INPUT INFORMATION AND THE LEARNING BOUNDARY

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning can fail because the input is insufficient, even when representation, optimization, and computation pose no obstacle. We study the corresponding information-theoretic learning boundary, which is the threshold at which the observations themselves become sufficient. Exact recovery makes this boundary observable by requiring an optimal learner to identify a finite structured target exactly. We show that exact recovery depends on two distinct properties of the input. The observations must distinguish the candidate set in total, and decisive evidence must reach every hidden coordinate. In forced-satisfiable planted k-SAT, we characterize this boundary at window precision. At $m = \frac{n(2^k-1)}{k}(\ln n + x + o(1))$ clauses, the complete version space is with high probability a random subcube, and its number of unresolved coordinates converges to a Poisson random variable with mean e^{-x}. This geometry separates two notions that are often conflated in learning. The data may permit an optimal learner to recover the planted target before they determine it uniquely. The limiting uniqueness probability is $\exp(-e^{-x})$, whereas Bayes-optimal and minimax exact recovery converge to $\exp(-e^{-x}/2)$. At every fixed success level, uniqueness therefore requires asymptotically $\ln 2\,(2^k - 1)n/k$ more clauses. The same window law persists under bounded non-uniform coordinate coverage. Equal marginal coverage minimizes the occupancy clock exactly at every finite sample size, and uniform weights attain it within the conditional product family. Calibrations across five further learning families and finite-size experiments show when total information predicts the boundary and when coordinate coverage becomes decisive. Together, the theory and experiments show that the learning boundary for exact recovery depends on both the amount of input information and its allocation across hidden coordinates.

1 INTRODUCTION

Learning depends not only on the learner but also on the information supplied by the input. When that information leaves essential alternatives indistinguishable, even an optimal procedure with unlimited computation cannot recover the target. Fano's inequality translates such indistinguishability into information-theoretic lower bounds (Fano & Hawkins, 1961). Recent recovery theory examines how noise and heterogeneity in the observation channel affect whether the input contains enough information to identify the target (Nikolakakis et al., 2021; Pilliat, 2025). These works establish sharp limits in particular models, but sample size or aggregate information alone does not reveal which property of the input makes learning possible. In structured problems, exact recovery can also depend on whether decisive evidence reaches every part of the target. We therefore study which properties of input information determine the information-theoretic learning boundary. This boundary describes when the observations themselves become sufficient for an optimal learner to recover the target exactly.

In constraint-based learning problems, observational sufficiency can be formalized through the version space left by the input. For a hidden target chosen from a finite structured family, this space contains all candidates compatible with the observations. This view is classical in concept learning and query-based identification (Mitchell, 1982; Angluin, 1988). Recovery asks whether an optimal learner can identify the target, whereas uniqueness asks whether the version space contains only one candidate. The boundaries can share the same first-order location while differing inside the critical window.

In locally observed structured problems, locating the recovery and uniqueness boundaries requires separating the amount of input information from its allocation across the target. *Total information* asks whether the information supplied per observation, accumulated over the sample, resolves the

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

entropy of the candidate set, as formalized by Fano's inequality. *Coverage* asks whether every hidden coordinate receives decisive evidence. When each observation provides evidence about only a small group of coordinates, achieving full coverage becomes a coupon-collection problem (Schilling & Henze, 2021). Consequently, channels with equal per-observation information can require different numbers of observations for exact recovery when they allocate evidence differently across coordinates.

Planted k-SAT makes the difference between total information and coverage sharp. A hidden assignment $\sigma \in \{0,1\}^n$ is planted, and every observed clause is sampled uniformly from those satisfied by σ. Earlier work studies identifiability, algorithms, and solution-space geometry in planted satisfiability and related planted-CSP and random-subcube models (Feldman et al., 2015; Flaxman, 2008; Mora & Zdeborová, 2008; Xu & Zhang, 2022). Total-information accounting on the nearest alternatives yields only $\Omega(n)$ clauses, whereas covering all n single-coordinate flips introduces an additional $\ln n$ factor. For $k = 3$, Motoki & Uehara (1999) established the first-order $\frac{7}{3} n \ln n$ uniqueness threshold. The corresponding critical-window law, complete solution-space geometry, and relation between recovery and uniqueness remained unresolved. App. A gives the full comparison with planted-CSP and sharp exact-recovery work.

Our main theorem characterizes the full critical window of the planted k-SAT boundary. At $m = \frac{n(2^k-1)}{k}(\ln n + x + o(1))$ clauses, all satisfying assignments form exactly the subcube generated by the surviving single flips with probability $1 - O_{k,x}(\ln n/n)$. The observations determine a unique assignment with limiting probability $\exp(-e^{-x})$, whereas Bayes-optimal and minimax exact recovery converge to $\exp(-e^{-x}/2)$. At every fixed success level, uniqueness therefore requires asymptotically $\ln 2\,(2^k-1)n/k$ more clauses. The same window law persists under bounded non-uniform coordinate coverage. It is governed by the occupancy clock $\Lambda_m = \sum_i (1-q_i)^m$. This extension connects non-uniform planted formulas (Friedrich et al., 2021) with group coupon-collection limits (Schilling & Henze, 2021), while also identifying the complete posterior geometry and both recovery curves.

Finite-size experiments support the theoretical predictions without fitted constants. In planted k-SAT, the empirical W-statistics track their exact finite-n values and the predicted Poisson law, while non-uniform allocations collapse onto the same curves when indexed by Λ_m. In the measured CIFAR-10H annotation channel, an item-specific prediction tracks the persistent error floor, whereas a homogeneous summary based on the pooled 5.14% error rate misses it by four orders of magnitude (Peterson et al., 2019). Adding image evidence lowers the error but leaves the coordinate-coverage signature visible. These experiments test finite-size consequences of the SAT theorem and calibrate how total information and coverage behave outside its assumptions. Our main results are summarized as follows.

- We characterize the learning boundary in planted k-SAT at critical-window precision. The posterior forms a Poisson-dimensional subcube, whose geometry separates the uniqueness and optimal-recovery curves by an explicit clause budget. The window law extends to non-uniform coordinate coverage, where equal marginal coverage exactly minimizes the occupancy clock at every finite n.
- We formalize total information and coordinate coverage as two necessity principles built from classical ingredients. Together they diagnose which input requirement is binding at the exact-recovery boundary.
- We calibrate the two requirements across five additional learning families. Finite-size experiments test the resulting predictions, while measured-channel diagnostics show when scalar and coordinate-aware summaries disagree.

2 Information Requirements for Exact Recovery

Exact recovery can fail because the observations contain too little information in total or because that information does not reach every part of the target, and we formulate these two obstructions as complementary necessity principles. A finite-target problem consists of a finite candidate set Θ, observation laws $\{P_\theta : \theta \in \Theta\}$, and a prior π. A target $\theta^* \sim \pi$ is drawn, and the learner observes $Z_{1:m}$ and returns $\hat{\theta}$. Its prior-average success is $\mathbb{P}(\hat{\theta} = \theta^*)$, Bayes success is the supremum of this quantity over estimators, and minimax success is $\sup_{\hat{\theta}} \min_\theta \mathbb{P}_\theta(\hat{\theta} = \theta)$. Our converses are prior-average statements unless minimax is stated explicitly. In families below the observation model is equivariant under a transitive group action and the prior is invariant, so symmetrizing any estimator

makes its per-target success constant and equal to its prior-average success, and Bayes and minimax values coincide. An arbitrary unsymmetrized estimator need not have constant per-target error.

2.1 N1: Total Information

The first requirement compares the number of distinguishable alternatives with the information supplied by one observation. Applying Fano's inequality to alternatives at a fixed Hamming distance isolates the locally difficult candidates from distant candidates that are easier to distinguish.

Theorem 2.1 (Localized Fano bound). *Fix θ^* and let $\mathcal{A}_r$ contain N_r alternatives, each at distance r from θ^*. Let Y be uniform on $\{\theta^*\} \cup \mathcal{A}_r$, and let $I_r = I(Y; Z_1)$ for observations that are i.i.d. conditional on Y. If an estimator identifies Y from $Z_{1:m}$ with error at most P_e, then*

$$mI_r \geq (1 - P_e) \ln(N_r + 1) - \ln 2.$$

This is a Bayes statement under the prior supported on $\{\theta^*\} \cup \mathcal{A}_r$, so a minimax procedure whose error is at most P_e on every candidate must satisfy it for every choice of $\mathcal{A}_r$, while a per-target conclusion holds for the symmetrized estimator under the transitive-equivariance assumption above.

The complete proof is in App. C.1. N1 is order-tight when difficulty is a single global event, as in the erasure family, where one observation reveals the entire target with probability ε and otherwise nothing, so recovery needs $\Theta(1/\varepsilon)$ samples whatever the number of candidates, but the same scalar accounting can be incomplete when the evidence must be spread over coordinates.

In the planted k-SAT model of Sec. 4, let Y be uniform on the planted assignment and its n single-coordinate flips. Then $I(Y; Z_1) = (1 + o(1))q_1 \ln n$, where $q_1 = k/[n(2^k - 1)]$ is the probability that a given single flip violates one clause (Lem. C.1 in App. C.2), so Thm. 2.1 certifies only $m = \Omega(1/q_1) = \Omega(n)$, whereas the true exact-recovery scale is $\Theta(n \log n)$. The missing logarithm is not missing entropy. It is the price of separately excluding all n single-coordinate alternatives.

2.2 N2: Coordinate Coverage

Coordinate coverage addresses the obstruction left unresolved by N1. It yields a converse either when an uncovered coordinate creates an equal-likelihood alternative, as in planted k-SAT, or when the posterior factorizes across coordinates, as in noisy indexing.

Theorem 2.2 (Coordinate-coverage necessity). *For every candidate θ and coordinate i, let $E_i(\theta)$ be the event that coordinate i is uncovered, receiving no decisive evidence. The event may depend on θ, the observations, and an auxiliary variable V that is ancillary given the observations, meaning the conditional law of V given $Z_{1:m}$ is the same under every θ. Set $K_\theta = \sum_i \mathbf{1}\{E_i(\theta)\}$. Suppose that*

$$\mathbb{E}[K_{\theta^*} \mid V] \xrightarrow{p} \infty, \qquad \frac{\mathrm{Var}(K_{\theta^*} \mid V)}{\mathbb{E}[K_{\theta^*} \mid V]^2} \xrightarrow{p} 0.$$

Then $K_{\theta^*} \xrightarrow{p} \infty$. Suppose additionally that, for every candidate θ and every realization, its K_θ uncovered coordinates supply K_θ pairwise distinct alternatives $\theta' \neq \theta$, each with likelihood equal to that of θ. Under the uniform prior, Bayes success, and hence the average success of every estimator, is at most $\mathbb{E}[1/(K_{\theta^*} + 1)] \to 0$.

The complete proof, including the every-candidate linkage required by the posterior argument, is in App. C.3. App. C.4 verifies the hypotheses of Thm. 2.2 in planted k-SAT, and App. C.5 records two boundary cases of these hypotheses.

When the prior is a product over coordinates, an auxiliary variable V has a law that does not depend on the target, and conditional on V the observations split into coordinate blocks that are independent given the target, with block i depending on the target only through its coordinate i, the posterior given V factorizes by coordinate. With $p_i(V)$ the conditional error probability given V of the Bayes decision on coordinate i, Bayes success is then at most $\mathbb{E}_V \exp(-\sum_i p_i(V))$, so $\sum_i p_i(V) \xrightarrow{p} \infty$ forces vanishing Bayes success (Cor. C.4 in App. C.6). Together, N1 compares the information supplied by the observations with the uncertainty over the candidate set, and N2 shows how unresolved coordinates can block exact recovery after that bound is met, which is why total information sets the erasure scale while planted k-SAT needs the additional $\ln n$ factor that Lem. C.1 cannot see.

Table 1: Input-information requirements against achievable exact-recovery scales, with lower and upper bounds matching at the order or window precision shown (App. C.7).

Family	Single-sample information	Recovery requirement	Exact-recovery scale
Hidden pixel	$\le \ln 2 - h_2(\eta)$ (MI)	Total information	$\Theta(\ln d/(\ln 2 - h_2(\eta)))$
One-edge Ising	$\Theta(\theta^2)$ (KL)	Total information	$\Theta(\theta^{-2} \log p)$
Planted k-SAT	$(1 + o(1))q_1 \ln n$ (MI)	Coordinate coverage	$(2^k - 1)n(\ln n + O(1))/k$
One-hidden-unit RBM	$\Theta(w^4)$ (KL)	Total information	$\Theta(w^{-4} \log p)$
Noisy indexing	$\Theta(n^{-1})$ (MI)	Coordinate coverage	$(1 + o(1))n \ln n/a$
Erasure	$\varepsilon \ln \lvert\Theta\rvert$ (MI)	Total information	$\Theta(1/\varepsilon)$

Here d is number of candidate pixel locations, p number of visible variables, η BSC crossover probability, θ Ising coupling, w RBM weight, ε erasure reveal probability, h_2 binary entropy in nats, and $a = 1 - 2\sqrt{\eta(1-\eta)}$ noisy-indexing rate. MI is $I(Y; Z_1)$ under the N1 candidate prior, uniform on the target and its n single-coordinate alternatives for SAT and indexing, while KL is the divergence to the independent visible law, which bounds MI for Ising and RBM under weak coupling. Noisy indexing assumes fixed $0 < \eta < \frac{1}{2}$.

3 Exact-Recovery Regimes Across Learning Problems

To locate the learning boundary, a necessary information requirement must be compared with an achievable recovery guarantee, which Tab. 1 does for hidden pixel, one-edge Ising, planted k-SAT, one-hidden-unit RBM, and noisy indexing, with erasure as a boundary case. The matching bounds separate three sources of sample cost, the number of alternatives to be distinguished, the evidence one sample supplies, and whether that evidence reaches every coordinate.

The total-information bound N1 matches the recovery order for hidden pixel, one-edge Ising, one-hidden-unit RBM, and erasure. Coordinate coverage becomes informative when observations must resolve separate coordinates, but it need not change the order, since in noisy indexing N1 already captures $n \ln n$ and the coordinate-level analysis fixes only the constant, whereas in planted k-SAT N1 on the single-coordinate alternatives gives $\Omega(n)$ against the $\Theta(n \ln n)$ clauses needed to exclude all single flips. A logarithmic factor alone therefore does not identify its source.

The second source, per-sample evidence, is isolated by comparing direct and hidden interactions over the same candidate pairs of visible variables. Marginalizing a hidden unit changes the observable evidence without changing which pairs must be distinguished. For $p \to \infty$ visible binary units and one hidden unit connected with weight $w = w_p \to 0$ to an unknown pair, under the uniform prior over the $\binom{p}{2}$ pairs, every exact-recovery procedure requires $m = \Omega(w^{-4} \log p)$ samples and correlation thresholding succeeds with $m = O(w^{-4} \log p)$. The reason is that marginalizing the hidden unit produces the exact visible coupling $J(w) = \frac{1}{2} \ln \cosh(2w) = \Theta(w^2)$, whose per-sample KL supply is $\Theta(J(w)^2) = \Theta(w^4)$ (Prop. C.6 in App. C.8, with constants uniform for $|w|$ below a small bound). The increased sample cost comes from weaker observable evidence rather than greater uncertainty over candidate pairs, and the conclusion concerns this minimal family, not general RBM learning.

Per-sample evidence can also be concentrated in rare decisive observations, as in the erasure family, where every sample reveals the entire target with probability ε or nothing. There the uncovered events of all coordinates coincide, N2 does not apply, and N1 alone gives the scale $\Theta(1/\varepsilon)$ for arbitrarily many candidates (App. C.5). Planted k-SAT instead eliminates single-coordinate alternatives through distinct observations, and its coverage scale does not yet describe all assignments compatible with the input inside the transition, so we turn to the full version-space geometry to determine the uniqueness and recovery curves and to quantify how reallocating evidence moves the boundary.

4 A Sharp Exact-Recovery Window in Planted k-SAT

4.1 Model and Window

Fix $k \ge 2$ and draw a hidden assignment σ uniformly from $\{0,1\}^n$. A clause is a scope $S \in \binom{[n]}{k}$ and a forbidden pattern $f \ne \sigma|_S$. We observe m clauses drawn uniformly with replacement from the $N = (2^k - 1)\binom{n}{k}$ clauses satisfied by σ, the forced-satisfiable uniform model of Motoki & Uehara (1999). The uniform draw is the prior for the Bayes statements, and the equivariance of Thm. 4.2 identifies the minimax value and lets every structural statement be proved conditional on a fixed σ.

For $A \subseteq [n]$, let $\tau_A = \sigma \oplus \mathbf{1}_A$. Encoding a clause by its difference pattern $D = \{j \in S : f_j \neq \sigma_j\}$ determines exactly which alternatives violate it,

$$\tau_A \text{ violates } (S, f) \quad \Longleftrightarrow \quad A \cap S = D. \tag{1}$$

Let $\mathcal{S}$ be the complete set of satisfying assignments, let F contain the coordinates whose single flips survive all clauses, and write $W = |F|$ for the number of unresolved coordinates. A given single flip violates one clause with probability $q_1 = k/[n(2^k - 1)]$. Let $\widetilde{q}_1 = -\ln(1 - q_1)$ and choose any integer sequence satisfying

$$m\widetilde{q}_1 - \ln n \longrightarrow x \in \mathbb{R}.$$

Each single flip survives one clause with probability $1 - q_1$, so $\mathbb{P}(i \in F) = (1 - q_1)^m$ and $\mathbb{E}W = n(1 - q_1)^m = \exp(-(m\widetilde{q}_1 - \ln n)) \to e^{-x}$, which centres the window where the expected number of unresolved coordinates is order one. Every assignment consistent with the observations has the same likelihood, so the posterior is uniform on $\mathcal{S}$. The exact joint survival probabilities of several flips and of alternatives at any Hamming distance are stated in Lem. C.7 in App. C.9.

4.2 Subcube Geometry, Uniqueness, and Recovery

Sec. 4.1 located the window through W, but uniqueness and optimal recovery concern the version space rather than this count. Uniqueness is the event $|\mathcal{S}| = 1$, and since the posterior is uniform on $\mathcal{S}$, the Bayes rule guesses uniformly on $\mathcal{S}$ and succeeds with probability $\mathbb{E}[1/|\mathcal{S}|]$. The count W does not determine $\mathcal{S}$, because a satisfying assignment at Hamming distance two or larger need not have any surviving single flip, and the first result shows that inside the window no such assignment exists.

Theorem 4.1 (Subcube structure)**.** *For fixed $k \geq 2$ and $x \in \mathbb{R}$,*

$$\mathbb{P}(\mathcal{S} = \{\sigma \oplus \mathbf{1}_B : B \subseteq F\}) = 1 - O_{k,x}\left(\frac{\ln n}{n}\right), \qquad W \Rightarrow \text{Poisson}(e^{-x}).$$

Thus the complete version space is, with high probability, exactly a random subcube generated by the coordinates that the input did not individually rule out.

Thm. 4.1 excludes all satisfying assignments outside this subcube simultaneously and therefore identifies the complete posterior geometry, not only the number of unresolved coordinates (proof in App. C.11). On the structure event $|\mathcal{S}| = 2^W$, so uniqueness is the event $W = 0$, Bayes success is $\mathbb{E}[2^{-W}]$ up to the probability that the structure event fails, and the Poisson law for W gives the uniqueness window $\mathbb{P}(|\mathcal{S}| = 1) \to \exp(-e^{-x})$ with $\log_2 |\mathcal{S}| \Rightarrow \text{Poisson}(e^{-x})$ (Cor. C.11).

Theorem 4.2 (Bayes and minimax recovery)**.** *Under the uniform prior on σ, Bayes exact-recovery success satisfies*

$$\text{Succ}_{\text{Bayes}} \longrightarrow \exp(-e^{-x}/2).$$

The model is equivariant under a transitive $(\mathbb{Z}_2)^n$ action, so minimax success has the same limit. Every estimator's uniform-prior average success is at most this value plus $o(1)$. More generally, if $m\widetilde{q}_1 - \ln n \to -\infty$, every estimator's prior-average success tends to zero, and Bayes/minimax success tends to one if and only if $m\widetilde{q}_1 - \ln n \to +\infty$.

The complete proof, including all quantifier and symmetrization statements, is in App. C.13. The first-order exact-recovery boundary is therefore $m_\star = (1 + o(1))\frac{n(2^k-1)}{k} \ln n$, whose constant for $k = 3$ is the uniqueness constant of Motoki & Uehara (1999), the contribution here being the critical-window law and the geometry that produces it. N1 certifies only $\Omega(n)$ clauses here (Sec. 2.1), so the factor $\ln n$ in $m_\star$ is the cost of covering every coordinate, and below the window $\mathbb{E}W$ diverges although that scale has long been passed, the threshold being set by where the information lands rather than by the entropy of the candidate set.

The two limiting curves answer different questions. Uniqueness discards every realization with $W > 0$, whereas recovery assigns those same realizations the graded value 2^{-W}. Hence

$$U(x) = e^{-e^{-x}}, \qquad R(x) = e^{-e^{-x}/2} = U(x + \ln 2), \tag{2}$$

and their horizontal separation is

$$\Delta m = \frac{\ln 2}{\widetilde{q}_1} = (1 + o(1)) \ln 2 \, \frac{(2^k - 1)n}{k}.$$

Requiring observations to determine a unique answer therefore costs $\Theta(n)$ more clauses than requiring an optimal learner to recover the planted answer, though both boundaries share the $\Theta(n \log n)$ first-order scale, a separation visible only at window precision. The gap is not an optimization advantage but the success the posterior permits after the input has left a finite number of coordinates unresolved.

4.3 Non-Uniform Coordinate Coverage

We now let the input allocate its evidence unevenly over scopes and ask which parts of the boundary move. Scopes are drawn from a conditional product law, $\mathbb{P}(S) \propto \prod_{i \in S} w_i$ over $S \in \binom{[n]}{k}$ with fixed weights $0 < c \le w_i \le C$, the forbidden pattern stays uniform among the $2^k - 1$ non-planted ones, and $w_i \equiv 1$ recovers the uniform model. The scope inclusion marginal $\xi_i = \mathbb{P}(i \in S)$ has a closed form in elementary symmetric polynomials (App. C.14), and $q_i = \xi_i/(2^k - 1)$ is the probability that the single flip of coordinate i violates one clause. The per-clause single-flip discrimination budget is conserved exactly, with $\sum_i q_i = k/(2^k - 1)$ under every scope law, and so is the per-clause information $I(\sigma; Z_1)$, so a non-uniform allocation redistributes a fixed budget and moves the joint information $I(\sigma; Z_{1:m}) = n \ln 2 - \mathbb{E} \ln |\mathcal{S}|$ instead (App. C.15). The expected number of unresolved coordinates is now $\Lambda_m = \sum_i (1 - q_i)^m$, which replaces $n(1 - q_1)^m$ as the occupancy clock.

Theorem 4.3 (Heterogeneous coverage window)**.** *Fix $k \ge 2$ and $0 < c \le C$, assume $c \le w_i \le C$ for all i, and let $m = m(n)$ satisfy $\Lambda_m \to \mu \in (0, \infty)$. Then $W \Rightarrow \mathrm{Poisson}(\mu)$, the complete solution set equals the subcube generated by the surviving single flips with probability $1 - o(1)$, uniqueness satisfies $\mathbb{P}(|\mathcal{S}| = 1) \to e^{-\mu}$, and Bayes and minimax exact recovery converge to $e^{-\mu/2}$. If $\Lambda_m \to \infty$, every estimator's prior-average success tends to 0, and if $\Lambda_m \to 0$, uniqueness and recovery tend to 1.*

Every layer of the uniform theorem survives, and the only change is the clock, since Λ_m replaces $m\widetilde{q}_1 - \ln n$ as the window coordinate and the boundary is still written by coordinate coverage even when the coordinates are covered at different rates (proof in App. C.17). The lower bound on the weights cannot be dropped, because a coordinate with zero weight is never resolved, so $W \ge 1$ and recovery is capped at one half however many clauses are drawn (App. C.18). Within the bounded-weight regime, Λ_m alone locates the window, so comparing allocations reduces to comparing their clocks.

Corollary 4.4 (Equal coverage and the cost of skew)**.**

(a) For every scope law and every $m \ge 2$, $\mathbb{E}W = \Lambda_m \ge n(1 - \bar{q})^m$ with $\bar{q} = k/[n(2^k - 1)]$, with equality if and only if all q_i are equal. Consequently the uniform allocation reaches any window position with the fewest clauses, exactly and at every finite n.

(b) Within the conditional product family with $n > k$, all marginals are equal if and only if all weights are equal.

(c) For a two-block law with a fraction α of coordinates at weight $\beta \in (0, 1)$ and the rest at weight one, the window location and the uniqueness premium Δm are both inflated by the factor $(\bar{w}/\beta)(1 + o(1))$ with $\bar{w} = \alpha\beta + 1 - \alpha$.

Part (a) is exact, law-free, and finite-sample, and is the first-moment form of a stochastic-order optimality of equal coverage (Anceaume et al., 2016), part (b) makes the equality case identifiable inside our family, and part (c) quantifies the deficit where it has a closed form. The window-scale skew $\beta_n = 1 - \zeta/\ln n$, with its quadratic inflation of Λ_m, is treated in App. C.19 with the complete proof. Fig. 2 tests both the invariance and the price without fitted constants.

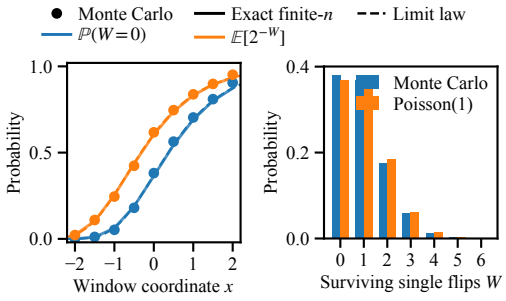
5 Experiments

Sec. 1 asked which properties of the input determine the learning boundary, and Secs. 2 to 4 answered with two requirements, total information and coordinate coverage, and with one window law in which coverage binds. The experiments put four questions to those answers.

RQ1. How accurately does the window geometry of Sec. 4 describe planted k-SAT at finite size?

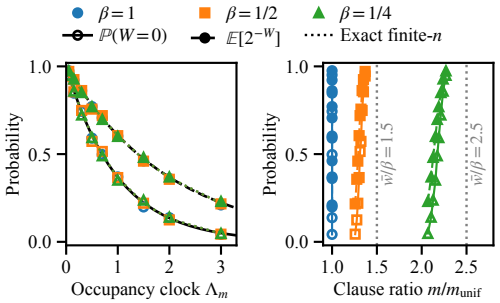
RQ2. How does the allocation of a conserved discrimination budget across coordinates change the number of clauses needed to reach a given window position?

RQ3. Across the learning families of Sec. 3, to what precision does total information alone locate the boundary?



(a) W-statistics. (b) W at $x = 0$.

Figure 1: The planted 3-SAT window at $n = 200$, 800 trials per point. Panel (a) compares Monte Carlo estimates of the uniqueness statistic $\mathbb{P}(W = 0)$ and the recovery statistic $\mathbb{E}[2^{-W}]$ with exact finite-n values and limit laws $\exp(-e^{-x})$ and $\exp(-e^{-x}/2)$. Panel (b) compares the law of W at $x = 0$ with Poisson(1).



(a) Clock collapse. (b) Clause premium.

Figure 2: The heterogeneous coverage window at $n = 200$ with two-block weights ($k = 3$, 800 trials per point, no fitted constants). Panel (a) plots both statistics against the exact Λ_m with the limits $e^{-\Lambda_m}$ and $e^{-\Lambda_m/2}$ in black. Panel (b) replots them against m/m_{unif} at matched Λ_m with the asymptotic ratios $\bar{w}/\beta$ marked.

RQ4. On a measured channel outside the theorems, does a coordinate-aware summary of the input predict the failures of measured decoders that a scalar summary misses?

5.1 RQ1: The SAT Window at Finite Size

RQ1 has two parts, how closely the two W-statistics and their separation follow the window law and how closely the full solution set matches the predicted subcube. For the first part, Fig. 1 evaluates the two statistics that become uniqueness and Bayes recovery on the subcube event. They track their exact finite-n formulas and the two distinct limit curves, while the empirical distribution of W matches the Poisson law, with the finite-n verification tabulated in App. D.1. Binary clauses with one forbidden pattern at $k = 2, 3, 4$ collapse onto the same curves after reparameterization by x (Fig. A.3 in App. F). These W-statistics alone would not establish that the entire solution set is a subcube.

For the second part, because the plotted statistics are only proxies off the structure event, we also enumerate all 2^{16} assignments. At $x = 0$ the true solution set differs from the predicted subcube in 39.5% of trials, and at $x = 2$ the structure-failure rate falls to 0.8%. Thus the exhaustive check exposes the finite-size gap instead of using the proxy to conceal it. Full true-curve results through $n = 22$ are in Apps. D.2 and F.

The rate sweep from $n = 50$ to 800 measures, at every size, a fixed detectable subset of the structure failures, namely two explicit modes defined in App. D. Their frequency is compared with a Poisson-approximation prediction assembled from their exact first moments, which for $k = 3$ at the window centre is asymptotic to $3 \ln n/n$. From $n = 100$ upward the observed frequencies agree with the prediction within Monte Carlo error, while $n = 50$ shows a finite-size discrepancy (Tab. A.6 and Fig. A.2 in App. F). The statistic is a detectable subset of the structure failures rather than their total probability, so the sweep supports the predicted $\ln n/n$ approach without being an empirical proof of the full $O_{k,x}(\ln n/n)$ theorem. The answer to RQ1 therefore has two layers. The W-statistics follow the window law at every size and every k tested, and the separation of their two curves is within one clause of the predicted budget Δm already at $n = 200$ (App. D.2). The true solution set matches the predicted subcube only up to a structure failure whose detectable part decays at the predicted rate, and at the sizes where the whole cube can be enumerated the true separation is still about a sixth short of Δm.

5.2 RQ2: Reallocating Clauses Across Coordinates

With the window law supported at finite size, RQ2 asks how the clause count needed to reach a given window position changes when the conserved per-clause discrimination budget of Sec. 4.3 is spread unevenly over the coordinates. Thm. 4.3 makes two finite-size predictions, that the W-statistics of different allocations collapse onto the same curves once plotted against Λ_m, and that a skewed allocation pays the predicted clause premium at matched window position. Fig. 2 tests both at

$n = 200$ with two-block weights, with m_{unif} the clause count of the uniform allocation at the same clock value. The collapse holds across the three skews tested with the exact finite-n curves overlaid, and an exhaustive $n = 16$ run of the full weighted model repeats the solution-set check of Sec. 5.1 rather than trusting the proxies (App. D.3). The answer to RQ2 separates an exact statement from a limit statement. The occupancy clock, and with it the clause count needed to reach any window position, is computed exactly at finite n, and at the target $\Lambda_m = 1$ the two skews tested pay 1.302 and 2.153 times the uniform clause count, against the asymptotic factors $\bar{w}/\beta = 1.5$ and 2.5 of Cor. 4.4, and by part (a) no unequal allocation needs fewer clauses. That matched clock values give matched recovery probabilities is the limit law of Thm. 4.3, and the exhaustive run shows the finite-size structure failure still differing between allocations at the same Λ_m.

5.3 RQ3: Total Information and Coverage Across Learning Families

RQ2 isolated the allocation of evidence inside one observation structure. For RQ3, Tab. 2 varies the observation structure itself across the families of Sec. 3 and pairs each prediction with a measurement from its own decoder, changing the target entropy at fixed channel quality (hidden pixel), the per-observation supply (Ising and RBM), the entropy under a whole-target reveal (erasure), the clause width (SAT), and the number of coordinates that must each be resolved (indexing).

Bounded supply. For hidden pixel, every observation is one noisy binary probe, so increasing the number of candidate positions raises the information demand while leaving the per-sample supply fixed, and the measured decoder cost follows $\log d$. The Ising and RBM experiments instead fix the candidate count and vary the channel strength. The visible Ising family is consistent with the predicted inverse-square dependence on θ (Santhanam & Wainwright, 2012) up to the search discretization and Monte Carlo fluctuation, while marginalizing one hidden unit replaces the visible coupling by $J(w) \asymp w^2$ and produces the inverse-fourth dependence of Prop. C.6, with the measured ratio against the exact identity $J(w) = \frac{1}{2} \ln \cosh 2w$ in Tab. 2.

Whole-target supply. The erasure family changes the number of candidates over four orders of magnitude while the recovery cost stays essentially flat, because a successful observation supplies the complete target and the learner pays for the waiting time to the first reveal rather than for $\log |\Theta|$ separate bits. This is the case where global information is sufficient and a coordinate-coverage formulation is inappropriate.

Table 2: Relative sample cost across the families of Sec. 3, the ratio of the reference scaling between the two endpoints of each parameter change against the ratio of the measured sample costs.

Family and change	Scaling	Ratio	
		theory	data
Total information			
Hidden pixel d, 8 to 1024	$\log d$	3.33	2.36
One-edge Ising θ, 0.35 to 0.14	θ^{-2}	6.25	5.98
One-hidden-unit RBM w, 0.5 to 0.25	$J(w)^{-2}$	13.0	13.2
Erasure $\lvert\Theta\rvert$, 10 to 10^5	constant	1.00	1.10
Coordinate coverage			
Planted k-SAT k, 2 to 4	$(2^k - 1)/k$	2.50	2.52
Noisy indexing n, 300 to 10^4	$n \ln n$	53.8	58.9

Ratios divide the final by the initial setting. Empirical costs use 90% decoder success for the first four rows, 50% exact recovery for noisy indexing, and the 50% crossing of the uniqueness proxy $\mathbb{P}(W = 0)$ for planted k-SAT at $n = 200$ (details in App. F).

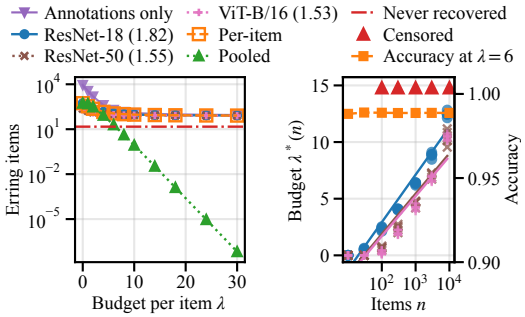
Coordinate coverage. Planted SAT and noisy indexing both require all coordinates to be resolved, but the observation metadata differ. An indexing sample names the coordinate it concerns, so the mixture information already contains the $n \ln n$ order and N2 sharpens its constant to $1/a$. The measured load at which exact recovery reaches one half rises from 0.73 to 0.80 of $n \ln n/a$ between $n = 300$ and 10^4, and the conjectural $\frac{1}{2} \ln \ln n$ correction of App. G accounts for more than half of the shortfall, its window constant for the rest (App. F). A SAT clause couples k coordinates without naming the alternative it rules out, so the localized Fano bound stops at $\Omega(n)$ and N2 supplies the missing logarithm, while the window location scales with k as the coverage calculation predicts. A scalar "information per sample" is therefore not a channel-invariant description of the learning boundary. The window law itself is tied to the one-forbidden-pattern hypothesis, and in planted Model RB (Xu & Li, 2000; Xu et al., 2007), where each constraint forbids $q = pd^k$ tuples over an alphabet of size $d = n^\alpha$, the same statistics are only proxies, approximately following the window law at weak clustering but stalling near $\mathbb{P}(W = 0) = 0.486$ at $(1+\alpha)p = 1.20$

instead of approaching e^{-1} (Fig. A.3 in App. F). That sweep is centred at the isolated-solution transition located by Xu & Zhang (2022), so nothing about its first-order location is ours.

The answer to RQ3 is that the precision of total information is read off the observation structure. In the bounded-supply families it fixes how the cost scales with the input, the demand $\log d$ or the supply θ^{-2} and w^{-4}, and the Ising and RBM ratios follow that scaling within the search-grid resolution, while the hidden-pixel cost is affine in $\log d$ and its endpoint ratio falls below the $\log d$ ratio (App. F). In erasure it predicts a cost independent of the candidate count, and the measured cost is flat. In planted k-SAT it falls short by the logarithm that coverage supplies, and in noisy indexing it fixes the order while coverage fixes the constant.

5.4 RQ4: A Measured Image-and-Annotation Channel

The families of RQ3 have known generating mechanisms, so the scaling of their recovery cost with the input could be predicted before it was measured. For RQ4, CIFAR-10H (Peterson et al., 2019) supplies a measured annotation channel whose mechanism is not identified, on which scalar and coordinate-aware summaries can still be compared on observed data, as a descriptive diagnostic rather than a theorem instance. Here λ denotes the number of annotations supplied per item, and $\lambda^*(n)$ the load at which the measured exact-recovery probability reaches one half on an n-item subset.



(a) Full pool. (b) Banded budget.

Figure 3: CIFAR-10H crowd annotations with image-assisted decoders. Panel (a) compares measured erring items for four decoders on the full pool at $n = 10^4$ with per-item and pooled predictions and marks items ResNet-18 never recovers. Panel (b) shows $\lambda^*(n)$ on the near-homogeneous band with fitted slopes per $\ln n$ in parentheses and whole-pool budgets censored at $\lambda = 30$.

Fig. 3(a) shows that a homogeneous channel built from the pooled 5.14% annotation error predicts essentially no errors at a large annotation budget, yet the measured error count plateaus because a small set of images has item-specific channels that point away from the gold label. A coordinate-aware prediction built on one half of the annotations tracks the other half, including this floor. Adding image evidence from ResNet-18, ImageNet-pretrained ResNet-50, or ImageNet-pretrained ViT-B/16 lowers error but does not remove the coverage signature of the annotation-only channel (Fig. A.8 in App. H). On the near-homogeneous band, the budget for exact recovery remains affine in $\ln n$ (Fig. 3(b)). The models correctly classify only 47 to 57 of the 86 annotation-hard items despite much higher accuracy on the remainder (Tab. A.7). The answer to RQ4 is that the per-item summary predicts where the measured decoders fail and the pooled summary does not, on the full pool through the error floor and on the near-homogeneous band through the $\ln n$ growth of the budget. Both describe the measured decoders rather than an algorithm-independent lower bound, and data splits, training details, controls, and limitations are in App. H.

6 Conclusion

We characterized a sharp exact-recovery window in planted k-SAT and used it to separate two demands on the input, enough total information to distinguish candidates and enough coverage to resolve every hidden coordinate. Inside the window the unresolved coordinates generate a Poisson-dimensional posterior subcube, and uniqueness costs $\Theta(n)$ clauses beyond Bayes/minimax recovery. Under bounded non-uniform allocation the same laws follow the occupancy clock Λ_m, so equal coverage reaches any clock value with the fewest clauses at every finite n. Exact-recovery sample complexity depends on how the input allocates its information across coordinates, not only on a scalar aggregate, as the calibrations across five further families and a measured annotation channel show. An exact-recovery boundary isolates failures of the input from those of the learner. Limitations and future work are in App. I.

AI USE STATEMENT

Generative AI tools were used to revise text.

ETHICS STATEMENT

Most experiments are synthetic. The real-channel study uses the public CIFAR-10H aggregate annotation release (Peterson et al., 2019), analyzes no annotator identity, and collects no new human-subject data. We are aware of no additional deployment-specific ethical concerns.

REPRODUCIBILITY STATEMENT

All formal claims are proved in App. C, and the exact finite-n SAT formulas are stated in App. D. App. E specifies every experiment at the level needed to rerun it, namely the generative model, the decoder, the parameter grid, the trial count, the seed, and the rule that turns trials into each reported number, and Apps. F and H report the outputs.

REFERENCES

Emmanuel Abbe, Afonso S Bandeira, and Georgina Hall. Exact recovery in the stochastic block model. *IEEE Transactions on information theory*, 62(1):471–487, 2015.

Ali Ahmadi, Kiarash Banihashem, Iman Gholami, Mohammad Taghi Hajiaghayi, and Jan Olkowski. Quiet planting for k-SAT, multiple solutions of arbitrary geometry. In *Proceedings of the 39th Annual Conference on Learning Theory (COLT)*, volume 336 of *Proceedings of Machine Learning Research*, pp. 75–105, 2026.

Animashree Anandkumar and Ragupathyraj Valluvan. Learning loopy graphical models with latent variables: Efficient methods and guarantees. *The Annals of Statistics*, pp. 401–435, 2013.

Emmanuelle Anceaume, Yann Busnel, and Bruno Sericola. New results on a generalized coupon collector problem using Markov chains. *Journal of Applied Probability*, 52(2):405–418, 2015.

Emmanuelle Anceaume, Yann Busnel, Ernst Schulte-Geers, and Bruno Sericola. Optimization results for a generalized coupon collector problem. *Journal of Applied Probability*, 53(2):622–629, 2016.

Dana Angluin. Queries and concept learning. *Machine learning*, 2(4):319–342, 1988.

George K Atia and Venkatesh Saligrama. Boolean compressed sensing and noisy group testing. *IEEE Transactions on Information Theory*, 58(3):1880–1901, 2012.

Raghu Raj Bahadur and R Ranga Rao. On deviations of the sample mean. *The Annals of Mathematical Statistics*, 31(4):1015–1027, 1960.

Guy Bresler and Rares-Darius Buhai. Learning restricted boltzmann machines with sparse latent variables. *Advances in Neural Information Processing Systems*, 33:7020–7030, 2020.

Luca Corinzia, Paolo Penna, Luca Mondada, and Joachim M Buhmann. Exact recovery for a family of community-detection generative models. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pp. 415–419. IEEE, 2019.

Robert M Fano and David Hawkins. Transmission of information: A statistical theory of communications. *American Journal of Physics*, 29(11):793–794, 1961.

Vitaly Feldman, Will Perkins, and Santosh Vempala. On the complexity of random satisfiability problems with planted solutions. In *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, pp. 77–86, 2015.

Abraham Flaxman. A spectral technique for random satisfiable 3cnf formulas. *Random Structures & Algorithms*, 32(4):519–534, 2008.

Tobias Friedrich, Frank Neumann, Ralf Rothenberger, and Andrew M Sutton. Solving non-uniform planted and filtered random sat formulas greedily. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 188–206. Springer, 2021.

Chao Gao, Yu Lu, and Dengyong Zhou. Exact exponent in optimal rates for crowdsourcing. In *International conference on machine learning*, pp. 603–611. PMLR, 2016.

Julia Gaudio and Nirmit Joshi. Community detection in the hypergraph sbm: Exact recovery given the similarity matrix. In *The Thirty Sixth Annual Conference on Learning Theory*, pp. 469–510. PMLR, 2023.

Hyeonsu Jeong and Hye Won Chung. Recovering top-two answers and confusion probability in multi-choice crowdsourcing. In *International Conference on Machine Learning*, pp. 14836–14868. PMLR, 2023.

Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pp. 286–295, 1983.

David R Karger, Sewoong Oh, and Devavrat Shah. Budget-optimal task allocation for reliable crowdsourcing systems. *Operations Research*, 62(1):1–24, 2014.

Tom M Mitchell. Generalization as search. *Artificial intelligence*, 18(2):203–226, 1982.

Thierry Mora and Lenka Zdeborová. Random subcubes as a toy model for constraint satisfaction problems. *Journal of Statistical Physics*, 131(6):1121–1138, 2008.

Mitsuo Motoki and Ryuhei Uehara. *Unique solution instance generation for the 3-Satisfiability (3SAT) problem*. Tokyo Institute of Technology. Department of Information Sciences, 1999.

Konstantinos E Nikolakakis, Dionysios S Kalogerias, and Anand D Sarwate. Predictive learning on hidden tree-structured ising models. *Journal of Machine Learning Research*, 22(59):1–82, 2021.

Joshua C Peterson, Ruairidh M Battleday, Thomas L Griffiths, and Olga Russakovsky. Human uncertainty makes classification more robust. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9617–9626, 2019.

Emmanuel Pilliat. Recovering labels from crowdsourced data: an optimal and polynomial-time method. In *Proceedings of the Thirty Eighth Conference on Learning Theory (COLT)*, volume 291 of *Proceedings of Machine Learning Research*, pp. 4565–4595, 2025.

Narayana P Santhanam and Martin J Wainwright. Information-theoretic limits of selecting binary graphical models in high dimensions. *IEEE Transactions on Information Theory*, 58(7):4117–4134, 2012.

Judith Schilling and Norbert Henze. Two poisson limit theorems for the coupon collector's problem with group drawings. *Journal of Applied Probability*, 58(4):966–977, 2021.

Ke Xu and Wei Li. Exact phase transitions in random constraint satisfaction problems. *Journal of Artificial Intelligence Research*, 12:93–103, 2000.

Ke Xu, Frédéric Boussemart, Fred Hemery, and Christophe Lecoutre. Random constraint satisfaction: Easy generation of hard (satisfiable) instances. *Artificial intelligence*, 171(8-9):514–534, 2007.

Wei Xu and Zhe Zhang. The solution space structure of planted constraint satisfaction problems with growing domains. *Journal of Statistical Mechanics: Theory and Experiment*, 2022(3):033401, 2022.

Chunyan Zhao, Ke Xu, and Zhiming Zheng. On the scaling window of model rb. *arXiv preprint arXiv:0801.3871*, 2008.

APPENDIX OVERVIEW

- **A. Related Work** 13
- **B. Notation and Conventions** 14
- **C. Proofs** 14
  - *C.1. Proof of Theorem 2.1 (Localized Fano Bound)* 14
  - *C.2. Mixture Information in Planted k-SAT (Lemma C.1)* 15
  - *C.3. Proof of Theorem 2.2 (Coordinate-Coverage Necessity)* 15
  - *C.4. Verification of the Hypotheses of Theorem 2.2 in Planted k-SAT* 16
  - *C.6. Exact Recovery with a Factorized Posterior (Corollary C.4)* 16
  - *C.5. Boundary Cases for the Hypotheses of N2* 16
  - *C.7. Proofs for the Rows of Table 1* 17
  - *C.8. Information Price of One Hidden Unit (Proposition C.6)* 18
  - *C.9. Exact Survival Probabilities and Likelihood Flatness (Lemma C.7)* 19
  - *C.10. Spurious-Survivor and Multi-Pattern Lemmas* 19
  - *C.11. Proof of Theorem 4.1 (Subcube Structure)* 20
  - *C.12. Uniqueness Window (Corollary C.11)* 21
  - *C.13. Proof of Theorem 4.2 (Bayes and Minimax Recovery)* 21
  - *C.14. Setting and Notation for Non-Uniform Coverage* 22
  - *C.15. Properties Preserved Under Arbitrary Scope Laws* 22
  - *C.16. Scope Probabilities and Survival Estimates* 23
  - *C.17. Proof of Theorem 4.3 (Heterogeneous Coverage Window)* 24
  - *C.18. A Counterexample Without Uniform Rarity* 25
  - *C.19. Proof of Corollary 4.4 (Equal Coverage and the Cost of Skew)* 26
- **D. Exact Finite-n Formulas and Numerical Protocol** 27
  - *D.1. Verification of the W-Statistic Laws at Finite n* 28
  - *D.2. Separation of Uniqueness and Recovery at Finite n* 28
  - *D.3. Two-Block Exact Formulas and Protocol* 29
- **E. Experimental Settings** 30
- **F. Additional Experimental Results** 32
- **G. A Conjectural Second-Order Window for Noisy Indexing** 36
  - *G.1. Supporting Calculation for Conjecture G.1* 37
- **H. Real-Channel and Vision Details** 38
- **I. Discussion and Limitations** 39

A Related Work

Planted constraint satisfaction. The direct predecessor is Motoki & Uehara (1999), who study the same forced-satisfiable uniform model and establish matching first-order bounds at $\frac{7}{3}n \ln n$ for $k = 3$. Their result locates the uniqueness threshold but does not identify a scaling window, a limit law inside it, or the geometry of the complete solution set. Those three layers are the SAT contribution here, and the first-order $k = 3$ constant is theirs. Feldman et al. (2015) obtain order-level identifiability results for planted CSPs and emphasize information-computation gaps. Algorithmic results for related planted formula distributions include Flaxman (2008). For non-uniform variable distributions, Friedrich et al. (2021) analyze greedy solvers on planted and filtered random formulas, and their planted-filtered equivalence theorem contains a first-order uniqueness sufficiency at $m \geq (1+\delta)(2^k - 1) \ln n / p_1$, where p_1 is the smallest variable probability and $\delta > 0$ is a fixed slack. The window-level limit law, the solution-set geometry, the converse direction, and the exact finite-n minimization of the occupancy clock in Sec. 4.3 are not addressed there. At the occupancy level, Schilling & Henze (2021) prove a Poisson limit for the number of uncovered cells in the coupon collector's problem with group drawings, under vanishing deviations of order $1/\log n$ from the uniform drawing law. This is a directly relevant precedent for the Poisson count in Thm. 4.3, which instead allows constant-factor weights and delivers the posterior geometry and both recovery curves. The optimality of equal coverage is likewise an occupancy statement with a precedent. For the coupon collector with a null coupon of fixed probability, Anceaume et al. (2015) prove that equal drawing probabilities stochastically minimize the time to complete the collection, and Anceaume et al. (2016) extend this to the time to collect any fixed number of coupons. The single-flip process of Sec. 4.3 is an instance, with $p_i = q_i$ and null probability $1 - k/(2^k - 1)$ fixed by Lem. C.13, so Cor. 4.4(a) is the first-moment form of that result (App. C.19). What is new here is the connection of the covering process to the complete SAT posterior and to the recovery window, not the optimality of equal coverage itself. Random subcubes were proposed by Mora & Zdeborová (2008) as a model of CSP solution geometry, and Thm. 4.1 instead derives a Poisson-dimensional subcube from the planted mechanism. Work on prescribed planted geometry (Ahmadi et al., 2026) addresses different questions. The Model RB calibration of Sec. 5.3 runs in the planted model of Xu & Li (2000), whose forced-satisfiable version is a standard generator of hard instances (Xu et al., 2007), and sweeps around the isolated-solution transition located by Xu & Zhang (2022). The satisfiability window of the un-planted RB model, a different object, is analyzed by Zhao et al. (2008).

Sharp exact recovery. Sharp exact-recovery thresholds with algorithm-independent converses appear in stochastic block models (Abbe et al., 2015) and community-detection models with refined maximum-likelihood asymptotics (Corinzia et al., 2019). Crowdsourcing has related sharp guarantees. These include exact error exponents under the Dawid-Skene model (Gao et al., 2016) and budget-optimal reliability bounds (Karger et al., 2014). Recent extensions study exact recovery from hypergraph similarity information (Gaudio & Joshi, 2023). These works establish sharp first-order thresholds, error exponents, or recovery guarantees. Our SAT result adds window-level posterior geometry. It distinguishes the point at which the observations have a unique explanation from the earlier point at which a Bayes learner can identify the planted one with a prescribed probability. Item- or task-dependent crowdsourcing models (Jeong & Chung, 2023; Pilliat, 2025) are relevant to the heterogeneity measured in App. H, while our experiment fits no generative model of this kind and instead treats per-item response distributions as measured channels.

Graphical-model selection and latent variables. For fully observed Ising selection, the weak-coupling scale $\Theta(\theta^{-2} \log p)$ is classical (Santhanam & Wainwright, 2012). Marginalizing hidden units is known to weaken effective interactions, and algorithms can sometimes avoid this loss for prediction objectives (Bresler & Buhai, 2020). Fourth-power sufficiency conditions occur in related latent graphical models (Anandkumar & Valluvan, 2013), while necessity questions under noisy observations remain delicate (Nikolakakis et al., 2021). Prop. C.6 supplies a necessity exponent only for the smallest one-hidden-unit family where the latent price can be computed exactly, and it is not a lower bound for general RBM learning.

Information-theoretic tools. N1 applies Fano's inequality (Fano & Hawkins, 1961) to alternatives at a fixed distance, in the tradition of conditional-mutual-information converses for sparse recovery and group testing (Atia & Saligrama, 2012). N2 packages coupon-collector and second-moment arguments, and negative association is used where loads are multinomial (Joag-Dev & Proschan, 1983). The large-deviation prefactor supporting Conj. G.1 is classical (Bahadur & Rao, 1960). Neither

N1 nor N2 is presented as a new inequality. Their contribution is organizational. Together they separate total discrimination from the distribution of evidence across the coordinates required by exact recovery.

B Notation and Conventions

Tab. A.1 lists the symbols used repeatedly in the main text and appendices. All logarithms are natural unless a base is shown. Asymptotics keep k fixed. Probability statements average over every source of randomness explicitly included in the model, and estimator-specific statements include the estimator's randomization.

Table A.1: Symbols used repeatedly in the main text and appendices. Local proof notation is introduced where needed.

Symbol	Meaning
General exact recovery	
$\theta^*, \hat{\theta}$	hidden target and estimator
m	number of target-dependent observations
Θ, P_θ, π	candidate set, observation law, and prior
$\mathrm{Succ}_{\mathrm{Bayes}}$	optimal prior-average exact-recovery probability
n	number of hidden coordinates or items, and Boolean variables in SAT
h_2	binary entropy function in nats
Planted k-SAT	
k	SAT clause width
σ, τ_A	planted assignment and the assignment flipped on A
$\mathcal{S}$	complete version space / satisfying-assignment set
F, W	surviving single-flip coordinates and their count
ξ_i	scope inclusion marginal $\mathbb{P}(i \in S)$
$q^{(r)}$	per-clause violation probability of a distance-r assignment in the uniform model
$q_i, \widetilde{q}_i$	single-flip violation probability and $-\ln(1-q_i)$
w_i, c, C	coordinate weight and its fixed lower and upper bounds
$\alpha, \beta, \bar{w}, \zeta$	two-block fraction, low weight, mean weight, and window-scale skew
x	centered coordinate in the SAT critical window
$U(x), R(x)$	limiting uniqueness and optimal-recovery curves
μ	Poisson intensity e^{-x} or the limit of Λ_m
Λ_m	occupancy functional $\sum_i (1-q_i)^m$ of Sec. 4.3
Noisy indexing and measured channels	
λ	mean observations per coordinate m/n
$\lambda^*(n)$	per-item load giving one-half exact-recovery probability
a	$1 - 2\sqrt{\eta(1-\eta)}$ for noisy indexing
C_η	$-\ln(2\sqrt{\eta(1-\eta)})$ for noisy indexing

C Proofs

Each subsection below proves one result of the main text, in order of appearance, or states and proves the supporting lemmas that the next proof needs.

C.1 Proof of Theorem 2.1 (Localized Fano Bound)

Proof. Let Y be uniform on the $N_r + 1$ hypotheses. Fano's inequality gives $H(Y \mid Z_{1:m}) \le h_2(P_e) + P_e \ln N_r$, hence $I(Y; Z_{1:m}) \ge (1 - P_e)\ln(N_r + 1) - \ln 2$. The chain rule and the

conditional i.i.d. structure give

$$I(Y; Z_{1:m}) = \sum_{t=1}^{m} I(Y; Z_t \mid Z_{1:t-1}) = \sum_t \big[H(Z_t \mid Z_{1:t-1}) - H(Z_t \mid Y)\big]$$
$$\leq \sum_t \big[H(Z_t) - H(Z_t \mid Y)\big] = m\, I(Y; Z_1),$$

using $H(Z_t \mid Z_{1:t-1}) \leq H(Z_t)$ and $H(Z_t \mid Y, Z_{1:t-1}) = H(Z_t \mid Y)$. □

C.2 Mixture Information in Planted k-SAT (Lemma C.1)

Lemma C.1 (Mixture information in planted k-SAT)**.** *Let Y be uniform on a planted assignment σ and its n single-coordinate flips. In the planted k-SAT model of Sec. 4,*

$$I(Y; Z_1) = (1 + o(1)) q_1 \ln n, \qquad q_1 = \frac{k}{n(2^k - 1)}.$$

Consequently, Thm. 2.1 certifies only $m = \Omega(1/q_1) = \Omega(n)$, whereas the true exact-recovery scale is $\Theta(n \log n)$.

Proof. Let Y be uniform on $\{\sigma, \tau_1, \ldots, \tau_n\}$ where $\tau_i := \tau_{\{i\}} = \sigma \oplus \mathbf{1}_{\{i\}}$. Given $Y = \theta$, the clause Z is uniform on the N clauses satisfied by θ. Every assignment satisfies exactly N clauses, so $P_\theta(z) = \mathbf{1}\{z \in \mathrm{Legal}(\theta)\}/N$ and the mixture is $\bar{P}(z) = c(z)/((n+1)N)$ with $c(z) = \#\{\theta : z \in \mathrm{Legal}(\theta)\}$. Hence

$$I(Y; Z_1) = \frac{1}{n+1} \sum_\theta \sum_{z \in \mathrm{Legal}(\theta)} \frac{1}{N} \ln \frac{n+1}{c(z)}.$$

Classify $z = (S, f)$ by its difference pattern D relative to σ:

- $f = \sigma|_S$ (the "distance-0" class), illegal under σ and under every τ_i with $i \notin S$, and legal exactly under the k flips τ_i, $i \in S$. So $c(z) = k$ and the log-term is $\ln \frac{n+1}{k}$.
- $|D| = 1$, say $D = \{j\}$, illegal only under τ_j, so $c(z) = n$ and the log-term is $\ln \frac{n+1}{n} \leq \frac{1}{n}$.
- $|D| \geq 2$, legal under all $n + 1$ hypotheses, with log-term 0.

Under $Y = \tau_i$, the first class occurs iff $i \in S$ and $f = \sigma|_S$, an event of probability exactly $q_1 = \frac{k}{n(2^k-1)}$ under the uniform legal clause of τ_i. Summing over the n flip hypotheses contributes $\frac{n}{n+1} q_1 \ln \frac{n+1}{k} = (1 + o(1)) q_1 \ln n$. The second class contributes at most $\frac{k}{2^k-1} \cdot \frac{1}{n} = O(q_1)$ per hypothesis, i.e. $O(q_1)$ in total, and the third contributes 0. Hence $I(Y; Z_1) = (1 + o(1)) q_1 \ln n$. Substituting into Thm. 2.1 with $\ln(N_1 + 1) = \ln(n + 1)$ gives $m \geq (1 - o(1))(1 - P_e)/q_1 = \Omega(n)$ and no more. □

C.3 Proof of Theorem 2.2 (Coordinate-Coverage Necessity)

Any measurable function of the observations qualifies as the auxiliary variable V in Thm. 2.2, and conditional independence or pairwise negative correlation of the E_i suffices for its variance condition.

Proof. Divergence of K. Write $\mu_V = \mathbb{E}[K \mid V]$ and $v_V = \mathrm{Var}(K \mid V)$. Chebyshev gives $\mathbb{P}(K \leq \frac{1}{2}\mu_V \mid V) \leq 4v_V/\mu_V^2$. Fix κ. On $\{\mu_V > 2\kappa\}$ we have $\mathbb{P}(K \leq \kappa \mid V) \leq 4v_V/\mu_V^2$, so

$$\mathbb{P}(K \leq \kappa) \leq \mathbb{E}\left[\min\left(1, \frac{4v_V}{\mu_V^2}\right)\right] + \mathbb{P}(\mu_V \leq 2\kappa) \longrightarrow 0$$

by the two hypotheses and bounded convergence. Hence $K \xrightarrow{p} \infty$. Conditional independence or pairwise negative correlation gives $v_V \leq \sum_i \mathbb{P}(E_i \mid V)(1 - \mathbb{P}(E_i \mid V)) \leq \mu_V$, so either condition implies the required variance ratio when $\mu_V \xrightarrow{p} \infty$.

Linkage. Fix an observation z and condition on it throughout. Define the posterior $\omega_z(\theta) = \mathbb{P}(\theta^* = \theta \mid Z_{1:m} = z)$. Under the uniform prior, $\omega_z(\theta) \propto P_\theta(z)$, so the level sets $L(\theta) = \{\theta' : P_{\theta'}(z) = P_\theta(z)\}$ partition Θ into classes of constant posterior mass, say ω_ℓ for each member of class ℓ. The linkage hypothesis, applied at θ rather than only at the truth, supplies K_θ pairwise distinct $\theta' \neq \theta$ of the same likelihood, so $|L(\theta)| \geq K_\theta + 1$ and therefore $1/(K_\theta + 1) \geq 1/|L(\theta)|$ for every θ. Summing over Θ and grouping by class gives

$$\mathbb{E}\Big[\frac{1}{K_{\theta^*}+1} \,\Big|\, Z_{1:m} = z\Big] \geq \sum_\theta \frac{\omega_z(\theta)}{|L(\theta)|} = \sum_\ell \omega_\ell \geq \max_\ell \omega_\ell = \max_\theta \omega_z(\theta).$$

The last quantity is the Bayes success given z. Averaging over z shows that Bayes success is at most $\mathbb{E}[1/(K_{\theta^*}+1)]$, which tends to 0 by $K_{\theta^*} \xrightarrow{p} \infty$ and bounded convergence. Since V is ancillary given the observations, conditioning on it does not change the posterior of the target. Every estimator's uniform-prior average success is at most the Bayes success. □

The level-set argument is needed because the posterior mass of the truth being at most $1/(\kappa+1)$ does not by itself bound an estimator's success. An estimator returns the posterior mode, while K is indexed by the truth, which the estimator does not know. The inequality above closes that gap, and this is why the hypothesis is imposed at every candidate.

The counterexample of Rem. C.2 in full. Requiring only $\mathbb{E}\big[\sum_i \mathbb{P}(E_i \mid V)\big] \to \infty$ together with conditional independence is not enough. Let $V \sim \mathrm{Bernoulli}(n^{-1/2})$, and given $V = 1$ let all n events occur independently with probability one, while given $V = 0$ none can occur. Then the expected conditional sum is $n^{1/2} \to \infty$, yet $K = 0$ with probability $1 - n^{-1/2} \to 1$. The convergence in probability of the conditional mean is essential.

C.4 Verification of the Hypotheses of Theorem 2.2 in Planted k-SAT

Take $E_i(\theta) = \{\text{no clause has difference pattern } \{i\} \text{ relative to } \theta\}$, so that $K_\sigma = W$. The every-candidate form of the linkage holds by $(\mathbb{Z}_2)^n$-equivariance, since the model looks the same from every assignment, and the K_θ alternatives are the distinct single flips of θ. Flatness is Lem. C.7(c), since each surviving single flip is a full-likelihood alternative. The per-clause violation events for distinct coordinates are disjoint (Lem. C.7(a)), so the survival indicators are negatively correlated. The theorem applies with V trivial, and $\mathbb{E}W = n(1-q_1)^m \to \infty$ whenever $m\widetilde{q}_1 - \ln n \to -\infty$. Thm. 4.2 gives the corresponding exact tail statement.

C.5 Boundary Cases for the Hypotheses of N2

Remark C.2 (Conditional rather than average divergence). It is insufficient to assume only $\mathbb{E}K \to \infty$. A rare mixture component can make K large in expectation while $K = 0$ with probability tending to one. Thm. 2.2 therefore requires conditional concentration in probability, and the complete counterexample is given with the proof of Thm. 2.2 above.

Remark C.3 (Erasure as a failure boundary). In the erasure family, all natural coordinate-level uncovered events are the same event that no sample revealed the target. They are perfectly positively correlated, so N2 does not apply, and N1 alone gives the scale $\Theta(1/\varepsilon)$, tight in order though not in its constant, as computed in App. C.7.

C.6 Exact Recovery with a Factorized Posterior (Corollary C.4)

Corollary C.4 (Exact recovery with a factorized posterior)**.** *Let the prior be a product over coordinates and let V be an auxiliary variable whose law does not depend on the target. Suppose that conditional on V the observations split into coordinate blocks that are independent given the target, block i depending on the target only through its coordinate i. Then the posterior given the observations and V factorizes by coordinate, and with $p_i(V)$ the conditional error probability given V of the Bayes decision on coordinate i, the Bayes success of a learner that observes $Z_{1:m}$ and V, and hence also of a learner that observes $Z_{1:m}$ alone, satisfies*

$$\mathrm{Succ}_{\mathrm{Bayes}} \leq \mathbb{E}_V \exp\left(-\sum_i p_i(V)\right).$$

In particular, $\sum_i p_i(V) \xrightarrow{p} \infty$ *implies vanishing Bayes success.*

Proof. Write $Z = Z_{1:m}$, let $\theta^* = (\theta_1^*, \dots, \theta_n^*)$ have the product prior $\pi = \bigotimes_i \pi_i$, and let $Z^{(1)}, \dots, Z^{(n)}$ be the coordinate blocks given V. By hypothesis $P(Z \mid \theta, V) = \prod_i P(Z^{(i)} \mid \theta_i, V)$ and $P(V \mid \theta) = P(V)$, so Bayes' rule gives

$$\mathbb{P}(\theta^* = \theta \mid Z, V) = \prod_i \pi_i(\theta_i \mid Z^{(i)}, V), \qquad \pi_i(a \mid Z^{(i)}, V) = \frac{\pi_i(a)\, P(Z^{(i)} \mid a, V)}{\sum_{a'} \pi_i(a')\, P(Z^{(i)} \mid a', V)},$$

which is the factorization of the posterior given (Z, V). When V is a function of Z, as the load vector is in noisy indexing, this is the posterior given Z. Otherwise the two posteriors differ, and the bound is proved below for the Bayes estimator that sees (Z, V), which dominates the Bayes estimator from Z alone, because $\mathbb{P}(\theta^* = \theta \mid Z) = \mathbb{E}[\mathbb{P}(\theta^* = \theta \mid Z, V) \mid Z]$ and a maximum of averages is at most the average of the maxima. The Bayes estimator maximizes the posterior, and a product of nonnegative factors indexed by disjoint coordinates is maximized coordinate by coordinate, so the Bayes estimator is the vector $(\hat\theta_1, \dots, \hat\theta_n)$ of local Bayes decisions $\hat\theta_i = \arg\max_a \pi_i(a \mid Z^{(i)}, V)$, and its success given (Z, V) is

$$\max_\theta \mathbb{P}(\theta^* = \theta \mid Z, V) = \prod_i M_i, \qquad M_i = \max_a \pi_i(a \mid Z^{(i)}, V) = \mathbb{P}(\hat\theta_i = \theta_i^* \mid Z, V).$$

Given V, the blocks $Z^{(1)}, \dots, Z^{(n)}$ are independent, because they are independent given (θ^*, V), block i depends on θ^* only through θ_i^*, and the coordinates of θ^* are independent under the product prior. Each M_i is a function of $(Z^{(i)}, V)$, so the M_i are conditionally independent given V, and the tower property gives

$$\mathrm{Succ}_{\mathrm{Bayes}} = \mathbb{E}\Big[\mathbb{E}\Big[\prod_i M_i \,\Big|\, V\Big]\Big] = \mathbb{E}\Big[\prod_i \mathbb{E}[M_i \mid V]\Big] = \mathbb{E}\Big[\prod_i \big(1 - p_i(V)\big)\Big],$$

since $\mathbb{E}[M_i \mid V] = \mathbb{P}(\hat\theta_i = \theta_i^* \mid V) = 1 - p_i(V)$. Finally $1 - p \le e^{-p}$ for $p \in [0, 1]$ gives $\prod_i (1 - p_i(V)) \le \exp(-\sum_i p_i(V))$, which is the claimed bound. If $\sum_i p_i(V) \xrightarrow{p} \infty$, then $\exp(-\sum_i p_i(V)) \xrightarrow{p} 0$ and is bounded by one, so its expectation tends to zero by bounded convergence, and every estimator's prior-average success is at most the Bayes success. □

In noisy indexing, let V be the load vector. Its law does not depend on σ, because each observation names a uniform coordinate, the prior on σ is a product, and given V the votes on distinct coordinates are independent given σ, each passing its own coordinate through $\mathrm{BSC}(\eta)$, so the hypotheses of the corollary hold, and since every observation names its coordinate, V is a function of the observations and the factorized posterior is the posterior given the observations themselves. Each local Bayes decision is the majority vote on that coordinate's noisy votes, and $\sum_i p_i(V)$ diverges in probability below the first-order threshold by the majority-vote error asymptotics, so the corollary gives vanishing Bayes success. The flatness linkage of Thm. 2.2 is not used here and does not hold in this family. The multinomial loads are negatively associated (Joag-Dev & Proschan, 1983), which is what transfers concentration statements about V itself.

C.7 Proofs for the Rows of Table 1

The rows of Tab. 1 are treated in order after the noisy-indexing information entry is computed. The one-hidden-unit RBM row is Prop. C.6, proved in App. C.8.

Example C.5 (Noisy indexing with an order-tight N1 bound)**.** Let Y be uniform on the target and its n single-coordinate alternatives. Conditional on the sampled coordinate, exactly one of these $n + 1$ candidates differs from the target there. The noisy bit therefore disagrees with the target with mixture probability $\eta + (1 - 2\eta)/(n + 1)$, while its conditional entropy given Y and the sampled coordinate is $h_2(\eta)$. Since the sampled coordinate is independent of Y, for fixed $0 < \eta < \frac{1}{2}$,

$$I_1 = h_2\left(\eta + \frac{1 - 2\eta}{n + 1}\right) - h_2(\eta) = \frac{1 - 2\eta}{n + 1} \ln \frac{1 - \eta}{\eta} + O_\eta(n^{-2}) = \Theta(n^{-1}).$$

This gives the information entry in Tab. 1. In the weak-signal regime, $I_1 \asymp \varepsilon_\eta^2/n$ with $\varepsilon_\eta = 1-2\eta$, so N1 gives $m = \Omega(n \ln n/\varepsilon_\eta^2)$, which is the correct order *including* the logarithm, because here the per-coordinate structure is visible to the mixture information. The constant is off by a factor of about 4 in the weak-signal limit (KL $\approx 2\varepsilon_\eta^2$ against the true per-coordinate rate $a \approx \varepsilon_\eta^2/2$).

Proof of the bounds in Table 1. **Hidden pixel.** Let Y be the uniform index of the hidden location. The demand is $\ln d$ nats and the supply per sample is $I(Y;Z_1) \le \ln 2 - h_2(\eta)$ nats by the channel capacity, so N1 gives $m = \Omega(\ln d/(\ln 2 - h_2(\eta)))$. A correlation estimator based on matched filtering achieves the same order for all fixed $\eta \in [0, \frac{1}{2})$ by a union bound over d positions.

One-edge Ising. Per-sample KL between the one-edge measure and the independent measure is $\theta \tanh\theta - \ln\cosh\theta = \frac{\theta^2}{2}(1+o(1))$ as $\theta \to 0$. N1 over the $\binom{p}{2}$ candidates gives $m = \Omega(\theta^{-2}\log p)$, and thresholding empirical pair correlations achieves $O(\theta^{-2}\log p)$ by Hoeffding plus a union bound. (Weak-coupling regime only, with no claim at large θ.)

Planted k-SAT. Necessity comes from Lem. C.1 (N1, order n) and Thm. 4.2 (full constant and window). Sufficiency follows from the layered union bound in Thm. 3.1 of Motoki & Uehara (1999), where $\mathbb{P}(\exists \text{ survivor at distance } r) \le \binom{n}{r}(1-q^{(r)})^m$, summed over $r \ge 1$, vanishes once $m \ge (1+\delta)\frac{n(2^k-1)}{k}\ln n$ for any fixed $\delta > 0$. For general k the same conclusion is contained in the tail statement of Thm. 4.2, proved in part (iv) of App. C.13, since such m has $m\widetilde{q}_1 - \ln n \to +\infty$, and then the uniqueness probability tends to one by the monotone coupling there together with Cor. C.11.

Noisy indexing. Necessity comes from Ex. C.5 (N1, order-tight) and the factorized-posterior Cor. C.4 (constant $1/a$), which applies because below $m = (1-\delta)n\ln n/a$ for fixed $\delta > 0$ the expected number of erring coordinates diverges. For sufficiency, a coordinate that received t votes is decoded wrongly by majority vote, ties broken at random, with probability at most $\mathbb{P}(\mathrm{Bin}(t,\eta) \ge t/2) \le (4\eta(1-\eta))^{t/2} = e^{-tC_\eta}$ by the Chernoff bound, with $C_\eta = -\ln(2\sqrt{\eta(1-\eta)})$ as in Conj. G.1. The load of a fixed coordinate is $\mathrm{Bin}(m, 1/n)$, so its error probability is at most $\mathbb{E}[e^{-tC_\eta}] = (1-a/n)^m \le e^{-a\lambda}$ with $a = 1 - e^{-C_\eta} = 1 - 2\sqrt{\eta(1-\eta)}$, and a union bound over the n coordinates gives an error probability at most $ne^{-a\lambda} \to 0$ once $a\lambda \ge (1+\delta)\ln n$. No Poissonization of the loads is needed at first order. (Window-level de-Poissonization is exactly what Conj. G.1 still requires.)

Erasure. In the erasure row of Tab. 1, every sample reveals the entire target with probability ε and is otherwise an erasure. For the mixture information, a sample is an erasure with probability $1-\varepsilon$ (carrying no information about Y) and otherwise reveals Y (carrying $\ln M'$ nats where $M' = N_r + 1$), so $I(Y;Z_1) = \varepsilon \ln M'$. Fano gives $m \ge [(1-P_e)\ln M' - \ln 2]/(\varepsilon \ln M') = (1 - P_e - o(1))/\varepsilon$. Conversely, the estimator that waits for the first revelation and otherwise guesses uniformly succeeds with probability $1 - (1-\varepsilon)^m + (1-\varepsilon)^m/M'$, so $m = c/\varepsilon$ achieves success $1 - e^{-c} + o(1)$. Both sides are $\Theta(1/\varepsilon)$ uniformly in M'. The bound is tight in order and not in its constant. At success level s the Fano side gives $m \ge (s - o(1))/\varepsilon$, whereas the optimal rule, which outputs the revealed target when a revelation has occurred and guesses otherwise, succeeds with probability exactly $1 - (1-\varepsilon)^m(1 - 1/M')$ and therefore needs $m = (-\ln(1-s) + o(1))/\varepsilon$ for large M'. At $s = 0.9$ the two constants are 0.9 and 2.30, so N1 is order-tight here and not constant-exact. □

C.8 Information Price of One Hidden Unit (Proposition C.6)

Proposition C.6 (Information price of one hidden unit)**.** *Consider a sequence of problems with $p \to \infty$ visible binary units and one hidden unit connected with weight $w = w_p \to 0$ to an unknown pair. Under the uniform prior over the $\binom{p}{2}$ pairs, every exact-recovery procedure requires*

$$m = \Omega(w^{-4}\log p),$$

and a correlation-thresholding estimator succeeds with $m = O(w^{-4}\log p)$. Marginalizing the hidden unit produces the exact visible coupling $J(w) = \frac{1}{2}\ln\cosh(2w) = \Theta(w^2)$, whose per-sample KL supply is $\Theta(J(w)^2) = \Theta(w^4)$.

Proof. The marginal over visibles is the one-edge Ising measure at coupling $J(w) = \frac{1}{2}\ln\cosh 2w$, by direct computation, since summing the hidden unit contributes $\cosh(w(v_i + v_j))$ and $\ln\cosh 2w$

is the resulting pair interaction after symmetrization. Applying the one-edge Ising row of App. C.7 with $\theta = J(w) \approx w^2$ gives necessity $\Omega(J^{-2}\log p) = \Omega(w^{-4}\log p)$ and matching sufficiency by pair-correlation thresholding. □

C.9 Exact Survival Probabilities and Likelihood Flatness (Lemma C.7)

Lemma C.7 (Exact survival probabilities and likelihood flatness)**.**

(a) For distinct coordinates $i_1, \dots, i_j$, the events that one clause is violated by their single flips are disjoint, and $\mathbb{P}(i_1, \dots, i_j \in F) = (1 - jq_1)^m$.

(b) For $|A| = r$, $\mathbb{P}(\tau_A \in \mathcal{S}) = (1 - q^{(r)})^m$, where $q^{(r)} = [1 - \binom{n-r}{k}/\binom{n}{k}]/(2^k - 1)$ and $q^{(1)} = q_1$.

(c) Every assignment consistent with the observations has the same likelihood, so the posterior is uniform on $\mathcal{S}$.

Proof. (a) The difference pattern D determines which single flip violates a clause, and $\{D = \{i_l\}\}$ are disjoint across distinct i_l, each of probability q_1, so the probability that none of j specified flips violates one clause is exactly $1 - jq_1$ (and $jq_1 \le nq_1 = \frac{k}{2^k-1} \le \frac{2}{3} < 1$). Independence across clauses gives the claim. (b) $\mathbb{P}(A \cap S = D) = q^{(r)}$ by counting. (c) $|\mathrm{Legal}(\tau)| = N$ for every τ, so the likelihood of the observed clause sequence is N^{-m} under every consistent hypothesis, hence the posterior is uniform on $\mathcal{S}$. □

C.10 Spurious-Survivor and Multi-Pattern Lemmas

Throughout Apps. C.10 to C.13, $k \ge 2$ and $x \in \mathbb{R}$ are fixed, m is any integer sequence with $m\widetilde{q}_1 - \ln n \to x$, $L = mq_1 = \ln n + x + o(1)$, and $\mu = e^{-x}$. Constants in $O_{k,x}$ may depend on k and x. A *spurious survivor* is an assignment τ_A with $|A| \ge 2$ that survives all clauses although some $i \in A$ is resolved, meaning that its single flip $\tau_{\{i\}}$ violates some clause, and it is *complete* if every $i \in A$ is resolved. The lemmas below bound the two ways in which the structure event of Thm. 4.1 can fail, a spurious survivor or a multi-pattern violation.

For $|A| = r \ge 2$ define per-clause categories (disjoint, since D identifies them):

$$D_i(A) = \{D = \{i\},\, S \cap A = \{i\}\}, \quad i \in A, \text{ each of probability } d_r = \frac{\binom{n-r}{k-1}}{\binom{n}{k}(2^k-1)};$$

$$G_i(A) = \{D = \{i\},\, S \cap A \supsetneq \{i\}\}, \quad \text{each of probability } g_r = q_1 - d_r.$$

Every $D_i(A)$ clause is violated by τ_A as well (by (1), $A \cap S = \{i\} = D$). With $\rho_r = d_r/q_1 = \prod_{j=0}^{k-2} \frac{n-r-j}{n-1-j}$ one has $1 - \rho_r \le \frac{(k-1)(r-1)}{n-k+1} \le \frac{2(k-1)r}{n}$ for $n \ge 2k$, hence $g_r \le \frac{2(k-1)r\,q_1}{n}$.

Lemma C.8 (Complete spurious survivors)**.** *Let $O_A = \{\tau_A \in \mathcal{S}\} \cap \{$every $i \in A$ is resolved$\}$. Then*

$$\sum_{r=2}^{n} \sum_{|A|=r} \mathbb{P}(O_A) = O_{k,x}\Big(\frac{(\ln n)^2}{n^2}\Big),$$

with region-I leading constant of order $36e^{10}(k-1)^2e^{-2x}$.

Lemma C.9 (Partial spurious survivors)**.** $\mathbb{P}\big(\exists\ \tau_A \in \mathcal{S},\ |A| \ge 2,$ *some $i \in A$ resolved*$\big) = O_{k,x}\big(\frac{\ln n}{n}\big)$.

Proof of both. Write $\Sigma_r = \sum_{|A|=r} \mathbb{P}(O_A)$ for the complete spurious-survivor contribution at distance r. On O_A, survival of τ_A forces all $D_i(A)$ counts to vanish, and resolving coordinate i then requires a $G_i(A)$ clause. Hence

$$O_A \subseteq \{X_{D_i} = 0\ \forall i \in A\} \cap \{X_{G_i} \ge 1\ \forall i \in A\}.$$

$\mathbb{P}(\text{no } D) = (1 - rd_r)^m$, and conditioned on that the m clauses remain i.i.d. and $(X_{G_i})_{i\in A}$ are multinomial counts with cell probabilities $g'_r = g_r/(1 - rd_r)$. Multinomial counts are negatively

associated (Joag-Dev & Proschan, 1983), and the events $\{X_{G_i} \geq 1\}$ are increasing in disjoint coordinates, so

$$\mathbb{P}(O_A) \leq (1-rd_r)^m \cdot \min\Big\{1,\ \prod_{i\in A} \mathbb{P}(X_{G_i} \geq 1 \mid \text{no } D)\Big\} \leq (1-rd_r)^m \min\{1, (mg'_r)^r\}. \quad (3)$$

Region I ($2 \leq r \leq r_0 = \lfloor n/((k-1)\ln n)\rfloor$). Here $rd_r \leq rq_1 \leq \frac{k}{2^k-1} \leq \frac{2}{3}$ gives $g'_r \leq 3g_r$, then $(1-rd_r)^m \leq e^{-r\rho_r L} \leq e^{-rL}e^{2(k-1)r^2L/n} \leq e^{-rL}e^{4r}$ (since $(k-1)rL/n \leq L/\ln n \leq 2$), and $mg'_r \leq \frac{6(k-1)rL}{n}$. Substituting into (3) and summing over A with $\binom{n}{r} \leq (en/r)^r$ gives

$$\Sigma_r \leq \Big[\frac{en}{r}\cdot e^{-L}e^4 \cdot \frac{6(k-1)rL}{n}\Big]^r = b_n^r,$$
$$b_n = 6e^5(k-1)Le^{-L} = 6e^5(k-1)e^{-x}\frac{\ln n}{n}(1+o(1)) \to 0,$$

so $\sum_{r=2}^{r_0}\Sigma_r \leq 2b_n^2 = O(e^{-2x}(\ln n)^2/n^2)$, an explicit bound in this region.

Region II ($r_0 < r \leq n/2$). Drop the G factor and use $\rho_r \geq \big(\frac{n/2-k}{n}\big)^{k-1} \geq 2^{-k}$ for large n, so

$$\Sigma_r \leq \Big[\frac{en}{r}e^{-2^{-k}L}\Big]^r \leq \big[2e(k-1)(\ln n)\, n^{-2^{-k}} e^{2^{-k}|x|}\big]^r,$$

whose bracket tends to 0, so the sum is dominated by $(\text{bracket})^{r_0}$ and is superpolynomially small.

Region III ($n/2 < r \leq n$). By Lem. C.7(b), $q^{(r)} \geq \frac{1-\binom{\lfloor n/2\rfloor}{k}/\binom{n}{k}}{2^k-1} \to \frac{1-2^{-k}}{2^k-1} > 0$, so for some $\gamma_k > 0$, $\sum_{r>n/2}\Sigma_r \leq 2^n(1-\gamma_k)^m = \exp(n\ln 2 - \Theta(n\ln n)) \to 0$.

Summing the three regions proves Lem. C.8 (region I dominates). For Lem. C.9, the event is contained in $\{X_{D_j} = 0\ \forall j \in A\} \cap \{\exists i \in A : X_{G_i} \geq 1\}$, and a union bound over the resolved coordinate replaces $(mg'_r)^r$ by $r\, mg'_r$ in (3). In region I the per-r term is at most $(\gamma/r)^r \cdot \frac{6(k-1)r^2L}{n}$ with $\gamma = e^5e^{-x}(1+o(1))$, and the series $C(\gamma) = \sum_{r\geq 2}(\gamma/r)^r r^2 < \infty$ (super-exponential decay once $r > \gamma$), so the region-I total is at most $6(k-1)C(\gamma)L/n = O_{k,x}(\ln n/n)$. The $r = 2$ term need not dominate, since for very negative x the dominant term sits near $r \asymp \gamma$, which is why the constant is allowed to depend on k and x. Regions II and III absorb the extra polynomial factors into their superpolynomial decay. □

Lemma C.10 (Multi-pattern violations)**.** $\mathbb{P}\big(\exists\ \text{clause } c : |D(c)| \geq 2,\ D(c) \subseteq F\big) \leq m\sum_{d=2}^{k}\frac{\binom{k}{d}}{2^k-1}(1-dq_1)^{m-1} = O\Big(e^{-2x}\frac{(k-1)\ln n}{n}\Big)$.

Proof. A clause with $|D| \geq 2$ is violated by no single flip, so the event that its D lies inside F depends only on the other $m-1$ independent clauses, with probability $(1-dq_1)^{m-1}$ by Lem. C.7(a), and a first moment over clauses and d finishes, the $d = 2$ term dominating. (This constant is explicit because the computation is a clean first moment.) □

C.11 Proof of Theorem 4.1 (Subcube Structure)

The proof uses the three lemmas of App. C.10.

Proof. The Poisson limit. For fixed j, Lem. C.7(a) gives $\mathbb{E}[(W)_j] = (n)_j(1-jq_1)^m$ exactly. Expanding $\ln(1-jq_1) = -j\widetilde{q}_1 - \frac{j^2-j}{2}q_1^2 + O(j^3q_1^3)$ and using $m\widetilde{q}_1 = \ln n + x + o(1)$ (the sequence definition guarantees only $o(1)$, which suffices) together with $mq_1^2 = O(\ln n/n)$:

$$\mathbb{E}[(W)_j] = (n)_j\, n^{-j}e^{-jx}(1+o(1)) \to e^{-jx} = \mu^j.$$

All factorial moments converge to those of $\mathrm{Poisson}(\mu)$, which is moment-determined, hence $W \Rightarrow \mathrm{Poisson}(\mu)$.

Assembly. Let $\mathcal{C} = \{\sigma \oplus \mathbf{1}_B : B \subseteq F\}$.

$\mathcal{C} \subseteq \mathcal{S}$ *whp.* If τ_B ($B \subseteq F$) violates a clause (S, f) then $B \cap S = D$. If $|D| = 1$ then $D = \{i\} \subseteq B \subseteq F$, contradicting $i \in F$, and if $|D| \geq 2$ then $D \subseteq F$ triggers Lem. C.10. So $\mathbb{P}(\mathcal{C} \not\subseteq \mathcal{S}) = O(e^{-2x}(k-1)\ln n/n)$.

$\mathcal{S} \subseteq \mathcal{C}$ *whp.* A survivor τ_A with $A \not\subseteq F$ has $|A| \geq 2$ (a surviving $\tau_{\{i\}}$ with $i \notin F$ is impossible, since a $D = \{i\}$ clause is violated by exactly the single flip $\tau_{\{i\}}$ by (1)) and some $i \in A$ resolved, which is the event of Lem. C.9.

Combining with the Poisson limit proves Thm. 4.1 with error $O_{k,x}(\ln n/n)$, since both lemmas are of that order, and the partial spurious-survivor constant is the non-explicit one. □

C.12 Uniqueness Window (Corollary C.11)

Corollary C.11 (Uniqueness window)**.** *Under the same assumptions,*

$$\mathbb{P}(|\mathcal{S}| = 1) \longrightarrow \exp(-e^{-x}), \qquad \log_2 |\mathcal{S}| \Rightarrow \text{Poisson}(e^{-x}).$$

Proof. Let $\mu = e^{-x}$, let $\mathcal{C} = \{\sigma \oplus \mathbf{1}_B : B \subseteq F\}$, and let $\mathcal{E}_n = \{\mathcal{S} = \mathcal{C}\}$ be the structure event. Thm. 4.1 gives $\mathbb{P}(\mathcal{E}_n^c) = O_{k,x}(\ln n/n) \to 0$ and $W \Rightarrow \text{Poisson}(\mu)$. The map $B \mapsto \sigma \oplus \mathbf{1}_B$ is injective, so $|\mathcal{C}| = 2^W$ and $\log_2 |\mathcal{S}| = W$ on $\mathcal{E}_n$.

Uniqueness. On $\mathcal{E}_n$, $|\mathcal{S}| = 1$ if and only if $W = 0$, so the events $\{|\mathcal{S}| = 1\}$ and $\{W = 0\}$ differ only inside $\mathcal{E}_n^c$, and

$$\big|\mathbb{P}(|\mathcal{S}| = 1) - \mathbb{P}(W = 0)\big| \leq \mathbb{P}(\mathcal{E}_n^c) \longrightarrow 0.$$

Since W is integer valued, $W \Rightarrow \text{Poisson}(\mu)$ gives $\mathbb{P}(W = 0) \to e^{-\mu}$, hence $\mathbb{P}(|\mathcal{S}| = 1) \to e^{-\mu} = \exp(-e^{-x})$.

Solution count. $\mathbb{P}(\log_2 |\mathcal{S}| \neq W) \leq \mathbb{P}(\mathcal{E}_n^c) \to 0$, so $\log_2 |\mathcal{S}| - W \xrightarrow{p} 0$, and Slutsky's theorem together with $W \Rightarrow \text{Poisson}(\mu)$ gives $\log_2 |\mathcal{S}| \Rightarrow \text{Poisson}(\mu)$.

The uniqueness statement also has a direct route with a sharper error. Every $i \in F$ yields the survivor $\tau_{\{i\}}$, so $\{|\mathcal{S}| = 1\} \subseteq \{W = 0\}$ for every n and m. On $\{W = 0\}$ every coordinate is resolved, so a survivor τ_A with $|A| \geq 2$ is a complete spurious survivor, and Lem. C.8 gives

$$0 \leq \mathbb{P}(W = 0) - \mathbb{P}(|\mathcal{S}| = 1) = \mathbb{P}(W = 0,\ |\mathcal{S}| > 1) \leq \sum_{r \geq 2} \sum_{|A| = r} \mathbb{P}(O_A) = O_{k,x}\Big(\frac{(\ln n)^2}{n^2}\Big).$$

□

C.13 Proof of Theorem 4.2 (Bayes and Minimax Recovery)

Proof. (i) By Lem. C.7(c) the posterior is uniform on $\mathcal{S}$, so the Bayes rule guesses uniformly on $\mathcal{S}$ and its success is $\mathbb{E}[1/|\mathcal{S}|]$. On the structure event $1/|\mathcal{S}| = 2^{-W}$, hence $|\mathbb{E}[1/|\mathcal{S}|] - \mathbb{E}[2^{-W}]| \leq \mathbb{P}(\text{structure fails}) \to 0$, and $\mathbb{E}[2^{-W}] \to \mathbb{E}[2^{-\text{Poisson}(\mu)}] = e^{-\mu/2}$ by the probability generating function, since $z \mapsto z^W$ is bounded.

(ii) Any estimator's prior-average success is at most the Bayes success.

(iii) The $(\mathbb{Z}_2)^n$ action $s \cdot (\sigma, \{(S_t, f_t)\}) = (\sigma \oplus s, \{(S_t, f_t \oplus s|_{S_t})\})$ preserves the model (it maps legal clauses of σ bijectively to legal clauses of $\sigma \oplus s$ and preserves uniformity) and acts transitively on targets, fixing the uniform prior. By the standard invariance reduction, the uniform prior is least favorable and minimax success equals Bayes success. For any $\hat{A}$, the randomized symmetrization $\hat{A}^s$ (conjugate by a uniform s) has per-target success constant in σ and equal to $\hat{A}$'s prior average. If some algorithm had per-target success $\geq$ curve $+ \delta$ for all targets, its prior average would also exceed the curve by δ, contradicting (ii).

(iv) Fixed-x limits do not automatically cover sequences $x_n \to \pm\infty$, so we use a monotone coupling. Realize the clauses as the first m elements of a single i.i.d. clause stream. Then $m' \geq m$ implies $\mathcal{S}_{m'} \subseteq \mathcal{S}_m$ pointwise, so the Bayes success $\mathbb{E}[1/|\mathcal{S}_m|]$ and the uniqueness probability $\mathbb{P}(|\mathcal{S}_m| = 1)$ are non-decreasing in m. For fixed $c \in \mathbb{R}$, define $m(n, c) = \lfloor (\ln n + c)/\widetilde{q}_1 \rfloor$, which satisfies

$m(n,c)\widetilde{q}_1 - \ln n \to c$. Let $x_n = m_n\widetilde{q}_1 - \ln n \to -\infty$. For any fixed c, eventually $m_n \le m(n,c)$, so

$$\limsup_n \ \mathrm{Succ}_{\mathrm{Bayes}}(m_n) \ \le \ \lim_n \ \mathrm{Succ}_{\mathrm{Bayes}}(m(n,c)) \ = \ e^{-e^{-c}/2},$$

and $c \to -\infty$ gives $\mathrm{Succ}_{\mathrm{Bayes}}(m_n) \to 0$, and any algorithm is dominated by Bayes success via (ii). (No per-target claim is made.) The side $x_n \to +\infty$ is symmetric, using $m_n \ge m(n,c)$ and $c \to +\infty$, and the achievability half is witnessed by the rule "output any surviving assignment," whose success tends to 1 because uniqueness does. For the "only if" direction, suppose $x_n \not\to +\infty$ and pick B and a subsequence with $x_{n_j} \le B$. Since $m(n_j, B+1)\widetilde{q}_1 = \ln n_j + B + 1 + o(1)$, eventually $m_{n_j} \le m(n_j, B+1)$, so by monotonicity the Bayes success along the subsequence has $\limsup \le e^{-e^{-(B+1)}/2} < 1$. Hence the optimal success does not tend to 1, and by (ii) neither does any algorithm's prior average. □

Remark C.12 (Sampling variants). All claims are for sampling with replacement. For $k \ge 3$, the without-replacement variant is equivalent up to total variation $O((\ln n)^2/n^{k-2}) = o(1)$ (a birthday bound, since with $m = \Theta(n \log n)$ draws from $N = \Theta(n^k)$ clauses the probability of any repeat vanishes), so the theorems transfer without change. For $k = 2$ the two samplings are *not* asymptotically equivalent in total variation, because the expected number of repeated clauses in the window is $\frac{3}{4}(\ln n + x)^2 \to \infty$. For $k = 2$ without replacement and for the binomial variant at every k, the theorems are therefore stated only for sampling with replacement. We expect the limit laws to hold for these variants as well, since discarding duplicates yields a uniform without-replacement sample whose random size fluctuates by $O_p((\ln n)^2)$, far below the window sensitivity of $\Theta(n)$ clauses, and the monotone coupling in part (iv) above applies to it, but this is a conjecture and not a proved statement.

C.14 Setting and Notation for Non-Uniform Coverage

Scopes follow the conditional product law $\nu(S) \propto \prod_{i \in S} w_i$ with $0 < c \le w_i \le C$ and c, C, k fixed, patterns are uniform among the $2^k - 1$ non-planted ones given the scope, and WLOG $\sigma = 0$. Write e_j for the jth elementary symmetric polynomial, w_{-A} for the weight vector with the indices in A removed, $\xi_i = \mathbb{P}(i \in S) = w_i e_{k-1}(w_{-i})/e_k(w)$, $q_i = \xi_i/(2^k - 1)$ for the single-flip violation probabilities, $q_A = [1 - \mathbb{P}(S \cap A = \varnothing)]/(2^k - 1)$, $\widetilde{q}_i = -\ln(1 - q_i)$, and $\Lambda_m = \sum_i (1 - q_i)^m$. Constants $C_1, C_2, \ldots$ and δ, γ, c_0 depend only on (c, C, k), and a displayed μ marks dependence on the window position as well. The proofs of Thm. 4.3 and Cor. 4.4 follow Apps. C.11 and C.13 step by step, and App. C.17 states exactly what changes.

C.15 Properties Preserved Under Arbitrary Scope Laws

Lemma C.13 (Budget conservation)**.** *For every scope law,* $\mathbb{P}(|D| = j) = \binom{k}{j}/(2^k - 1)$ *for* $j = 1, \ldots, k$*, and in particular* $\sum_i q_i = k/(2^k - 1)$.

Proof. Given any scope, exactly $\binom{k}{j}$ of the $2^k - 1$ forbidden patterns have $|D| = j$, so $\mathbb{P}(|D| = j \mid S)$ does not depend on S, and $\sum_i \mathbb{P}(D = \{i\}) = \mathbb{P}(|D| = 1)$. □

What is conserved is the per-clause single-flip discrimination mass, and only the $j = 1$ layer enters the proofs. The per-clause mutual information is conserved as well, which is worth recording because it is easy to guess otherwise. Under the uniform prior on σ the scope S is independent of σ, the pattern f is uniform on $\{0,1\}^k$ given S once σ is averaged out, and uniform on the $2^k - 1$ patterns other than $\sigma|_S$ given (S, σ), so

$$I(\sigma; Z_1) = I(\sigma; f \mid S) = k \ln 2 - \ln(2^k - 1)$$

under every scope law. The mixture information $I(Y; Z_1)$ of Lem. C.1 is conserved too. Its proof uses only the counts $c(z)$, which depend on (S, f) and not on the scope law, and the probabilities of the three clause classes, which under $Y = \tau_i$ are q_i, $k/(2^k - 1) - q_i$, and the remainder, so summing over the $n + 1$ hypotheses leaves only $\sum_i q_i = k/(2^k - 1)$. What the allocation changes is therefore not the information carried by one clause but the joint information $I(\sigma; Z_{1:m}) = n \ln 2 - \mathbb{E} \ln |\mathcal{S}|$, since the posterior is uniform on $\mathcal{S}$ (Lem. C.14) and on the structure event $\ln |\mathcal{S}| = W \ln 2$ with

$\mathbb{E}W = \Lambda_m$. Two allocations with identical per-clause information can thus sit at different window positions after the same number of clauses, and neither $I(\sigma; Z_1)$ nor the product $m\, I(\sigma; Z_1)$ should be read as the total information the observations carry about the target.

Lemma C.14 (Flatness and equivariance). *Lem. C.7(c) holds without change, and the model remains equivariant under the transitive $(\mathbb{Z}_2)^n$ action of App. C.13.*

Proof. For a fixed candidate τ and scope S, exactly $2^k - 1$ patterns are satisfied by τ, each of probability $\nu(S)/(2^k - 1)$, so the per-clause normalization $\sum_S \nu(S) = 1$ is the same for every candidate, and every consistent assignment has likelihood $\prod_t [\nu(S_t)/(2^k - 1)]$. The global flip s maps (S, f) to $(S, f \oplus s|_S)$, leaves scopes and pattern uniformity untouched, and acts transitively, so the invariance reduction and the symmetrization argument of App. C.13 apply unchanged. Coordinate-permutation symmetry is lost and is used nowhere in that argument. □

Lemma C.15 (Exact survival and negative association). *For distinct $i_1, \dots, i_j$, $\mathbb{P}(i_1, \dots, i_j \in F) = (1 - \sum_l q_{i_l})^m$, and the survival indicators are negatively associated.*

Proof. Per clause the events $\{D = \{i\}\}$ are disjoint, so the single flip that violates a clause is one categorical draw with cell probabilities $(q_1, \dots, q_n, 1 - \sum_i q_i)$, i.i.d. across clauses, and the per-coordinate violation counts are multinomial. The product formula is immediate. Multinomial counts are negatively associated (Joag-Dev & Proschan, 1983), and the survival indicators are decreasing functions of disjoint blocks. □

C.16 Scope Probabilities and Survival Estimates

Lemma C.16 (Marginals, pinning, and tail scopes). *With $\bar{w} = \frac{1}{n}\sum_j w_j$, the following hold.*

(a) $\xi_i = \frac{k w_i}{n \bar{w}}\big(1 + O(1/n)\big)$ uniformly in i, so $C_1^{-1}/n \le q_{\min} \le q_{\max} \le C_1/n$.

(b) Conditioned on $i \in S$, the remaining $k - 1$ scope coordinates follow the conditional product law with weights w_{-i}.

(c) For $i \in A$ with $|A| = r$, $\rho_i(A) = \mathbb{P}(S \cap (A \setminus \{i\}) = \varnothing \mid i \in S)$ satisfies $\rho_i(A) \ge 1 - C_3 r/n$, and $\rho_i(A) \ge \delta > 0$ for all $r \le n/2$.

(d) For $|A| > n/2$, $q_A \ge \gamma/(2^k - 1)$ for some $\gamma > 0$.

(e) For $|B| = d \le k$, $\mathbb{P}(B \subseteq S) \le C_4^d \prod_{i \in B} \xi_i$.

Proof. (a) Comparing $k!\, e_k(w)$, a sum over distinct k-tuples, with $(\sum_j w_j)^k$, a sum over all k-tuples, the tuples containing a repeat contribute at most $\binom{k}{2}(\sum_j w_j^2)(\sum_j w_j)^{k-2}$, and $\sum_j w_j^2 \le \frac{C^2}{c^2 n}(\sum_j w_j)^2$, so $e_k(w) = \frac{(n\bar{w})^k}{k!}(1 + O(1/n))$ and $e_{k-1}(w_{-i}) = \frac{(n\bar{w})^{k-1}}{(k-1)!}(1 + O(1/n))$, giving the display. (b) The factor w_i cancels from the conditional law. (c) The first bound is a union bound over $A \setminus \{i\}$ using the marginals of the pinned law from (b), and the second is the ratio bound $e_{k-1}(w_{-i} \setminus A)/e_{k-1}(w_{-i}) \ge (c/C)^{k-1}\big((n - r - k)/n\big)^{k-1}$, bounded below for $r \le n/2$. (d) The law equals k i.i.d. draws from $p_i \propto w_i$ conditioned on distinctness, a $1 + O(1/n)$ correction, so $\mathbb{P}(S \cap A = \varnothing)$ is at most $(1 + o(1))$ times the k-th power of the weight-mass fraction of A^c, and $|A| > n/2$ with $c \le w_i \le C$ caps that fraction at $1 - \frac{c}{2C}$, so $\gamma = \frac{1}{2}\big[1 - (1 - \frac{c}{2C})^k\big]$ works for n large and depends only on (c, C, k). A count ratio in place of the mass ratio fails here once $C \ge 2c$. (e) Direct from $\mathbb{P}(B \subseteq S) = (\prod_{i \in B} w_i)\, e_{k-d}(w_{-B})/e_k(w)$ and the termwise bounds $\binom{n}{j}c^j \le e_j \le \binom{n}{j}C^j$. □

Lemma C.17 (Exponential approximation and uniform rarity). *If $\Lambda_m \to \mu \in (0, \infty)$, then $m \asymp n \ln n$, $m q_{\min} \ge c_9 \ln n$ for some $c_9 > 0$, $\max_i (1 - q_i)^m \to 0$, and $e^{-m q_i} = (1 + O(\ln n/n))(1 - q_i)^m$ uniformly in i. Consequently*

$$A_m := \sum_i e^{-m q_i} = (1 + o(1))\, \Lambda_m, \qquad \sum_i m q_i e^{-m q_i} \;\le\; m q_{\max} A_m \;=\; O(\mu \ln n).$$

Proof. From $ne^{-m\widetilde{q}_{\max}} \le \Lambda_m \le ne^{-m\widetilde{q}_{\min}}$ and $\Lambda_m \to \mu$, we get $m\widetilde{q}_{\max} \ge \ln n - O(1)$ and $m\widetilde{q}_{\min} \le \ln n + O(1)$, and with the two-sided marginal bound of Lem. C.16(a) this forces $m \asymp n \ln n$ and $m\widetilde{q}_{\min} \ge C_1^{-2}(\ln n - O(1))$, so $\max_i (1-q_i)^m \le n^{-C_1^{-2}+o(1)} \to 0$. The bridge follows from $mq_{\max}^2 = O(\ln n/n)$, since $e^{-mq_i} = (1-q_i)^m e^{O(mq_i^2)}$. The last display then follows by summing, and this route through A_m is necessary, because $(1-q)^m < e^{-mq}$ makes the direct comparison with Λ_m run in the wrong direction. □

C.17 Proof of Theorem 4.3 (Heterogeneous Coverage Window)

The steps are those of the proofs of Thms. 4.1 and 4.2 in Apps. C.11 and C.13, and only the bookkeeping changes.

Proof. The Poisson limit. By Lem. C.15, $\mathbb{E}[(W)_j] = \sum^{\neq}(1-\sum_l q_{i_l})^m$ over ordered distinct j-tuples. By the bridge, each term equals $\prod_l e^{-mq_{i_l}}\big(1+O(j^2 \ln n/n)\big)$, and expanding the unrestricted product sum and removing tuples with a coincidence gives $\sum^{\neq}\prod_l e^{-mq_{i_l}} = A_m^j + O\big(j^2\delta_n A_m^{j-1}\big)$ with $\delta_n = \max_i e^{-mq_i} \to 0$. Combining the two,

$$\mathbb{E}[(W)_j] = A_m^j + O\Big(\frac{j^2 \ln n}{n} A_m^j\Big) + O\big(j^2\delta_n A_m^{j-1}\big),$$

both error terms vanish for fixed j, so $\mathbb{E}[(W)_j] \to \mu^j$, and $W \Rightarrow \mathrm{Poisson}(\mu)$ by moment determinacy.

Spurious survivors. The spurious-survivor lemmas of App. C.10 apply with the per-coordinate categories

$$D_i(A) = \{D = \{i\},\, S \cap A = \{i\}\}, \quad d_i(A) = q_i\rho_i(A),$$
$$G_i(A) = \{D = \{i\},\, S \cap A \supsetneq \{i\}\}, \quad g_i(A) = q_i\big(1-\rho_i(A)\big),$$

replacing d_r and g_r, with $g_i(A) \le C_3 q_i r/n$ by Lem. C.16(c), and the bound (3) becomes

$$\mathbb{P}(O_A) \;\le\; \Big(1-\sum_{i\in A} d_i(A)\Big)^m \min\Big\{1,\, \prod_{i\in A} m g_i'(A)\Big\}, \qquad g_i'(A) = \frac{g_i(A)}{1-\sum_{j\in A} d_j(A)},$$

by the same negative-association step, which is unchanged because it uses only the structure of one clause. For fixed A the categories $D_i(A)$ and $G_i(A)$, $i \in A$, are pairwise disjoint events of a single clause under every scope law, since D names the coordinate. On O_A every $D_i(A)$ count vanishes and every $i \in A$ needs a $G_i(A)$ clause. Conditioned on the absence of $D_i(A)$ clauses the m clauses remain i.i.d., the counts $(X_{G_i})_{i\in A}$ are multinomial with cell probabilities $g_i'(A)$ and hence negatively associated (Joag-Dev & Proschan, 1983), and the events $\{X_{G_i} \ge 1\}$ are increasing in disjoint coordinates, which gives the product bound with $\mathbb{P}(X_{G_i} \ge 1 \mid \text{no } D) \le m g_i'(A)$. Exactly two bookkeeping devices change. First, in Region I ($2 \le r \le r_0 = \lfloor c_0 n/\ln n \rfloor$) the uniform proof sums over A as $\binom{n}{r}$ times a common factor, which is unavailable here because the factor depends on A. The per-A bound is therefore made explicit before summing. By Lem. C.16(c), $d_i(A) \ge q_i(1-C_3 r/n)$ and $g_i(A) \le C_3 q_i r/n$. By Lem. C.16(a) and the proof of Lem. C.17, $q_{\max} \le C_1/n$ and $mq_{\max} \le C_5 \ln n$ with $C_5 = 2C_1^2$ for n large, since $q_{\max} \le C_1^2 q_{\min} \le C_1^2 \widetilde{q}_{\min}$ and $m\widetilde{q}_{\min} \le \ln n + O(1)$. Hence for $r \le r_0$ and n large, $\sum_{j\in A} d_j(A) \le r q_{\max} \le C_1 c_0/\ln n \le \frac{1}{3}$, so $g_i'(A) \le \frac{3}{2} g_i(A) \le \frac{3}{2} C_3 q_i r/n$, while

$$\Big(1-\sum_{i\in A} d_i(A)\Big)^m \le \exp\Big(-m\sum_{i\in A} q_i + C_3\frac{r}{n}\, r\, mq_{\max}\Big) \le e^{C_6 r}\prod_{i\in A} e^{-mq_i}, \qquad C_6 = C_3C_5c_0,$$

using $r \ln n/n \le c_0$ in Region I. Inserting both estimates into the bound on $\mathbb{P}(O_A)$ gives

$$\mathbb{P}(O_A) \le e^{C_6 r}\prod_{i\in A} e^{-mq_i}\prod_{i\in A}\frac{3C_3 r}{2n}\, mq_i = e^{C_6 r}\Big(\frac{3C_3 r}{2n}\Big)^r \prod_{i\in A} b_i, \qquad b_i = mq_i e^{-mq_i}.$$

Write $\Sigma_r = \sum_{|A|=r}\mathbb{P}(O_A)$ for the complete spurious-survivor sum at distance r. Summing the per-A bound over all A of size r produces the elementary symmetric polynomial of the vector b,

$$\Sigma_r \;\le\; e^{C_6 r}\Big(\frac{3C_3 r}{2n}\Big)^r e_r(b), \qquad e_r(b) \le \frac{(\sum_i b_i)^r}{r!} \le \Big(\frac{e\sum_i b_i}{r}\Big)^r,$$

where a worst-member bound times $\binom{n}{r}$ would diverge once $c < C$. Second, the estimate $\sum_i b_i = O(\mu \ln n)$ must be routed through the bridge of Lem. C.17, whose last display gives $\sum_i b_i \le m q_{\max} A_m \le C_5(1+o(1))\mu \ln n$, so that Λ_m controls the Region-I error through $A_m = (1+o(1))\Lambda_m$. Combining the two displays,

$$\Sigma_r \;\le\; \Big(\tfrac{3}{2}e^{C_6+1}C_3\,\frac{\sum_i b_i}{n}\Big)^r = b_n^r, \qquad b_n \;\le\; \tfrac{3}{2}e^{C_6+1}C_3C_5(1+o(1))\,\frac{\mu\ln n}{n} \;\to 0,$$

so $\sum_{r\ge 2}\Sigma_r \le 2b_n^2$, matching the uniform Region-I bound with μ in place of e^{-x}. For partial spurious survivors the union bound over the resolved coordinate replaces the product $\prod_{i\in A} m g_i'(A)$ by the sum $\sum_{i\in A} m g_i'(A)$. Set $a_i = e^{C_6}e^{-mq_i}$ and $c_i = \frac{3}{2}C_3\frac{r}{n}\,mq_i$ at the current r, so that the per-A bound is at most $(\prod_{i\in A} a_i)\sum_{i\in A} c_i$, and the sum over A is organized as

$$\sum_{|A|=r}\Big(\prod_{i\in A}a_i\Big)\sum_{i\in A}c_i = \sum_i c_i a_i e_{r-1}(a_{-i}) \le \Big(\sum_i c_i a_i\Big)e_{r-1}(a).$$

By the last display of Lem. C.17, $\sum_i c_i a_i \le \frac{3}{2}C_3 e^{C_6}\frac{r}{n}\sum_i mq_i e^{-mq_i} = O(r\mu\ln n/n)$, and $e_{r-1}(a) \le \big(e\sum_i a_i/(r-1)\big)^{r-1} \le \big(C_7\mu/(r-1)\big)^{r-1}$ since $\sum_i a_i = e^{C_6}A_m = O(\mu)$. The per-$r$ total is therefore $O(\mu\ln n/n)\cdot r\big(C_7\mu/(r-1)\big)^{r-1}$, whose series over $r \ge 2$ converges, with super-exponential decay once $r > C_7\mu$, so the Region-I partial spurious-survivor total is $O_{c,C,k,\mu}(\ln n/n)$, and the constant depends on μ exactly as the uniform constant depends on x. In Region II ($r_0 < r \le n/2$), Lem. C.16(c) gives $\sum_{i\in A} d_i(A) \ge \delta r q_{\min}$ and Lem. C.17 gives $mq_{\min} \ge c_9 \ln n$, so $\Sigma_r \le [(e/c_0)(\ln n)\, n^{-\delta c_9}]^r$, superpolynomially small in total. In Region III ($r > n/2$), Lem. C.16(d) replaces the count-ratio bound of Lem. C.7(b) by a mass-ratio bound with the explicit $\gamma = \frac{1}{2}[1-(1-\frac{c}{2C})^k]$ given in its proof. Dropping the G factor, $\mathbb{P}(O_A) \le \mathbb{P}(\tau_A \in \mathcal{S}) = (1-q_A)^m \le (1-\gamma/(2^k-1))^m$, and since $m \ge c_{10} n\ln n$ for some $c_{10} > 0$ by Lem. C.17, $\sum_{r>n/2}\Sigma_r \le 2^n(1-\gamma/(2^k-1))^m \le \exp\big(n\ln 2 - c_{10}\gamma\, n\ln n/(2^k-1)\big) \to 0$, with every constant depending only on (c, C, k). The partial spurious-survivor events in Regions II and III are contained in the corresponding no-D and survival events, so they pay only an extra factor $r \le n$, which the superpolynomial decay absorbs. For multi-pattern violations no negative association is needed, since the bound is a first moment over clauses. Lem. C.16(e) gives $\mathbb{P}(D = B) \le C_4^d(2^k-1)^{d-1}\prod_{i\in B} q_i$ for $|B| = d$, and the first moment over clauses and $d \ge 2$ is $O(\mu^2\ln n/n)$ with the $d = 2$ term dominating, using $\sum_i q_i e^{-mq_i} \le q_{\max}A_m = O(\mu/n)$.

Assembly. The assembly of App. C.11 is purely logical and applies without change, since it uses only the kill rule (1), the multi-pattern bound, and the partial spurious-survivor bound, all established above under the general scope law, so the structure probability is $1 - O_{c,C,k,\mu}(\ln n/n)$, and on the structure event $|\mathcal{S}| = 2^W$, giving $\mathbb{P}(|\mathcal{S}| = 1) \to e^{-\mu}$.

Recovery curves. Parts (i) to (iii) of App. C.13 apply without change, using Lem. C.14 for flatness and equivariance. Part (i) uses only that the posterior is uniform on $\mathcal{S}$ and that $|\mathcal{S}| = 2^W$ on the structure event, part (ii) is the definition of Bayes success, and part (iii) uses only the transitive $(\mathbb{Z}_2)^n$ action, which Lem. C.14 preserves, while the coordinate-permutation symmetry of the uniform model enters nowhere. Hence Bayes and minimax recovery tend to $e^{-\mu/2}$, and if $\Lambda_m \to \infty$ then $W \xrightarrow{p} \infty$ (mean Λ_m, variance at most the mean by negative association), so every estimator's prior-average success is at most $\mathbb{E}[1/(W+1)] \to 0$. For $\Lambda_m \to 0$, the monotone coupling of App. C.13(iv) closes the tail without any spurious-survivor sums. On a single clause stream, the uniqueness probability and the Bayes success are nondecreasing in m. Fix $\ell \in (0,1)$ and let $m_n(\ell) = \min\{m : \Lambda_m \le \ell\}$. Since

$$\Lambda_{m-1} - \Lambda_m = \sum_i q_i(1-q_i)^{m-1} \;\le\; q_{\max}\Lambda_{m-1}, \qquad q_{\max} = O(1/n),$$

we get $\Lambda_{m_n(\ell)} \to \ell$, so the fixed-$\mu$ theorem applies along $m_n(\ell)$ with $\mu = \ell$. Any sequence with $\Lambda_{m_n} \to 0$ eventually has $m_n \ge m_n(\ell)$, hence $\liminf_n \mathbb{P}(|\mathcal{S}| = 1) \ge e^{-\ell}$ and $\liminf_n$ of the Bayes success is at least $e^{-\ell/2}$, and $\ell \downarrow 0$ gives both limits equal to one. □

C.18 A Counterexample Without Uniform Rarity

The bounded-weight condition supplies, through Lem. C.17, the uniform rarity $\max_i(1-q_i)^m \to 0$ that the Poisson limit needs, and the following example shows the pure Poisson form can fail without

it. Take $w_1 = 0$ and $w_2, \dots, w_n \in [c, C]$. No scope contains coordinate 1, the induced law on the other coordinates is again conditional product with the same k, and the violation criterion (1) concerns $A \cap \{2, \dots, n\}$ only, so exactly

$$\mathcal{S} = \mathcal{S}' \times \{0, 1\}, \qquad |\mathcal{S}| = 2|\mathcal{S}'|, \qquad W = 1 + W',$$

with primes for the $(n-1)$-coordinate system. If $\Lambda'_m \to \mu_0 \in (0, \infty)$, Thm. 4.3 applies to the primed system, so $\Lambda_m = 1 + \Lambda'_m \to 1 + \mu_0$ while $W \Rightarrow 1 + \mathrm{Poisson}(\mu_0)$, the uniqueness probability is zero at every m, and Bayes recovery tends to $\frac{1}{2}e^{-\mu_0/2}$, capped at one half by the fair bit the observations never touch. The factorial moments converge to $\mathbb{E}[(W)_j] \to \mu_0^j + j\mu_0^{j-1}$, which differs from $(1 + \mu_0)^j$ already at $j = 2$. We do not claim that uniform rarity is logically necessary for a Poisson limit.

C.19 Proof of Corollary 4.4 (Equal Coverage and the Cost of Skew)

Part (c) extends to window-scale skew. For fixed $\zeta > 0$, the two-block law with $\beta_n = 1 - \zeta/\ln n$ inflates Λ_m at the uniform window centre by $\alpha e^{(1-\alpha)\zeta} + (1-\alpha)e^{-\alpha\zeta}$, which is $1 + \frac{\alpha(1-\alpha)}{2}\zeta^2 + O(\zeta^3)$ as $\zeta \to 0$. This is derived at the end of part (c) below.

Part (a) is the first-moment form of an optimality that the coupon-collector literature states in stochastic order. Per clause, the resolved single flip is one categorical draw with cells $(q_1, \dots, q_n, 1 - \sum_i q_i)$ (Lem. C.15), and the null cell has probability $1 - k/(2^k - 1)$ under every scope law (Lem. C.13). This is the generalized coupon collector with a null coupon of Anceaume et al. (2015), who prove that the almost-uniform law, with all non-null probabilities equal, stochastically minimizes the time to complete the collection, so that $\mathbb{P}(W = 0)$ is maximized by equal q_i at every n and m, and Anceaume et al. (2016) extend this to the time to collect any fixed number of coupons, which makes W stochastically smallest under equal q_i at every m. The inequality of part (a) follows from that ordering, the equality case is Jensen's, and the direct proof below is kept because the clock needs only the mean. The ordering concerns the single-flip count W and does not by itself order the true curves $\mathbb{P}(|\mathcal{S}| = 1)$ and $\mathbb{E}[1/|\mathcal{S}|]$, which depend on the structure event of Thm. 4.3, and this is where the SAT geometry enters.

Proof. *(a)* By Lem. C.15, $\mathbb{E}W = \sum_i \mathbb{P}(i \in F) = \sum_i (1 - q_i)^m = \Lambda_m$. Lem. C.13 fixes $\sum_i q_i = k/(2^k - 1)$ for every scope law, each $q_i \le 1/(2^k - 1) < 1$, and $t \mapsto (1 - t)^m$ is strictly convex on $[0, 1)$ for $m \ge 2$, so Jensen's inequality gives $\Lambda_m \ge n(1 - \bar{q})^m$ with equality iff all q_i are equal, a value the uniform law attains exactly. Since $\Lambda_m^{\mathrm{law}} \ge \Lambda_m^{\mathrm{unif}}$ pointwise in m, the hitting time $\min\{m : \Lambda_m \le \mu\}$ can only increase, at every finite n. The increase is weak. The strict statement is the clock inequality, and integer hitting times can coincide for a small skew, because one clause moves Λ_m by a factor $1 - O(1/n)$ only. At $n = 200$, $k = 3$, $\alpha = \frac{1}{2}$ and $\mu = 1$, the uniform law and the two-block law with $\beta = 0.999$ both first reach $\Lambda_m \le 1$ at $m = 2470$, while $\beta = 0.99$ needs 2471, and the clause premium of part (c) is the asymptotic size of the increase for a fixed skew. At $m = 1$, every law gives the same $\Lambda_1 = n - k/(2^k - 1)$, so one clause cannot distinguish allocations through $\mathbb{E}W$.

(b) Splitting off coordinate j via $e_{k-1}(w_{-i}) = e_{k-1}(w_{-ij}) + w_j e_{k-2}(w_{-ij})$, and symmetrically for i, makes the cross terms cancel in $w_i e_{k-1}(w_{-i}) - w_j e_{k-1}(w_{-j})$, giving the identity $\xi_i - \xi_j = (w_i - w_j)\, e_{k-1}(w_{-ij})/e_k(w)$. For positive weights, $e_{k-1}(w_{-ij}) > 0$ requires $n - 2 \ge k - 1$. At $n = k$ the scope is deterministic, every marginal equals one for any weights, and identifiability genuinely fails. Weights are identified up to the common scaling $w \mapsto tw$, which leaves the law unchanged.

(c) Put $n_1 = \lceil \alpha n \rceil$ coordinates at weight β and $n_2 = n - n_1$ at weight one. Write q_ℓ, q_{h}, and $q_{\mathrm{u}} = q_1$ for the single-flip violation probabilities of a low-weight coordinate, a high-weight coordinate, and a coordinate under uniform weights. Also write $\widetilde{q}_s = -\ln(1 - q_s)$ for $s \in \{\ell, \mathrm{h}, \mathrm{u}\}$. Lem. C.16(a) gives $q_\ell = \frac{k\beta}{n\bar{w}(2^k - 1)}(1 + O(1/n))$ and $q_{\mathrm{h}} = \frac{k}{n\bar{w}(2^k - 1)}(1 + O(1/n))$, so $\widetilde{q}_{\mathrm{h}}/\widetilde{q}_\ell = \beta^{-1}(1 + O(1/n))$. Whenever $\Lambda_m \le B_0$, the bound $n_1 e^{-m\widetilde{q}_\ell} \le B_0$ forces $m\widetilde{q}_\ell \ge \ln n - O(1)$, and then

$$n_2 e^{-m\widetilde{q}_{\mathrm{h}}} \;\le\; n_1 e^{-m\widetilde{q}_\ell} \cdot \frac{1 - \alpha}{\alpha}\, n^{-(\beta^{-1} - 1)(1 - o(1))}, \qquad \text{so} \qquad \Lambda_m = n_1(1 - q_\ell)^m\big(1 + O(n^{-\kappa})\big)$$

with $\kappa = \frac{1}{2}(\beta^{-1} - 1)$, and the window is written by the thin layer, at $m\widetilde{q}_\ell = \ln(\alpha n) - \ln \mu + o(1)$. With the uniform window at $m\widetilde{q}_{\mathrm{u}} = \ln n - \ln \mu + O(\widetilde{q}_{\mathrm{u}})$, the hitting times compare as

$$\frac{m^{(\beta)}(\mu)}{m^{\mathrm{unif}}(\mu)} = \frac{\bar{w}}{\beta} \cdot \frac{\ln(\alpha n) - \ln \mu + o(1)}{\ln n - \ln \mu + o(1)} \longrightarrow \frac{\bar{w}}{\beta},$$

and halving Λ_m costs $\ln 2/\widetilde{q}_\ell$ clauses by the dominance display, so the uniqueness premium is $\Delta m = (\bar{w}/\beta)\ln 2/\widetilde{q}_{\mathrm{u}}\,(1 + o(1))$, inflated by the same factor as the window. At finite n the leading constant can be sharpened by replacing $\bar{w}/\beta$ with the exact $\widetilde{q}_{\mathrm{u}}/\widetilde{q}_\ell$, a refinement rather than part of the theorem. At $n = 200$, $\alpha = \frac{1}{2}$, $\mu = 1$, the displayed formula with its $o(1)$ terms dropped gives 1.304 and 2.173 for $\beta = \frac{1}{2}$ and $\frac{1}{4}$, the refinement gives 1.298 and 2.153, and the ratios computed from the exact Λ_m hitting times are 1.302 and 2.153. For fixed $\zeta > 0$, consider the window-scale skew $\beta_n = 1 - \zeta/\ln n$. The same marginal estimates give $\beta_n/\bar{w}_n = 1 - (1-\alpha)\zeta/\ln n + O(1/\ln^2 n)$ and $1/\bar{w}_n = 1 + \alpha\zeta/\ln n + O(1/\ln^2 n)$. The exponential bridge gives $m(\widetilde{q}_s - q_s) = O(\ln n/n)$ uniformly for $s \in \{\ell, \mathrm{h}, \mathrm{u}\}$. At the uniform window centre $m\widetilde{q}_{\mathrm{u}} = \ln n + x + o(1)$ we get $m\widetilde{q}_\ell = \ln n + x - (1-\alpha)\zeta + o(1)$ and $m\widetilde{q}_{\mathrm{h}} = \ln n + x + \alpha\zeta + o(1)$, whence

$$\Lambda_m \longrightarrow e^{-x}\big[\alpha e^{(1-\alpha)\zeta} + (1-\alpha)e^{-\alpha\zeta}\big],$$

and the bracket equals $1 + \frac{\alpha(1-\alpha)}{2}\zeta^2 + O(\zeta^3)$ as $\zeta \to 0$. □

D Exact Finite-n Formulas and Numerical Protocol

The experimental settings of this appendix use $k = 3$, and the exact formulas below hold for general k. With $q_1, \widetilde{q}_1, m$ as in the main text and $p_{\mathrm{pair}} = \frac{k(k-1)}{n(n-1)(2^k-1)}$ (the probability that a clause has $D = \{i\}$ and $j \in S$, which equals the probability of $D = \{i, j\}$), and with $q^{(2)}$ as in Lem. C.7(b) of App. C.9, the exact quantities are as follows.

- The factorial-moment generating function is $\mathbb{E}[z^W] = \sum_{j=0}^{n} \binom{n}{j} (z-1)^j (1 - jq_1)^m$, which gives $\mathbb{E}W = n(1-q_1)^m$, $\mathbb{P}(W = 0) = \mathbb{E}[0^W]$ at $z = 0$, and $\mathbb{E}[2^{-W}]$ at $z = \frac{1}{2}$.
- Spurious pairs with one resolved coordinate, ordered (i, j) with i resolved, j unresolved and $\tau_{\{i,j\}}$ surviving, $2\binom{n}{2}\big[(1 - q^{(2)} - p_{\mathrm{pair}})^m - (1 - q^{(2)} - 2p_{\mathrm{pair}})^m\big]$.
- Spurious pairs with both coordinates resolved: $\binom{n}{2}\big[(1-q^{(2)})^m - 2(1 - q^{(2)} - p_{\mathrm{pair}})^m + (1 - q^{(2)} - 2p_{\mathrm{pair}})^m\big]$.
- Subcube-violation *pairs* (both flips surviving, at least one $D = \{i, j\}$ clause): $\binom{n}{2}\big[(1 - 2q_1)^m - (1 - 2q_1 - p_{\mathrm{pair}})^m\big]$.
- Subcube-violation *clause occurrences* (with multiplicity): $\binom{n}{2}\, m\, p_{\mathrm{pair}}\, (1 - 2q_1)^{m-1}$.

The last two are different statistics (at $n = 60$, $x = 0$ their expectations over 6000 trials are 375.71 and 402.63), and both are reported, each against its own formula. The simulation encodes clauses as scope triples with a difference-pattern index, computes W by marking resolved coordinates, detects spurious pairs by intersecting the scopes of the violated clauses, and uses fixed seeds. The configurations $(n, x, \mathrm{trials}, \mathrm{seed})$ are $(60, 0, 6000, 1)$, $(200, -1, 1500, 2)$, $(200, 0, 1500, 3)$, $(200, 1, 1500, 4)$, $(800, 0, 600, 5)$, $(800, 1, 600, 6)$. The simulation implements $k = 3$ only, while the exact formulas cover general k. Tab. A.2 reports the pair-level counts.

The prediction column of Tab. A.6 in App. F is $1 - \exp[-(\mathrm{EN}_A + \mathrm{EN}_B)]$, where EN_A is the subcube-violation clause-occurrence expectation above and EN_B is the sum of the two spurious-pair expectations, which telescopes to $\binom{n}{2}\big[(1-q^{(2)})^m - (1 - q^{(2)} - 2p_{\mathrm{pair}})^m\big]$. The composition is a Poisson approximation rather than an identity, and for $k = 3$ at $x = 0$ it is asymptotic to $3\ln n/n$, one unit from the size-two patterns and two from the distance-two survivors.

Table A.2: Pair-level diagnostics, empirical (e) against exact finite-n (x) counts, totals over trials.

(n, x), trials	one resolved e/x	both resolved e/x	violation pairs e/x	violation clauses e/x
$(60, 0)$, 6000	726 / 751.4	50 / 56.5	365 / 375.7	383 / 402.6
$(200, -1)$, 1500	457 / 463.0	12 / 10.3	243 / 231.5	246 / 236.5
$(200, 0)$, 1500	83 / 76.7	0 / 2.1	41 / 38.4	42 / 39.4
$(200, 1)$, 1500	12 / 12.3	0 / 0.4	8 / 6.1	9 / 6.3
$(800, 0)$, 600	11 / 9.9	0 / 0.08	5 / 5.0	5 / 5.0
$(800, 1)$, 600	2 / 1.5	0 / 0.01	2 / 0.8	2 / 0.8

Table A.3: Empirical vs. exact finite-n vs. limiting values of the three window statistics ($k = 3$). "exact" uses the closed-form finite-n formulas, and "limit" is the theorem curve at the same x.

(n, x), trials	$\mathbb{E}W$ emp / exact / limit	$\mathbb{P}(W{=}0)$ emp / exact / limit	$\mathbb{E}2^{-W}$ emp / exact / limit
$(60, 0)$, 6000	0.980 / 1.001 / 1	0.359 / 0.359 / 0.368	0.607 / 0.603 / 0.607
$(200, -1)$, 1500	2.673 / 2.717 / 2.718	0.068 / 0.063 / 0.066	0.262 / 0.254 / 0.257
$(200, 0)$, 1500	1.019 / 1.000 / 1	0.356 / 0.365 / 0.368	0.601 / 0.605 / 0.607
$(200, 1)$, 1500	0.342 / 0.368 / 0.368	0.705 / 0.691 / 0.692	0.842 / 0.832 / 0.832
$(800, 0)$, 600	1.025 / 1.000 / 1	0.363 / 0.367 / 0.368	0.601 / 0.606 / 0.607
$(800, 1)$, 600	0.358 / 0.368 / 0.368	0.705 / 0.692 / 0.692	0.838 / 0.832 / 0.832

D.1 Verification of the W-Statistic Laws at Finite n

Because every quantity in Lem. C.7(a) is exact at finite n, the limit laws for the W-statistics can be checked against closed-form finite-n values rather than only against asymptotics. This is a check of the dimension laws and not of the solution-set statements, which have no closed form and are checked instead by the exhaustive enumeration of Sec. 5.1 and App. D.2. The factorial-moment generating function gives $\mathbb{E}[z^W] = \sum_j \binom{n}{j}(z-1)^j(1-jq_1)^m$, and the spurious-survivor and subcube-violation events admit exact pair-level inclusion-exclusion formulas (stated earlier in this appendix). We simulated the model at $k = 3$ for six configurations (n, x) with $n \in \{60, 200, 800\}$ and $x \in \{-1, 0, 1\}$, with fixed seeds (600 to 6000 trials per configuration). Tab. A.3 reports the three window statistics. All empirical values match the exact finite-n formulas within Monte Carlo error, and the finite-n values visibly converge to the limit curves.

Two further checks are diagnostic. First, the spurious-pair counts (the error events of Lem. C.8 at distance two) match their exact pair-level formulas across all configurations, and their scaling decisively distinguishes the correct asymptotic constant $\propto (k-1)^2$ from the incorrect $\propto k^2$. At $(n, x) = (60, 0)$ over 6000 trials the observed both-resolved count is 50 against a $(k-1)^2$-prediction of 55.9 and a k^2-prediction of 125.7. Second, we track *two distinct* subcube-violation statistics, the number of violating unordered pairs and the number of violating clause occurrences, each against its own exact formula. The two differ by 7% in expectation already at $n = 60$, a difference that a single simulation run may fail to expose at finite sample sizes, so both columns are reported separately (App. D).

D.2 Separation of Uniqueness and Recovery at Finite n

Identity (2) is between the two limit curves, so the finite-n question is how far the finite-n curves are from being translates of each other. Two different pairs of curves can be asked about and they must not be conflated.

The first pair is the W-statistics $\mathbb{P}(W = 0)$ and $\mathbb{E}[2^{-W}]$, which are exactly the uniqueness and optimal-recovery quantities on the structure event of Thm. 4.1 and proxies for them off it. Both are exactly computable at every finite n from the factorial-moment generating function above, at $z = 0$ and $z = \frac{1}{2}$. Solving $\mathbb{E}[0^W] = s$ and $\mathbb{E}[2^{-W}] = s$ for the window coordinate at matched success level s and differencing gives Tab. A.4. This proxy separation exceeds $\ln 2$ everywhere, by an amount largest at low success levels that halves with each doubling of n, reaching 0.25% of $\ln 2$ at $n = 200$, under one clause out of the 323 that Δm amounts to there. That is a property of the two proxy curves.

It is not a proved error bound on the true ones, and the $O(1/n)$ rate it exhibits is not the $O_{k,x}(\ln n/n)$ of Thm. 4.1, which bounds a different quantity.

The second pair is the true curves $\mathbb{P}(|\mathcal{S}| = 1)$ and $\mathbb{E}[1/|\mathcal{S}|]$. These have no closed form. One direction is provable and holds at every n. A surviving single flip is itself a solution, so $\{|\mathcal{S}| = 1\} \subseteq \{W = 0\}$ and the true uniqueness curve lies pointwise below its proxy. For recovery only a one-sided bound with an error term is available. On the event $\mathcal{C} \subseteq \mathcal{S}$ we have $|\mathcal{S}| \geq 2^W$, so $\mathbb{E}[1/|\mathcal{S}|] \leq \mathbb{E}[2^{-W}] + \mathbb{P}(\mathcal{C} \not\subseteq \mathcal{S})$ with the last term bounded by Lem. C.10. This does not place the true recovery curve below its proxy at every n. A clause with $|D| \geq 2$ inside F removes points of $\mathcal{C}$ from $\mathcal{S}$ and can leave $|\mathcal{S}| < 2^W$, which pushes $1/|\mathcal{S}|$ above 2^{-W}, while spurious survivors push it below, and the bound does not order the two effects. In the enumerations reported below the true recovery curve does lie below its proxy at every grid point (Fig. A.1), which is a measurement and not a theorem. What these facts do *not* determine is the effect on the *difference* of the two horizontal shifts, since when both curves sit below their proxies the biases partly cancel. That has to be measured, and measuring it needs the whole cube.

Both tables read a horizontal separation off curves that are known only on a grid, so the extraction has to be stated. The window coordinate is swept over the 49 points $x \in \{-2.5, -2.375, \ldots, 3.5\}$, and the clause count at each is the nearest integer to $(\ln n + x)/\widetilde{q}_1$. Each trial draws one i.i.d. clause stream of the maximum length and reads it at every prefix, so within a trial the solution set only shrinks as x grows and all four empirical curves are nondecreasing in x. For a success level s, the crossing $x_U(s)$ is the linear interpolant between the two adjacent grid points that bracket s on the uniqueness curve, $x_R(s)$ likewise on the recovery curve, and both are undefined if s falls outside the swept range. Bootstrap intervals resample trials, so they carry the Monte Carlo error and not the extraction error, and we calibrate that separately in two ways. Halving the grid step on identical clause streams moves the separations by at most 0.020, and by under 0.003 at six of the nine levels. Running the same recipe on the W-proxy pair, whose crossings are exactly computable so that the answer is known, gives a bias of at most 0.019, with no systematic sign and mean absolute value 0.005. Both are well below the effect reported next.

Tab. A.5 measures the true pair by enumeration at $n = 16$, 18, 20 and 22, which is as far as 2^n candidates per trial can be pushed, over 20,000 trials at each size. At every size the true separation is well below $\ln 2$, running from 0.437 to 0.660 at $n = 16$ and from 0.449 to 0.683 at $n = 22$, short of $\ln 2$ by 0.135 and by 0.113 on average over the sweep. The proxy separation measured on the same clause streams runs from 0.682 to 0.716 at $n = 22$ and reproduces the exact proxy values 0.694 to 0.709 at that size within the bootstrap width, which is the evidence that the pipeline is measuring what it claims to. The distance between the two readings is the point. Taking Tab. A.4 for a statement about the true curves would put them within a percent of (2) already at $n = 200$, whereas at the sizes where the true curves are available at all they are still short of it by about a sixth.

Three quantities move in the same direction as n grows. The true separation rises, and significantly so in the middle of the sweep, where the $n = 16$ and $n = 22$ intervals are disjoint at levels 0.2, 0.4 and 0.5. The average shortfall falls from 0.135 to 0.113. And the structure event fails less often, in 38.6% of trials at $n = 16$ against 30.2% at $n = 22$, measured at $x = 0$, with the first of these matching the 39.5% obtained independently for Fig. A.1. Four sizes spanning a factor 1.375 in n establish that direction but do not support a fitted rate. At the top of the sweep the intervals widen and the two largest sizes do reach $\ln 2$, which is where the true pair should agree with the proxy first, since a higher success level means more clauses and a structure event that holds more often. Tab. A.4 therefore reports the separation of two exactly computable proxy curves, and we make no finite-n claim about the true pair beyond the bounds above and this measurement. Tab. A.4 is exact arithmetic with no simulation, and the two extraction calibrations quoted above are specified in App. E.

D.3 Two-Block Exact Formulas and Protocol

For the two-block law the elementary symmetric polynomials have closed forms, $e_j = \sum_u \binom{n_1}{u}\binom{n_2}{j-u}\beta^u$, so the marginals $\xi_i = w_i e_{k-1}(w_{-i})/e_k(w)$, and hence q_ℓ and q_h, are computed exactly, with the budget identity of Lem. C.13 verified to 250 decimal digits. The factorial-moment generating function

$$\mathbb{E}[t^W] = \sum_{j_1, j_2} \binom{n_1}{j_1}\binom{n_2}{j_2}(t-1)^{j_1+j_2}\big(1 - j_1 q_\ell - j_2 q_\mathrm{h}\big)^m$$

Table A.4: Finite-n separation of the two W-*proxy* curves $\mathbb{P}(W = 0)$ and $\mathbb{E}[2^{-W}]$ at $k = 3$, computed from their exact finite-n formulas at matched success level s, against the limiting $\ln 2 = 0.69315$. The last column is the largest deviation over the sweep $s \in [0.1, 0.9]$.

	horizontal separation $x_U(s) - x_R(s)$ at success level s						
n	$s = 0.1$	$s = 0.3$	$s = 0.5$	$s = 0.7$	$s = 0.9$	max dev. from $\ln 2$	as % of $\ln 2$
50	0.70009	0.69669	0.69516	0.69417	0.69344	0.00695	1.00%
100	0.69660	0.69492	0.69416	0.69367	0.69330	0.00345	0.50%
200	0.69486	0.69404	0.69366	0.69341	0.69322	0.00172	0.25%
400	0.69400	0.69359	0.69340	0.69328	0.69319	0.00086	0.12%
800	0.69358	0.69337	0.69328	0.69321	0.69317	0.00043	0.06%

Table A.5: The *true* two-curve separation by exhaustive enumeration at $k = 3$, from $\mathbb{P}(|\mathcal{S}| = 1)$ and $\mathbb{E}[1/|\mathcal{S}|]$ over all 2^n assignments, against the W-proxy separation measured on the same clause streams. Each entry is the point estimate and half the 95% percentile-bootstrap interval over 20,000 trials. The mean row is over the full nine-level sweep, a negative entry means the separation sits above $\ln 2$, and the last row is the structure-failure rate at $x = 0$.

	true separation $x_U(s) - x_R(s)$, by enumeration				W-proxy
level s	$n = 16$	$n = 18$	$n = 20$	$n = 22$	$n = 22$
0.1	0.437 ± 0.013	0.442 ± 0.013	0.443 ± 0.011	0.449 ± 0.011	0.716 ± 0.014
0.3	0.513 ± 0.011	0.500 ± 0.013	0.519 ± 0.013	0.535 ± 0.013	0.704 ± 0.013
0.5	0.556 ± 0.015	0.564 ± 0.013	0.582 ± 0.016	0.594 ± 0.011	0.698 ± 0.013
0.7	0.615 ± 0.018	0.614 ± 0.015	0.640 ± 0.018	0.624 ± 0.020	0.682 ± 0.019
0.9	0.660 ± 0.028	0.659 ± 0.027	0.674 ± 0.030	0.683 ± 0.026	0.708 ± 0.029
mean gap to $\ln 2$	0.135	0.135	0.121	0.113	-0.010
$\lvert\mathcal{S}\rvert \neq 2^W$ at $x = 0$	38.6%	35.1%	32.5%	30.2%	

follows from Lem. C.15 and is evaluated at 300 decimal digits, because the alternating double sum cancels catastrophically at $n = 200$ in double precision. Monte Carlo draws use the exact categorical reduction of Lem. C.15, the exhaustive $n = 16$ runs sample the full weighted model by rejection and enumerate all 2^{16} assignments, and the worst deviation between Monte Carlo and exact values over the 54 comparisons behind Fig. 2 is 2.2 standard errors, with the variance of 2^{-W} computed exactly as $\mathbb{E}[4^{-W}] - \mathbb{E}[2^{-W}]^2$. In the exhaustive runs, with $n_1 = n_2 = 8$, 2000 trials, and m chosen so that $\Lambda_m = 1$, the true solution set fails to be the predicted subcube in 0.280 of trials at $\beta = \frac{1}{2}$ against 0.392 for the uniform allocation, so at this size the finite-size structure failure still depends on the allocation at matched clock value.

E Experimental Settings

This appendix specifies every experiment of Sec. 5 and Apps. D, F, and H at the level needed to rerun it, namely the generative model, the decoder, the parameter grid, the trial count, the seed, and the rule that turns trials into each reported number. The exact finite-n formulas that the simulations are compared with are in App. D, and the outputs are reported in Apps. F and H.

Common conventions. Every simulation draws its randomness from one fixed seed, stated below, and every Monte Carlo estimate is reported together with its trial count. A sample cost m^* at a stated success level is the smallest sample size on a search grid at which the empirical success over the stated number of trials reaches that level, with the grid searched upward from a stated start, either geometrically with a stated ratio or in unit steps, and nothing is fitted. Exact finite-n curves are evaluated from the formulas of App. D, in multiple precision with `mpmath` wherever an alternating sum cancels catastrophically. Level crossings of a Monte Carlo curve are read by linear interpolation between the two grid points that bracket the level. The synthetic experiments run on one CPU core and train nothing, and the only trained models are the fifteen image models described in the last paragraph, trained once with a pinned recipe.

Planted k-SAT window (Fig. 1, Tabs. A.2 and A.3). The model of Sec. 4.1 is simulated at $n = 200$ and $k = 3$ with clauses drawn with replacement, $m = \mathrm{round}((\ln n + x)/\widetilde{q}_1)$, $x \in \{-2, -1.5, \ldots, 2\}$, 800 trials per point, and seed 11. Each trial encodes clauses as scope triples with a difference-pattern index, computes W by marking resolved coordinates, and evaluates the detected-modes flag of App. D, namely a distance-two spurious survivor or a size-two difference pattern inside F. The empirical $\mathbb{P}(W = 0)$, $\mathbb{E}[2^{-W}]$, and the law of W at $x = 0$ are compared with the exact factorial-moment formulas. The six-configuration verification of Tabs. A.2 and A.3 uses the $(n, x, \text{trials}, \text{seed})$ list of App. D.

Structure failure at small and growing n (Sec. 5.1, Fig. A.1, Tab. A.6). The exhaustive check runs at $n = 16$ and $k = 3$ for $x \in \{-1, -0.5, \ldots, 2\}$ with 600 trials per point and seed 12, enumerating all 2^{16} assignments in every trial to obtain $\mathcal{S}$, $\mathbb{P}(|\mathcal{S}| = 1)$, $\mathbb{E}[1/|\mathcal{S}|]$, and the indicator of $\mathcal{S} \neq \mathcal{C}$, and recomputing the detected-modes flag on the same trials. The rate sweep runs at $x = 0$ for $n \in \{50, 100, 200, 400, 800\}$ with 1500 trials per size and seed 91, and compares the detected-modes frequency with $1 - \exp[-(\mathrm{EN}_A + \mathrm{EN}_B)]$ computed from the first moments of App. D.

Separation of the two curves (Tabs. A.4 and A.5). Tab. A.4 is exact arithmetic at 400 decimal digits for $n \in \{50, 100, 200, 400, 800\}$ and levels $s \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$, solving each crossing on the continuous sample size. Tab. A.5 enumerates the cube at $n \in \{16, 18, 20, 22\}$ and $k = 3$ with 20,000 trials per size and seed 20260828. Each trial draws one clause stream and reads it at the prefix lengths $m(x) = \mathrm{round}((\ln n + x)/\widetilde{q}_1)$ for x from -2.5 to 3.5 in steps of 0.125, crossings are interpolated at the nine levels $s \in \{0.1, \ldots, 0.9\}$, and intervals are percentile bootstraps over 400 resamples of the trials. The two extraction calibrations of App. D.2 rerun $n = 16$ at half the grid step on identical streams and run the same recipe on the exactly computable proxy pair.

Non-uniform coverage (Fig. 2, App. D.3). The two-block law is simulated at $n = 200$, $k = 3$, and $\alpha = \frac{1}{2}$, so 100 coordinates carry weight $\beta \in \{1, \frac{1}{2}, \frac{1}{4}\}$ and 100 carry weight one. For each of the target clock values $\Lambda \in \{3, 2, 1.5, 1, 0.7, 0.5, 0.3, 0.15, 0.05\}$ the clause count is the first integer m with exact $\Lambda_m \leq \Lambda$, where q_ℓ and q_h come from the two-block elementary symmetric sums and $\mathbb{E}[t^W]$ is evaluated at 300 digits, and 800 trials per point with seed 20260831 use the categorical reduction of Lem. C.15. The weighted exhaustive run has $n = 16$ with $n_1 = n_2 = 8$, 2000 trials, m chosen so that $\Lambda_m = 1$, $\beta \in \{1, \frac{1}{2}\}$, rejection sampling of the weighted scope law, and enumeration of all 2^{16} assignments.

General k and Model RB (Fig. A.3, planted k-SAT row of Tab. 2). Panel (a) simulates one-forbidden-pattern clauses at $n = 200$ for $k \in \{2, 3, 4\}$ on the grid $x \in \{-2, -1.5, \ldots, 2\}$ with 3000 trials per point and seed $4103 + k$, drawing explicit scopes. The Tab. 2 row is the ratio between $k = 4$ and $k = 2$ of the clause count at which $\mathbb{P}(W = 0)$ crosses one half. Panel (b) simulates planted Model RB with domain size $d = \mathrm{round}(n^\alpha)$ and $q = \mathrm{round}(pd^k)$ forbidden tuples per constraint, in regime I with $k = 2$, $\alpha = 0.8$, $p = 0.4$, $n = 6400$ and in regime II with $k = 3$, $\alpha = 1$, $p = 0.6$, $n \in \{400, 6400\}$, using $\max(3000, \min(20000, 4 \cdot 10^6/n))$ trials per point, that is 10,000 at $n = 400$ and 3000 at $n = 6400$, with seed $4103 + n$. The Model RB runs use a multinomial coverage-count sampler, which lets one constraint touch a coordinate twice, and the script validates it against the explicit-scope sampler at 32,000 trials per cell on three settings before use.

Hidden pixel (Tab. 2 row, Fig. A.4). Sample t draws a public uniform $R_t \in \{0, 1\}^d$ and reveals $R_t[\theta^*]$ through $\mathrm{BSC}(0.1)$, for $d \in \{8, 16, \ldots, 1024\}$. The agreement decoder returns the index whose probe bits agree with the most observations. The empirical $m^*(d)$ at 90% success is searched in unit steps from $m = 4$, with 2000 trials per candidate and seed 21, and the second seed quoted in App. F is 22. The floor is Thm. 2.1 at $P_e = 0.1$, and the guarantee is a union bound over the $d - 1$ alternatives with the exact pairwise error of the decoder.

One-edge Ising and one-hidden-unit RBM (Tab. 2 rows, Fig. A.5). Both families are sampled through the exact one-edge marginal on $p = 30$ visible ± 1 spins, in which the coupled pair agrees with probability $e^J/(e^J + e^{-J})$ and every other spin is an independent fair sign. The Ising family uses $J = \theta \in \{0.35, 0.28, 0.22, 0.18, 0.14\}$, and the RBM family uses $J = J(w) = \frac{1}{2} \ln \cosh 2w$ with $w \in \{0.5, 0.42, 0.35, 0.3, 0.25\}$. The decoder returns the pair with the largest absolute empirical correlation and breaks ties uniformly at random, which lets the coupled pair be fixed to the first two spins by exchangeability. The cost m^* at 90% success is searched by $m \leftarrow \lfloor 1.03m \rfloor + 1$ from $m = 16$ with 2000 trials per candidate and seed 31, the second run quoted in App. F uses seed 32, and the Tab. 2 ratios divide the cost at the weakest coupling by the cost at the strongest.

Erasure (Tab. 2 row, Fig. A.7). Each sample reveals the target with probability $\varepsilon = 0.05$ and is otherwise empty, for $M \in \{10, 10^2, \dots, 10^5\}$ candidates. The rule outputs the revealed target if any sample revealed it and otherwise guesses uniformly. The cost m^* at 90% success is searched in unit steps from $m = 2$ with 400 trials per candidate and seed 51, the closed form is $\ln 10/(-\ln(1-\varepsilon))$, and the floor is the N1 bound of App. C.7 at $P_e = 0.1$.

Noisy indexing (Tab. 2 row, Fig. A.6). With $\eta = 0.1$ and $n \in \{300, 1000, 3000, 10^4\}$, the $m = \lambda n$ observations are allocated to coordinates by a multinomial draw, each coordinate is decoded by majority vote with ties broken by a fair coin, and exact recovery means every coordinate is correct. The sweep fixes $y = a\lambda - \ln n + \frac{1}{2}\ln\ln n \in \{-1, -0.33, \dots, 3\}$ in steps of 0.667 and solves for λ, with 160 trials per point and seed 41. The reference curve is $\exp(-\mathrm{EN})$ with EN computed exactly from binomial loads, and the Tab. 2 row and the loads 0.73 and 0.80 are read at the one-half crossing. The constants of Conj. G.1 are checked at 400 digits, the parity constants at $\eta \in \{0.05, 0.1, 0.2, 0.35\}$ and $t \in \{400, 1600\}$, and c_η at $\eta \in \{0.05, 0.1, 0.2\}$ and $\lambda \in \{40, 80, \dots, 640\}$.

Measured annotation channel, annotations only (Fig. A.8, App. H). The raw CIFAR-10H release is parsed with attention-check trials removed, giving $514{,}200$ annotations over the $10{,}000$ test images. Each image's vote counts are thinned binomially with probability one half into halves A and B, with seed 20260826. The per-item channel is the half-A vote distribution of the item, and the homogeneous baseline places probability $1 - \bar{\eta}$ on the gold class and spreads $\bar{\eta}$ evenly over the other nine classes, with $\bar{\eta} = 5.14\%$ the pooled half-A error. A run at (n, λ) draws a uniform n-subset of items, allocates $m = \lambda n$ votes to them multinomially, draws each item's votes from its channel, half B for measurements and half A or the baseline for predictions, and decodes by plurality with uniform tie-breaking. The full-pool sweep uses $n \in \{100, 1000, 10^4\}$, $\lambda \in \{2, 4, 6, 8, 10, 14, 18, 24, 30\}$, and 120 trials per cell. The near-homogeneous band consists of the $8{,}624$ items whose half-A plurality equals the gold label and whose half-A error is at most 0.10, swept at $n \in \{30, 100, 300, 1000, 3000\}$ and $\lambda = 1, \dots, 40$ with 200 trials per cell, and $\lambda^*(n)$ is the interpolated one-half crossing with its slope fitted by least squares against $\ln n$.

Image-assisted decoders (Fig. 3, Tab. A.7, App. H). The thinning seed is 0, the measurement seed 90210, and the prediction seed 31337. Three architectures are trained with seeds 0 to 4 on the full $50{,}000$-image CIFAR-10 training set. The CIFAR ResNet-18 has BasicBlock stages $[2, 2, 2, 2]$, a 3×3 stride-one stem, and no max-pool, and is trained from scratch for 60 epochs at batch 128, learning rate 0.1, and input size 32. The ResNet-50 and ViT-B/16 start from the `torchvision` ImageNet weights `IMAGENET1K_V2` and `IMAGENET1K_V1`, replace the classifier head, and are fine-tuned for 10 epochs at batch 64, learning rate 0.01, and input size 224, the ViT with micro-batches of 32 and gradient accumulation. All models use SGD with momentum 0.9, weight decay $5 \cdot 10^{-4}$, a cosine schedule to zero, random crops of size 32 with padding 4 and horizontal flips, and per-channel normalization with CIFAR statistics for the scratch model and ImageNet statistics for the pretrained ones. The recipe is pinned in advance, there is no validation split and no model selection, and the test logits are cached once. The pooled confusion matrix $C[v \mid y]$ is estimated from half A against the gold labels with Laplace smoothing one. The combined decoder returns $\arg\max_y [\log p_{\mathrm{model}}(y \mid \text{image}) + \sum_v n_v \log C[v \mid y]]$ with n_v the votes for class v, the annotations-only decoder drops the image term, and both break ties uniformly. Sweeps use $\lambda \in \{0, 1, 2, 4, 6, 8, 10, 14, 18, 24, 30\}$, $n \in \{10, 30, 100, 300, 1000, 3000\}$ and the full pool of $10{,}000$, and 500 trials per cell on half-B votes. The per-item prediction $q_i(k)$, the probability that item i is decoded correctly from k votes, is estimated from half A with 256 Monte Carlo draws for every $k \le 72$ and combined over the binomial law of the load. The band, $\lambda^*(n)$, and the slopes are as in the previous paragraph. The annotation-hard items are the 86 whose half-A plurality differs from the gold label, the model-only column of Tab. A.7 is the combined decoder at $\lambda = 0$, the never-recovered count at $\lambda = 30$ is the number of items that no trial decodes correctly, and every entry is the mean over the five seeds.

F Additional Experimental Results

The figures below and the rate-sweep table of Sec. 5.1 accompany Secs. 5.1 and 5.3, and all run under the settings of App. E. The planted k-SAT and noisy indexing rows of Tab. 2 are read from the curves of Figs. A.3 and A.6 by linear interpolation of their one-half crossings, so they inherit the trial counts and grids of those figures.

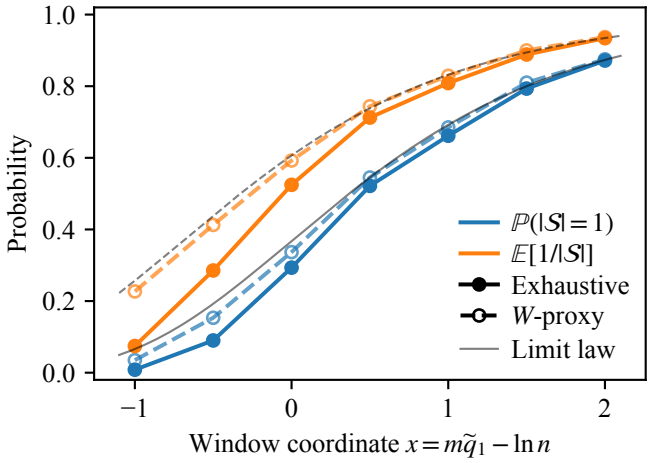


Figure A.1: Exhaustive validation at $n = 16$, showing the true uniqueness probability $\mathbb{P}(|\mathcal{S}| = 1)$ and the true Bayes success $\mathbb{E}[1/|\mathcal{S}|]$, computed by enumerating all assignments, against the W-proxies $\mathbb{P}(W = 0)$ and $\mathbb{E}[2^{-W}]$ and the limit curves.

Table A.6: Detectable subcube-structure violations at $x = 0$ across a sixteen-fold range of n. Observed values are Monte Carlo frequencies over 1500 trials with 95% confidence intervals. The statistic is a fixed detectable subset of the structure failures, not their total probability.

n	m	Predicted	Observed
50	454	0.207	0.181 ± 0.020
100	1072	0.128	0.116 ± 0.016
200	2470	0.076	0.068 ± 0.013
400	5589	0.044	0.044 ± 0.010
800	12475	0.025	0.027 ± 0.008

Fixed-k sweep and Model RB (Fig. A.3). Panel (a) checks the fixed-k prediction beyond the historically studied $k = 3$ case. At $n = 200$, the window-centre clause counts are 1587, 2470, and 3971 for $k = 2, 3, 4$, yet the reparameterized statistics, from 3000 trials per point, have maximum deviation 0.017 from their corresponding curves. Panel (b) is the Model RB sweep discussed in Sec. 5.3, with domain size $d = n^{\alpha}$ for a growth exponent α, forbidden-tuple density p, and $q = pd^k$ forbidden tuples per constraint. Regimes I and II have $(1+\alpha)p = 0.72$ and 1.20. When the clustering strength is small, the coordinate count behaves approximately as in the theorem, and when it is large, $\mathbb{P}(W = 0)$ remains near 0.486 at $n = 6400$ instead of approaching e^{-1}.

Hidden pixel (Fig. A.4). We instantiate the family as random one-bit probes, in which sample t draws a public uniform $R_t \in \{0, 1\}^d$ and reveals $R_t[\theta^*]$ through $\mathrm{BSC}(\eta)$, so the per-sample supply is capped by one channel use. The three lines are the N1 floor (no algorithm below it, by Thm. 2.1), the guarantee of the agreement decoder via a union bound whose pairwise error is computed exactly, and the empirical $m^*(d)$ at 90% success over 2000 trials, searched in unit steps. The empirical cost rises from 11 at $d = 8$ to 26 at $d = 1024$, by two samples per doubling of d except for one step of three, and a second seed gives 11 to 27. This is the $\log d$ law with an additive term, since at a fixed success level the union bound asks the pairwise error exponent to reach $\ln(d - 1) + \ln 10$, and at $d = 8$ the second term is as large as the first. The endpoint ratio 2.36 in Tab. 2 is therefore below the ratio 3.33 of $\log d$ itself, while the computed guarantee, which carries the same term, rises from 12 to 32, a ratio of 2.67, and the empirical points sit between 8% and 19% below it and well above the floor, as they must. They do not, and cannot, *verify* the floor.

One-edge Ising and one-hidden-unit RBM (Fig. A.5). We measure the empirical $m^*(\theta)$ of the maximal-correlation pair decoder for the visible one-edge Ising family at $p = 30$, and $m^*(w)$ for the one-hidden-unit RBM family, whose visible marginal is sampled exactly as the one-edge Ising

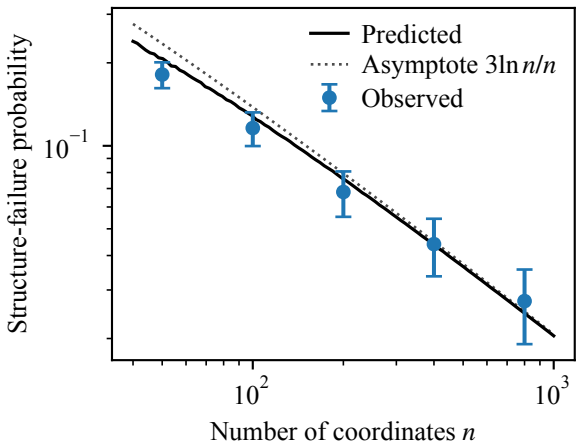
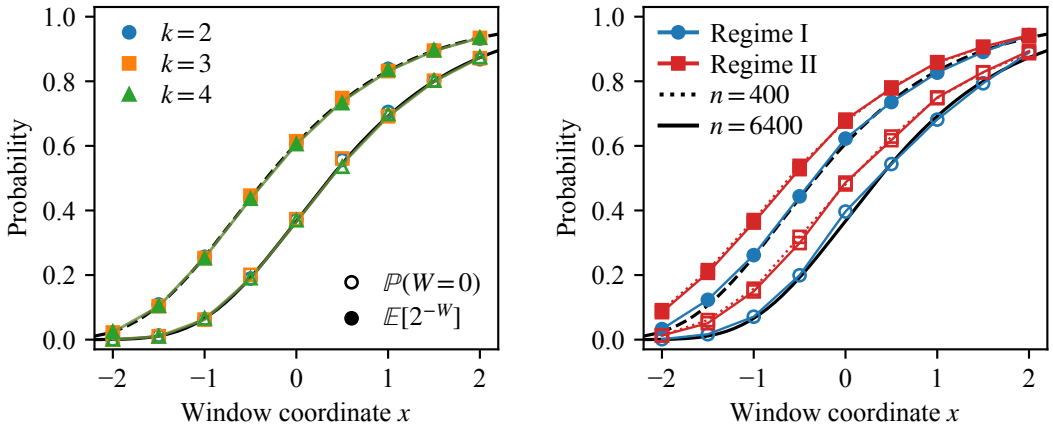


Figure A.2: The data of Tab. A.6, the detected two-mode violation statistic at the window centre against n, with the exactly computed first-moment prediction and its asymptote $3 \ln n/n$. Error bars are 95% Monte Carlo intervals.



(a) Binary clauses with one forbidden pattern, $k = 2, 3, 4$.

(b) Planted Model RB.

Figure A.3: Testing the mechanism behind the SAT window. Panel (a) shows the window statistics at $n = 200$ for binary clauses with one forbidden pattern and $k = 2, 3, 4$, reparameterized by x. Panel (b) shows the same statistics in planted Model RB (Xu & Li, 2000), where no structure theorem is available and they are only proxies for uniqueness and recovery, with $(1 + \alpha)p = 0.72$ in regime I and 1.20 in regime II.

measure at coupling $J(w) = \frac{1}{2} \ln \cosh 2w$. The search grid has ratio 1.03, the success target is 90%, and each point uses 2000 trials. On log-log axes the Ising points are compared with the reference slope -2 and the RBM points with the exact-identity curve $m^* \propto J(w)^{-2}$ and its asymptote of slope -4, each anchored at the weakest coupling of its family and none of them fitted. Between $\theta = 0.35$ and 0.14 the measured cost ratio is 5.98 against 6.25 from θ^{-2}, and between $w = 0.5$ and 0.25 it is 13.2 against the identity prediction 13.0, an end-to-end slope of 3.7. Both measured ratios carry the 3% discretization of the grid at each endpoint and the Monte Carlo fluctuation of 2000 trials per point, a second seed gives 6.37 and 13.2, and the decoder sees the pair correlation $\tanh J$ rather than J, so its finite-coupling predictions are 5.85 and 12.7. The two slopes are Prop. C.6 made visible, since one marginalized unit squares the sample price.

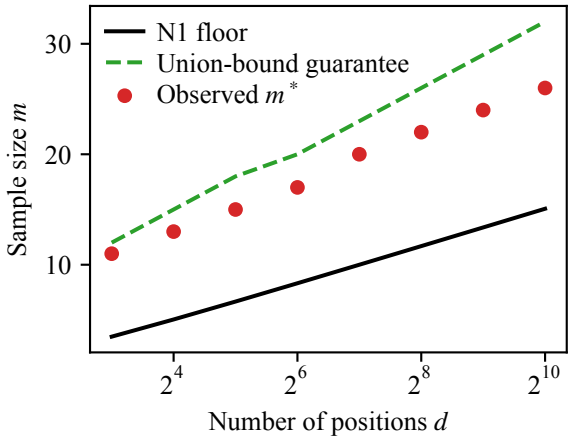


Figure A.4: Hidden pixel at $\eta = 0.1$: N1 floor, computed union-bound guarantee, and empirical sample cost of the agreement decoder.

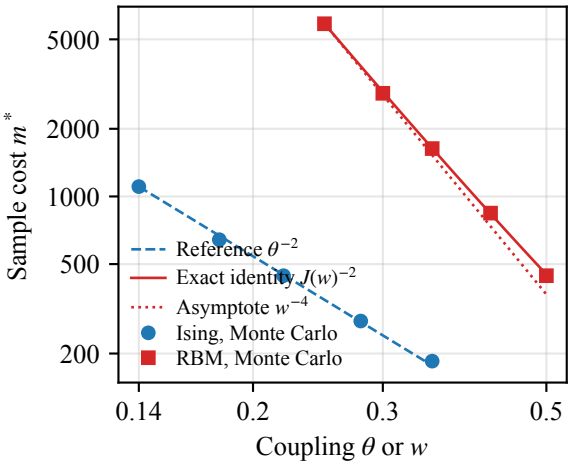
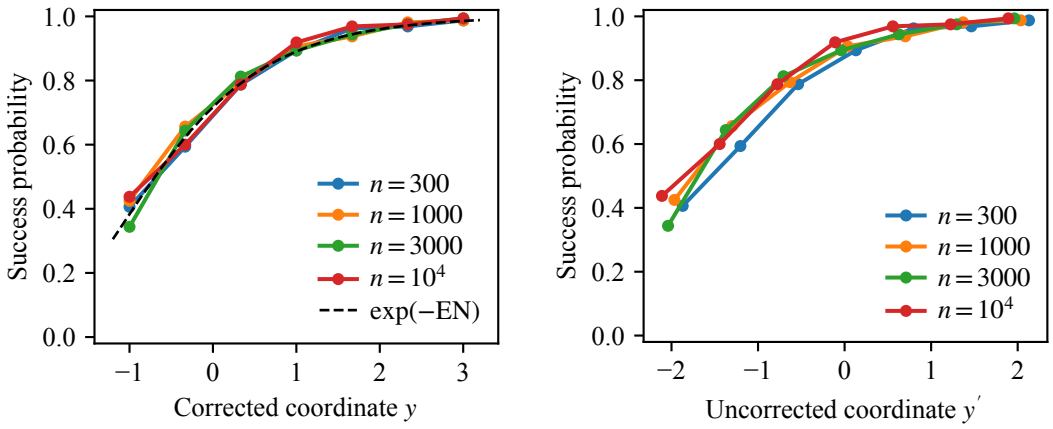


Figure A.5: Sample cost against coupling strength at $p = 30$ and 90% success for the visible one-edge Ising and the one-hidden-unit RBM family, with the theoretical curves anchored at the weakest coupling of each family.

Numerical Evidence for the Noisy Indexing Conjecture (Fig. A.6). For n from 300 to 10^4 we plot the empirical exact-recovery success of per-coordinate majority voting over 160 trials per point against two horizontal axes, $y = a\lambda - \ln n + \frac{1}{2}\ln\ln n$ (with the conjectured correction) and $y' = a\lambda - \ln n$ (without). Under the corrected axis the curves for different n collapse onto a common profile, and without the correction they drift with n in the predicted direction. The size of this effect is small by construction. The predicted horizontal drift between $n = 300$ and $n = 10^4$ is $\frac{1}{2}\ln(\ln 10^4/\ln 300) \approx 0.24$, slow because it is doubly logarithmic, so the drift in the right panel is modest though systematic across the four sizes, and the sharper diagnostic is the left panel's collapse onto the exactly computed reference. The overlaid reference is $\exp(-\mathrm{EN})$ where EN, the expected number of erring coordinates, is computed *exactly* from binomial loads. The exponential form is the conjectured Poissonization, while the mean is not conjectural. Read at the one-half crossing by linear interpolation on the grid of step 0.667, the load is 0.73 and 0.80 of $n\ln n/a$ at $n = 300$ and $n = 10^4$. The $\frac{1}{2}\ln\ln n$ term accounts for 0.15 and 0.12 of the shortfalls 0.27 and 0.20 from the first-order threshold, and the crossings sit at $y = -0.67$ and -0.74 on the corrected axis against the conjectured $\ln(c_\eta/\ln 2) = -0.755$, so the window constant c_η accounts for the remainder within the Monte Carlo resolution. The figure is labeled as evidence for a conjecture, not as a verified law.



(a) Corrected coordinate $y = a\lambda - \ln n + \frac{1}{2}\ln\ln n$.

(b) Uncorrected coordinate $y' = a\lambda - \ln n$.

Figure A.6: Noisy indexing at $\eta = 0.1$. Exact-recovery success of per-coordinate majority voting for n from 300 to 10^4, against the corrected window coordinate in panel (a), where the dashed reference is $\exp(-\mathrm{EN})$ with the mean computed exactly, and against the uncorrected coordinate in panel (b).

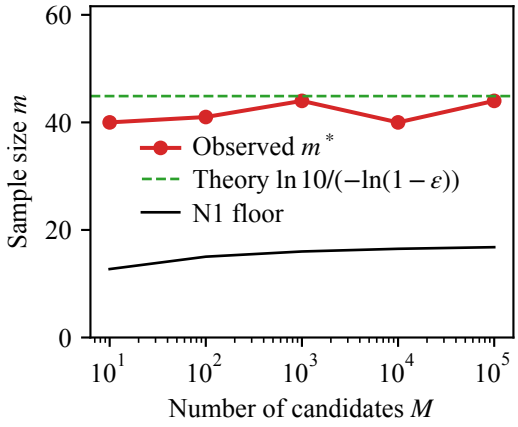


Figure A.7: Erasure family at $\varepsilon = 0.05$, where the empirical sample cost is flat in M and matches the closed-form cost of the wait-for-reveal rule, with the order-tight N1 floor flat as well.

Erasure family (Fig. A.7). Write $M = |\Theta|$ for the number of candidates. The empirical sample cost of the wait-for-reveal strategy, at 90% success over 400 trials with unit search steps, is flat in M across four orders of magnitude (m^* between 40 and 44 for M from 10 to 10^5, against the closed form $\ln 10/(-\ln(1-\varepsilon)) = 44.9$), and the N1 floor, which is tight in order for this family but sits a constant factor below the true cost (its constant at success level s is s against the true $-\ln(1-s)$, App. C.7), is likewise flat. The flatness is what Rem. C.3 predicts, since a family whose difficulty is one global event has no per-coordinate structure for N2 to use, and no $\log M$ in its sample cost.

G A Conjectural Second-Order Window for Noisy Indexing

Noisy indexing has a proved first-order exact-recovery boundary $m = (1 + o(1))n\ln n/a$, where each observation chooses a coordinate uniformly and returns its sign through $\mathrm{BSC}(\eta)$. The following window refinement is supported by the calculation in App. G.1, but is not used by any main theorem.

Conjecture G.1 (Noisy-indexing window). *Fix $\eta \in (0, \frac{1}{2})$, let $a = 1 - 2\sqrt{\eta(1-\eta)}$ and $C_\eta = -\ln(2\sqrt{\eta(1-\eta)})$. If*

$$a\lambda = \ln n - \tfrac{1}{2}\ln\ln n + y + o(1), \qquad \lambda = m/n,$$

then the number of incorrectly decoded coordinates under per-coordinate Bayes decoding converges to $\mathrm{Poisson}(c_\eta e^{-y})$, where

$$c_\eta = \sqrt{\frac{e^{C_\eta}+1}{8\pi}},$$

and optimal exact-recovery success converges to $\exp(-c_\eta e^{-y})$.

The unresolved steps are window-precision de-Poissonization of multinomial loads and an upgrade from the mean error count to a Poisson limit. App. G.1 derives the $-\frac{1}{2}\ln\ln n$ correction and the constant, and App. F reports the simulation evidence.

G.1 Supporting Calculation for Conjecture G.1

Conditionally on its load t (the number of samples that touched it), a coordinate is decoded by majority vote with ties broken uniformly. With $C_\eta = -\ln(2\sqrt{\eta(1-\eta)})$ as in Conj. G.1, the exact large-deviation prefactor of Bahadur & Rao (1960) gives

$$\mathbb{P}(\text{coordinate errs} \mid t) \;=\; c_{\mathrm{par}(t)}\, t^{-1/2} e^{-tC_\eta}\,(1 + O(1/t)),$$

where the prefactor splits by the parity of t, because the binomial is a lattice variable and an even load can tie,

$$c_{\mathrm{even}} = \frac{1}{(1-2\eta)\sqrt{2\pi}}, \qquad c_{\mathrm{odd}} = e^{-C_\eta}\, c_{\mathrm{even}}. \tag{4}$$

These are the lattice Bahadur-Rao constant $2/\big((1-e^{-\tau})\sqrt{2\pi}\big)$ at tilt $\tau = \ln\frac{1-\eta}{\eta}$, whose tilted standard deviation is $\frac{1}{2}$. For even t, the constant is reduced by half the local-limit mass $2e^{-tC_\eta}/\sqrt{2\pi t}$ sitting exactly on the tie. For odd t, it is multiplied by $e^{-\tau/2}$ because the threshold sits half a lattice step higher. The $t^{-1/2}$ and the constants in (4) are classical and are not claimed as a contribution.

Loads are asymptotically $\mathrm{Poisson}(\lambda)$, and the Poisson tilting identity

$$\mathbb{E}\big[e^{-tC_\eta} g(t)\big] \;=\; e^{-\lambda a}\,\mathbb{E}\big[g(\mathrm{Poisson}(\lambda e^{-C_\eta}))\big], \qquad t \sim \mathrm{Poisson}(\lambda),$$

moves the computation to a Poisson of mean $\lambda_{\mathrm{tilt}} = \lambda e^{-C_\eta}$. Its parity is equidistributed up to $e^{-2\lambda_{\mathrm{tilt}}}$, so the effective prefactor is the average $\bar{c} = \frac{1}{2}(c_{\mathrm{even}} + c_{\mathrm{odd}})$ and $\mathbb{E}\big[c_{\mathrm{par}(t)} t^{-1/2}\big] = \bar{c}\,\lambda_{\mathrm{tilt}}^{-1/2}(1 + O(1/\lambda_{\mathrm{tilt}}))$. The expected number of erring coordinates is therefore

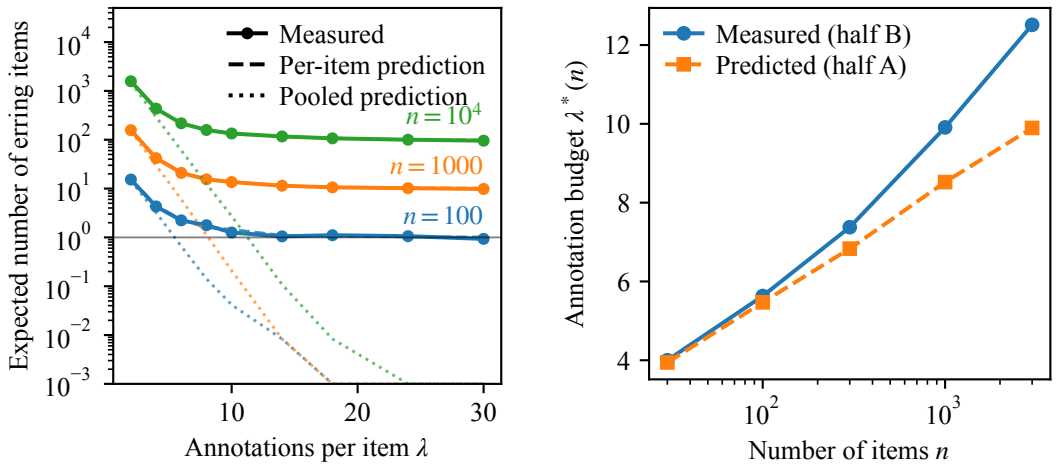
$$\mathrm{EN} \;=\; n\, e^{-\lambda a}\,\bar{c}\, e^{C_\eta/2}\,\lambda^{-1/2}\,(1 + o(1)),$$

in which $e^{-\lambda a}$ produces the first-order threshold $a\lambda = \ln n$ and $\lambda^{-1/2}$, at $\lambda = (1+o(1))\ln n/a$, produces the $-\frac{1}{2}\ln\ln n$ correction. Substituting $a\lambda = \ln n - \frac{1}{2}\ln\ln n + y$ leaves $\mathrm{EN} \to c_\eta e^{-y}$ with $c_\eta = \bar{c}\, e^{C_\eta/2}\sqrt{a}$, and since $(1-2\eta)^2 = 1 - 4\eta(1-\eta) = 1 - e^{-2C_\eta}$ this collapses to the closed form of Conj. G.1,

$$c_\eta \;=\; \frac{\sqrt{1+e^{-C_\eta}}}{2\sqrt{2\pi e^{-C_\eta}}} \;=\; \sqrt{\frac{e^{C_\eta}+1}{8\pi}}.$$

Both (4) and c_η are checked numerically at high precision under the settings of App. E. The parity constants match the exact binomial to four digits by $t = 1600$ at $\eta \in \{0.05, 0.1, 0.2, 0.35\}$, and the exactly computed $\mathbb{E}[\mathbb{P}(\text{errs} \mid \mathrm{Poisson}(\lambda))]\, e^{a\lambda}\sqrt{\lambda a}$ reaches c_η to within 0.09%, 0.16% and 0.36% at $\eta = 0.05, 0.1, 0.2$ by $\lambda = 640$, with the residual halving each time λ doubles, as the $O(1/\lambda_{\mathrm{tilt}})$ term predicts.

Two things are still missing for a theorem. The first is de-Poissonization of the multinomial loads at window precision, where the $O(1)$ window and the constant must survive rather than merely the first order, and the second is an independence or negative-dependence argument upgrading the mean of the error count to its Poisson limit law. The evaluation of c_η uses one further unproved step, the



(a) Erring items against the annotation load.

(b) The exact-recovery budget on the near-homogeneous band.

Figure A.8: The measured CIFAR-10H annotation channel. Panel (a) compares the measured number of erring items on subsets of $n = 100$, 1000, and 10^4 items with the per-item and pooled predictions. Panel (b) shows that the budget $\lambda^*(n)$ on the band of items with per-item error at most 0.10 grows with $\ln n$, measured on half B and predicted from the half-A channel.

interchange of the $t \to \infty$ prefactor asymptotics with the tilted expectation, which requires uniform integrability of the $O(1/t)$ remainder under $\mathrm{Poisson}(\lambda_{\mathrm{tilt}})$. The numerics above support it at the observed $O(1/\lambda_{\mathrm{tilt}})$ rate, and it is not written out here. This also explains why the constant is hard to read off the simulations of Fig. A.6 directly. At $n = 10^4$ the tilted mean is only $\lambda_{\mathrm{tilt}} \approx 17$, so $\mathrm{EN}\, e^y$ still sits about 2% above c_η, and it approaches it like $1/\ln n$. Dense analogues (deterministic load $t = m$) exhibit the same $-\frac{1}{2}\ln\ln n$ correction, so the correction itself is not specific to random assignment, and the literature comparison for this family is in Sec. A.

H Real-Channel and Vision Details

Measured annotation channel. CIFAR-10H contains 514,200 crowd annotations of the CIFAR-10 test images (Peterson et al., 2019). We split annotations within each image into two halves. Half A estimates an item-specific categorical response distribution, and half B supplies independent resampling trials. A homogeneous baseline instead pools all items into a single confusion channel. We write λ for annotations per item and $\lambda^*(n)$ for the value at which exact-recovery probability reaches one half on an n-item subset. The pooled error rate is 5.14%, but 86 images have a half-A plurality away from the gold class. At $n = 10^4$ and $\lambda = 18$, the homogeneous calculation predicts 0.008 erring items while the measured count is 107, and the per-item prediction tracks the measurement, nonzero floor included, at every budget (Fig. A.8(a)). On a near-homogeneous subset capped at 10% estimated item error, $\lambda^*(n)$ grows affinely in $\ln n$ with slope 1.85 (Fig. A.8(b)).

Adding the image channel. We train a CIFAR ResNet-18 from scratch and fine-tune ImageNet-pretrained ResNet-50 and ViT-B/16 models, with five seeds per architecture. Their test accuracies are 0.9454, 0.9767, and 0.9794. The image log posterior is combined with annotation likelihoods from the pooled confusion matrix, which is estimated from half A against the gold labels, and no test label is used for training or model selection. A budget of zero annotations is a valid point on every sweep, and for the annotations-only decoder that point is a prior guess, which is what the 8.7 of Tab. A.7 records against a chance value of 8.6. On the near-homogeneous band, the fitted $\lambda^*(n)$ slopes are 1.82, 1.55, and 1.53. The right axis of Fig. 3(b) shows the mean per-item accuracy of the ResNet-18 combined decoder at $\lambda = 6$ on the full-pool subsets, which stays between 0.988 and 0.989 from $n = 10$ to 10^4 while the exact-recovery budget grows with $\ln n$, since exact recovery asks for every item at once and per-item accuracy does not. On the 86 annotation-hard images, the

Table A.7: Annotation-only and image-assisted decoders. ResNet-50 and ViT-B/16 use external ImageNet pretraining. Values are means over five seeds.

	annotations only	ResNet-18	ResNet-50	ViT-B/16
test accuracy	–	0.9454	0.9767	0.9794
slope of $\lambda^*(n)$ vs. $\ln n$	1.85	1.82	1.55	1.53
hard items correct, model only	8.7	47.0	56.8	52.2
never recovered at $\lambda = 30$	14.6	14.8	15.0	15.6

models alone recover 47.0, 56.8, and 52.2 images on average, far below what their accuracy on the other $9{,}914$ images predicts. A model treating these items as ordinary would recover 81 to 84 of them. ViT-B/16 is the most accurate model overall and the weaker of the two pretrained models on this set, 52.2 against 56.8, an inversion for which an exact permutation test over the five seeds of each architecture gives $p = 0.04$, so difficulty on these items does not track overall model strength. These measurements show shared concentration of difficulty, not an impossibility result for arbitrary vision learners.

The combined decoder and the floor. Combining the channels does not move the floor. Relative to the annotations-only 14.6, the never-recovered count at $\lambda = 30$ rises to 14.8, 15.0, and 15.6, and ViT-B/16 adds 1.4 items of its own that no trial recovers against 0.4 it removes. The mechanism is a property of the combination rule. Under the single global confusion matrix, the model log posterior is worth a median of 1.8 to 2.3 crowd votes, so at thirty votes per item it is outvoted by a factor of 13 to 17, and the number of hard items recovered falls monotonically from the model-alone 47 to 57 at zero annotations to between 34.5 and 35.8 at thirty. A decoder that assumes a homogeneous channel is outvoted by a channel that is not homogeneous. The failure belongs to this misspecified combination rule, not to the information in either channel, and the model-only row of Tab. A.7 is therefore the informative measurement.

Items Where Both Channels Disagree with the Gold Label. Thirteen images are labelled away from the CIFAR gold class both by the roughly 51 human annotations and by both pretrained models at all five seeds. Four of them receive zero gold votes, the models side with the crowd on all four, and on image 2405 every one of 51 annotators and 10 of 10 model runs answer frog against a gold label of cat. These four are gold-label errors rather than information boundaries. A fifth enters only through a 25 to 25 tie broken uniformly. On the remaining eight, the gold class holds 10% to 33% of the votes without ever holding a plurality, and the models' modal answer coincides with the crowd's plurality on all eight. Whether these stimuli are semantically ambiguous or merely underdetermined at 32×32 resolution is a question our measurements do not decide and our claims do not need. On these eight coordinates both measured channels leave the gold label short of a plurality, and that is the sense in which the floor is a property of the two measured channels rather than a labelling artifact. It is a statement about the pluralities of two channels measured at the gold labels, not an algorithm-independent bound, since neither channel is a fully identified generative model.

Protocol and scope. The complete settings of every experiment, including seeds, trial counts, search grids, and the training recipe of the image models, are in App. E, and the vision analysis is deterministic given the cached test outputs of the fifteen models. The pooled confusion matrix of the homogeneous baseline and the per-item error rates are both calibrated against the gold labels of the evaluated test images, using half A of their annotations, so the diagnostic is retrospective rather than a recovery experiment under unknown labels. No gold label of an evaluated item enters model training, model selection, or hyperparameter choice, and the image models see no test label at any point. Simulations are evidence about finite-size predictions, not verification of a necessity theorem. The annotation channel is measured only at the gold labels and is not a fully identified counterfactual generative model.

I Discussion and Limitations

Interpretation. The two input requirements answer different questions. Total information asks whether the observations collectively distinguish enough hypotheses. Coverage asks whether each hidden degree of freedom receives evidence. In planted k-SAT, the latter creates both the missing

$\ln n$ factor and the critical-window posterior geometry. The distinction also clarifies why uniqueness is stronger than learning, since an input can leave several symmetric explanations while still allowing the planted one to be recovered with nonzero Bayes probability.

Limitations. First, exact recovery is stricter than approximate recovery or prediction, and the thresholds here need not govern those objectives. Second, information-theoretic necessity does not imply computational attainability, and planted CSPs are a standard setting for information-computation gaps (Feldman et al., 2015). Third, the SAT window holds for fixed k and fixed x, with constants not uniform over the window tails, and the without-replacement transfer of Rem. C.12 in App. C.13 is proved only for $k \geq 3$. Fourth, the second-order noisy-indexing window remains a conjecture. Fifth, the five families are calibration points, not an exhaustive taxonomy of learning boundaries. Finally, the real-channel study establishes a mismatch between scalar and item-specific summaries for the measured decoders, not a universal lower bound for vision or crowdsourcing.

Future work. Natural next steps are to prove the noisy-indexing window, characterize analogous coverage windows under clustered observations, and study how observation design can deliberately move the binding coordinate layer. A broader challenge is to connect these exact-identification boundaries to approximate recovery without collapsing the distinction between information, computation, and representation.