DECISION TRANSPARENCY: WHEN CAN A BLACK-BOX DECISION REVEAL ITS HIDDEN INFORMATION?

Anonymous authors
Paper under double-blind review

ABSTRACT

A good decision can be correct for the wrong reason. Decision quality asks whether an action performs well. Decision transparency asks what the released action reveals about the information that shaped it. This is a missing evaluation dimension for black-box decision systems. We call this method *FanOpt*, a public regret-to-evidence compiler. It makes decision transparency testable from one action under a structured alternative. If the null and alternative induce the same action law, no action-only test can have power above its Type-I error. FanOpt breaks this barrier when low regret meets a public geometric contract. Before private observations arrive, the public evaluator fixes a source-independent baseline and a public distance. Under the null, low regret keeps the returned action near the baseline. Under a structured alternative, low regret forces the action away because remaining near the baseline gives up the information-dependent utility gain. A fixed distance threshold can then turn one released action into statistical evidence. We prove finite-sample Type I and Type II error bounds and an explicit certified regret radius. A marginal-load bound carries estimation and rounding error into the regret condition required by the test. We give certificates for 2-factors and the stated spanning-tree policy. Matching also fits the framework when its family certificate holds. In controlled experiments, power rises from 16.0% at $q = 1$ to 100.0% at $q = 2$, while none of 2,196 low-regret outputs falls on the wrong side of the public threshold. These results show when one black-box combinatorial decision provides valid decision-transparency evidence under a structured alternative. The test does not need the private model, objective, or solver trace.

1 INTRODUCTION

A good decision can be correct for the wrong reason. Consider two routing services that return the same low-cost tree. One uses a private real-time congestion forecast. The other reuses a route learned from historical traffic. The two decisions may have the same utility or even be identical. Yet only the first responds to the current source. Decision quality alone cannot reveal the difference to an observer who sees only the released tree.

Modern AI services, including foundation-model agents, expose decisions but hide their sources and computation. Evaluation still centers on whether an action is good. It rarely asks whether the action reveals why it is good. We call this missing dimension *decision transparency*. It asks what an action can reveal when the data, model, and objective remain private.

Existing approaches cross this information boundary. They inspect internal computation, observe repeated behavior, or use downstream outcomes. None answers the one-shot question when the objective and computation remain hidden. Interpretability explains a model or mechanism. Decision transparency tests whether the released action carries evidence about a stated source.

Fundamental limitation. Let the null omit a structured private source and the alternative include it. If $\mathsf{P}_0^A = \mathsf{P}_1^A$, then every level-$\alpha$ action-only test has power at most α. Theorem 4.1 establishes the identifiability boundary of decision transparency. Repeated decisions, outcomes, or solver access could resolve the ambiguity, but this setting provides none. The routing pair shows why equal utility and equal output leave quality-based evaluation silent.

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.

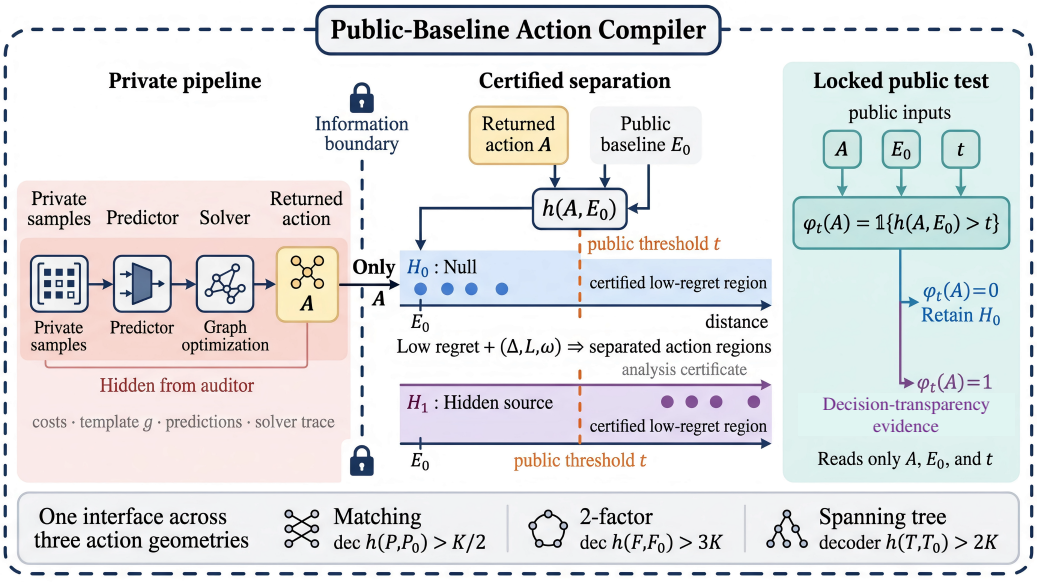

Figure 1. Decision quality alone does not establish decision transparency. FanOpt keeps the private pipeline closed. Only A crosses the information boundary. The public test compares A with a baseline E_0 fixed before private observations and rejects when $h(A, E_0) > t$. Low regret places null actions and actions shaped by structured information on opposite sides of t, turning the action into public evidence.

FanOpt uses low regret to separate the action laws. The public evaluator fixes a source-independent baseline before private observations arrive. Low-regret actions stay near it under the null and cross a public threshold under a structured alternative. The action then provides evidence while the private objective remains hidden.

Existing methods use richer information. Decision-focused learning improves actions through task losses or solver feedback (Donti et al., 2017; Wilder et al., 2019; Elmachtoub & Grigas, 2022; Vlastelica Pogančić et al., 2020; Mandi et al., 2022). Inverse optimization uses repeated decisions (Bärmann et al., 2017; Ren et al., 2025). Decision-risk certification uses calibrated outcome distributions (Zhou et al., 2026). FanOpt evaluates one action after optimization. The prediction, objective, solver, and outcomes remain private. Its target is structured source influence under the stated regret geometry.

We pair the impossibility boundary with a public geometry certificate that turns low regret into separated action regions and finite-sample error bounds. A marginal-load theorem controls estimation and rounding error. We give 2-factor and spanning-tree constructions, a conditional matching extension, and experiments in which power rises from 16.0% at $q = 1$ to 100.0% at $q = 2$.

The *FanOpt compiler* fixes the public baseline, distance, threshold, and test. *FanOpt-Spectral* is one private pipeline. Any pipeline that meets the regret geometry can use the same rule. The evaluator sees only the action, baseline, and threshold.

Our contributions are fourfold:

- **A missing evaluation dimension.** We formalize decision transparency for black-box AI systems and characterize its statistical limits. Quality asks whether an action works. Transparency asks what shaped it.
- **An identifiability boundary.** Identical action laws rule out source-influence evidence from one released action. Every level-α action-only test then has power at most α.
- **A regret-to-evidence principle.** A public baseline and a uniform certificate separate low-regret actions under the null and structured alternative. The compiler turns regret-tail bounds into finite-sample Type-I and Type-II error bounds from the released action.
- **Combinatorial realizations and evidence.** Marginal-load bounds link estimation error to regret for 2-factors and spanning trees. Matching requires its family certificate. Experiments verify the regret–evidence transition across five private generators and test transfer.

2 RELATED WORK

Decision quality versus decision transparency. Decision-quality methods ask how predictions can improve a downstream action. Decision-focused learning pursues this goal through task losses, differentiable solvers, ranking objectives, and learned optimization maps (Donti et al., 2017; Elmachtoub & Grigas, 2022; Vlastelica Pogančić et al., 2020; Mandi et al., 2022). Recent methods reduce solver calls, handle uncertain constraints, or protect regret under measurement error and shift (Berden et al., 2025; Mandi et al., 2025; Kong et al., 2025; Yamao et al., 2026). FanOpt asks what one released action can show when the objective remains private. Model interpretability explains the hidden mechanism (Rudin, 2019). Decision transparency asks whether the released action carries statistical evidence about a stated source. FanOpt leaves the private predictor and solver unchanged. Appendix comparisons test whether the public interface preserves measured decision utility.

Optimization interfaces. Black-box differentiation, score-function estimators, learned surrogates, and search policies change how a private action is computed (Vlastelica Pogančić et al., 2020; Silvestri et al., 2026; Tang & Khalil, 2024; Balcan et al., 2018). FanOpt runs after that computation. It needs only the released action and a public baseline, so it can sit on top of any engine whose output meets the stated regret and geometry conditions.

Repeated behavior and outcome-based checks. Inverse optimization recovers objectives, constraints, or system states from repeated decisions (Bärmann et al., 2017; Ren et al., 2025; Bao & Sakaue, 2025). Decision-risk certification evaluates a candidate action using calibrated outcome distributions (Zhou et al., 2026). Both settings provide information that is unavailable here. FanOpt tests one solver output against a source-independent baseline fixed before private observations. Its distinct information set defines the problem. The method does not recover the objective.

Planted inference and low-regret reductions. Planted inference tests or recovers a hidden template from noisy observations (Bagaria et al., 2020; Chen & Xu, 2016; Even & Ganassali, 2026). Active testing chooses observations (Yang & Xing, 2026). Low-regret reductions connect hidden structure to feasibility or selection (Awasthi et al., 2015; Gangrade et al., 2024; Nayebi, 2026). FanOpt changes the endpoint of the reduction. The final statistic is computed from a released combinatorial action. The geometry certificate gives the test, and the marginal-load result states when it survives estimation and rounding for the stated 2-factor and spanning-tree constructions. Matching follows conditionally when its family certificate is verified.

3 PROBLEM FORMULATION

Let E be a finite ground set. We represent each combinatorial action as a subset of E. Let $\mathcal{A} \subseteq 2^E$ be the finite feasible family. We observe cost samples $C_1, \dots, C_m \in \mathbb{R}^E$. A budget-$m$ rule returns

$$A_m = \mathsf{Alg}_m(C_1, \dots, C_m; U), \tag{1}$$

where U is its internal randomness. For a population utility $U_g : \mathcal{A} \to \mathbb{R}$, we measure population regret as

$$\mathsf{Reg}_g(A_m) = \max_{A \in \mathcal{A}} U_g(A) - U_g(A_m). \tag{2}$$

For a population cost mean $\mu = \mathbb{E}C_i$, put $c_\mu(A) = \langle \mu, \mathbf{1}_A \rangle$, $U_\mu(A) = -c_\mu(A)$, and $\mathsf{Reg}_\mu(A) = c_\mu(A) - \min_{B \in \mathcal{A}} c_\mu(B)$. We write $\mathfrak{A}_{\mathrm{all}}$ for all such rules. The class $\mathfrak{A}_{\mathrm{poly}}$ contains uniform randomized algorithms with runtime at most Cs^c. Here C and c do not depend on dimension, and s is the input-array size. We use a real-Gaussian computation model. We count real arithmetic, comparisons, Gaussian draws, and every combinatorial-oracle call. Information bounds optimize over $\mathfrak{A}_{\mathrm{all}}$. Computational bounds use $\mathfrak{A}_{\mathrm{poly}}$. The budget also counts sample generation, perturbations, calibration, recovery, and the final oracle call. The reduction cannot hide work in preprocessing.

Public transparency contract. We fix the public baseline before private observations. It can be a historical assignment, a static cycle cover, or a fixed routing tree. Independent calibration may choose from a fixed public threshold menu $\mathcal{T}$, without using the final evaluation suffix. This prevents final-suffix tuning. It does not provide distribution-free Type-I control under temporal drift. We publish $(E_0, h, \mathcal{T})$ and lock t before evaluation. At deployment, anyone can evaluate the test from

(A, E_0, t). Its error guarantee is conditional on the declared regret-tail bounds and the uniform geometry event in Theorem 4.2. Anyone can evaluate the rule from the public inputs. Its guarantee still depends on the stated validity conditions.

The private service may use any estimator and solver. We measure evasion against its declared private objective. Theorem 4.2 gives regret greater than ρ to every wrong-side action. This includes a null action with $h(A, E_0) > t$. On $\mathcal{G}$, it also includes an alternative action with $h(A, E_0) \le t$. Forcing an informed action below the threshold gives up more than ρ of private utility. This result measures the cost of evasion. It does not give a general incentive-compatibility guarantee. If no calibrated threshold protects the null, the service returns E_0.

The test targets source influence. It does not recover the source. Rejecting the null means that the released action conflicts with low-regret null behavior under the locked baseline and threshold. The result does not identify the hidden template or certify every property of the private system. A failure to reject also does not prove that no private signal exists. The guarantee applies when the returned action has low regret and the action family satisfies the stated geometry certificate.

3.1 THE STATISTICAL TEST

The source-level statistical problem is

$$H_0 : \mathsf{P} = \mathsf{P}_0, \qquad H_1 : \mathsf{P} = \mathsf{P}_1 := \int \mathsf{P}_g \, d\Pi(g), \tag{3}$$

where P_0 and P_1 are laws of the private observations and Π is the mixing law over hidden templates fixed in advance. Let P_0^A and P_1^A denote the induced laws of the returned action under the decision rule in (1), including its internal randomness. Under P_0, the public baseline E_0 is optimal. The variable g indexes a hidden alternative. From one returned action A, FanOpt computes

$$T(A) := h(A, E_0), \qquad \phi_t(A) := \mathbf{1}\{T(A) > t\}. \tag{4}$$

Its Type-I error, Type-II error, power, and testing risk are

$$\begin{aligned} \mathsf{Err}_{\mathrm{I}}(t) &:= \mathsf{P}_0^A\{\phi_t(A) = 1\}, & \mathsf{Err}_{\mathrm{II}}(t) &:= \mathsf{P}_1^A\{\phi_t(A) = 0\}, \\ \mathsf{Power}(t) &:= 1 - \mathsf{Err}_{\mathrm{II}}(t), & \mathsf{Risk}_{\mathrm{test}}(t) &:= \mathsf{Err}_{\mathrm{I}}(t) + \mathsf{Err}_{\mathrm{II}}(t). \end{aligned} \tag{5}$$

At deployment, the statistic uses only A, E_0, and t. Public size parameters or independent calibration fix t before evaluation. The hidden template, private costs, certificate constants, solver state, and trace remain inside the optimization service. Theorem 4.2 gives the condition under which low regret makes both errors in (5) small.

The analysis certificate proves geometric separation. Deployment publishes the baseline, distance, and threshold needed to use that separation. The source premise enters only when deriving source-specific recovery conditions. Appendix Table 3 records which quantities are public and which remain private. Appendix Table 4 separates sufficient compiler assumptions from family-specific construction conditions. The counterexamples identify failure cases for uniform guarantees. They do not show that this contract is necessary for every action-only test. Optimization quality alone is therefore not evidence. It becomes evidence only when the public geometry places low-regret actions on opposite sides of a fixed threshold.

4 FROM REGRET TO PUBLIC EVIDENCE

FanOpt is a post-solver interface. It publishes a feasible baseline E_0, a distance h, and a threshold t. A private predictor and solver return a feasible action A. The public rule then computes $\phi_t(A) = \mathbf{1}\{h(A, E_0) > t\}$. The compiler theorem places low-regret null and mixture actions on opposite sides of the same threshold. Marginal-load stability bounds the regret caused by estimation under the stated load conditions. The 2-factor and spanning-tree constructions give their baselines, thresholds, and error constants. Matching uses the same interface when its geometry certificate is verified. The private predictor and solver can change. The public interface does not. Figure 1 shows this separation between private computation and public testing.

4.1 FUNDAMENTAL LIMITS AND THE REGRET-TO-EVIDENCE PRINCIPLE

We begin with the boundary of any action-only method.

Theorem 4.1 (Decision-transparency impossibility)**.** Let $\phi(A) \in \{0, 1\}$ be any action-only test with $\mathsf{P}_0^A\{\phi(A) = 1\} \le \alpha$. If the induced action laws satisfy $\mathsf{P}_0^A = \mathsf{P}_1^A$, then

$$\mathsf{P}_1^A\{\phi(A) = 1\} = \mathsf{P}_0^A\{\phi(A) = 1\} \le \alpha. \tag{6}$$

Distinct returned-action laws are necessary for power above the Type-I level.

Proof. The equality of the two laws makes the rejection probabilities equal for every measurable function of A, including ϕ. □

The theorem defines the boundary. The compiler below gives a constructive condition that separates the two action laws. It needs one public statistic, and low regret must push this statistic in opposite directions under the two hypotheses. Equal-cardinality action families provide it through distance from a public baseline. Let $\mathcal{A}$ be a finite family of edge sets. Let $E_0 \in \mathcal{A}$ be fixed publicly. Assume that every action has cardinality M. Define

$$h(A) = |A \setminus E_0| = |E_0 \setminus A|. \tag{7}$$

The proof uses standard regret identities, deterministic score perturbation, and ordinary error accounting. FanOpt adds a public link from these regret inequalities to an action-only statistical test. The combinatorial families give concrete instances of this link. For a hidden template $g : E \to [-G, G]$, write

$$s_g(A) = \sum_{e \in A} g(e), \qquad S_g = \max_{A \in \mathcal{A}} s_g(A), \qquad b_g = s_g(E_0). \tag{8}$$

We say that (Δ, L, ω) is a baseline-geometry certificate if

$$\begin{gathered} S_g - b_g \ge \Delta, \\ \exists Q_g \in \mathcal{A} : \; s_g(Q_g) = S_g, \qquad h(Q_g) \le L, \\ s_g(A) - s_g(E_0) \le \omega h(A) \quad (A \in \mathcal{A}). \end{gathered} \tag{9}$$

The last inequality always holds with $\omega = 2G$. If $g(e) = 0$ for every $e \in E_0$, then $\omega = G$ suffices. The triple is an analysis certificate. It may depend on the hidden template. The deployed test reads only $h(A)$ and a fixed threshold. It does not observe g, Δ, L, or ω.

Certificate-compatible structured alternatives. For public (E_0, t, ρ), a measurable class $\mathcal{G}$ is *certificate-compatible* when every $g \in \mathcal{G}$ has the displayed geometry certificate and satisfies the same threshold inequalities below. A mixture may place at most η_{geom} mass outside $\mathcal{G}$. Alternatives without this public separation are outside the theorem's scope.

The baseline creates the separating coordinate. A useful baseline is feasible, source independent, and fixed before private observations. The gap $S_g - b_g$ measures gain beyond this reference. Small L provides a nearby alternative witness. Bounded ω prevents hidden score from growing faster than public distance. Together, these properties make a single threshold possible. Changing the baseline requires checking these constants again. Table 1 records the constructed families and the conditional matching extension.

For $a > 0$, $0 < \beta < a\omega$, and a regret radius $\rho \ge 0$, consider

$$U_0(A) = \beta|A \cap E_0|, \qquad U_g(A) = \beta|A \cap E_0| + a s_g(A). \tag{10}$$

Theorem 4.2 (regret-to-test compiler)**.** Under (9), E_0 is the unique null maximizer and

$$\mathrm{Reg}_0(A) = \beta h(A). \tag{11}$$

Under the hidden template,

$$\mathrm{Reg}_g(A) \ge a\Delta - \beta L - (a\omega - \beta)h(A). \tag{12}$$

Let $\mathcal{G}$ be a certificate-compatible structured alternative in the sense stated above. Its common public threshold must obey

$$\rho < \beta t, \qquad \rho < a\Delta - \beta L - (a\omega - \beta)t. \tag{13}$$

Then $\mathbf{1}\{h(A) > t\}$ separates all outputs with regret at most ρ. For a pre-specified mixing probability measure Π on the template class, draw $g \sim \Pi$ first and then draw the data and returned action conditionally on g. The same public t must work for every g in $\mathcal{G}$. The mixture error then gains $\Pi(\mathcal{G}^c)$ under this joint law. A pointwise certificate gives only pointwise separation. It does not make a template-dependent threshold public.

Now let the returned action be randomized. Write P_0^A for the induced null law of A. For a fixed template g, write $\mathsf{P}_{1,g}^A$ for the alternative law of the returned action, and let $\mathsf{P}_1^A = \int \mathsf{P}_{1,g}^A \, d\Pi(g)$ be the mixture action law. Suppose

$$\begin{gathered} \mathsf{P}_0^A\{\mathrm{Reg}_0(A) > \rho\} \le \eta_0, \qquad \sup_{g\in\mathcal{G}} \mathsf{P}_{1,g}^A\{\mathrm{Reg}_g(A) > \rho\} \le \eta_1, \\ \Pi(\mathcal{G}^c) \le \eta_{\mathrm{geom}}, \end{gathered} \tag{14}$$

and that (13) holds uniformly on $\mathcal{G}$. Then the test $\phi_t(A) = \mathbf{1}\{h(A) > t\}$ satisfies

$$\begin{aligned} \mathsf{Err}_{\mathrm{I}}(t) &\le \eta_0, & \mathsf{Err}_{\mathrm{II}}(t) &\le \eta_1 + \eta_{\mathrm{geom}}, \\ \mathsf{Risk}_{\mathrm{test}}(t) &\le \eta_0 + \eta_1 + \eta_{\mathrm{geom}}, & \mathsf{Power}(t) &\ge 1 - \eta_1 - \eta_{\mathrm{geom}}. \end{aligned} \tag{15}$$

Proof sketch. Equal cardinality gives $|A \cap E_0| = M - h(A)$, so the null regret is exactly $\beta h(A)$. Under the alternative, compare A with the witness Q_g. The witness condition contributes at least $a\Delta$, moving from E_0 to Q_g costs at most βL, and the oscillation condition bounds the remaining score change by $a\omega h(A)$. Combining these terms gives (12). The two inequalities in (13) then imply the deterministic event inclusions for $\{\mathrm{Reg}_0(A) \le \rho\}$ and $\{\mathrm{Reg}_g(A) \le \rho\}$. Taking probabilities under the induced action laws yields (15). The geometric failure event contributes η_{geom}. □

Limits of a uniform guarantee. Constant policies and overlapping low-regret sets rule out uniform action-only guarantees. A post-hoc baseline can also erase the distance statistic. Appendix A distinguishes these two obstructions.

Reading the guarantee. The exact null identity converts public distance into regret. The alternative bound reverses the implication and forces a low-regret action beyond the same threshold. Equation (15) turns regret tails into Type-I and Type-II error bounds. The deployed rule never identifies g. The private procedure recovers the information needed to construct A. The public test then reads only $h(A, E_0)$. Different predictors and optimizers can share this test without exposing their internal representations.

Theorem 4.1 gives the identifiability boundary. Without an observable action-law separation, no one-shot action-only procedure can have power above its Type-I level.

Continuous envelope and deployed radius. Set $D := a\Delta - \beta L > 0$ and assume $0 < \beta < a\omega$. The two bounds meet at $t_* = D/(a\omega)$. They certify every $\rho < \rho_* := \beta D/(a\omega)$, where

$$\rho_* = \max_{t \ge 0} \min\{\beta t, D - (a\omega - \beta)t\}. \tag{16}$$

This is a continuous envelope. Because $h(A, E_0)$ is integer-valued, the deployed test uses a fixed public threshold from Table 1. The resulting discrete guarantee can be stronger than the continuous value ρ_*. Appendix D records the endpoint construction.

Envelope endpoint sharpness. The two affine terms in (16) are not artifacts of the proof. The appendix constructs an equal-cardinality family with only three feasible actions that reaches both bounds at the same distance. Neither affine term can be increased uniformly at the stated integer distance using only (Δ, L, ω). The continuous envelope should not be read as a universal optimum for the integer-valued deployed test. A larger discrete radius may follow from additional action-space structure.

4.2 TRANSPORTING ESTIMATION ERROR THROUGH MARGINAL LOAD

The compiler needs a low-regret action, but a private pipeline usually optimizes an estimated source. The next result identifies the error quantity that matters. It counts how often a mistaken entity can affect the returned action. This avoids charging every coordinate equally.

Theorem 4.3 (marginal-load plug-in stability)**.** Let E be a finite coordinate set and V be a finite entity set. Let $\partial_s : E \to 2^V$ be a public sensitivity map. Invariant coordinates have $\partial_s(e) = \varnothing$. Fix a finite $\mathcal{A} \subseteq 2^E$. Define $U_g(A) = B(A) + a\sum_{e\in A} g(e)$, where $a > 0$, $G \geq 0$, $g, \widehat{g} : E \to [-G, G]$, and B is public. For a policy q, define $\lambda_v(q) = \mathbb{E}_{A\sim q}\sum_{e\in A} \mathbf{1}\{v \in \partial_s(e)\}$. Let $Z \subseteq V$ be a finite set of entities on which the two score maps may differ. Fix a nonempty compact public subset $\mathcal{Q}_\Lambda$ of the probability simplex on $\mathcal{A}$. Assume $\lambda_v(q) \leq \Lambda$ for every $q \in \mathcal{Q}_\Lambda$ and $v \in V$. Here $\Lambda \geq 0$. Let $\varepsilon_\Lambda(g) = \max_A U_g(A) - \max_{q\in\mathcal{Q}_\Lambda} \mathbb{E}_q U_g(A)$. If g and $\widehat{g}$ agree when $\partial_s(e) \cap Z = \varnothing$, choose $\widehat{q} \in \arg\max_{q\in\mathcal{Q}_\Lambda} \mathbb{E}_q U_{\widehat{g}}(A)$. The perturbation bound is uniform over $\mathcal{Q}_\Lambda$ for each realized $(g, \widehat{g}, Z)$. It also applies to a data-dependent $\widehat{q}$ when these conditions hold for that realization. Every such plug-in optimizer satisfies

$$\mathbb{E}_{\widehat{q}} \operatorname{Reg}_g(A) \leq \varepsilon_\Lambda(g) + 4aG\Lambda|Z|. \tag{17}$$

If $\widehat{q}$ is supported on actions with common values $B(A) = B_0$, $s_{\widehat{g}}(A) = S_0$, and $\max_A U_g(A) = B_0 + aS_0$, a direct bound with factor two holds:

$$\mathbb{E}_{\widehat{q}} \operatorname{Reg}_g(A) \leq 2aG\Lambda|Z|. \tag{18}$$

Proof sketch. For any action in the support of q, the two score mismatches are nonzero only on coordinates whose sensitivity set meets Z. Each such coordinate changes by at most $2G$, and the load condition charges each affected entity at most Λ in expectation under q. We bound the mismatch once for an optimal feasible policy and once for the plug-in policy. This gives the factor-four term in (17). Under all three additional equalities, regret equals one score mismatch, giving the factor-two bound directly. Appendix B gives both proofs and the conversion from expected regret to the tail probability required by the compiler. □

Marginal load turns local estimation errors into regret control. An entrywise error bound treats all action coordinates equally. The load bound uses only coordinates that can change the action score. It then counts how often each affected entity can appear. This gives deterministic constants for matching and 2-factors. Trees replace worst-case degree with expected load.

A source estimator must control errors that enter the action score. The load bound and (20) then control the alternative regret tail. The null regret tail requires a separate bound, such as control of a gate's false activations. The public statistic and threshold stay fixed when the private estimator or optimizer changes.

For matching and 2-factors, bounded sensitive degree makes the load deterministic. The factor-four bound in Theorem 4.3 is the default guarantee. Its factor-two variant requires the additional common-score and common-public-term conditions stated in the theorem. Trees have no uniform degree bound. They use a uniform-spanning-tree construction in Appendix E.2. For balanced estimated signs, its component-tree policy has expected load $2 - 1/K$ and optimizes over the public load-constrained class with zero approximation penalty (Lemma E.2). Arbitrary tree optimizers need not satisfy this load cap.

4.3 PUBLIC STATISTICAL TESTS ACROSS THREE ACTION SPACES

The analysis certificate may use the latent template. The deployed test does not. Each family fixes its baseline and threshold before private samples arrive. The test reads only the action and public baseline. In the table, d_0 is the deterministic load cap and δ is matching-certificate slack.

The three rows are instances of the same compiler. The bounded-degree and spanning-tree appendices construct the 2-factor and tree witnesses and verify their constants. The matching row is conditional on supplying the displayed uniform certificate and sufficient compiler parameters for the chosen public family. The general theorem alone does not give its threshold for arbitrary δ. Matching and 2-factors use deterministic degree bounds. The tree uses expected load from a balanced uniform-spanning-tree policy.

Table 1. One public-test form covers three different action spaces. The baseline, threshold, and error-transport argument change, but every deployed test reads only the returned action and public data.

Family	Public baseline	Certificate and load	Public statistical test
Matching	source-independent permutation	verified uniform matching certificate with $\delta \le 1/16$, $\Delta \ge 2(1-\delta)K$, $L \le 4K$, $\omega = 2$, $d_0 = 1$, and the sufficient compiler choice $\beta = a/64$, $\rho = aK/256$	$\mathbf{1}\{h(P, P_0) > K/2\}$
2-factor	two fixed same-side cycles	$\Delta = 4K$, $L \le 8K$, $\omega = 1$, $d_0 = 2$	$\mathbf{1}\{h(F, F_0) > 3K\}$
Spanning tree	public root star	$\Delta = 4K - 2$, $L = 4K - 2$, $\omega = 1$, $\Lambda = 2 - 1/K$	$\mathbf{1}\{h(T, E_0) > 2K\}$

Deployment procedure. Deployment publishes a feasible baseline, a distance, and a threshold fixed from public size parameters or independent calibration. After the optimizer returns A, the test computes $h(A, E_0)$ and rejects the null when the distance exceeds the threshold. It neither reconstructs g nor reads costs or solver state. A third party can reproduce the decision from the action and public instance alone.

Geometry changes the constants. The test remains public. Every row of Table 1 evaluates a thresholded distance from a source-independent action. Matching and 2-factors use deterministic degree to carry estimation error. Trees use expected load. Their baselines and thresholds differ because their feasible sets differ. Every deployed test uses the same inputs. One action, one public baseline, and public size parameters suffice in every family.

Computational scope. The public certificate is finite-sample and does not require a particular recovery algorithm. Recovery and exact optimization are private implementation choices. This paper reports the public testing guarantee and the marginal-load stability condition. Source-specific recovery limits are outside the main claim.

The compiler works through geometric separation. Under the null, distance from the public baseline incurs regret. Under a structured alternative, staying near it forfeits the hidden-source advantage. These opposing envelopes create a fixed public threshold, and marginal-load stability bounds estimation loss for the specified private policies.

5 Experiments

5.1 When does one action become evidence?

The experiments ask when decision quality can support decision transparency. The main measures are detection power and Type-I error. Regret checks the premise of the public test. We apply one locked rule to private generators with different failure modes. Useful or strongly displaced actions remain invalid evidence when their null error is uncontrolled. We then test whether transparency appears in the certified regret region and evaluate robustness across action families. The road studies are semi-synthetic topology checks, not evaluations of independently operated services.

Protocol. The scaling experiment varies the normalized budget. It does not change the decision problem after outcomes are observed. We use $q_{\mathrm{comp}} = ma^2K^2/n$ for the spectral construction and $q_{\mathrm{info}} = ma^2K/\log n$ for exhaustive recovery. Alternative costs have mean $-\beta I - axy^\top$. We draw x and y independently, each with K positive and K negative entries. The null mean is $-\beta I$. We set $K = \mathrm{round}(n^\alpha/2)$, $a = (2K)^{-1/2}$, $\beta = a/32$, and report $\mathrm{Reg}_\mu(A)/(2aK)$ without clipping. The internal gate thresholds the leading singular value of $-\sqrt{m}(\bar{C} + \beta I)$ at $2.6\sqrt{n}$. The 6,720-task grid spans four values each of n and α. It uses eight values of q under both hypotheses. A disjoint 6,720-task grid repeats the experiment with seed 2027080302. Paired intervals resample tasks. Proportions use Wilson bounds. Appendix F gives the full grid and additional results.

Table 2. When does a decision become evidence? Every generator uses the same public rule. Regret values are normalized points on the 840-task ledger; intervals and hashes appear in Appendix H.1.

Action generator	Type-I	Power	Null regret	Alt. regret	Evidence reading
FanOpt-Spectral (gated)	0.0%	100.0%	0.00	2.09	Valid, powerful
Population oracle	0.0%	100.0%	0.00	0.00	Upper control
Sample ERM	100.0%	100.0%	17.13	56.73	Invalid null
Spectral (no gate)	100.0%	100.0%	5.02	2.09	Invalid null
Public baseline	0.0%	0.0%	0.00	94.95	No response

The same rule exposes two failures. Table 2 gives the gated pipeline and population oracle full power with no observed null errors. Ungated spectral and sample ERM reject every null. The fixed baseline never rejects. Movement alone is not evidence.

All 2,196 low-regret outputs stay on the certified side. Without the gate, 720 null actions activate. Rounding and oracle knockouts add 2.522 $[2.345, 2.696]$ and 93.99 $[93.54, 94.47]$ regret points. The equal-energy control never activates.

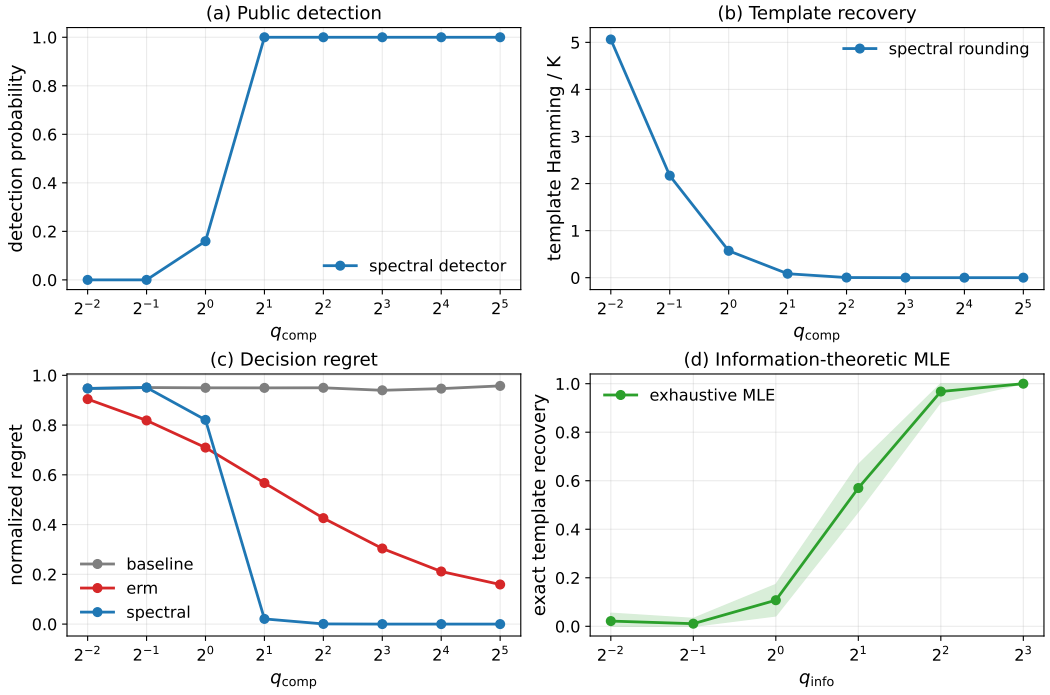

Figure 2. Regret and detection across budgets. In (a), gate activation matches the reconstructed public test. Separate replay certifies 420 nulls and 258/420 alternatives (Appendix H.1).

As regret falls, power rises from 16.0% at $q = 1$ to 100.0% at $q = 2$. None of 3,360 null trials activates (0.115% Wilson upper bound). Three noise laws retain 100% power and 2.70–3.01% regret. Road transfer gives AUC 0.970–1.000; public decoding takes 11.65 microseconds at $n = 10^5$ (Appendix F.4).

6 Conclusion

Decision quality asks whether an action works. Decision transparency asks what it reveals about the information that shaped it. Equal action laws give no power above the Type-I error. FanOpt shows when one action can provide evidence: a fixed public baseline and a geometric margin turn low regret into action separation. An outside evaluator can apply the test without seeing the private model or objective.

The experiments show this transition as regret falls. They also show why movement alone is weak evidence: ungated actions move under the null. Across the tested action families, the same public-test form applies to different private solvers. Decision transparency thus adds a measurable check to decision quality under structured alternatives that satisfy the public contract.

AI USE STATEMENT

We used generative AI tools for language editing, LaTeX and figure formatting, and the retrieval and discovery of related work.

REPRODUCIBILITY STATEMENT

Section 5 specifies the evaluation settings, baselines, inference units, and statistical procedures. The appendix gives the assumptions, complete proofs, implementation details, protocols, random seeds, and unit-level results. The anonymous supplementary package contains the FanOpt implementation, fixed configurations, public-data preparation instructions, archived result ledgers, and scripts that reproduce or verify the key experiments and certificate figures. Bundled checks replay or reconstruct the stored evidence. Fresh raw-data reruns require the documented public datasets and external solver dependencies, which are not redistributed when subject to third-party terms. Upon acceptance, we will release the submitted first-party code, configurations, and documentation publicly.

REFERENCES

Pranjal Awasthi, Moses Charikar, Kevin A. Lai, and Andrej Risteski. Label optimal regret bounds for online local learning. In *Proceedings of the 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pp. 150–166, 2015.

Vivek Bagaria, Jian Ding, David Tse, Yihong Wu, and Jiaming Xu. Hidden Hamiltonian cycle recovery via linear programming. *Operations Research*, 68(1):53–70, 2020.

Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, and Ellen Vitercik. Learning to branch. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 344–353, 2018.

Han Bao and Shinsaku Sakaue. Inverse optimization with prediction market: A characterization of scoring rules for eliciting system states. In *Proceedings of the 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 451–459, 2025.

Andreas Bärmann, Sebastian Pokutta, and Oskar Schneider. Emulating the expert: Inverse optimization through online learning. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 400–410, 2017.

Senne Berden, Ali Mahmutoğulları, Dimos Tsouros, and Tias Guns. Solver-free decision-focused learning for linear optimization problems. In *Advances in Neural Information Processing Systems*, volume 38, pp. 158127–158145, 2025.

Yudong Chen and Jiaming Xu. Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices. *The Journal of Machine Learning Research*, 17(27):1–57, 2016.

Priya L. Donti, Brandon Amos, and J. Zico Kolter. Task-based end-to-end model learning in stochastic optimization. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Adam N. Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1): 9–26, 2022.

Bertrand Even and Luca Ganassali. Statistical-computational gap in multiple Gaussian graph alignment. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 3133–3141, 2026.

Harold N. Gabow. Data structures for weighted matching and extensions to b-matching and f-factors. *ACM Transactions on Algorithms*, 14(3):39:1–39:80, 2018.

Aditya Gangrade, Aditya Gopalan, Venkatesh Saligrama, and Clayton Scott. Testing the feasibility of linear programs with bandit feedback. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 14513–14539, 2024.

Shengnan Guo, Youfang Lin, Ning Feng, Chao Song, and Huaiyu Wan. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 922–929, 2019.

Vishal Gupta and Michael Huang. Decision-focused learning with directional gradients. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Lingkai Kong, Wenhao Mu, Jiaming Cui, Yuchen Zhuang, B. Aditya Prakash, Bo Dai, and Chao Zhang. DF^2: Distribution-free decision-focused learning. In *Proceedings of the Forty-first Conference on Uncertainty in Artificial Intelligence*, volume 286 of *Proceedings of Machine Learning Research*, pp. 2269–2290, 2025.

Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *International Conference on Learning Representations*, 2018.

Jayanta Mandi, Víctor Bucarey, Maxime Mulamba Ke Tchomba, and Tias Guns. Decision-focused learning: Through the lens of learning to rank. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 14935–14947, 2022.

Jayanta Mandi, Marianne Defresne, Senne Berden, and Tias Guns. Feasibility-aware decision-focused learning for predicting parameters in the constraints. In *Advances in Neural Information Processing Systems*, volume 38, pp. 163285–163309, 2025.

Maxime Mulamba, Jayanta Mandi, Michelangelo Diligenti, Michele Lombardi, Victor Bucarey, and Tias Guns. Contrastive losses and solution caching for predict-and-optimize. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pp. 2833–2840, 2021.

Aran Nayebi. What capable agents must know: Selection theorems for robust decision-making under uncertainty. In *Proceedings of the 42nd Conference on Uncertainty in Artificial Intelligence*, volume 337 of *Proceedings of Machine Learning Research*, pp. 4773–4795, 2026.

Ke Ren, Peyman Mohajerin Esfahani, and Angelos Georghiou. Inverse optimization via learning feasible regions. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 51471–51488, 2025.

Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1:206–215, 2019.

Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.

Mattia Silvestri, Senne Berden, Gaetano Signorelli, Ali Irfan Mahmutogullari, Jayanta Mandi, Brandon Amos, Tias Guns, and Michele Lombardi. Score function gradient estimation to widen the applicability of decision-focused learning. *Journal of Artificial Intelligence Research*, 85(10): 1–34, 2026.

Bo Tang and Elias B. Khalil. CaVE: A cone-aligned approach for fast predict-then-optimize with binary linear programs. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pp. 193–210, 2024.

W. T. Tutte. A short proof of the factor theorem for finite graphs. *Canadian Journal of Mathematics*, 6:347–352, 1954.

Marin Vlastelica Pogančić, Anselm Paulus, Vit Musil, Georg Martius, and Michal Rolinek. Differentiation of blackbox combinatorial solvers. In *International Conference on Learning Representations*, 2020.

Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1658–1665, 2019.

David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pp. 296–303, 1996.

Shoki Yamao, Ken Kobayashi, Ryo Matsui, Shota Nagai, Naoki Nishimura, and Kazuhide Nakata. Robust decision-focused learning via worst-case regret minimization. In *Proceedings of the 42nd Conference on Uncertainty in Artificial Intelligence*, volume 337 of *Proceedings of Machine Learning Research*, pp. 7701–7736, 2026.

Jiaqi Yang, Enming Liang, Zicheng Su, Zhichao Zou, Peng Zhen, Jiecheng Guo, Wanjing Ma, and Kun An. DFF: Decision-focused fine-tuning for smarter predict-then-optimize with limited data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 26868–26876, 2025.

Zhilin Yang, William W. Cohen, and Ruslan Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 40–48, 2016.

Zichuan Yang and Yiming Xing. Active hypothesis testing for correlated combinatorial anomaly detection. *arXiv preprint arXiv:2601.17430*, 2026.

Wenbin Zhou, Agni Orfanoudaki, and Shixiang Zhu. Conformalized decision risk assessment. In *International Conference on Learning Representations*, 2026.

APPENDIX. PROOFS, CONSTRUCTIONS, AND EXPERIMENTAL DETAILS

A IDENTIFIABILITY LIMITS AND COUNTEREXAMPLES

The following counterexamples rule out uniform guarantees over classes that contain them. They motivate the contract used by FanOpt. They do not establish that its particular baseline, equal-cardinality condition, or certificate is necessary for every action-only test.

Proposition A.1 (Counterexamples to a uniform guarantee)**.** Consider procedures that observe only the returned action. If the policy class contains a source-blind rule that returns a fixed action $A^\star$ under both hypotheses, then without a low-regret restriction the induced action laws are identical and no level-α test has power above α. If a feasible action $A^\star$ satisfies

$$\mathrm{Reg}_0(A^\star) \le \rho \qquad \text{and} \qquad \sup_{g\in\mathcal{G}} \mathrm{Reg}_g(A^\star) \le \rho, \tag{19}$$

then the low-regret premise alone cannot separate the hypotheses. Finally, if the baseline is selected after the private action, choosing $E_0 = A$ makes $h(A, E_0) = 0$ for every output. A fixed-threshold distance test is then constant. The first two cases rule out a uniform action-only power guarantee above the Type-I level. The last case rules out such a guarantee for this distance test.

Proof. The first claim follows from the equality of the induced laws for a constant policy. For the second, the point mass at $A^\star$ satisfies both regret premises and has no action-law separation. For the last claim, post-hoc choice $E_0 = A$ makes every fixed-threshold distance test constant. Other statistics of A may still separate the hypotheses. For example, let the null always return $\{e_0\}$ and the alternative always return $\{e_1\}$, with $e_0 \neq e_1$. Choosing $E_0 = A$ erases the distance, yet the action-only test $\mathbf{1}\{e_1 \in A\}$ has zero Type-I error and power one. □

B PROOFS OF THE MAIN COMPILER RESULTS

Full proof of Theorem 4.2. Because every feasible action and E_0 have cardinality M, $|A \cap E_0| = M - h(A)$. Hence

$$U_0(E_0) - U_0(A) = \beta h(A),$$

which proves the null identity and uniqueness of E_0 when $\beta > 0$. For the alternative, the witness gives

$$\begin{aligned} U_g(Q_g) - U_g(A) &= \beta\big(h(A) - h(Q_g)\big) + a\big(S_g - s_g(A)\big) \\ &\ge \beta h(A) - \beta L + a\big(\Delta - \omega h(A)\big) \\ &= a\Delta - \beta L - (a\omega - \beta)h(A), \end{aligned}$$

where the last step uses $h(Q_g) \le L$, $S_g - s_g(E_0) \ge \Delta$, and $s_g(A) - s_g(E_0) \le \omega h(A)$. Since $\max_B U_g(B) \ge U_g(Q_g)$, this proves (12). If $\mathrm{Reg}_0(A) \le \rho$, then $\beta h(A) \le \rho < \beta t$, so $h(A) < t$. If $\mathrm{Reg}_g(A) \le \rho$ and $h(A) \le t$, the alternative lower bound is at least $a\Delta - \beta L - (a\omega - \beta)t > \rho$, a contradiction. Thus every low-regret null action lies on the non-rejection side and every low-regret alternative action lies on the rejection side. Taking probabilities under the induced action laws gives $\mathsf{Err}_{\mathrm{I}} \le \eta_0$. For each $g \in \mathcal{G}$, it gives $\mathsf{P}^A_{1,g}\{\phi_t(A) = 0\} \le \eta_1$. Integrating over Π and adding the mass $\Pi(\mathcal{G}^c)$ gives $\mathsf{Err}_{\mathrm{II}} \le \eta_1 + \eta_{\mathrm{geom}}$. The remaining bounds follow by addition. □

Full proof of Theorem 4.3. Let $D_Z(A) = \sum_{e\in A} \mathbf{1}\{\partial_s(e) \cap Z \neq \varnothing\}$. The agreement condition and the range bounds imply

$$|U_g(A) - U_{\widehat{g}}(A)| \le 2aGD_Z(A).$$

For every $q \in \mathcal{Q}_\Lambda$,

$$\mathbb{E}_{A\sim q} D_Z(A) \le \sum_{v\in Z} \lambda_v(q) \le \Lambda|Z|.$$

Let $q^\star_g$ maximize $\mathbb{E}_q U_g(A)$ over $\mathcal{Q}_\Lambda$. Since $\widehat{q}$ maximizes $\mathbb{E}_q U_{\widehat{g}}(A)$,

$$\begin{aligned} \max_A U_g(A) - \mathbb{E}_{\widehat{q}} U_g(A) &= \varepsilon_\Lambda(g) + \mathbb{E}_{q^\star_g} U_g(A) - \mathbb{E}_{\widehat{q}} U_g(A) \\ &= \varepsilon_\Lambda(g) + \mathbb{E}_{q^\star_g}(U_g - U_{\widehat{g}}) + \big(\mathbb{E}_{q^\star_g} U_{\widehat{g}} - \mathbb{E}_{\widehat{q}} U_{\widehat{g}}\big) + \mathbb{E}_{\widehat{q}}(U_{\widehat{g}} - U_g) \\ &\le \varepsilon_\Lambda(g) + 4aG\Lambda|Z|. \end{aligned}$$

The left side equals $\mathbb{E}_{\widehat{q}}\,\mathrm{Reg}_g(A)$. If the estimated optimizer has objective error at most $\tau \geq 0$ within $\mathcal{Q}_\Lambda$, the middle term is at most τ. It need not be zero. The factor-four bound then gains τ. Any rounding step must preserve the load constraint and account for its estimated-objective loss in this term. Under the common-support conditions, the public term and estimated score are constant on the support, while $\max_A U_g(A) = B_0 + aS_0$. Therefore

$$\mathbb{E}_{\widehat{q}}\,\mathrm{Reg}_g(A) = a\mathbb{E}_{\widehat{q}}\big[s_{\widehat{g}}(A) - s_g(A)\big] \leq 2aG\Lambda|Z|,$$

which proves both bounds. □

From expected regret to a regret tail. The expectation in the plug-in theorem is over the returned action, conditional on the data used to construct $\widehat{q}$. Suppose a data event $\mathcal{E}$ has probability at least $1 - \xi$ and, on that event, the theorem applies with $|Z| \leq z$ and $\varepsilon_\Lambda(g) \leq \varepsilon$. Conditional Markov's inequality and then total probability give, for $\rho > 0$,

$$\Pr\{\mathrm{Reg}_g(A) > \rho\} \leq \min\left\{1,\ \xi + \frac{\varepsilon + 4aG\Lambda z}{\rho}\right\}. \tag{20}$$

Under the factor-two conditions, replace the numerator by $2aG\Lambda z$. For a τ-approximate optimizer, add τ to the factor-four numerator, provided that this error bound holds on $\mathcal{E}$. A small expected regret alone does not imply that every sampled action lies inside the certified radius. To apply the compiler's mixture bound, these event and regret bounds must hold uniformly over its declared template class. This argument bounds an alternative regret tail. The null regret tail, or a null bound for a gate that returns E_0 when inactive, needs a separate proof.

C Public compiler definitions and test inputs

This appendix defines the notation and public inputs, records one sharpness example, and gives minimal instantiations for bounded-degree, 2-factor, and spanning-tree actions. The remaining sections report additional experiments.

Proof provenance. The FanOpt bridges are the regret-to-test compiler and the marginal-load stability result in the main text. Weighted f-factor optimization supplies the 2-factor oracle, and Wilson's algorithm samples the component trees. The proofs use finite counting, score perturbation, symmetry, linearity of expectation, and Markov's inequality. The family constructions verify the conditions needed to apply the two main results.

The objects below are shared by the compiler, the family constructions, and the experiments. A template-dependent certificate proves separation. Each deployed decoder receives only the returned action, the public baseline, and public size parameters. Table 3 identifies the quantities available to the test.

Table 3. The statistical test uses only public inputs. The remaining quantities appear only in the analysis.

Object	Meaning	Deployment status
$E, \mathcal{A}$	Ground set and feasible action family	public
$E_0, h(A)$	Public baseline and half symmetric-difference distance	public
g, a, G	Hidden template, signal scale, and entry bound	hidden/analysis
Δ, L, ω	Score margin, witness travel, and oscillation bound	analysis certificate
Z, Λ	Error entities and their load cap	analysis certificate
d_0, δ	Family-specific load cap and matching slack	construction metadata

Continued on the next page

Table 3. Continued from the previous page.

Object	Meaning	Deployment status
$K, n\ (N), m$	Per-sign count, side length, and observation budget	public instance metadata
$q_{\mathrm{comp}}, q_{\mathrm{info}}$	Polynomial and information normalized axes	experiment/analysis scales
$\phi_{\mathrm{match}}, \phi_{\mathrm{2fac}}, \phi_{\mathrm{tree}}$	Action-only tests	public

Table 4. Assumptions used by the compiler and by each action-family construction. A family row does not enlarge the general theorem's scope.

Layer	Required condition	Role
Compiler	Public baseline fixed before private observations	Public test contract
Compiler	Null and fixed-template regret tails	Type-I and Type-II bounds
Compiler	Uniform certificate margin and geometry event	Common threshold
Structured alternative	Templates in $\mathcal{G}$ are certificate-compatible, with $\Pi(\mathcal{G}^c) \le \eta_{\mathrm{geom}}$	Scope of source-influence guarantee
Plug-in	$q \in \mathcal{Q}_\Lambda$ with the stated load cap	Data-dependent estimator
2-factor	$K \ge 2$, $n - 2K \ge 3$ and bounded sensitive degree	Explicit witness construction
Tree	Balanced true and estimated signs, with the component UST policy in the public load class	Expected-load transport
Matching	A separately verified family certificate	Conditional extension

The family-specific baselines, certificate constants, and tests are collected once in Table 1. When a closed-form threshold is not available, an independent calibration split selects from a finite public menu. The selected baseline and threshold remain fixed during evaluation. Intervals use independent tasks for the synthetic grid and task blocks for native graphs.

D CERTIFICATE SHARPNESS

Sharpness example. Fix $a > 0$, $0 < \beta < a\omega$, integers $1 \le t \le L$, and scalars $0 < \omega \le 2G$ and $\omega t \le \Delta \le \omega L$. If $a\Delta - \beta L - (a\omega - \beta)t \ge 0$, then an equal-cardinality family with three actions and a template bounded by G satisfies both equalities

$$\mathrm{Reg}_0(A_t) = \beta t, \qquad \mathrm{Reg}_g(A_t) = a\Delta - \beta L - (a\omega - \beta)t. \tag{21}$$

Hence neither affine term in (16) can be improved uniformly from (Δ, L, ω) alone.

The construction is explicit. Let the public baseline be

$$E_0 = \{u_1, \dots, u_t\} \cup \{v_1, \dots, v_L\},$$

and introduce one replacement u_i' for each u_i and v_j' for each v_j. The other two feasible actions are

$$A_t = \{u_1', \dots, u_t'\} \cup \{v_1, \dots, v_L\}, \qquad Q = \{u_1, \dots, u_t\} \cup \{v_1', \dots, v_L'\}.$$

Thus all actions have cardinality $t + L$, while $h(A_t) = t$ and $h(Q) = L$. Assign

$$g(u_i) = -\omega/2, \quad g(u_i') = \omega/2, \qquad g(v_j) = -\Delta/(2L), \quad g(v_j') = \Delta/(2L).$$

The assumptions $\Delta/L \le \omega \le 2G$ keep every entry in $[-G, G]$. Writing $b_g = s_g(E_0)$ gives $s_g(A_t) - b_g = \omega t$ and $s_g(Q) - b_g = \Delta$. Since $\Delta \ge \omega t$, Q maximizes hidden score, and all three actions satisfy the oscillation inequality with equality for A_t. Therefore the geometry certificate is exactly (Δ, L, ω). Under the null, $\mathrm{Reg}_0(A_t) = \beta t$. Under the alternative,

$$U_g(Q) - U_g(A_t) = a\Delta - \beta L - (a\omega - \beta)t.$$

The assumed nonnegativity makes Q an alternative utility maximizer as well, so this difference is exactly $\mathrm{Reg}_g(A_t)$ and proves the claim.

E FAMILY-SPECIFIC CONSTRUCTIONS

The next two subsections apply the same public compiler to 2-factors and spanning trees. They supply family-specific baselines, witnesses, load bounds, and exact or randomized decoding rules. Matching uses the same interface only when its separate family certificate is verified, as summarized in Table 1.

E.1 BOUNDED-DEGREE TRANSPORT AND 2-FACTOR CONSTRUCTION

Bounded-degree load. Suppose $g, \widehat{g} \in [-1,1]^E$ differ only on edges incident to row and column error sets of sizes e_R and e_C. If every feasible action has maximum degree at most d, at most $d(e_R + e_C)$ selected edges touch an erroneous endpoint. Each affected score changes by at most two. Hence

$$\sup_{A \in \mathcal{A}} |s_g(A) - s_{\widehat{g}}(A)| \le 2d(e_R + e_C). \tag{22}$$

For $U_g(A) = B(A) + a s_g(A)$ with the same public B in both objectives, let $\widehat{A}$ maximize $U_{\widehat{g}}$ over the full family $\mathcal{A}$. The usual two-mismatch comparison gives

$$\mathrm{Reg}_g(\widehat{A}) \le 4ad(e_R + e_C). \tag{23}$$

All actions satisfy the load cap d. Thus the full probability simplex is a valid $\mathcal{Q}_d$ and $\varepsilon_d(g) = 0$. An optimizer with estimated objective error at most τ adds τ to the bound. A smaller feasible policy class also needs the approximation term from Theorem 4.3.

A fixed public baseline. Let $R = \{r_1, \dots, r_n\}$ and $C = \{c_1, \dots, c_n\}$ be disjoint vertex sets, and let $\mathcal{F}_{2n}$ contain all simple spanning 2-factors of the complete graph on $R \cup C$. Assume $K \ge 2$ and $n - 2K \ge 3$. Before selecting any hidden signs, fix the Hamiltonian cycles

$$H_R = (r_1, r_2, \dots, r_n, r_1), \qquad H_C = (c_1, c_2, \dots, c_n, c_1), \qquad T_0 = H_R \cup H_C.$$

For arbitrary $x, y \in \{0, +1, -1\}^n$ having K entries of each nonzero sign, set $g_{x,y}(r_i, c_j) = x_i y_j$ and set all same-side scores to zero. The supports can vary without changing T_0. Optimization over this family is the standard $f \equiv 2$ case of weighted f-factor optimization (Tutte, 1954; Gabow, 2018).

Lemma E.1 (Uniform 2-factor witness)**.** For every such x, y, the baseline T_0 has a score-maximizing witness $Q_{x,y}$ with

$$S_{x,y} = s_{x,y}(Q_{x,y}) = 4K, \qquad h(Q_{x,y}) \le 8K. \tag{24}$$

Thus $(\Delta, L, \omega) = (4K, 8K, 1)$ is a uniform certificate.

Proof. Delete the $2K$ active vertices from H_R. The remaining $n - 2K \ge 3$ neutral vertices form c_R nonempty paths, including possible singletons. Each path follows a distinct nonempty run of deleted vertices in the original cyclic order, so $1 \le c_R \le 2K$. Join the paths in that order, connecting the last vertex of each to the first vertex of the next. This adds exactly c_R nonbaseline edges and forms a simple neutral cycle. For $c_R = 1$ its path has at least three vertices. For $c_R = 2$ the two paths cannot both be singletons, so the two connecting edges are distinct. For $c_R \ge 3$, cyclic joining is simple even with singleton paths. Apply the same construction on C, adding $c_C \le 2K$ nonbaseline edges.

For each sign, the positive biclique $K_{K,K}$ contains an alternating $2K$-cycle, since $K \ge 2$. The two positive cycles and the two neutral cycles form a simple spanning 2-factor $Q_{x,y}$. Its $4K$ cross-side edges have score one, and every neutral edge has score zero. No 2-factor has higher score. Its total degree at the $4K$ active vertices is $8K$. Each positive edge uses two of these incidences. Thus it has at most $4K$ positive edges, while negative edges can only reduce the score.

All positive-cycle edges lie outside T_0. The only other new edges are the neutral connectors. Hence

$$h(Q_{x,y}) = 4K + c_R + c_C \le 8K.$$

Finally, every baseline edge has score zero and every other edge has score at most one. Hence $s_{x,y}(T_0) = 0$ and $s_{x,y}(T) - s_{x,y}(T_0) \le h(T)$ for all feasible T, proving the certificate. □

Public separation. The compiler gives

$$\mathrm{Reg}_{x,y}(T) \geq 4aK - 8\beta K - (a-\beta)h(T). \tag{25}$$

With $\beta = a/32$ and $\rho_0 = aK/32$, this implies

$$\mathrm{Reg}_0(T) \leq \rho_0 \;\Rightarrow\; h(T) \leq K, \qquad \mathrm{Reg}_{x,y}(T) \leq \rho_0 \;\Rightarrow\; h(T) \geq \frac{119}{31}K > 3K. \tag{26}$$

The public action-only rule is $h(T) > 3K$. For an exact estimated optimizer, (23) bounds regret by $8a(e_R + e_C)$. A sufficient condition to enter the radius $aK/32$ is thus $e_R + e_C \leq K/256$. This strict sufficient condition requires exact sign recovery when $K < 256$. A small nonzero error rate alone does not certify that radius. The proof uses the same baseline for every hidden support. Its complete-graph and same-side-cycle assumptions are distinct from the empirical tests on native road topologies.

E.2 Spanning-tree construction

Public geometry. Let $V = \{r\} \dot\cup R \dot\cup C$, with $|R| = |C| = N$, and let $\mathcal{T}$ contain all spanning trees of the complete graph on V. The public baseline E_0 is the star centered at r, and every tree has $M = 2N$ edges. For $2 \leq K \leq N/2$, define

$$\mathcal{X}_{N,K} = \{z \in \{-1, 0, +1\}^N : |\{i : z_i = +1\}| = |\{i : z_i = -1\}| = K\}.$$

For $x, y \in \mathcal{X}_{N,K}$, put $g(r_i, c_j) = x_i y_j$ and set all other edge scores to zero. The positive graph consists of two copies of $K_{K,K}$. Any tree has at most $2K - 1$ positive edges in each copy. Choose a tree in each copy, attach each component to the root with one edge, and connect each inactive vertex directly to the root. This spanning tree has both hidden score and baseline distance $S = 4K - 2$. The score bound is attained, and $s_g(T) \leq h(T)$ for every tree because all baseline edges have score zero. Thus

$$(\Delta, L, \omega) = (S, S, 1), \qquad \mathrm{Reg}_g(T) \geq (a-\beta)(S - h(T)). \tag{27}$$

For $t = 2K$, choose $\beta = a/32$ and $\rho = aK/32$. Then $\rho < \beta t = aK/16$ and $\rho < (a-\beta)(S - t) = 31a(K-1)/16$ for $K \geq 2$. Both strict compiler conditions hold, giving the public rule $h(T, E_0) > 2K$.

A public sensitivity map. Use $\partial_s(\{r_i, c_j\}) = \{r_i, c_j\}$ on every edge in $R \times C$, including edges with zero true score. Set $\partial_s(e) = \varnothing$ only for root and same-side edges, whose scores stay zero for every template. For estimates $\widehat{x}, \widehat{y} \in \mathcal{X}_{N,K}$, let Z contain every row or column vertex with an incorrectly estimated sign, including support errors. Then g and $\widehat{g} = \widehat{x}\widehat{y}^\top$ agree whenever $\partial_s(e) \cap Z = \varnothing$. This map is fixed publicly and does not depend on either hidden or estimated support.

Lemma E.2 (Balanced tree plug-in policy)**.** Let $0 < \beta < a$ and $\Lambda = 2 - 1/K$. Define $\mathcal{Q}_\Lambda$ to be all distributions on $\mathcal{T}$ whose expected sensitive load at every vertex is at most Λ. For balanced estimates $\widehat{x}, \widehat{y} \in \mathcal{X}_{N,K}$, sample independently a uniform spanning tree in each of the two estimated positive bicliques. Attach each component to the root once and connect each estimated-inactive vertex directly to the root. The resulting policy $\widehat{q}$ is an estimated-utility maximizer in the public class $\mathcal{Q}_\Lambda$. Also, $\varepsilon_\Lambda(g) = 0$ and

$$\mathbb{E}_{\widehat{q}}\,\mathrm{Reg}_g(T) \leq 2a(2 - 1/K)|Z|. \tag{28}$$

Proof. A uniform spanning tree of $K_{K,K}$ has $2K - 1$ edges. Vertex symmetry gives expected degree $2(2K-1)/(2K) = \Lambda$ at each active vertex. Root attachments have empty sensitivity. All other vertices have sensitive load zero, so the constructed policy belongs to $\mathcal{Q}_\Lambda$. The class is nonempty and compact. It is the intersection of the finite probability simplex with closed linear load constraints.

For any balanced template, $s_g(T) \leq \min\{h(T), S\}$. It follows that

$$U_g(T) = \beta(M - h(T)) + a s_g(T) \leq \beta M + (a-\beta)S.$$

For $h(T) \leq S$, use $s_g(T) \leq h(T)$ and $a > \beta$. For $h(T) \geq S$, use $s_g(T) \leq S$ and $\beta > 0$. Every component-tree construction above reaches this upper bound for its own template because

$h(T) = s_g(T) = S$. The estimated policy thus maximizes estimated utility even over all trees. The analogous true-template policy belongs to the same public $\mathcal{Q}_\Lambda$ and reaches the true global optimum, so $\varepsilon_\Lambda(g) = 0$.

Every tree in the support of $\widehat{q}$ has $B(T) = B_0 = \beta(M - S)$ and $s_{\widehat{g}}(T) = S$. Also $\max_T U_g(T) = B_0 + aS$. All three factor-two conditions of Theorem 4.3 hold with $G = 1$, which proves the claim. □

Wilson's algorithm samples each component tree exactly with expected polynomial runtime on these unweighted complete bipartite graphs (Wilson, 1996). The load calculation and regret proof need only symmetry and linearity of expectation.

Regret tails and the role of the gate. For balanced estimates under a fixed alternative template g, suppose that $|Z| \le z$ with probability at least $1 - \xi$. The component policy and conditional Markov's inequality imply, at $\rho = aK/32$,

$$\mathsf{P}_{1,g}\{\mathrm{Reg}_g(T) > \rho\} \le \min\{1,\ \xi + 64(2 - 1/K)z/K\}. \tag{29}$$

This bound is informative only when the second term is below one. For small K, a nonzero error bound may give no useful tail guarantee. An exact-recovery event ($z = 0$) gives the bound ξ.

The component policy alone does not control Type-I error. Every sampled tree has $h(T, E_0) = S = 4K - 2 > 2K$, even under the null. Let a private gate $J \in \{0, 1\}$ return $A = E_0$ when $J = 0$ and use the balanced component policy when $J = 1$. Then, for every realization,

$$\phi_{2K}(A) = J, \qquad \mathsf{Err}_{\mathrm{I}} = \mathsf{P}_0(J = 1), \qquad \mathsf{Err}_{\mathrm{II}} = \mathsf{P}_1(J = 0). \tag{30}$$

Thus the gate needs its own error bounds. The load lemma controls the utility of the activated policy. Applying (29) to the gated action must also account for missed gates under the alternative. These identities concern this balanced construction on the complete graph. The native-road exact-MST experiments test the same public distance interface. They do not establish the conditions of this component-policy construction.

F Empirical validation of decision transparency

This block asks whether the public transparency rule is robust and transferable. Non-Gaussian noise tests stability, road networks test two action families, and generic action-only scores test whether movement alone is valid evidence. The last unit records the cost of the public interface.

F.1 Non-Gaussian robustness with the decoder held fixed

The confirmatory grid changes only the entry-noise distribution. It uses a disjoint seed namespace and preserves the original $2.6\sqrt{n}$ spectral gate, action decoder, dimensions, sparsity exponents, and $q_{\mathrm{comp}} = 2$ budget. Each row aggregates 270 null and 270 alternative tasks. Proportions use Wilson intervals. Normalized-regret intervals use 10,000 task-level bootstrap draws.

Table 5. The public action test keeps 100% power and zero observed null activation under three variance-matched non-Gaussian noise laws. The Wilson interval is $[98.60, 100.00]$ for every 100% power entry and $[0.00, 1.40]$ for every zero-activation entry.

Noise	Null activation	Power	Normalized regret
Gaussian reference	0.00%	100.00%	2.88% [2.45, 3.32]
Rademacher	0.00%	100.00%	2.70% [2.30, 3.10]
Student-t_5	0.00%	100.00%	3.01% [2.55, 3.48]
5% contaminated Gaussian	0.00%	100.00%	2.99% [2.58, 3.42]

These results test noise robustness without retuning. No threshold or noise parameter was chosen from these outcomes. At $q_{\mathrm{comp}} = 1$, power remains between 15.6% and 18.5% across the four laws. The non-Gaussian replication therefore preserves both sides of the finite-size transition.

F.2 Action-family transfer across road networks

The traffic studies use chronological splits and public graph structures to test tree and 2-factor transfer. They preserve the information boundary but do not claim independently operated hidden-event labels. Road-study intervals are descriptive. Calendar-day bootstrap intervals account for within-day dependence. Wilson and timestamp-bootstrap intervals do not account for temporal dependence, so they do not establish 95% coverage under that dependence.

F.2.1 Spanning-tree construction and frozen transfer

METR-LA defines the spanning-tree rule and public statistic. PeMS-BAY then applies that fixed rule to a different network and time period without model reselection.

METR-LA construction. METR-LA records five-minute speeds on a 206-sensor road network (Li et al., 2018). The public baseline T_0 is the minimum spanning tree of the published graph. An exact oracle returns T_s after private edge costs are adjusted by time-of-day congestion. The public statistic is $h(T_s, T_0)$. Development fixes the congestion rule, and a chronological split calibrates $t = 38$ before evaluating 242 timestamps from 12 days (196 null and 46 alternative). The same traffic profiles define events and private costs. This study tests protocol transfer. It does not use independent event labels.

Endpoint	Estimate	95% interval
Action-only AUC	1.000	$[1.000, 1.000]$
Power	100.00%	$[92.29, 100.00]$
Type-I error	2.55%	$[1.09, 5.83]$
Tree-distance gap	49.88 edges	$[41.84, 55.40]$
Private-cost reduction	10.75%	$[7.01, 15.89]$

Table 6. A threshold fixed on the preceding calibration period separates natural congestion events from one returned METR-LA spanning tree. The cost reduction compares the returned tree with the public baseline tree.

Intervals use calendar-day cluster bootstrap draws or Wilson bounds. The test observes only T_s, T_0, and t. Speeds, costs, and solver state remain private.

Frozen cross-network transfer to PeMS-BAY. PeMS-BAY provides an independent 325-sensor network (Li et al., 2018). We use its 319-sensor connected component. The METR-LA rule and coefficient are frozen, and the exact oracle returns one tree every 30 minutes. A 70/15/15 chronological split fits profiles, calibrates $t = 63$, and evaluates one final suffix.

Endpoint	Estimate	95% interval
Action-only AUC	1.000	$[1.000, 1.000]$
Power	100.00%	$[97.50, 100.00]$
Type-I error	3.92%	$[1.81, 8.29]$
Tree-distance gap	93.36 edges	$[88.18, 97.57]$
Private-cost reduction	19.29%	$[16.15, 22.68]$

Table 7. The METR-LA rule transfers without model reselection to the independent 319-sensor PeMS-BAY component. The final suffix contains 303 returned actions from 28 calendar days.

Intervals use calendar-day cluster bootstrap draws or Wilson bounds. The fixed rule reaches 100% power at 3.92% point Type-I error, a 93.36-edge distance gap, and 19.29% registered cost reduction. This tests transfer across topology and time without another round of model selection.

F.2.2 2-factor transfer on PEMS08

PEMS08 provides five-minute measurements and a public graph for 170 sensors (Guo et al., 2019). The public baseline F_0 is the minimum-cost spanning 2-factor of the symmetrized shortest-path

metric. An exact mixed-integer oracle returns F_s after private costs are adjusted from time-of-day profiles, while the test reads only $h(F_s, F_0)$. A 60/20/20 chronological split fits profiles, calibrates $t = 34$, and evaluates 72 actions (56 alternative and 16 null). The public data source is fresh, but labels and costs share the traffic profiles, so this remains a semi-synthetic transfer check.

Endpoint	Estimate	95% interval
Action-only AUC	0.970	$[0.947, 0.993]$
Power	80.36%	$[68.16, 88.66]$
Type-I error	0.00% (0/16)	$[0.00, 19.36]$
2-factor distance gap	14.66 edges	$[13.02, 17.40]$
Private-cost reduction	2.77%	$[2.48, 3.12]$

Table 8. Fresh public-source rerun on the 170-sensor PEMS08 graph. The exact spanning 2-factor separates the final traffic suffix from its public cycle-cover baseline. The threshold is fixed on the preceding calibration period.

Intervals use calendar-day cluster bootstrap draws or Wilson bounds. Every returned action has degree two, and the test reads only F_s, F_0, and t. This is the real-topology counterpart to the controlled 2-factor study.

F.3 Generic one-shot action-only baselines

Each PEMS08 evaluation receives one returned action. Generic scores use that action, the baseline, and 64 null calibration actions. Hamming fits a scalar null threshold. Energy and RBF-MMD compare the action with the null set. FanOpt retains $t = 34$, selected on the full calibration period of 64 null and 46 alternative actions. The methods therefore use different calibration information, but share the final held-out suffix. Deployment reads no private cost, regret, or solver state. A naive inverse residual requires the private objective.

Table 9. One-shot action-only baseline comparison on the fresh public-source PEMS08 suffix. The final set has 56 alternatives and 16 nulls. Brackets are task-bootstrap 95% intervals. Energy and RBF-MMD have higher observed power and null FPR of $1/16$. The 5% calibration target is nominal, not a finite-sample error guarantee.

Method	Threshold	AUC	Power	Null FPR
FanOpt distance	34.000	0.970 $[0.932, 0.996]$	80.36% $[69.35, 90.16]$	0/16
Calibrated Hamming distance	34.000	0.970 $[0.932, 0.996]$	80.36% $[69.35, 90.16]$	0/16
One-shot energy distance	71.079	0.994 $[0.978, 1.000]$	100.00% $[100, 100]$	1/16
One-shot RBF-MMD	0.723	0.994 $[0.978, 1.000]$	100.00% $[100, 100]$	1/16
Naive inverse residual	–	–	–	Requires private objective

Retrospective chronological diagnostic. Generic thresholds use the first 42 null actions. FanOpt retains $t = 34$ from the original full calibration period. The expanded evaluation combines the last 55 calibration actions with the 72 final actions, giving 38 null and 89 alternative actions. It is not held out from FanOpt's threshold selection. FanOpt distance reached AUC 0.950 $[0.912, 0.979]$, power 80.9% $[72.2, 88.7]$, and an observed null FPR of $2/38 = 5.26\%$. Calibrated Hamming reached the same AUC and 84.3% power, but its FPR was $4/38 = 10.5\%$. Energy and RBF-MMD reached 100% point power, while both had FPR $7/38 = 18.4\%$. These are descriptive results with unequal calibration information and partial calibration reuse. They do not establish a power advantage at fixed population Type-I error.

Table 10. Retrospective action-only diagnostic on PEMS08. Generic thresholds use 42 null actions. FanOpt retains its full-calibration threshold. The 127-action evaluation includes 55 original calibration actions. Brackets are descriptive bootstrap 95% intervals.

Method	AUC	Power	Null FPR	Rule
FanOpt distance	0.950 $[0.912, 0.979]$	80.9% $[72.2, 88.7]$	2/38 (5.26%)	retained t
Calibrated Hamming	0.950 $[0.912, 0.979]$	84.3% $[76.3, 91.4]$	4/38 (10.5%)	null scalar
One-shot energy distance	0.996 $[0.989, 1.000]$	100.0%	7/38 (18.4%)	null energy
One-shot RBF-MMD	0.996 $[0.988, 1.000]$	100.0%	7/38 (18.4%)	null kernel

F.4 Public-decoder cost

Public decoding is linear in the action representation. It does not invoke the private predictor or the optimizer. The released implementation has median latency $11.65\,\mu$s at $n = 10^5$.

G Decision-quality compatibility

This block asks whether decision transparency can be added without reducing decision quality. It first checks native citation utility, then compares matched decision-focused references on Cora, and finally transfers the same quality check to PubMed and two co-purchase graphs.

G.1 Native citation-utility confirmation

The Cora experiment uses tasks 0 through 21 for rule selection and tasks 22 through 26 for evaluation. Task-level resampling avoids treating censoring seeds as independent graph instances. At 20% and 40% visibility, feature-baseline-minus-selected normalized-regret gains are 23.60 points [20.06, 26.47] at 20% and 37.13 points [33.94, 39.82] at 40%. Direct-reveal-minus-selected gains are 1.06 points [0.65, 1.48] and 6.89 points [4.70, 9.23].

CiteSeer's deterministic construction yields 16 selection tasks and five evaluation tasks. At the 20% and 40% budgets, feature-baseline-minus-selected gains are 13.44 points [7.27, 19.62] and 25.16 points [16.32, 35.33], positive on every task. Direct-reveal-minus-selected is 2.32 points $[-0.17, 5.23]$ and 9.80 points [5.56, 14.94]. The adverse uncertainty at 20% is retained. Brackets are 95% task-cluster bootstrap intervals over within-graph tasks. Hierarchical intervals lead to the same gate decision. The public binary citation outcomes are native utility, while independent thinning creates controlled observation budgets. These studies test native decision value.

G.2 Matched decision-focused references on Cora

These experiments test whether decision transparency can be added without reducing decision quality. Here *FanOpt-Spectral* denotes the private experimental pipeline used for this compatibility check. *FanOpt* is the public compiler. We use representative decision-focused methods as compatibility references under matched observations, tuning budgets, and exact oracle access. They are not comparisons of transparency guarantees. The references are SPO+ (Elmachtoub & Grigas, 2022), PG Loss (Gupta & Huang, 2024), black-box differentiation (Vlastelica Pogančić et al., 2020), Listwise DFL (Mandi et al., 2022), NCE with solution caching (Mulamba et al., 2021), LAVA-2swap (Berden et al., 2025), zero-QP CaVE-Hybrid (Tang & Khalil, 2024), SFGE (Silvestri et al., 2026), DFF (Yang et al., 2025), and Robust DFL (Yamao et al., 2026). Every row receives the same Cora features, revealed edges, tuning budget, and exact assignment oracle. Published update rules are used when directly applicable. Assignment adaptations and the public zero-QP configuration are labeled in the table. Throughout this unit, a reported difference is comparator regret minus FanOpt-Spectral regret, so a positive value indicates higher comparator regret. Gen-DFL, Diffusion-DFL, R-DFL,

Online DFL, and CREDO require distributional targets, sequential or recursive feedback, or outcome samples absent from this static assignment task. No table entry is imported from an incompatible published benchmark.

Method	20%: regret / difference	40%: regret / difference
FanOpt-Spectral	37.37 / $+0.00$ $[+0.00, +0.00]$	23.76 / $+0.00$ $[+0.00, +0.00]$
PG Loss	37.62 / $+0.25$ $[-0.77, +1.20]$	23.96 / $+0.20$ $[-0.08, +0.48]$
Black-box	37.43 / $+0.06$ $[-0.63, +0.58]$	24.13 / $+0.37$ $[+0.18, +0.56]$
NCE cache	38.08 / $+0.71$ $[+0.03, +1.25]$	24.34 / $+0.58$ $[+0.28, +0.95]$
SFGE (adapt.)	37.58 / $+0.20$ $[-0.51, +0.78]$	26.56 / $+2.79$ $[+2.22, +3.28]$
DFF (adapt.)	38.09 / $+0.72$ $[+0.44, +1.00]$	30.74 / $+6.98$ $[+4.57, +9.48]$
LAVA-2swap	43.85 / $+6.48$ $[+4.68, +8.59]$	31.73 / $+7.97$ $[+5.72, +10.12]$
Robust DFL (adapt.)	49.33 / $+11.95$ $[+9.31, +15.01]$	37.58 / $+13.82$ $[+11.13, +16.50]$
CaVE-Hybrid (zero-QP)	64.12 / $+26.74$ $[+24.80, +28.47]$	65.51 / $+41.75$ $[+38.83, +44.75]$
SPO+	48.78 / $+11.41$ $[+9.05, +14.39]$	36.74 / $+12.98$ $[+10.37, +15.51]$
Listwise DFL	49.81 / $+12.43$ $[+9.78, +15.80]$	38.79 / $+15.03$ $[+12.25, +17.81]$

Table 11. Decision-quality compatibility under matched optimization settings. FanOpt-Spectral remains at the same regret level as PG Loss and SFGE at 20% and keeps competitive regret across the full comparison. Entries are normalized-regret points. Positive differences indicate higher regret than FanOpt-Spectral. Intervals resample the five task blocks after averaging observation and fit seeds within each block.

FanOpt-Spectral's differences from PG Loss cross zero at both budgets. Its 20% differences from black-box differentiation and 2026 SFGE also cross zero. At 40%, SFGE has 2.79 more regret points $[2.22, 3.28]$. The differences are 6.48–15.03 points for LAVA-2swap, SPO+, and Listwise DFL, 0.72–6.98 points for DFF, 11.95–13.82 points for Robust DFL, and 26.74–41.75 points for the zero-QP CaVE-Hybrid configuration. The smaller NCE gaps are 0.71 and 0.58 points. LAVA-2swap evaluates the complete 2-swap neighborhood of each assignment, not the full adjacency construction used in the original solver-free benchmarks.

G.3 Graph-transfer compatibility

The remaining datasets test whether the decision-quality conclusion persists under a frozen or task-matched protocol beyond Cora.

PubMed transfer. The Planetoid PubMed construction (Yang et al., 2016) has 30 disjoint 50×50 assignment tasks. Tasks 0–19 train, 20–24 select hyperparameters, and 25–29 form five evaluation blocks. Every method receives the same features, revealed citations, exact oracle, and refit budget. FanOpt transfers the two raw-edge rules selected on Cora without tuning. Intervals average seeds within blocks and use 10,000 task-block bootstrap resamples. The † rows add zero-QP CaVE-Hybrid, SFGE, DFF, and Robust DFL.

Table 12. Decision-quality compatibility on PubMed under matched optimization settings. FanOpt-Spectral remains at the same regret level as PG Loss, black-box differentiation, and NCE. The † rows extend the check to CaVE-Hybrid, SFGE, DFF, and Robust DFL. Positive differences indicate higher regret than FanOpt-Spectral. Intervals resample five held-out task blocks.

Method	20% regret, difference	40% regret, difference
FanOpt-Spectral	43.16, reference $[0.00, 0.00]$	26.25, reference $[0.00, 0.00]$
PG Loss	43.16, tie $[0.00, 0.00]$	26.26, $+0.005$ $[0.000, +0.010]$

Continued on the next page

Table 12 (continued)

Method	20% regret, difference	40% regret, difference
Black-box	43.16, tie $[0.00, 0.00]$	26.26, $+0.005$ $[0.000, +0.010]$
NCE cache	43.16, tie $[0.00, 0.00]$	26.27, $+0.02$ $[+0.00, +0.05]$
SFGE (adapt.)†	44.29, $+1.13$ $[+0.45, +1.80]$	26.27, $+0.02$ $[+0.00, +0.03]$
DFF (adapt.)†	43.43, $+0.27$ $[+0.07, +0.54]$	33.59, $+7.34$ $[+2.87, +11.80]$
LAVA-2swap	49.45, $+6.29$ $[+2.73, +9.94]$	33.91, $+7.65$ $[+3.01, +12.30]$
Robust DFL (adapt.)†	56.81, $+13.64$ $[+8.38, +18.85]$	42.35, $+16.10$ $[+7.85, +24.35]$
CaVE-Hybrid (zero-QP)†	71.83, $+28.67$ $[+24.96, +31.51]$	71.77, $+45.52$ $[+40.71, +49.83]$
SPO+	55.38, $+12.22$ $[+7.14, +17.42]$	39.74, $+13.49$ $[+6.43, +20.60]$
Listwise DFL	55.11, $+11.95$ $[+6.89, +16.52]$	42.92, $+16.67$ $[+8.19, +25.16]$

The closest task-matched methods tie FanOpt at 20% reveal and differ by at most 0.02 regret points at 40%. SFGE has 1.13 more regret points $[0.45, 1.80]$ at 20% and 0.02 more points $[0.00, 0.03]$ at 40%. The differences are 6.29 and 7.65 points for LAVA-2swap, 12.22 and 13.49 points for SPO+, and 11.95 and 16.67 points for Listwise DFL. The DFF differences are 0.27 and 7.34 points, and the Robust DFL differences are 13.64 and 16.10 points. This pre-outcome cross-dataset result supports the compatibility conclusion from Cora. FanOpt-Spectral preserves competitive decision quality. The public compiler adds the statistical test.

Amazon Photo. Amazon Photo is a public co-purchase graph with 7,650 nodes and 745 features (Shchur et al., 2018). We form 30 disjoint 50×50 tasks. We use 20 for training, five for selection, and five for evaluation. FanOpt-Spectral, PG Loss, black-box differentiation, and NCE receive matched observations, fitting budget, and oracle access. Intervals bootstrap the five held-out task blocks after averaging observation and fit seeds within task.

Method	20% regret, difference	40% regret, difference
FanOpt-Spectral	33.84, reference $[0.00, 0.00]$	19.04, reference $[0.00, 0.00]$
PG Loss	33.69, -0.15 $[-0.60, +0.24]$	19.16, $+0.12$ $[-0.41, +0.51]$
Black-box differentiation	34.71, $+0.87$ $[+0.13, +1.56]$	19.54, $+0.51$ $[+0.20, +0.80]$
NCE solution cache	33.83, -0.01 $[-0.56, +0.53]$	18.98, -0.06 $[-0.61, +0.41]$
Direct reveal	33.89, $+0.05$ $[-0.03, +0.15]$	20.25, $+1.21$ $[+0.21, +2.29]$
Feature-only	70.51, $+36.67$ $[+29.24, +43.36]$	70.51, $+51.47$ $[+41.34, +60.42]$

Table 13. Decision-quality compatibility on Amazon Photo. Brackets are 95% task-block bootstrap intervals over five held-out blocks. The differences from PG Loss and NCE include zero, while black-box differentiation has higher regret at both budgets.

FanOpt-Spectral reduces regret relative to the feature-only rule by 52.0% at 20% reveal and 73.0% at 40% reveal. Its differences from PG Loss and NCE are at most 0.15 regret points and their intervals include zero. The two positive black-box differences exclude zero, while direct reveal has 1.21 more regret points at 40%. The added test remains compatible with the closest task-matched training objectives on a fresh graph source. This experiment measures the decision-quality part of the method on a natural graph.

Amazon Computers. The source-independent constructor produces 60 disjoint 50×50 tasks with 767-dimensional node features. We assigned 30 tasks to training, 10 to development, and 20 to evaluation. FanOpt transfers its raw-edge rules without reselection. PG Loss, black-box differentiation, and NCE use the same finite development menus, fitting budget, observations, and exact oracle. Intervals bootstrap the 20 held-out task blocks after averaging observation and fit seeds within task.

Table 14. Decision-quality compatibility on 20 held-out Amazon Computers blocks. Brackets are 95% task-block bootstrap intervals. Black-box differentiation has higher regret at both budgets, while PG Loss and NCE remain within the one-point non-inferiority margin.

Method	20% regret, difference	40% regret, difference
FanOpt-Spectral	44.32, reference $[0.00, 0.00]$	26.97, reference $[0.00, 0.00]$
PG Loss	44.85, $+0.52$ $[-0.09, +1.10]$	27.20, $+0.23$ $[-0.14, +0.60]$
Black-box differentiation	45.70, $+1.37$ $[+0.53, +2.23]$	27.67, $+0.70$ $[+0.14, +1.25]$
NCE solution cache	44.09, -0.23 $[-0.80, +0.36]$	26.74, -0.23 $[-0.55, +0.11]$
Direct reveal	44.39, $+0.06$ $[-0.01, +0.16]$	27.93, $+0.95$ $[+0.51, +1.45]$
Feature-only	78.42, $+34.10$ $[+30.87, +37.20]$	78.42, $+51.45$ $[+48.09, +54.68]$

Relative to feature-only, FanOpt-Spectral reduces regret by 43.5% at 20% reveal and 65.6% at 40%. Both black-box intervals favor FanOpt-Spectral. PG Loss and NCE remain within the one-point non-inferiority margin, and direct reveal has 0.95 more regret points at 40%. This second co-purchase graph confirms that the decision quality conclusion is not specific to Amazon Photo.

H CERTIFICATE REPLAY AND CROSS-FAMILY CONFIRMATION

This block checks the theorem's observable predictions. The locked replay tests the public threshold on every stored action, and the cross-family study tests the same compiler under 2-factor and spanning-tree geometries.

H.1 LOCKED PUBLIC TEST AND EMPIRICAL CERTIFICATE

We test whether the public rule depends on the choice of action generator. The experiment fixes the identity baseline and the rule $\mathbf{1}\{h(A, E_0) > K/2\}$ before running any task. Five action generators return actions on the same 840 assignment tasks. The population oracle is an upper control, not a deployable learner. The public replay parses each returned permutation and uses no cost, source label, regret, or detector state.

Table 15. The same public rule applied to five action generators. Brackets are 95% Wilson or task-bootstrap intervals. The final column uses the strict finite-grid cutoff $\mathrm{Reg} \le aK/256$. A dash means that the row produced no action below this cutoff.

Action generator	Type-I error	Power	Null regret	Alt. regret	Wrong / low-regret
Public baseline	0.0% [0.0,0.91]	0.0% [0.0,0.91]	0.00%	94.95% [94.08,95.82]	0/420
Population oracle	0.0% [0.0,0.91]	100% [99.09,100]	0.00%	0.00%	0/840
Sample ERM	100% [99.09,100]	100% [99.09,100]	17.13% [16.05,18.20]	56.73% [55.49,57.93]	– / 0
Spectral, gated	0.0% [0.0,0.91]	100% [99.09,100]	0.00%	2.09% [1.79,2.41]	0/678
Spectral, no gate	100% [99.09,100]	100% [99.09,100]	5.02% [4.95,5.09]	2.09% [1.78,2.40]	0/258

The action-level ledger contains 4,200 serialized permutations. Recomputing every hash, feasibility check, distance, threshold decision, and regret gives zero failures. The public rule places none of 2,196 low-regret actions on the wrong side of the threshold. In contrast, sample ERM rejects every null task, and the fixed baseline misses every alternative. These controls violate the low-regret premise in exactly the hypothesis where they fail.

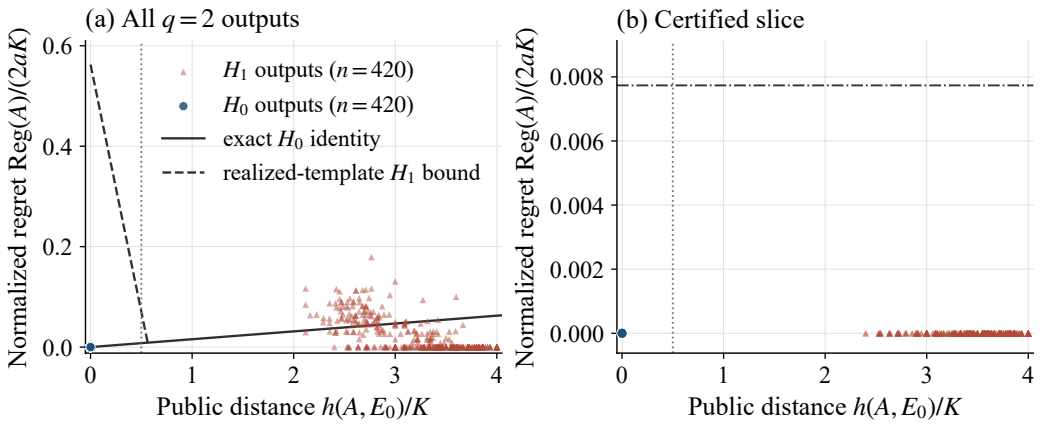

Figure 3. Regret–distance certificate on 840 original $q = 2$ tasks. Panel (a) shows the spectral outputs, the exact null identity, the uniform bound over the 420 realized alternative templates, and the fixed threshold $t/K = 0.5$. Panel (b) shows the strict certified slice. All 420 null and 258 alternative outputs in this slice lie on the predicted sides of the threshold (0/678 wrong side). The radius $0.99/128$ and the normalized Table 15 cutoff $1/512$ differ but select the same 678 outputs on this grid.

This replay also corrects an ambiguity in the earlier intervention receipt. Its stored `detected` field records the private spectral gate, not the public action test. Directly replaying the 720 frozen no-gate null actions shows 720 public threshold crossings, while the private gate fires 0 times. We keep these two events separate throughout the revised tables.

H.2 Cross-family controlled confirmation

The same synthetic protocol tests whether the assignment mechanism transfers to the other two theorem families. Every study uses the same $\tau = 2.6$, $a = (2K)^{-1/2}$, $\beta = a/32$, and $q_{\text{comp}} \in \{0.5, 1, 2, 4\}$, without family-specific tuning, and compares with the exact family oracle. The smallest 2-factor case is theorem-valid. Tree inference averages 32 UST draws within each observation before resampling the 128 independent observations. Normalized regret is $\text{Reg}/(4aK)$ for 2-factors and $\text{Reg}/[a(4K-2)]$ for trees.

Family and scale	Run	Trials	Power	Normalized regret	Null FPR
2-factor, $n = 7, K = 2$	Original	64	98.4%	3.48%	0%
	Replication	64	92.2%	10.21%	0%
2-factor, $n = 11, K = 3$	Scale check	512	98.6%	5.43%	0%
2-factor, $n = 15, K = 4$	Scale check	512	99.6%	4.30%	0%
Tree, $n = 32, K = 4$	Original	128	100%	0.091%	0%
	Replication	128	100%	0.305%	0%
Tree, $n = 64, K = 8$	Original	128	100%	0.419%	0%
	Replication	128	100%	0.677%	0%

Table 16. The fixed $q_{\text{comp}} = 2$ test maintains high power and zero observed null activation across 2-factor and tree geometries. Each run has its own row. Tree regret is averaged over UST draws within an observation.

Each trial or observation is the inference unit, with tree regret averaged over 32 UST draws. Observation-bootstrap 95% regret intervals are $[0.79, 7.25]\%$ / $[4.24, 17.26]\%$ for the original/replicated minimum 2-factor cell, $[4.34, 6.63]\%$ and $[3.56, 5.09]\%$ for the two scaled 2-factor cells, $[0, 0.27]\%$ / $[0, 0.71]\%$ for the smaller tree, and $[0.18, 0.70]\%$ / $[0.34, 1.05]\%$ for the larger tree. The matching Wilson power intervals are $[91.7, 99.7]\%$ / $[83.0, 96.6]\%$, $[97.2, 99.3]\%$, $[98.6, 99.9]\%$,

and $[97.1, 100]\%$ for each tree namespace. Null upper bounds are 5.67%, 0.745%, and 2.92% at 64, 512, and 128 units.

The gated rule has no observed null activation. The gate-free plug-in activates on every null observation. From $q_{\mathrm{comp}} = 1$ to 2, power rises sharply. The minimum replication reaches 92.2% power and 10.21% regret. The scale checks reach 98.6–99.6% power and 4.30–5.43% regret.

Runtime scaling. The median recorded task-suite time is 0.637 seconds at $n = 2048$, with a log-log slope of 1.95. This includes method execution, shared statistics, and the exact evaluation oracle.

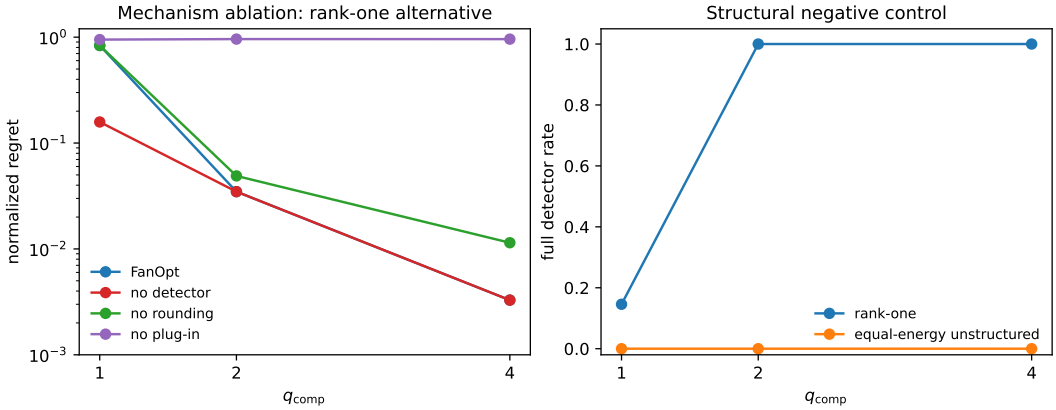

Figure 4. Component interventions identify the source of the transition. Removing the detector causes a non-baseline action on every null task, while removing balanced rounding increases paired alternative regret and removing the plug-in oracle nearly returns to the baseline. A same-energy unstructured sparse control has zero full-detector rate at $q_{\mathrm{comp}} \in \{1, 2, 4\}$, so the transition comes from the coherent rank-one source and not from entrywise energy.

The analysis combines template recovery, singular-vector perturbation, sign rounding, and standard weighted combinatorial oracles. FanOpt connects these private steps to a source-independent public baseline. The resulting distance test reads one low-regret action. Marginal load controls recovery error through deterministic degree for the 2-factor construction and expected degree for the uniform-spanning-tree decoder. Matching follows when its certificate is supplied.