JS-refined stability} \quad \longrightarrow \quad \text{exact pullback stability.}$$

This keeps the note's main message intact: approximate distributional closeness gives controlled tester degradation, while exact transport gives an unchanged testing problem.

## References


[1] Jonathan R. Landers. *Gaussian Halfspace Testing, Distributional Stability, and the Exact Lognormal Pullback*. Draft note, 2026. https://jonathan-r-landers.s3.us-east-1.amazonaws.com/landers_halfspace_query_complexity_note_arxiv.pdf

[2] Jonathan R. Landers. *Closed-Form Beta Distribution Estimation from Sparse Statistics with Random Forest Implicit Regularization*. arXiv:2507.23767, 2025. https://arxiv.org/abs/2507.23767

[3] Jianhua Lin. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information Theory*, 37(1):145–151, 1991.