Arithmetic blindness and its removal in energy-transfer channels

Supplementary proofs and computational methods

Ahmet Faruk Arapoğlu
Independent Researcher, Istanbul, Türkiye
ORCID: 0009-0005-9374-3368
ahmetarapoglu@gmail.com

31 October 2026

Abstract

This supplement specifies the calculations behind the main article and gives additional proofs for the channel representation, graph quotient, physical constraints and finite-shot decision rule. It also documents a source-analysis tool that accepts user-supplied rational probability tables. All saved numerical data are synthetic model calculations. The universal statements are established analytically; exact and floating-point checks serve different verification roles.

S1 Conventions and an exact channel derivation

The receiver basis is $(|d\rangle, |r\rangle)$. Source occupations are nonnegative integers, $\hbar = 1$, and $x = gt$. For each initial source basis state (n, m) write $u = \psi(n, m+1)$ and $v = \psi(n+1, m)$. In the interaction picture, tracing the evolved source gives three Kraus operators:

$$K_{nm,0} = \sqrt{P_{nm}}\,\mathrm{diag}(\cos xu, \cos xv), \tag{S1}$$

$$K_{nm,+} = -i\sqrt{P_{nm}}\,\sin(xu)|r\rangle\langle d|, \tag{S2}$$

$$K_{nm,-} = -i\sqrt{P_{nm}}\,\sin(xv)|d\rangle\langle r|. \tag{S3}$$

For either input receiver level the sum of $K^\dagger K$ is its source weight times $\cos^2 + \sin^2$. Summing P_{nm} gives the identity. Thus the map is completely positive and trace preserving. Its excited population is

$$q' = q + (1-q)\sum_{nm} P_{nm}\sin^2(xu) - q\sum_{nm} P_{nm}\sin^2(xv), \tag{S4}$$

and its off-diagonal entry is $c\sum P_{nm}\cos(xu)\cos(xv)$. These are the main-text formulas. Cross terms between the transition Kraus operators do not survive the source trace: they have distinct final number labels.

Symmetry of ψ interchanges u and v when n, m are swapped. For the transpose pair, the channel difference is consequently $X \mapsto \mathcal{A}_P\,\mathrm{Tr}(X)Z_R$. Tensoring with an identity on an ancilla gives $\mathcal{A}_P Z_R \otimes \mathrm{Tr}_R\,\rho_{RA}$, so its trace norm on every density operator is $2|\mathcal{A}_P|$. This proves both the ancillary statement and the single-use equal-prior success probability. It does not describe repeated interactions with the same disturbed source; the experimental protocol uses fresh preparations.

S2 One dark positive pair does not force matched blindness

The pair graph has vertices $\{a, b\}$ and edges joining $\{i, j+1\}$ and $\{i+1, j\}$. Along each edge the pair sum is fixed, and the smaller member increases by one. Thus it is a disjoint union of paths, possibly broken into shorter paths by gaps in the input alphabet.

First merge only the mandatory zero vertices $\{0, s\}$. Each is a leaf and occurs at most once in its component. The result F_0 is a forest. Suppose the positive unordered pair amplitudes are mutually distinct, with exactly one such pair v assigned zero. If the common zero vertex is absent, no extra identification occurs. If it is present, identifying v with it either joins different components or closes one path in a single tree. In the latter case the frequency graph has a one-dimensional cycle space, allowing a loop as a one-edge cycle.

That unique path starts on the fixed-sum branch containing v. It uses exactly one alphabet edge incident to occupation zero, namely $\{0, s-1\}$. Every nonzero frequency-cycle flow therefore has nonzero occupation divergence at zero. It cannot be a matched circulation. The simultaneous kernel of the two boundaries remains zero even though the frequency graph can contain a cycle.

The hypothesis concerns injectivity on *all* positive unordered pair labels, including the possible zero-valued label. Injectivity only among nonzero amplitudes would allow many dark interior pairs and would not justify this proof. The reproducibility script checks 105 single-dark-pair cases for contiguous alphabets of sizes two through eight, with all other pair labels distinct.

S3 Physical norm bounds and short-time response

For an edge $i < j$, positivity implies

$$|J_{ij}| = |P_{ij} - P_{ji}| \le P_{ij} + P_{ji}, \qquad \sum_{i<j} |J_{ij}| \le 1. \tag{S5}$$

Matching also gives $P_{ij}, P_{ji} \le \min(\pi_i, \pi_j)$, hence $|J_{ij}| \le \min(\pi_i, \pi_j)$. An incidence column has absolute sum at most two. The positive-frequency boundary therefore obeys

$$\|w_+\|_1 \le 2, \qquad \|w_+\|_\infty \le \Delta_+ \pi_{\max}. \tag{S6}$$

Using $8\nu_c^2 = \|w_+\|_2^2 \le \|w_+\|_1 \|w_+\|_\infty$ proves the marginal-spread bound. For a standard positive squared label k, an unequal unordered factor pair can contribute at most two incidences and an equal pair at most one. This is bounded by the number of divisors of k.

For completeness, the subpower divisor estimate needed here is elementary. Fix $\epsilon > 0$. If a prime $p \ge 2^{1/\epsilon}$, then $a+1 \le 2^a \le p^{\epsilon a}$ for all $a \ge 0$. For each of the finitely many smaller primes, $(a+1)/p^{\epsilon a}$ has a finite supremum. Multiplying these finite constants over the small primes gives $d(n) \le C_\epsilon n^\epsilon$. Since $k \le (N-1)^2$, the largest incidence degree is $N^{o(1)}$. This proves the required growth bound without relying on a theorem-number convention in a reference book.

The expansion $\sin^2 y = y^2 - y^4/3 + O(y^6)$ gives the fourth-order contrast in the main article. The uniform inequality follows from the Taylor remainder

$$|\sin^2 y - y^2| = \tfrac{1}{2}|1 - \cos(2y) - 2y^2| \le y^4/3. \tag{S7}$$

On an off-diagonal alphabet edge the two squared frequencies are at most $(N-1)^2$. After the matched quadratic terms cancel, summing the two remainder bounds against $|J_{ij}|$ gives $|\mathcal{A}_P| \le 2(N-1)^4 x^4/3$. The additional bound one is simply the maximum difference of probabilities.

S4 Calibration and finite-shot inference

For the minimum-gap certificate, use the set consisting of zero and all distinct positive nominal amplitude labels, even when the nominal graph has no zero-frequency vertex. If each positive-pair amplitude has absolute error below $\Delta_L/2$, different labels cannot meet. Exchange symmetry ensures that identical unordered pairs still have identical amplitudes. The mandatory zero boundary is preserved by assumption. These structural requirements are part of the certificate: independent, unconstrained errors on repeated pair occurrences would be a different channel model.

For a perturbed nominally blind source, each edge contrast is a difference of two $\sin^2$ terms. The derivative of $\sin^2 y$ has magnitude at most one, giving $2|x|\eta|J_{ij}|$ per edge and the summed perturbation bound. For a unit vector in the old blind subspace, Cauchy–Schwarz gives $\|J\|_{1,E} \leq \sqrt{\binom{N}{2}}$. Averaging x^2 over $[0,T]$ yields $T^2/3$; the min–max principle then bounds the indicated number of small Gramian eigenvalues.

At one measurement setting let X_s, Y_s be independent Bernoulli outcomes for the two sources. The estimator is $S^{-1}\sum_s (X_s - Y_s)$. Its $2S$ independent weighted summands each have range width $1/S$, so Hoeffding gives

$$\Pr(|\widehat{\mathcal{A}} - \mathbb{E}\widehat{\mathcal{A}}| > h) \leq 2e^{-Sh^2}. \tag{S8}$$

A union bound over L_t settings gives $h = \sqrt{\log(2L_t/\alpha)/S}$. Bias δ_{tot} is added outside this stochastic radius. Under the null, the true expected contrast has magnitude at most δ_{tot}, so rejecting above $h + \delta_{\mathrm{tot}}$ controls false positives. Under a specified nominal alternative of magnitude a, the estimate is at least $a - h - \delta_{\mathrm{tot}}$ on the coverage event. Thus $a > 2(h + \delta_{\mathrm{tot}})$ suffices for detection on that event.

Source matching alone does not calibrate the entire measurement. If $\widehat{P}$ approximates P with total-variation error ϵ_P, a bounded transition probability changes by at most ϵ_P. Separate source-preparation errors, readout offsets, timing uncertainty and dynamical mismatch can therefore be translated into conservative probability-bias contributions, then combined into δ_{tot}. A bound on matching error alone is not necessarily a bound on all these effects. The supplied 0.01 systematic budget is an example design target.

S5 Exact and numerical calculations

S5.1 Exact graph construction

`code/model.py` builds the occupation incidence D and frequency incidence B from the same oriented edge list. Standard labels are integer products. Cubic and bounded-cubic labels are unordered positive pairs; the analytic injectivity theorem justifies this representation. Mandatory zeros are one common label. No rounding threshold is used to decide exact collisions.

For the contiguous standard alphabet, `standard_rank` constructs the boundary of anchored triangles directly. It uses SymPy's rational-domain matrix rank rather than an approximate numerical rank. The saved $N = 9$ kernel vectors are primitive integer vectors checked against the relevant exact matrix and converted to edge circulations. The independent audit in `audit/` uses a separate fraction-based elimination implementation at sizes three through twelve.

For uniform marginals the set of allowed tables is the scaled Birkhoff polytope. If P is a convex combination of permutation sources, convexity of the RMS bounds it by the largest RMS in that combination. Therefore a permutation maximizes the RMS. Integer boundary norms are enumerated for all permutations through $N = 9$; the total is $6 + 24 + 120 + 720 + 5040 + 40320 + 362880 = 409110$. Dividing the largest squared norm by $8N^2$ gives the recorded exact squared RMS. Finite-time assignment calculations instead maximize a single-time linear functional and can use a different permutation at every time.

S5.2 Matrix evolution and numerical tolerance

The full matrix is built by joining each allowed basis pair with its stated amplitude. It is not obtained by inserting the reduced-channel formula. With ten Fock levels per mode its dimension is 200. Although outermost unpopulated basis blocks can be truncated, every initially populated state and its transition partner are included; no population reaches a truncated partner in these checks. Hermitian eigendecomposition supplies the unitary and a direct partial trace supplies the receiver state.

Numerical comparisons use explicit absolute tolerances and save their actual residuals. They verify floating-point implementations of the stated formulas, rather than prove arbitrary-precision equalities. Exact graph ranks, rational source normalization and integer kernel tests are recorded separately. Finite-dimensional norm and energy checks illustrate the bounded-Hamiltonian theorem; the full-space result follows from the analytic operator bound, not from a cutoff extrapolation.

S5.3 Conditioning at equal peak amplitude

Both dressed laws are normalized using all 36 populated source entries of the six-occupation example and both receiver transition endpoints. This includes diagonal entries even though they cancel from $P - P^{\mathsf{T}}$. On this support the raw and bounded maxima are, respectively, 17.0379090684 and 0.944561201844. Dividing each amplitude by its own maximum makes the largest populated Rabi rate equal.

Let $f(x)$ be the row of edge contrasts and Q an orthonormal basis for $\ker D$. For quadrature nodes x_j with positive weights a_j summing to one, the weighted response design has rows $\sqrt{a_j} f(x_j) Q$. Its singular values are the square roots of the eigenvalues of the positive-semidefinite quadrature Gramian. This avoids interpreting small negative eigenvalues from subtractive round-off as physical results.

The principal calculations use 20001 trapezoidal nodes. Repeating the $T = 50$ and $T = 300$ comparisons with 10001 and 40001 nodes changes the minimum gain by at most 7.2×10^{-7} relative in the recorded refinement. Tiny $T = 10$ gains are stored as order-of-magnitude diagnostics and are not used for the quantitative comparison in the main article. Refinement is numerical convergence evidence, not a certified enclosure of the continuum eigenvalues.

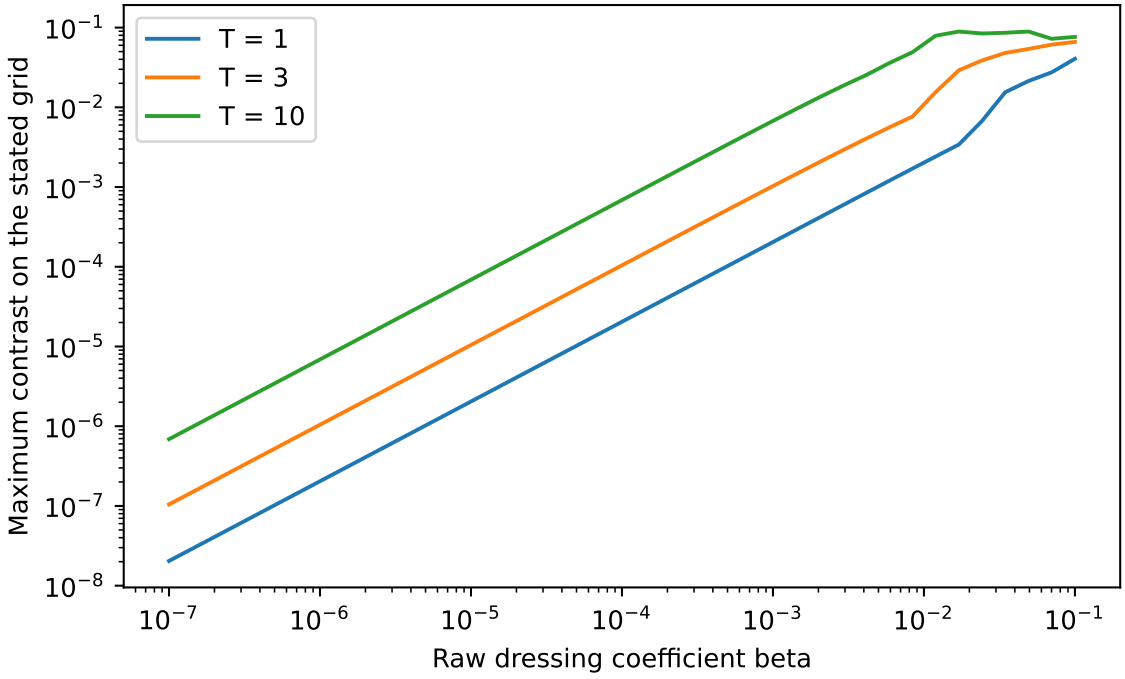


Figure 1: Weak raw dressing of the exactly hidden pair. Finite-window response vanishes continuously as the dressing tends to zero, despite generic exact-rank restoration. Curves show maxima over the documented finite time grids.

S6 How to reproduce or inspect another source

The release uses Python 3.12 with the versions listed in `requirements.txt`. From the package root:

```
python3 -m venv .venv
.venv/bin/python -m pip install -r requirements.txt
.venv/bin/python code/reproduce.py
.venv/bin/python code/analyze_source.py \
    examples/hidden_pair.json --output my_source_check
```

The reproduction script has no online data dependency after installation. It regenerates the numerical results and figures under `results/` and `figures/`. The source analyzer accepts rational strings such as `"1/72"`, validates normalization, nonnegativity and matching, then reports standard/cubic/bounded graph boundaries, signed means, squared RMS and finite-grid contrasts. It rejects an existing output prefix to avoid silently replacing a previous inspection.

The archive contains the exact witness table, source analyzer, three-channel matrix checks, rank and RMS tables, four figures in vector PDF and PNG formats, and a manifest of file hashes. The editable English article and methods document are supplied in a separate manuscript-source archive. The file map below covers both the manuscript sources and the computational supplement.

S7 File map and interpretation

File or directory	Purpose
`main.tex`, English body files	Editable journal-format article
`supplement.tex`	This methods and proof document
`code/model.py`	Exact graph and channel primitives
`code/reproduce.py`	Deterministic calculations and figure generation
`code/analyze_source.py`	Analysis of user-supplied matched sources
`results/results.json`	Parameters, exact results, numerical residuals
`results/test_log.json`	Individual executed checks
`results/kernel_N9.json`	Exact blind-direction certificate
`examples/hidden_pair.json`	Editable rational source example
`audit/`	Independent standard-library calculations and saved checks

Author. Ahmet Faruk Arapoğlu is the sole author. Funding and competing-interest statements are given in the main article.

Use of generative artificial intelligence. ChatGPT versions 5.6 and 6.0 were used for mathematical calculations, literature searches, translation, restructuring, language editing and manuscript formatting. The author takes full responsibility for the scientific reasoning, derivations, computational results and final text.