Sensor Decisions under Unknown Dependence: Interval Frontiers and Paid Calibration

B-Theory research programme

Technical addendum, version 0.2
9 November 2026

Prepared with AI assistance. Internal AI review is distinct from external peer review.

Abstract

We extend a finite sensor-reliability model in which a controller may purchase two imperfect reference readings or abstain. The references have known marginal error rates and an unknown shared dependence parameter. First, we derive the likelihood of the complete four-reading calibration record and reduce it to three sufficient categories. The primary readings add statistical information beyond reference disagreement, although a strict reduction in the calibration sample requirement for the present decision task is not established. Second, we characterize the exact risk–coverage–cost frontier for a closed interval of dependence parameters. Finally, we combine this decision geometry with standard concentration and change-of-measure arguments. Outside the no-training region, positive coverage and cost margins yield matching upper and lower sample-size orders: inverse squared margin times log inverse failure probability, for failure probabilities at most .1. No finite fixed calibration sample can uniformly certify the exact worst-law oracle budget or maximum coverage. These are conditional mathematical results for a stipulated finite model, with no claim of physical validation, historical priority or a consciousness criterion.

1. Purpose and relation to previous work

The sensor report [1] distinguishes state identification, sensor-law identification and reliable decisions. Its third iteration solves a single-episode robust decision problem with no inherited calibration evidence. This addendum adds a paid training stage and answers how accurately the shared sensor law must be learned before a deployment policy can meet a prescribed certificate.

For an unknown dependence parameter in $[0, T]$, the earlier report establishes a robust coverage and sensing-cost frontier. In particular, at $T = .04$, maximum guaranteed coverage is $11/16$, whereas a controller told the actual parameter can attain $3/4$. The earlier report also establishes the broader feasibility threshold $T = 3/40$. The new results below, including all statistical constants, are restricted to the parameter space $\Theta = [0, 1/25]$.

The specialized increments are the complete calibration likelihood, the arbitrary-interval decision frontier, and the resulting certification bounds and endpoint impossibilities. We attribute experiment comparison to classical statistical decision theory [2], concentration to Hoeffding [3], and the testing lower-bound method to established change-of-measure arguments [4]. Risk calibration, selective-error control and budgeted sequential calibration already have substantial antecedents [5–7]. We claim no new general concentration, calibration or testing principle.

2. Model and certification contract

2.1 One episode

An episode has a fair target bit θ, fixed throughout the episode. A report d is correct exactly when $d = \theta$. The controller may instead abstain.

The primary sensor has modes $\mathsf{H}, \mathsf{R}, \mathsf{U}$ with initial probabilities

$$\pi = (2/5, 2/5, 1/5).$$

On input x, mode H returns x through an independent binary symmetric channel with error $1/10$; mode R reverses x before that channel; mode U returns an independent fair bit.

A mandatory observation C tests a trusted physical zero through the initial mode L_0, at cost $1/50$. The mode then changes once according to

$$P(L_1 = m \mid L_0 = l) = \tfrac{3}{4}\, 1_{\{m=l\}} + \tfrac{1}{4}\pi_m.$$

A mandatory observation P tests θ through L_1, at cost $1/100$. The mode remains fixed thereafter. Primary channel draws are conditionally independent given the modes and inputs. The initial mode is independent of θ.

The two potential reference readings are $Q_i = \theta \oplus F_i$. Their joint error probabilities are

(F_1, F_2)	00	01	10	11
P_t	$3/5 + t$	$1/5 - t$	$1/5 - t$	$t.$

Both marginals have error $1/5$. The error pair is independent of θ, all primary modes and noise, and controller coins. The unknown t is fixed in $\Theta = [0, 1/25]$. Each purchased reference costs $2/25$. All charges are stipulated model units, not a monetary quotation.

After C, P, a deployment policy may abstain or report, purchase Q_1, and then stop or purchase Q_2. Unpurchased readings are not available. Policies may use all acquired readings and private randomization. Observations have no physical side effects in this model.

2.2 Calibration and deployment

Each of n calibration episodes purchases all four observations and exposes the complete record

$$Z_i = (C_i, P_i, Q_{1i}, Q_{2i}).$$

Its charge is $19/100$, including the compulsory primary pair. Calibration episodes are independent conditional on the same t, with fresh targets, modes and noise. Neither truth labels nor mode labels are revealed.

The integer n is fixed before sampling. A learner maps the records and any independent private coins to a frozen deployment policy. Deployment episodes are fresh independent episodes at the same t; they do not update the learner. Marginal accuracies, reference polarity and the primary model are supplied assumptions, not quantities learned here.

An episode policy can be represented by a distribution over the finite set of deterministic allowed history trees. Thus its performance is linear in its mixture weights. Arbitrary real probabilities are allowed mathematically; no hardware cost or exact rational implementability claim follows.

For a fixed policy φ, define

$$C_m(t;\varphi) = P_t(\text{accept} \mid L_1 = m), \qquad E_m(t;\varphi) = P_t(\text{accept}, d \neq \theta \mid L_1 = m),$$

$$N_m(t;\varphi) = \mathbb{E}_t[\text{number of purchased references} \mid L_1 = m].$$

These average over the target, old mode, readings and policy coins on a fresh episode. They are neither per-history nor per-target-state guarantees.

Fix $0 < \gamma < 3/4$, $0 < \delta < 1/2$ and $0 \le b \le 2$. Let

$$\mathcal{F}_t(\gamma, b) = \{\varphi : \ \forall m,\ C_m(t;\varphi) \ge \gamma,\ E_m(t;\varphi) \le C_m(t;\varphi)/10,\ N_m(t;\varphi) \le b\}.$$

The learner must satisfy

$$\forall t \in \Theta, \qquad P_t\{\widehat{\varphi} \in \mathcal{F}_t(\gamma, b)\} \ge 1 - \delta. \tag{PC}$$

The probability is over training and learner randomization. All three population criteria and all modes must hold on the same good-training event. Uniformity means one learner works for every actual t; it does not identify datasets drawn under different laws.

Define $n^*(\gamma, b, \delta)$ as the least fixed sample size for which such a learner exists, or infinity. This certificate concerns population performance conditional on training, not the realized error fraction in finitely many deployment trials.

3. Likelihood and information in the paid record

Proposition 1 (complete likelihood and sufficiency). Put

$$A(t) = 3/5 + 2t, \qquad \eta = 144/625.$$

For every four-bit record,

$$P_t(c, p, q_1, q_2) = \begin{cases} (1/5 - t)/4, & q_1 \neq q_2, \\ [A(t) + \eta(-1)^{c+p+q_1}]/8, & q_1 = q_2. \end{cases} \tag{1}$$

Define categories $D, +, -$ as disagreement, agreement with $c + p + q_1$ even, and agreement with that sum odd. Their probabilities are

$$p_D(t) = 2/5 - 2t, \quad p_+(t) = 519/1250 + t, \quad p_-(t) = 231/1250 + t. \tag{2}$$

For fixed n, the three category counts are minimal sufficient for t.

Proof. The conditional sign means of the primary modes are $s = (4/5, -4/5, 0)$. Both C and $P \oplus \theta$ have zero sign mean. The refresh kernel gives their sign correlation

$$a = \tfrac{3}{4} \sum_m \pi_m s_m^2 + \tfrac{1}{4} \left(\sum_m \pi_m s_m \right)^2 = 48/125.$$

Consequently

$$P(C = c, P = p \mid \theta) = [1 + a(-1)^{c+p+\theta}]/4.$$

Sum this law against the reference-error table and the fair θ. When references disagree, both possible target values give reference weight $1/5 - t$; when they agree, the weights are $3/5 + t$ and t. This yields (1), with $\eta = (3/5)a$.

There are eight records in category D and four in each agreement category. Within each category their conditional distribution is uniform and independent of t, proving sufficiency and (2). The count likelihood is proportional to

$$p_D(t)^{n_D} p_+(t)^{n_+} p_-(t)^{n_-}.$$

A ratio of two such likelihoods can be constant on the parameter interval only if every count difference is zero: the logarithmic derivative has three distinct simple poles, whose residues are those differences. This proves minimal sufficiency. Two counts suffice for storage because their sum is n. □

All observed-record probabilities are positive on Θ, including at $t = 0$. A zero probability for the latent error pair 11 at that endpoint does not create a zero observed-record probability.

Proposition 2 (strict information comparison). The complete experiment is strictly more informative than retaining only the disagreement bit. Its per-episode Fisher information exceeds that of disagreement by

$$I_{\text{full}}(t) - I_D(t) = \frac{4\eta^2}{A(t)[A(t)^2 - \eta^2]} > 0. \tag{3}$$

Proof. Disagreement is a deterministic function of the complete record. To exclude a parameter-independent reverse randomization, write $q = p_D(t)$. Then

$$p_-(t) = 481/1250 - q/2.$$

A kernel producing category minus from disagreement would require

$$P(- \mid D = 0) = 481/1250, \qquad P(- \mid D = 1) = -72/625,$$

which is impossible. Thus the experiments are not equivalent under parameter-independent randomization, in the comparison framework of [2].

Direct differentiation of the positive categorical probabilities gives

$$I_{\text{full}} = 4/p_D + 1/p_+ + 1/p_-, \qquad I_D = 4/p_D + 4/A.$$

Subtraction gives (3). □

Strict informativeness does not by itself establish a smaller n^* for any specified certification target. A full-record learner can always discard information, so it can do no worse than the restricted learner at the same sensing charge; strict finite-sample improvement for this task remains open.

4. Exact interval decision frontier

For $0 \le l \le u \le 1/25$, write

$$K_u = 3/5 - 8u, \quad M_u = 1 + K_u, \quad L_{l,u} = 6/5 + 2l - 8u, \quad R_u = 2 + K_u.$$

Theorem 3 (interval frontier). A single deployment policy can satisfy accepted error at most $1/10$ and coverage at least γ in every current mode, for every $t \in [l, u]$, if and only if

$$0 \le \gamma \le \Gamma(l, u) := \frac{L_{l,u}}{M_u}. \tag{4}$$

At feasible γ, its minimum worst-law, worst-current-mode expected reference count is

$$B(l, u; \gamma) = \frac{R_u}{L_{l,u}}\gamma. \tag{5}$$

At zero coverage the error condition is interpreted as $E_m \le C_m/10$; the conditional ratio need not be defined.

Proof. First condition on $L_1 = \mathsf{U}$. Stationarity and reversibility of the refresh kernel give

$$P(L_0 = l \mid L_1 = \mathsf{U}) = \tfrac{3}{4}1_{\{l=\mathsf{U}\}} + \tfrac{1}{4}\pi_l.$$

Thus C is fair and P is an independent fresh fair bit. The pair (C, P) is independent of θ, the reference errors and t. In this conditional experiment it supplies only auxiliary randomization.

For any policy, let a_0, x be its accepted masses before any reference and after exactly one reference. Let q_1, q_2 be its query probabilities. Let w, v be its acceptance realization weights following agreement and disagreement after two queries. These weights are independent of t, bounded above by q_2, and are multiplied by the corresponding agreement/disagreement probabilities to obtain accepted masses. Independence follows because the observable prefixes before the second reference have t-independent laws; the fair target also makes the first reference value independent of the potential error pair. Averaging over prefixes and private coins therefore gives

$$q_1 + a_0 \le 1, \quad q_2 + x \le q_1, \quad w, v \le q_2,$$

$$C_{\mathsf{U}}(t) = a_0 + x + (3/5 + 2t)w + (2/5 - 2t)v,$$

$$E_{\mathsf{U}}(t) \ge a_0/2 + x/5 + tw + (1/5 - t)v, \qquad N_{\mathsf{U}} \ge x + 2w. \tag{6}$$

The error inequality uses the smallest possible error at each leaf, so it bounds even policies reporting the wrong polarity.

The risk condition at u implies

$$4a_0 + x + M_u v \le K_u w.$$

Put $z = K_u w - x - 4a_0 - M_u v \ge 0$, and $s = R_u/L_{l,u}$. Then

$$C_{\mathsf{U}}(l) = L_{l,u}(w - v) - 3a_0 - z,$$

$$N_{\mathsf{U}} - sC_{\mathsf{U}}(l) \geq v + (3s - 4)a_0 + (s - 1)z \geq 0. \tag{7}$$

Here $s \geq 13/6$: at fixed u it is minimized at $l = u$, and $R_u/L_{u,u}$ increases with u. This proves the cost lower bound.

Also $a_0 + x + w \leq 1$, whence $M_u(w - v) \leq 1 + 3a_0 + z$. Therefore

$$C_{\mathsf{U}}(l) \leq \Gamma(l, u) + [\Gamma(l, u) - 1](3a_0 + z) \leq \Gamma(l, u),$$

because $\Gamma(l, u) \leq 3/4$. Conditioning on one mode supplies valid lower bounds for the full policy class.

For attainment, after compulsory C, P, buy Q_1 with probability $M_u\gamma/L_{l,u}$. After Q_1, accept its answer with probability K_u/M_u; otherwise buy Q_2, accepting only on agreement. The policy ignores primary readings. In every current mode,

$$C_m(t) = \gamma(K_u + 3/5 + 2t)/L_{l,u}, \quad E_m(t) = \gamma(K_u/5 + t)/L_{l,u}, \quad N_m(t) = R_u\gamma/L_{l,u}. \tag{8}$$

Coverage is minimized at l; $E_m - C_m/10$ is increasing in t and vanishes at u. All coin probabilities are valid exactly within (4). This proves both claims. □

The theorem recovers the earlier $[0, T]$ result on the present parameter range and the known-law formula at $l = u = t$:

$$\Gamma(t, t) = 3/4, \qquad B(t, t; \gamma) = \frac{13/5 - 8t}{6/5 - 6t}\gamma. \tag{9}$$

For interval width $h = u - l$,

$$\Gamma(l, u) = 3/4 - \frac{2h}{8/5 - 8u}. \tag{10}$$

The broader threshold 3/40 in [1] concerns its stated robust model; the calibration constants below are not extended to that larger domain.

5. Certification lower bounds and attainable rates

5.1 Incompatible decisions yield a testing problem

Lemma 4 (full-record testing bound). For $0 \leq l < u \leq T = 1/25$, if

$$\gamma > \Gamma(l, u) \quad \text{or} \quad b < B(l, u; \gamma), \tag{11}$$

then any learner satisfying (PC) with n samples obeys

$$n \geq \frac{\mathrm{kl}(1 - \delta, \delta)}{\mathrm{KL}(p(l), p(u))}, \tag{12}$$

and the corresponding inequality with the KL direction reversed. Here logarithms are natural and $p(t) = (p_D(t), p_+(t), p_-(t))$.

Proof. For a fixed policy its performance is affine in t, since the only varying episode law is the affine reference-error table. Its feasibility inequalities are linear in finite-tree mixture weights. Thus $\mathcal{F}_t$ is closed and measurable, and feasibility at l, u is equivalent to feasibility throughout their interval. Theorem 3 shows that (11) makes $\mathcal{F}_l$ and $\mathcal{F}_u$ disjoint.

The event that the learned policy lies in $\mathcal{F}_l$ has probability at least $1 - \delta$ under l, and at most δ under u. Data processing for relative entropy bounds the divergence between these two event probabilities by that of the training records. Independence gives n times the one-record KL. Independent learner coins add no divergence. The within-category distributions in Proposition 1 are parameter-independent, so the categorical KL equals the full four-reading KL. Monotonicity of binary KL for event probabilities on opposite sides of $1/2$ gives (12). Interchanging the laws gives the reverse-direction bound. □

This is a specialization of classical testing arguments, including the change-of-measure framework of [4, Lemma 1]. It does not replace the full-record experiment with its disagreement projection.

5.2 Complete feasibility classification

Put $T = 1/25$, so

$$M_T = 32/25, \qquad L_{T,T} = 24/25, \qquad R_T = 57/25.$$

Theorem 5 (calibration feasibility and matching-order bounds). Under (PC):

1. If $b \le (19/8)\gamma$, then $n^* = \infty$.

2. If $b > (19/8)\gamma$, define

$$w_{\text{cov}} = \frac{M_T}{2}(3/4 - \gamma), \qquad w_{\text{cost}} = \frac{L_{T,T} - R_T\gamma/b}{2}, \qquad w = \min(w_{\text{cov}}, w_{\text{cost}}). \tag{13}$$

Both widths are positive. If $w \ge T$, then $n^* = 0$. Equivalently this is the region

$$\gamma \le 11/16, \qquad b \ge (57/22)\gamma.$$

If $0 < w < T$, then

$$\frac{\text{kl}(1-\delta, \delta)}{84w^2} \le n^* \le \left\lceil \frac{\log(2/\delta)}{2w^2} \right\rceil. \tag{14}$$

Consequently, for $0 < \delta \le .1$ and $0 < w < T$,

$$n^* = \Theta\left(\frac{\log(1/\delta)}{w^2}\right).$$

3. If the coverage domain is extended to $\gamma = 3/4$, no finite n satisfies (PC), regardless of the permitted $b \le 2$. Coverage above $3/4$ is impossible even at a known law.

Proof. The known-law cost at T, from (9), is $(19/8)\gamma$. A smaller budget makes $\mathcal{F}_T$ empty. At equality, every pair $T - d, T$, $0 < d < T$, is cost-incompatible. Its full-record KL tends to zero as $d \downarrow 0$, whereas $\text{kl}(1-\delta, \delta) > 0$. Lemma 4 therefore excludes every finite n. This is a uniform-over-t

impossibility, not a prohibition on exact optimality at a particular known law. At $\gamma = 3/4$, (10) makes every positive-width pair incompatible, and the same reasoning applies.

Consider positive widths in (13). If $w \geq T$, Theorem 3 supplies a policy valid on all of Θ without training. Substitution into (13) gives the equivalent region stated above.

For $0 < w < T$, an upper bound can use only disagreement. Let $\widehat{q}$ be its sample mean and $\widehat{t} = (.4 - \widehat{q})/2$. By Hoeffding's inequality [3],

$$P_t(|\widehat{t} - t| > w/2) = P_t(|\widehat{q} - q(t)| > w) \leq 2e^{-2nw^2}. \tag{15}$$

Intersect $[\widehat{t} - w/2, \widehat{t} + w/2]$ with Θ. On its coverage event it is nonempty and has width at most w. If it is empty, choose a predeclared bounded fallback, such as abstention; this cannot occur on the good event.

Use the attaining policy of Theorem 3 for the interval obtained. At fixed width h, capacity is smallest at upper endpoint T, by (10). The cost slope is largest there because

$$\frac{d}{du}\frac{13/5 - 8u}{6/5 - 6u - 2h} = \frac{6 + 16h}{(6/5 - 6u - 2h)^2} > 0.$$

The two width conditions in (13) therefore imply all three criteria in every mode. No further union bound over modes is needed. Equation (15) yields the upper bound in (14).

For the lower bound, choose

$$d_0 = \min(2w, (T + w)/2).$$

Then $w < d_0 < T$ and $d_0 \leq 2w$. At $l = T - d_0, u = T$, either the coverage or cost threshold in (13) is exceeded, so (11) holds. The categorical probabilities obey

$$p_+ \geq 519/1250, \qquad p_- \geq 231/1250, \qquad p_D \geq 8/25.$$

Their derivatives are $1, 1, -2$, up to ordering. The inequality $\mathrm{KL}(P, Q) \leq \sum_j (P_j - Q_j)^2/Q_j$, obtained from $\log x \leq x - 1$, gives

$$\mathrm{KL}(p(l), p(u)) \leq d_0^2(1250/519 + 1250/231 + 25/2) < 21 d_0^2 \leq 84w^2.$$

Apply Lemma 4. The zero-sample case in this region is also excluded for randomized learners: their output law would be identical under two incompatible laws, but it cannot put probability greater than $1/2$ in each of two disjoint feasibility sets.

Finally $\mathrm{kl}(1 - \delta, \delta) = (1 - 2\delta)\log((1 - \delta)/\delta)$ is comparable to $\log(1/\delta)$ for $0 < \delta \leq .1$. This proves the rate statement. □

The upper bound discards the additional information identified in Proposition 2. Nevertheless it matches the full-record lower bound in order. Optimal constants and exact finite-sample savings from retaining the primary readings remain unresolved.

6. Cost interpretation and scope

For a horizon of H deployment episodes, on the good-training event the total expected sensing charge is at most

$$.19n + H(.03 + .08b). \tag{16}$$

This bounds conditional expected cost, not a deterministic realized bill. Since each episode buys at most two references, an unconditional bound is

$$.19n + H\{.03 + .08[(1 - \delta)b + 2\delta]\}. \tag{17}$$

Calibration failure also affects acceptance; one cannot generally add δ to the conditional accepted-error target 1/10.

For $\gamma \leq 11/16$, the certain no-training benchmark has reference count $b_0 = (57/22)\gamma$. If $b < b_0$, the certified upper bound (16) is below that benchmark whenever

$$H > \frac{.19n}{.08(b_0 - b)}. \tag{18}$$

This is sufficient for a certified good-event expected-cost saving. An actual learned policy can break even sooner if its expected reference count is below b. This accounting comparison is not a new value-of-information theorem. The confidence guarantee differs from the certain baseline and must be disclosed. Unconditional cost comparisons should instead use (17). For coverage exceeding 11/16, no training is infeasible; the comparison is then one of feasibility.

The results assume known marginal accuracies, known polarity, known primary dynamics, fresh independent episodes and a common stationary parameter. They do not validate those assumptions from the same unlabeled data. They do not treat persistent target states, unknown marginal errors, changing dependence, physical hazards, active stopping of training or deployment-time learning.

B-Theory here denotes the Being/Becoming research programme; no thesis about the metaphysics of time is assumed. The extension distinguishes decisions about a present state from information about a measurement law shared across episodes. It quantifies a finite resource limit on that learning problem. An explicit self-model and a compiled controller using the same evidence can implement the same policies. No architectural superiority, natural self/world boundary or phenomenal-awareness conclusion follows.

7. Attribution, reproducibility and review status

The relationship to [1] is explicit: its sensor law, single-episode objective, $[0, T]$ frontier and known-law benchmark are inherited. The current likelihood reduction, interval extension and their calibration consequences are the additions. The report's threshold 3/40 is retained as inherited context; the statistical constants here use $T = .04$ throughout.

The concentration and change-of-measure machinery is classical [3,4]. Learn then Test [5] already treats selective risk and simultaneous criteria, while [6] studies active budgeted anytime-valid calibration. The basic confidence-bound approach is also established in [7]. Their existence prevents treating paid calibration or confidence wrapping in general as a contribution. The present specialization involves an unlabeled latent-sensor record and its particular decision geometry. Historical priority for that specialization remains unconfirmed.

The likelihood is reproducible by summing the finite hidden-state law over θ, L_0, L_1 for each of sixteen observed records and comparing the resulting affine coefficients with (1). The frontier can be checked through the elimination identities (6)–(8), and the statistical bound through

the probability floor and the derivative identities in Theorem 5. These finite checks corroborate algebra; the proofs supply the continuum and all-learner claims.

The derivation and manuscript have undergone separate internal AI mathematical and hostile-review passes. The review corrected the sufficient-condition interpretation of the cost comparison and made the testing lemma's distinct-parameter domain explicit. This process is not external human peer review or proof-assistant certification. No simulation or physical experiment establishes performance in this addendum. Exact finite n^*, optimal leading constants, a demonstrated finite-sample advantage of full-record calibration, and external priority remain open.

References

[1] B-Theory research programme. *B-Theory Sensor Reliability: Final Three-Iteration Research Report.* Internal research report, 4 November 2026, especially iteration 3, sections 1–4. The source report is to accompany this addendum's reproducibility package.

[2] David Blackwell. "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics* 24(2), 265–272, 1953. https://doi.org/10.1214/aoms/1177729032

[3] Wassily Hoeffding. "Probability Inequalities for Sums of Bounded Random Variables." *Journal of the American Statistical Association* 58(301), 13–30, 1963. https://doi.org/10.1080/01621459.1963.10500830

[4] Emilie Kaufmann, Olivier Cappé and Aurélien Garivier. "On the Complexity of Best-Arm Identification in Multi-Armed Bandit Models." *Journal of Machine Learning Research* 17(1), 1–42, 2016. https://jmlr.org/papers/v17/kaufman16a.html

[5] Anastasios N. Angelopoulos, Stephen Bates, Emmanuel J. Candès, Michael I. Jordan and Lihua Lei. *Learn then Test: Calibrating Predictive Algorithms to Achieve Risk Control.* arXiv:2110.01052, version 5, 2022. https://arxiv.org/abs/2110.01052

[6] Ziyu Xu, Nikos Karampatziakis and Paul Mineiro. *Active, anytime-valid risk controlling prediction sets.* NeurIPS 2024; arXiv:2406.10490, version 2. https://arxiv.org/abs/2406.10490

[7] Stephen Bates, Anastasios Angelopoulos, Lihua Lei, Jitendra Malik and Michael I. Jordan. *Distribution-Free, Risk-Controlling Prediction Sets.* arXiv:2101.02703, 2021. https://arxiv.org/abs/2101.02703