Matroidal Bayes Bounds for General Quantum Process Discrimination

Canonical Compression, Support Congestion, and Exact Qubit Phase Families

Lluis Eriksson

August 2026

Abstract

Minimum-error discrimination of quantum processes is normally optimized over testers whose normalization may encode parallel, sequential, or indefinite-order access. We give an architecture-independent upper bound that retains the labeled dependence structure of the process hypotheses and arbitrary nonuniform priors. For a fixed deterministic tester normalization, Moore–Penrose compression maps the tester exactly to a POVM acting on normalized effective states; it preserves every conditional probability. We then associate to the support subspaces of arbitrary positive process operators a Rado matroid on the hypothesis labels. The vector of correct-label probabilities of every general tester lies in this matroid's independence polytope. Consequently the Bayes success probability is at most the prior weight of a maximum-weight independent transversal, computable by the matroid greedy algorithm. Equality for a proposed tester is audited by rank saturation only at the prior prefixes followed by a strict prior drop. To remove the discontinuity of exact supports, we also prove a robust spectral-tail version: replacing each hypothesis by any positive low-rank core changes the ceiling by at most the prior-weighted worst-case discarded tester mass. In particular, admixture of an arbitrary valid process with weight η weakens the certificate by at most η.

For rank-one processes the Rado matroid reduces to the vector matroid of the process vectors and strictly refines the sum-of-largest-priors dimension bound. We solve an infinite physical family exactly. Any finite collection of distinct qubit phase gates lies in a rank-two flat, while one or both off-diagonal Pauli gates are coloops. The Bayes ceiling is the sum of the two largest phase priors plus all coloop priors. Whenever I and Z carry the two largest phase priors, a fixed Bell tester attains it, yielding the exact general-tester optimum throughout an open prior chamber. For $I, Z, \mathrm{diag}(1, i), X, Y$ with priors $(.30, .25, .20, .15, .10)$, the exact optimum is .80 while the total-dimension relaxation is .90. We also give a tight trine-plus-coloop family in which every dependent phase hypothesis receives a nonzero effect: for three equatorial phase gates with priors (a, a, a) and an X gate of prior $b < a$, a trine POVM attains the matroid value $2a + b$, strictly below the top-dimension value $3a$. Perfect rank-one process identification is separately equivalent to a deterministic Gram feasibility condition. The support theorem is not claimed to be always tight; identical full-rank states provide a sharp counterexample to such an interpretation. The result combines standard canonical tester normalization and classical Rado–Edmonds theory; neither ingredient is claimed as new.

1 Introduction

Quantum channel and process discrimination asks which unknown transformation was inserted into an experiment. The optimizing experiment may use parallel entanglement, adaptive memory, or a more general process matrix. The quantum tester formalism places all such strategies, for a fixed access model, in one positive-operator optimization [1, 2]. This generality is useful but can obscure simple impossibility certificates. A full tester semidefinite program may be expensive, and an ambient-dimensional bound may discard most of the structure of the promised hypotheses.

For equiprobable pure states spanning a D-dimensional space, the elementary ceiling is D/M. Its nonuniform analogue is the sum of the D largest priors. Shah recently gave this latter statement explicitly for state discrimination, together with a spectral mixed-state extension [3]. In repeated unitary-channel discrimination, Bavaresco, Murao, and Quintino obtained a uniform general-strategy ceiling from the ambient symmetric-tensor dimension [4]. Such total-rank bounds do not see *which* labels cause the rank. This loss matters when high-prior hypotheses are dependent but a low-prior hypothesis supplies the last independent direction.

The missing object is a matroid. For rank-one process operators, every label determines a process vector, and linear independence defines a representable matroid. For mixed process operators, every label instead determines a support subspace. Rado's independent-transversal theorem associates a matroid to this family of subspaces [5]. Its independent sets are the label sets from whose supports one can choose linearly independent representatives. Edmonds' rank inequalities describe the corresponding independence polytope [6].

Our main observation is that a tester supplies a point of exactly this polytope. If s_i is the conditional probability of correctly guessing label i, then for every set B of labels,

$$\sum_{i\in B} s_i \leq \dim \sum_{i\in B} \operatorname{supp} \rho_i,$$

after canonical compression to effective states. Combining these support inequalities with $s_i \leq 1$ gives all Rado-rank inequalities. The common linear map induced by the tester normalization can only lower the Rado rank. Maximizing the Bayes functional over the resulting independence polytope is a maximum-weight independent-set problem, solved greedily.

This produces three layers of bounds:

$$P_{\text{succ}}^{\text{GEN}} \leq \beta(\mathcal{M}_E, p) \leq \Pi_D(p) \leq \Pi_{H_X(k)}(p).$$

The successive terms are the maximum independent prior mass, the mass of the D largest priors, and an optional algebraic ambient ceiling. The first inequality retains every dependent subfamily and may be strict in the second. Its proof does not assume a group orbit, a uniform prior, a compact parameter space, or a definite causal order. It does assume the usual deterministic tester normalization for the chosen process access model.

The contributions are as follows. First, we state an exact fixed-normalization equivalence between general testers and ordinary state discrimination, with singular normalizations handled on their support. Second, we prove the mixed Rado-matroid Bayes ceiling and an exact strict-prior-prefix equality audit. Third, we prove a robust low-rank-core extension controlled by discarded tester mass, so full-rank noise need not erase the certificate. Fourth, we give an explicit implementation route through linear matroid intersection, including arithmetic complexity and a mixed-support regression fixture. Fifth, we specialize the exact theorem to rank-one processes and characterize perfect identification by a deterministic Gram equation. Sixth, we solve a continuous family of one-query qubit phase ensembles in which the matroid ceiling is exact on an open prior chamber and strictly stronger than the total-rank ceiling. A second trine construction attains the ceiling while giving every dependent hypothesis a nonzero decision effect. A Hilbert-function ceiling is retained only as a brief comparison with global span bounds. Finally, we exhibit a mixed-state counterexample showing that support data alone do not determine the optimum.

The canonical tester construction, Rado's theorem, Edmonds' polytope theorem, and the pure-state top-d-priors bound are established ingredients. The claim here is the support-matroid organization of correct-label probabilities, its lift to the general tester setting, and the operational consequences. We do not claim that the resulting upper bound is universally attained.

2 General testers and canonical compression

All Hilbert spaces are finite-dimensional. Let $J_1, \dots, J_M \succeq 0$ be process operators with priors $p_i > 0$ and $\sum_i p_i = 1$. The process operators may be Choi operators of channels, tensor products of repeated channel uses, or more general combs. The access model fixes a convex set $\mathcal{D}$ of physical deterministic tester normalizations. Its partial-trace equations are not needed below; the only identity we use explicitly is

$$\operatorname{Tr}(WJ_i) = 1, \qquad W \in \mathcal{D}, \quad i = 1, \dots, M. \tag{1}$$

A tester with outcomes labeled by the hypotheses is a family $T_i \succeq 0$ satisfying

$$\sum_{i=1}^{M} T_i = W \in \mathcal{D}. \tag{2}$$

Its conditional probabilities and Bayes success are

$$q(j|i) = \operatorname{Tr}(T_j J_i), \qquad P_{\text{succ}}(T) = \sum_i p_i \operatorname{Tr}(T_i J_i). \tag{3}$$

This is the standard tester convention in which process Choi operators need not have unit trace. The deterministic normalization supplies the missing normalization through equation (1).

Write $S_W = \operatorname{supp} W$, let P_W be its orthogonal projector, and let $W^{+1/2}$ denote the Moore–Penrose inverse square root: it is $W^{-1/2}$ on S_W and zero on $S_W^{\perp}$.

Theorem 2.1 (Exact canonical compression)**.** *Fix $W \in \mathcal{D}$. Define normalized effective states*

$$\rho_i^W = W^{1/2} J_i W^{1/2}. \tag{4}$$

The maps

$$M_i = W^{+1/2} T_i W^{+1/2}, \qquad T_i = W^{1/2} M_i W^{1/2} \tag{5}$$

are mutually inverse between testers summing to W and POVMs on S_W. Moreover they preserve the full conditional-probability table:

$$\operatorname{Tr}(T_j J_i) = \operatorname{Tr}(M_j \rho_i^W). \tag{6}$$

Consequently

$$P_{\text{succ}}^{\text{GEN}}(\{p_i, J_i\}) = \sup_{W \in \mathcal{D}} P_{\text{state}}(\{p_i, \rho_i^W\}). \tag{7}$$

Proof. From $0 \preceq T_i \preceq W$, positivity implies $\operatorname{supp} T_i \subseteq S_W$. Hence

$$\sum_i M_i = W^{+1/2} W W^{+1/2} = P_W.$$

Conversely, if $M_i \succeq 0$ and $\sum_i M_i = P_W$, then the second map in equation (5) gives positive tester effects summing to W. The support inclusion shows that the maps are inverse. Cyclicity of trace and the support identities give equation (6). Finally, $\operatorname{Tr} \rho_i^W = \operatorname{Tr}(WJ_i) = 1$, so the effective operators are states. Optimizing the POVM at fixed W and then optimizing the physical deterministic normalization proves equation (7). □

This construction is the canonical POVM associated with a tester normalization [2]. The point of theorem 2.1 is not a new dilation theorem, but an exact route by which state-discrimination constraints can be transported to every admissible tester normalization. It is essential to use the Moore–Penrose inverse on S_W; an ordinary inverse does not exist for singular normalizations.

3 The Rado matroid of process supports

Let

$$\mathcal{S}_i = \operatorname{ran} J_i, \qquad d_E(B) = \dim \sum_{i \in B} \mathcal{S}_i \tag{8}$$

for $B \subseteq [M] = \{1, \ldots, M\}$, with $d_E(\varnothing) = 0$.

Definition 3.1 (Support Rado matroid)**.** A label set $I \subseteq [M]$ is independent if there are representatives $w_i \in \mathcal{S}_i$, $i \in I$, that are linearly independent. The resulting matroid on $[M]$ is denoted $\mathcal{M}_E$.

Rado's theorem is a matroidal Hall theorem. In the present linear-subspace form it says that a set I has a full independent transversal exactly when

$$\dim \sum_{i \in B} \mathcal{S}_i \geq |B| \quad \text{for every } B \subseteq I. \tag{9}$$

The corresponding rank formula is the following.

Lemma 3.2 (Rado rank formula)**.** *The rank of $A \subseteq [M]$ in $\mathcal{M}_E$ is*

$$r_E(A) = \min_{B \subseteq A} \left(|A \setminus B| + d_E(B) \right). \tag{10}$$

Proof. If $I \subseteq A$ admits independent representatives, then for every $B \subseteq A$ at most $d_E(B)$ representatives with labels in B can be independent, while at most $|A \setminus B|$ labels lie outside B. Thus $|I| \leq |A \setminus B| + d_E(B)$, proving the upper bound.

For the reverse inequality, choose a finite basis for each $\mathcal{S}_i$ and tag each basis vector by its label, even when two tagged vectors coincide. On the disjoint union of these tagged bases, intersect the ambient linear matroid with the partition matroid allowing at most one vector per label. The matroid-intersection min–max theorem gives the maximum size of a common independent set. For a label set B, taking the union of the tagged bases with labels in B gives the value $d_E(B) + |A \setminus B|$. Conversely, for an arbitrary subset Q of tagged vectors, let B contain precisely the labels whose complete tagged basis is contained in Q. The linear rank of Q is at least $d_E(B)$, and the partition rank of its complement is $|A \setminus B|$. Therefore the matroid-intersection minimum is exactly equation (10). ☐

The independence polytope of a matroid with rank r is

$$\mathcal{P}(\mathcal{M}) = \left\{ x \in \mathbb{R}^M_{\geq 0} : x(A) \leq r(A) \text{ for all } A \subseteq [M] \right\}, \quad x(A) = \sum_{i \in A} x_i. \tag{11}$$

Edmonds proved that this polytope is integral: its vertices are the incidence vectors of independent sets [6]. Linear optimization over it is equivalently solved by the matroid greedy algorithm.

4 The general-tester Bayes ceiling

Define the maximum independent prior mass

$$\beta(\mathcal{M}_E, p) = \max \left\{ \sum_{i \in I} p_i : I \text{ independent in } \mathcal{M}_E \right\}. \tag{12}$$

Theorem 4.1 (Rado-matroid Bayes ceiling)**.** *For arbitrary positive process hypotheses and every general tester in the chosen access model,*

$$P^{\mathrm{GEN}}_{\mathrm{succ}} \leq \beta(\mathcal{M}_E, p). \tag{13}$$

The right-hand side is computed by inspecting labels in nonincreasing prior order and retaining a label exactly when it increases the Rado rank equation (10).

Proof. Fix a deterministic normalization W and its canonical POVM. Put

$$s_i = \mathrm{Tr}(M_i \rho_i^W). \tag{14}$$

These are the conditional correct-label probabilities, so $0 \le s_i \le 1$. The effective support subspace is

$$\mathcal{S}_i^W = \operatorname{ran} \rho_i^W = W^{1/2} \mathcal{S}_i. \tag{15}$$

Indeed, $\rho_i^W = (W^{1/2} J_i^{1/2})(W^{1/2} J_i^{1/2})^\dagger$, whose range is $W^{1/2} \operatorname{ran} J_i^{1/2} = W^{1/2}\mathcal{S}_i$. Let $d_W(B) = \dim \sum_{i \in B} \mathcal{S}_i^W$, and let P_B^W project onto that sum. Because a density operator has operator norm at most one, $\rho_i^W \preceq P_B^W$ for $i \in B$. Therefore

$$\begin{aligned} s(B) &= \sum_{i \in B} \mathrm{Tr}(M_i \rho_i^W) \\ &\le \mathrm{Tr}\left(P_B^W \sum_{i \in B} M_i\right) \le \mathrm{Tr}\, P_B^W = d_W(B). \end{aligned} \tag{16}$$

For $B \subseteq A$, combine this with $s_i \le 1$:

$$s(A) = s(B) + s(A \setminus B) \le d_W(B) + |A \setminus B|. \tag{17}$$

Minimizing over $B \subseteq A$ and applying theorem 3.2 gives

$$s(A) \le r_W(A). \tag{18}$$

Thus $s \in \mathcal{P}(\mathcal{M}_W)$, where $\mathcal{M}_W$ is the Rado matroid of the effective support family.

The map $W^{1/2}$ is common to all labels and cannot increase the dimension of a sum of subspaces. Hence $d_W(B) \le d_E(B)$ for every B, and equation (10) gives

$$r_W(A) \le r_E(A). \tag{19}$$

Consequently $s \in \mathcal{P}(\mathcal{M}_E)$. Edmonds' integrality theorem now yields

$$\sum_i p_i s_i \le \max_{x \in \mathcal{P}(\mathcal{M}_E)} p \cdot x = \beta(\mathcal{M}_E, p).$$

This holds for every $W \in \mathcal{D}$, and theorem 2.1 completes the proof. The standard matroid greedy theorem gives the stated algorithm. □

The proof has a useful direct form that avoids enumerating all 2^M rank inequalities. Relabel the hypotheses so that $p_1 \ge \cdots \ge p_M$, put $A_j = \{1, \ldots, j\}$, and set $p_{M+1} = 0$. Abel summation gives

$$\sum_i p_i s_i = \sum_{j=1}^{M} (p_j - p_{j+1}) s(A_j) \le \sum_{j=1}^{M} (p_j - p_{j+1}) r_E(A_j). \tag{20}$$

After the first j labels, greedy has retained exactly $r_E(A_j)$ labels. The last expression in equation (20) is therefore the greedy independent set's total prior weight.

Corollary 4.2 (Exact prefix equality audit)**.** *A particular tester attains equation* (13) *if and only if, for every strict prior drop $p_j > p_{j+1}$,*

$$s(A_j) = r_W(A_j) = r_E(A_j). \tag{21}$$

Proof. Each coefficient $p_j - p_{j+1}$ in equation (20) is nonnegative, and the proof uses the termwise chain $s(A_j) \le r_W(A_j) \le r_E(A_j)$. Equality of the weighted sums is equivalent to equality at every prefix whose coefficient is positive. □

Ties are handled automatically: no condition is imposed inside a constant prior block. Because all priors are positive, the terminal drop $p_M > p_{M+1} = 0$ is included. If zero-prior labels are present, they should be removed before applying the statement; effects assigned to them are operationally irrelevant.

4.1 A robust low-rank-core certificate

Exact support is discontinuous: an arbitrarily small full-rank perturbation can make every $\mathcal{S}_i$ equal to the ambient space. The next statement replaces exact support by a freely chosen positive core and charges only the discarded tester mass. Choose operators

$$0 \preceq L_i \preceq J_i, \qquad R_i = J_i - L_i, \qquad \mathcal{S}_i^L = \operatorname{ran} L_i, \tag{22}$$

let $\mathcal{M}_L$ be the Rado matroid of the core supports, and define

$$\bar{\varepsilon}_i = sup_{W \in \mathcal{D}} \operatorname{Tr}(W R_i). \tag{23}$$

By equation (1), $0 \le \bar{\varepsilon}_i \le 1$.

Theorem 4.3 (Spectral-tail robust Rado ceiling)**.** *For every choice of positive cores equation* (22)*,*

$$P_{\text{succ}}^{\text{GEN}} \le \min \left\{ 1,\ \beta(\mathcal{M}_L, p) + \sum_{i=1}^{M} p_i \bar{\varepsilon}_i \right\}. \tag{24}$$

The exact theorem is recovered by $L_i = J_i$.

Proof. Fix $W \in \mathcal{D}$ and its canonical POVM $\{M_i\}$. Split the effective state as

$$\rho_i^W = A_i^W + E_i^W, \qquad A_i^W = W^{1/2} L_i W^{1/2}, \quad E_i^W = W^{1/2} R_i W^{1/2}.$$

Since $0 \preceq A_i^W \preceq \rho_i^W$ and $\operatorname{Tr} \rho_i^W = 1$, the operator A_i^W is bounded above by the projector onto $W^{1/2}\mathcal{S}_i^L$. Also $0 \preceq M_i \preceq P_W$, hence

$$\operatorname{Tr}(M_i E_i^W) \le \operatorname{Tr} E_i^W = \operatorname{Tr}(W R_i) \le \bar{\varepsilon}_i.$$

Consequently, for every label set B,

$$\sum_{i \in B} s_i \le \dim \sum_{i \in B} W^{1/2} \mathcal{S}_i^L + \sum_{i \in B} \bar{\varepsilon}_i. \tag{25}$$

For $A \subseteq [M]$, choose a minimizer $B \subseteq A$ in the Rado rank formula for the uncompressed core family. Combining equation (25) with $s_i \le 1$ and rank monotonicity under $W^{1/2}$ gives

$$s(A) \le |A \setminus B| + d_L(B) + \sum_{i \in B} \bar{\varepsilon}_i \le r_L(A) + \sum_{i \in A} \bar{\varepsilon}_i.$$

Apply this to the sorted-prior prefixes in equation (20). The residual terms telescope by Abel summation:

$$\sum_{j=1}^{M} (p_j - p_{j+1}) \sum_{i \le j} \bar{\varepsilon}_i = \sum_{i=1}^{M} p_i \bar{\varepsilon}_i.$$

The rank terms give $\beta(\mathcal{M}_L, p)$, proving equation (24); the cap by one is the elementary probability bound. ☐

Corollary 4.4 (Linear stability under valid process admixture)**.** *Suppose $J_i = (1-\eta) J_i^{(0)} + \eta N_i$ with $0 \le \eta \le 1$, where both $J_i^{(0)}$ and N_i obey the tester normalization equation* (1)*. Then*

$$P_{\text{succ}}^{\text{GEN}} \le \min\{1, \beta(\mathcal{M}_0, p) + \eta\}, \tag{26}$$

where $\mathcal{M}_0$ is the Rado matroid of $\operatorname{ran} J_i^{(0)}$.

Proof. If $\eta = 1$, equation (26) is just $P_{\text{succ}}^{\text{GEN}} \le 1$. Otherwise take $L_i = (1-\eta) J_i^{(0)}$ and $R_i = \eta N_i$; then $\bar{\varepsilon}_i = \eta$ and $\operatorname{ran} L_i = \operatorname{ran} J_i^{(0)}$. ☐

For ordinary state discrimination, a convenient core is the spectral truncation of each density operator; then $\bar{\varepsilon}_i$ is exactly its discarded spectral mass. For multi-time process models, the supremum in equation (23) retains the physical tester normalization rather than an ambient trace norm.

5 Rank-one processes, dimension relaxations, and equality

Suppose now

$$J_i = |v_i\rangle\langle v_i|. \tag{27}$$

Then $\mathcal{S}_i = \mathbb{C}v_i$. An independent transversal is simply a linearly independent subfamily of the process vectors. Thus $\mathcal{M}_E$ is their vector matroid and

$$\beta(\mathcal{M}_E, p) = \max\left\{\sum_{i\in I} p_i : \{v_i : i \in I\} \text{ is linearly independent}\right\}. \tag{28}$$

Let

$$D_E = \dim \operatorname{span}\{v_1, \dots, v_M\}, \qquad \Pi_d(p) = \sum_{j=1}^{\min\{d,M\}} p_j^{\downarrow}. \tag{29}$$

Corollary 5.1 (Vector-matroid and top-prior bounds)**.** *For rank-one process hypotheses,*

$$P_{\text{succ}}^{\text{GEN}} \le \beta(\mathcal{M}_E, p) \le \Pi_{D_E}(p). \tag{30}$$

For uniform priors this becomes $P_{\text{succ}}^{\text{GEN}} \le \min\{1, D_E/M\}$*.*

Proof. Every independent set in a rank-D_E matroid has at most D_E elements, so its prior weight is at most the sum of the D_E largest priors. Under a uniform prior, every basis has weight D_E/M when $D_E < M$. □

The second inequality in equation (30) is the pure-state sum-of-top-d-priors bound after canonical compression. It is not claimed as new; see Shah [3]. The first inequality can be strictly stronger because it remembers which high-prior labels lie in a common flat.

There is also a useful equality description for the coarser dimension-only bound. Fix W, write $\phi_i = W^{1/2}v_i$, let $R = \operatorname{span}\{\phi_i\}$ and $r = \dim R$, and compress the POVM to R: $M_i' = P_R M_i P_R$. Let $\tau = p_r^{\downarrow}$ and

$$H = \{i : p_i > \tau\}, \qquad E = \{i : p_i = \tau\}, \qquad L = \{i : p_i < \tau\}.$$

Proposition 5.2 (Dimension-face equality)**.** *The fixed-W strategy attains $\Pi_r(p)$ if and only if*

$$\begin{aligned} M_i' &= |\phi_i\rangle\langle\phi_i|, && i \in H, \\ M_i' &= c_i|\phi_i\rangle\langle\phi_i|, && i \in E, \\ M_i' &= 0, && i \in L, \end{aligned} \tag{31}$$

where $0 \le c_i \le 1$ and

$$\sum_{i\in H} |\phi_i\rangle\langle\phi_i| + \sum_{i\in E} c_i|\phi_i\rangle\langle\phi_i| = P_R. \tag{32}$$

Proof. Put $s_i = \langle\phi_i|M_i'|\phi_i\rangle$ and $x_i = \operatorname{Tr} M_i'$. Since $|\phi_i\rangle\langle\phi_i| \preceq P_R$, $0 \le s_i \le x_i$, while $\sum_i x_i = r$ and $s_i \le 1$. Equality in the linear program defining $\Pi_r(p)$ fills labels above the threshold, empties labels below it, and permits fractional filling on the tie block. As all priors are positive, equality also forces $s_i = x_i$ termwise. For a positive operator M_i', equality $\operatorname{Tr}[M_i'(P_R - |\phi_i\rangle\langle\phi_i|)] = 0$ forces its support into $\mathbb{C}\phi_i$. This gives equation (31); POVM completeness gives equation (32). The converse is immediate. □

For uniform priors, equation (32) is positive Parseval scalability. If $K = (\langle\phi_i, \phi_j\rangle)_{ij}$ and $C = \operatorname{diag}(c_i)$, it is equivalent to

$$KCK = K. \tag{33}$$

The tie block is essential. The simpler phrase "the top r states must be orthogonal" is valid only when the relevant threshold is separated from all discarded priors and no fractional tie choice remains.

6 Perfect rank-one process identification

Let $V = [v_1 \ \cdots \ v_M]$ be the process-vector matrix. The tester normalization controls the effective Gram matrix:

$$K_W = (\langle \phi_i, \phi_j \rangle)_{ij} = V^\dagger W V. \tag{34}$$

Theorem 6.1 (Deterministic Gram feasibility)**.** *The rank-one processes are perfectly distinguishable by a general tester if and only if there exists a physical deterministic normalization $W \in \mathcal{D}$ such that*

$$V^\dagger W V = I_M. \tag{35}$$

Proof. For fixed W, theorem 2.1 reduces the problem to deterministic discrimination of normalized pure states ϕ_i. Such states are perfectly distinguishable if and only if they are pairwise orthogonal. Their diagonal Gram entries already equal one by equation (1), so orthogonality is exactly equation (35). Conversely, when equation (35) holds, projective measurement onto the effective vectors, completed on the orthogonal complement, gives a perfect canonical POVM and hence a perfect tester. □

The condition is linear in W. Together with the positive-semidefinite and partial-trace constraints defining $\mathcal{D}$, it is an SDP feasibility problem. It implies $D_E \geq M$, but that rank condition is not sufficient: the physical deterministic affine slice may contain no W that orthogonalizes the process vectors. Replacing $\mathcal{D}$ by all positive matrices normalized only on the finite promise would change the operational problem and can create spurious certificates.

7 Brief comparison with global span bounds

A companion analysis develops exact Hilbert laws and harmonic spectra for structured unitary-oracle families [7]. We need only the following coarse comparison here; the new support-resolved statement of this paper is theorem 4.1.

For an ordinary d-dimensional unitary oracle, use the unnormalized Choi vector $u(U) = |U\rangle\!\rangle \in \mathbb{C}^{d^2}$. With k calls to the same unknown channel, the pure process vector is

$$v(U) = u(U)^{\otimes k}. \tag{36}$$

Let $X \subset \mathbb{P}(\mathbb{C}^{d^2})$ be a projective variety containing the projectivized promise $[u(U_i)]$. Write

$$H_X(k) = \dim\bigl(\mathbb{C}[X]_k\bigr) \tag{37}$$

for the degree-k Hilbert function of its homogeneous coordinate ring [8].

Corollary 7.1 (Algebraic oracle ceiling)**.** *For every physical general k-slot tester of a finite unitary promise in X,*

$$P_{\text{succ}}^{\text{GEN}} \leq \beta(\mathcal{M}_E, p) \leq \Pi_{D_E}(p) \leq \Pi_{H_X(k)}(p), \tag{38}$$

where $D_E = \dim \operatorname{span}\{u(U_i)^{\otimes k} : i = 1, \ldots, M\}$.

Proof. The tensors in equation (36) are the degree-k Veronese images of the projective oracle points. The linear span of the Veronese image of X has vector dimension $H_X(k)$. A finite subset can span no more, so $D_E \leq H_X(k)$. Apply theorem 5.1 and monotonicity of $\Pi_d(p)$. □

The exact finite-sample vector matroid is at least as informative as the Hilbert relaxation. Repeated points, finite collisions, or special flats can lower D_E or β without changing X. Conversely, $H_X(k) \geq M$ is not an attainability theorem. The access model also matters: controlled-U, calls to $U^\dagger$, postselected supermaps, or a different oracle primitive have different feature vectors and must be re-analyzed. We make no claim that the rank-one Hilbert argument extends unchanged to noisy mixed channels.

8 Exact qubit phase families

The Rado ceiling is not merely a finite isolated obstruction. We now solve a continuous-angle family with arbitrary ensemble size and an open chamber of priors. For $\theta \in \mathbb{R}/2\pi\mathbb{Z}$, let

$$D(\theta) = \mathrm{diag}(1, e^{i\theta}). \tag{39}$$

Choose $L \geq 2$ pairwise distinct phase channels $D(\theta_1), \ldots, D(\theta_L)$, including the anchors $D(0) = I$ and $D(\pi) = Z$. Add a set $C \subseteq \{X, Y\}$ of $c \in \{1, 2\}$ off-diagonal Pauli channels. All channels are distinct modulo global phase.

Every phase gate obeys

$$D(\theta) = \frac{1 + e^{i\theta}}{2} I + \frac{1 - e^{i\theta}}{2} Z. \tag{40}$$

Thus their Choi vectors lie in $\mathrm{span}\{|I\rangle\rangle, |Z\rangle\rangle\}$. Pairwise distinct phase channels determine distinct projective lines, so every two are independent and every three are dependent. The vectors $|X\rangle\rangle, |Y\rangle\rangle$ are mutually orthogonal and orthogonal to the diagonal plane.

Theorem 8.1 (Phase-flat family)**.** *The one-query vector matroid of the preceding ensemble is*

$$U_{2,L} \oplus U_{c,c}. \tag{41}$$

For arbitrary positive priors, write $p^{\mathrm{ph}}_{(1)} \geq p^{\mathrm{ph}}_{(2)} \geq \cdots$ for the sorted phase-gate priors. Every general one-slot tester obeys

$$P^{\mathrm{GEN}}_{\mathrm{succ}} \leq \beta = p^{\mathrm{ph}}_{(1)} + p^{\mathrm{ph}}_{(2)} + \sum_{Q \in C} p_Q. \tag{42}$$

If I and Z carry the two largest phase priors, the bound is attained by a fixed maximally entangled input and Pauli/Bell measurement. Therefore equation (42) *is the exact general-tester optimum throughout the open prior chamber*

$$p_I, p_Z > \max_{\theta \notin \{0,\pi\}} p_{D(\theta)}. \tag{43}$$

Proof. Equation (40) and pairwise distinctness give the uniform rank-two matroid $U_{2,L}$ on the phase labels. Each included off-diagonal Pauli vector adds an independent orthogonal coloop, proving equation (41). A maximum-weight basis must contain all coloops and the two largest-prior phase labels, which proves the upper bound through theorem 4.1.

For attainment, prepare $|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ and apply the unknown unitary to one half. The output is $|U\rangle\rangle/\sqrt{2}$. Measure in the normalized Pauli-Choi, or Bell, basis

$$|I\rangle\rangle/\sqrt{2}, \quad |Z\rangle\rangle/\sqrt{2}, \quad |X\rangle\rangle/\sqrt{2}, \quad |Y\rangle\rangle/\sqrt{2}.$$

Guess I, Z and every included coloop on their corresponding outcomes; never guess any other phase gate. If a Pauli outcome is not assigned to an included coloop, assign it arbitrarily to I or Z, making the decision rule complete. No promised state for a different assigned label has support on that unused outcome. The correct-label probabilities are one for I, Z and the coloops and zero for all other phase labels. Under equation (43), their prior mass is exactly the right side of equation (42). A physical parallel strategy has therefore attained the ceiling valid for every general tester. □

The strict chamber in equation (43) is sufficient, not necessary: ties can also be attained when a Bell-resolvable maximizing phase basis is chosen. The theorem deliberately states a simple open chamber with one fixed measurement. It does not assert that the same measurement is optimal when the two most likely phase gates are arbitrary nonorthogonal members of the flat.

The ceiling can also be attained without discarding any member of the dependent phase flat. The equal-prior trine measurement used below is a standard symmetric-state construction [9, 10]; trine discrimination with arbitrary priors is analyzed in [11]. Our claim is the exact match with the Rado–tester ceiling after adjoining a process direction orthogonal to the trine support, not novelty of the trine measurement itself.

Corollary 8.2 (Trine phases with a coloop)**.** *Let*

$$\theta_j = \frac{2\pi j}{3}, \quad j = 0, 1, 2, \qquad \textit{and include the channel } X. \tag{44}$$

Give the three phase gates equal prior a and X prior b, where

$$3a + b = 1, \qquad 0 < b < a. \tag{45}$$

Then

$$P_{\mathrm{succ}}^{\mathrm{GEN}} = 2a + b = 1 - a, \tag{46}$$

while the total-rank top-prior relaxation is $3a$. Their gap is $a - b > 0$. There is an attaining POVM with a nonzero effect for every hypothesis.

Proof. The matroid is $U_{2,3} \oplus U_{1,1}$, so theorem 8.1 gives the upper bound $2a + b$. Let

$$|\phi_j\rangle = \frac{|00\rangle + e^{2\pi i j/3}|11\rangle}{\sqrt{2}}, \qquad P_{\mathrm{diag}} = |00\rangle\langle 00| + |11\rangle\langle 11|. \tag{47}$$

The cubic roots of unity sum to zero, hence

$$\sum_{j=0}^{2} |\phi_j\rangle\langle\phi_j| = \frac{3}{2} P_{\mathrm{diag}}. \tag{48}$$

Define

$$M_j = \frac{2}{3}|\phi_j\rangle\langle\phi_j| \quad (j = 0, 1, 2), \qquad M_X = \frac{1}{2}|X\rangle\!\rangle\langle\!\langle X| + \frac{1}{2}|Y\rangle\!\rangle\langle\!\langle Y|. \tag{49}$$

The first three effects sum to P_{diag}; the last is the projector onto the off-diagonal subspace. Thus they form a complete POVM on the normalized Choi output. Each phase gate is guessed correctly with conditional probability $2/3$, while X is guessed correctly with probability one. The success probability is therefore $3a(2/3) + b = 2a + b$, attaining the general tester ceiling. Every effect in equation (49) is nonzero.

The total process-vector rank is three. Since $a > b$, its dimension-only relaxation selects the three phase priors and equals $3a$. Subtracting equation (46) gives $3a - (2a + b) = a - b > 0$. ☐

The matroid bound strictly improves the top-dimension relaxation whenever the globally largest $2 + c$ priors contain more than two phase labels and an omitted coloop prior is strictly smaller than the extra phase prior that displaces it. The smallest example using both coloops is

$$I, \; Z, \; S = \mathrm{diag}(1, i), \; X, \; Y, \qquad p = (.30, .25, .20, .15, .10). \tag{50}$$

The matroid is $U_{2,3} \oplus U_{2,2}$. Its maximum-weight basis is $\{I, Z, X, Y\}$, so $\beta = .80$. The total rank is four and the coarser top-four-prior relaxation is

$$\Pi_4(p) = .90. \tag{51}$$

The Bell tester in the proof has correct-label vector $s = (1, 1, 0, 1, 1)$ and succeeds with probability .80. Hence

$$P_{\mathrm{succ}}^{\mathrm{parallel}} = P_{\mathrm{succ}}^{\mathrm{GEN}} = .80 \quad \text{for equation (50)}. \tag{52}$$

This is exact one-slot GEN optimality by sandwiching, not evidence about a causal-order advantage or its absence in multi-slot tasks.

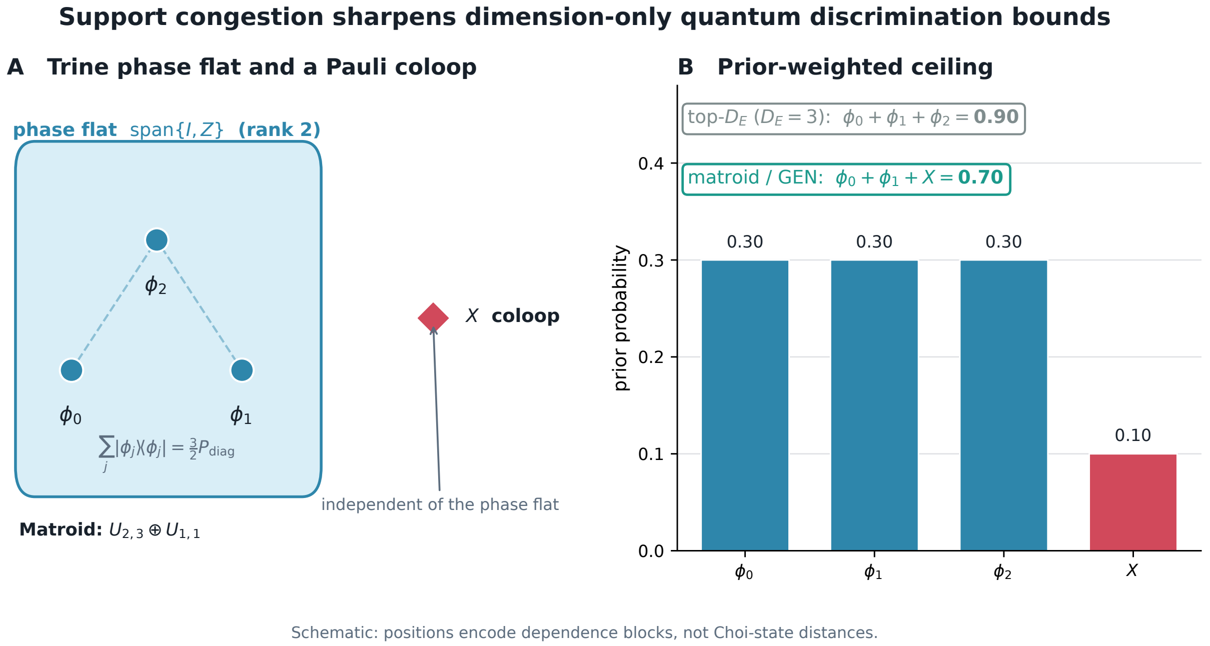

Figure 1: The tight trine-plus-coloop instance. Panel A records only the dependence blocks: three phase hypotheses form a rank-two flat and the Pauli X channel is a coloop. Panel B compares the total-span top-D_E relaxation with the exact Rado ceiling at $(a, a, a, b) = (.30, .30, .30, .10)$. The canonical trine POVM, completed on the orthogonal Choi direction and assigned to the X decision, attains .70; all four hypothesis effects are nonzero.

9 Mixed processes: scope and a counterexample to tightness

The main theorem includes mixed process operators through their support subspaces. This extension is exact as an upper bound, but support data can be too coarse to determine the optimum.

Example 9.1 (Identical full-rank states). Take three identical qubit states $\rho_1 = \rho_2 = \rho_3 = I_2/2$ with priors $(1/2, 3/10, 1/5)$. Every nonempty support is $\mathbb{C}^2$. Formula equation (10) gives the uniform matroid $U_{2,3}$, hence

$$\beta = 1/2 + 3/10 = 4/5.$$

However, for any POVM,

$$\sum_i p_i \operatorname{Tr}(M_i \rho_i) \le p_{\max} \sum_i \operatorname{Tr}(M_i I_2/2) = p_{\max} = 1/2,$$

and always guessing label 1 attains 1/2. Thus the Rado ceiling is not tight here.

The example also shows why a support-only theorem should not be advertised as an exact formula for mixed discrimination. Shah's mixed-state spectral bound [3] examines the eigenvalues of the weighted states $p_i \rho_i$. Here those eigenvalues are $1/4, 1/4, 3/20, 3/20, 1/10, 1/10$; the sum of the two largest is 1/2, exactly recovering the optimum. The two bounds retain different information: the Rado theorem records labeled support deficiencies and survives arbitrary tester normalization, whereas the spectral theorem can be sharper after a specific effective ensemble is known. Optimizing a spectral expression over all deterministic W is generally no longer a cheap promise-only screen.

Noisy full-rank processes often make the exact support matroid close to a uniform matroid, weakening equation (13). This is a limitation of the exact support certificate, not a failure of the proof. The robust theorem 4.3 avoids the resulting discontinuity whenever the process operators admit positive low-rank cores with small discarded tester mass. Conversely, low-rank mixed processes with heavily overlapping support subspaces can have a nonuniform Rado matroid and retain a useful labeled obstruction even without truncation.

10 Operational use, reproducibility, and limitations

The Rado ceiling is designed as a fail-closed pre-screen before a tester SDP. For rank-one hypotheses, sort labels by prior and perform incremental exact or certified rank tests on their process vectors. Mixed supports admit an equally explicit implementation. Let B_i be a basis matrix for $\mathcal{S}_i$, tag every column of B_i by label i, let $N = \sum_i \dim \mathcal{S}_i$, and let $R \leq \min\{N, \dim \mathcal{H}\}$ be the total linear rank. On the N tagged columns, intersect (i) their linear matroid with (ii) the partition matroid allowing at most one column of each label. By Rado's theorem, the maximum common-set size over labels A is exactly $r_E(A)$.

Greedy implementation. The following pseudocode computes the certificate without enumerating the 2^M subsets in equation (10).

1. Sort labels so that $p_1 \geq \cdots \geq p_M$; initialize $A \leftarrow \varnothing$, $r \leftarrow 0$, and $\beta \leftarrow 0$.
2. For $i = 1, \ldots, M$, compute by linear matroid intersection the largest common independent set among tagged columns with labels in $A \cup \{i\}$; call its size r'.
3. If $r' > r$, set $A \leftarrow A \cup \{i\}$, $r \leftarrow r'$, and $\beta \leftarrow \beta + p_i$; otherwise discard label i.
4. Return β and the retained independent transversal.

Cunningham's linear-matroid-intersection algorithm uses $O(NR^2 \log R)$ arithmetic operations for one such instance [12]. A straightforward prefix implementation therefore uses $O(MNR^2 \log R)$ arithmetic operations after sorting; warm starts can improve this practical cost but are not needed for the stated bound. In the rank-one case, incremental row reduction or QR gives the simpler $O(MR^2)$ arithmetic bound. Rational row reduction, finite-field lifting, or interval-certified singular-value tests can make the rank decisions fail-closed. The robust theorem is computed by the same routine on the chosen core supports, followed by the scalar correction in equation (24).

If the resulting ceiling is below a target success probability, no parallel, adaptive, or admissible indefinite-order tester in the chosen access model can meet the target. If it is high, no positive conclusion follows: a lower-bound strategy or an SDP is still required.

The accompanying exact-arithmetic script checks: canonical compression with a genuinely singular normalization and Moore–Penrose inverse; a mixed rank-two support family inducing $U_{2,3} \oplus U_{1,1}$; the robust $\beta + \eta$ fixture; positive Parseval scalability of a trine frame; all subset-rank inequalities and the strict separation equation (52); and the mixed-state non-tightness example. The script uses no floating-point thresholds. It is a regression fixture, not a substitute for the proofs or a general-purpose matroid-intersection implementation.

Several scope restrictions are essential.

1. W must belong to the physical deterministic tester or comb cone. An arbitrary positive operator normalized only on the listed hypotheses is not automatically implementable.
2. The Moore–Penrose support convention is required for singular W.
3. The Hilbert corollary assumes ordinary repeated access to the same unitary channel, with process tensor $J(U)^{\otimes k}$. Augmented oracle primitives change the feature matroid.
4. Matroid rank, total rank, and Hilbert function are upper-bound data, not attainability data. Perfect rank-one discrimination requires equation (35).

5. The mixed Rado theorem is support-based. Eigenvalues and overlaps can produce substantially stronger bounds. The robust theorem depends on a chosen positive core and on a valid upper bound for equation (23); it does not optimize that choice automatically.

6. The targeted priority search described below is not exhaustive, so no absolute first-discovery claim is made.

11 Related work and priority boundary

The optimization of quantum measurements for minimum-error discrimination goes back to Holevo and Yuen–Kennedy–Lax; see [13, 14, 15]. Quantum combs and testers provide the corresponding framework for transformations and multi-time networks [1]. Sedlak, Reitzner, Chiribella, and Ziman explicitly associate a canonical inverse-square-root POVM with a tester normalization [2]. The compression in theorem 2.1 uses that standard construction.

Bavaresco, Murao, and Quintino compare parallel, sequential, and general strategies for unitary channel discrimination and derive a uniform ambient dimension bound for general strategies [4]. Shah proves the sum-of-top-d-priors bound for pure-state communication and a sharp spectral-information bound for mixed states [3]. We therefore do not attribute novelty to the second inequality of equation (30).

Rado's independent-transversal theorem and Edmonds' submodular polyhedra are classical combinatorics [5, 6]. Their use here is not a new matroid theorem. The contribution specific to this paper is to recognize that the conditional correct-label vector generated by a quantum tester obeys the support inequalities equation (16), which become the Rado independence-polytope inequalities after singleton capping, and that the tester normalization acts as a common rank-nonincreasing linear map.

The symmetric trine POVM and square-root measurement are established constructions [9, 10], and analytic trine discrimination beyond equal priors is also known [11]. The trine in theorem 8.2 is therefore an exactness witness for the support-matroid bound, not a new state-discrimination protocol.

Targeted searches for combinations of "quantum state discrimination", "Rado matroid", "independence polytope", "subset rank bound", "quantum tester", and "nonuniform prior" did not locate this formulation. That negative search result is not proof of novelty; related work may use different language, especially in quantum decision theory or semidefinite packing. A safe priority statement is therefore that we derive the Rado-matroid tester ceiling by combining established tester normalization with established matroid theory.

12 Conclusion

A general quantum tester cannot assign correct-label probabilities arbitrarily. After exact canonical compression, every subset of hypotheses consumes at most the dimension of its joint effective support. Rado's rank formula turns these simultaneous support budgets and the elementary cap $s_i \leq 1$ into one matroid independence polytope. The Bayes objective is then bounded by a greedy maximum-weight independent transversal.

For rank-one processes, this retains the complete labeled linear-dependence pattern and can be strictly stronger than any ceiling based only on total dimension. The qubit phase-flat family demonstrates an exact open chamber, and its five-channel member has attained matroid value .80 while the total-rank relaxation is .90. Global algebraic spans supply a coarser Hilbert-function comparison. For mixed processes, the Rado construction remains valid but is deliberately presented as an upper bound rather than an exact formula. Positive core truncation makes that upper bound stable under full-rank admixture, with an explicit prior-weighted spectral-tail correction.

The practical outcome is a light, architecture-independent impossibility certificate that can precede causal-strategy optimization, together with a concrete matroid-intersection algorithm and an arithmetic complexity bound. The mathematical outcome is a bridge between quantum tester discrimination, independent transversals, and spectral-tail control. Further work could optimize the choice of positive cores, characterize when a physical deterministic normalization attains a matroid face, or combine the certificate with the full spectral geometry of effective states.

AI assistance. AI assistance was used for exploratory algebra, drafting, and consistency checks. The author is responsible for the statements, proofs, references, and final manuscript.

References

[1] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Theoretical framework for quantum networks. *Physical Review A*, 80:022339, 2009.

[2] Michal Sedlak, Daniel Reitzner, Giulio Chiribella, and Mario Ziman. Incompatible measurements on quantum causal networks. *Physical Review A*, 93:052323, 2016.

[3] Ronit Shah. Qudits offer no advantages over dits for sending random messages, 2025.

[4] Jessica Bavaresco, Mio Murao, and Marco Túlio Quintino. Unitary channel discrimination beyond group structures: Advantages of sequential and indefinite-causal-order strategies. *Journal of Mathematical Physics*, 63:042203, 2022.

[5] Richard Rado. A theorem on independence relations. *The Quarterly Journal of Mathematics*, os-13(1):83–89, 1942.

[6] Jack Edmonds. Submodular functions, matroids, and certain polyhedra. In Richard Guy, Haim Hanani, Norbert Sauer, and J. Schönheim, editors, *Combinatorial Structures and Their Applications*, pages 69–87, New York, 1970. Gordon and Breach.

[7] Lluis Eriksson. Algebraic query support for unitary oracles: Exact hilbert laws, harmonic spectra, and general-tester bounds. Archive for Rigorous Research, ARR-2026-6WX2JF38WE87GB2M, 2026.

[8] David Eisenbud. *Commutative Algebra with a View Toward Algebraic Geometry*, volume 150 of *Graduate Texts in Mathematics*. Springer, 1995.

[9] Masashi Ban, Keiko Kurokawa, Rei Momose, and Osamu Hirota. Optimum measurements for discrimination among symmetric quantum states and parameter estimation. *International Journal of Theoretical Physics*, 36(6):1269–1288, 1997.

[10] Yonina C. Eldar and G. David Forney, Jr. On quantum detection and the square-root measurement. *IEEE Transactions on Information Theory*, 47(3):858–872, 2001.

[11] Graeme Weir, Catherine Hughes, Stephen M. Barnett, and Sarah Croke. Optimal measurement strategies for the trine states with arbitrary prior probabilities. *Quantum Science and Technology*, 3(3):035003, 2018.

[12] William H. Cunningham. Improved bounds for matroid partition and intersection algorithms. *SIAM Journal on Computing*, 15(4):948–957, 1986.

[13] Horace P. Yuen, Robert S. Kennedy, and Melvin Lax. Optimum testing of multiple hypotheses in quantum detection theory. *IEEE Transactions on Information Theory*, 21(2):125–134, 1975.

[14] Carl W. Helstrom. *Quantum Detection and Estimation Theory*. Academic Press, New York, 1976.

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