PROMOTION IS RENT-OR-BUY: EXACT LEDGERS AND COMPETITIVE LIMITS FOR MEMORY CONSOLIDATION IN CONTINUAL AGENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

A continual learning agent must decide when to move an item from its memory into its model weights. Current systems decide by heuristic, with no guarantees. We cast the decision as a rent-or-buy problem: keeping an item in memory costs a small rent on every query, while promoting it into the weights costs a one-time price. The classical answer is the break-even rule (rent until the total rent equals the price, then buy). Promotion, however, has a side effect the classical problem lacks: crowding. Each promoted item makes the shared classifier slightly worse at separating nearby concepts. Our guarantees depend on how crowding is charged. If it is charged once, at the moment of promotion, the break-even rule stays within a constant factor of the optimal schedule. If it is charged on every query, as in deployed systems, no promote-only policy has a bounded competitive ratio (worst-case cost divided by the best possible cost), even with two items. Letting the agent demote items back out of the weights restores a bounded ratio. We test these predictions in six pre-registered experimental programs on synthetic data and features from real vision and text datasets, and report where the results went against our hypotheses.

1 INTRODUCTION

Every continual agent with a memory faces the same decision. Some item can stay in its memory, where it costs a little, or it can be moved into its weights (*consolidated*): paid for once. A recent position paper calls this move a basic building block of agent memory (Dorovatas et al., 2026). When to move the item is left to rules of thumb (move when the memory is full, move on a schedule, move when something happens) or to learned policies; deployed agent-memory systems decide the same way (Park et al., 2023; Chhikara et al., 2025). None of these rules comes with a bound on its competitive ratio, its worst-case total cost divided by the best possible schedule's cost. We give the move a price, and find that whether any such bound is possible depends on one bookkeeping choice (Figure 1).

Consolidation is a classic online problem in disguise: *ski rental*. A skier does not know in advance how many days they will ski, and each day chooses between renting skis and buying them once. Keeping an item in memory is renting; moving it into the weights (*promoting* it) is buying. The classical answer (Karlin et al., 1988; 1994) is the *break-even* rule: rent until the total rent equals the purchase price, then buy. Break-even never pays more than twice what a planner who knew the future would pay, no deterministic rule does better than break-even, and even the best randomized rule lowers the factor only to $e/(e-1) \approx 1.58$. Consolidation inherits these results, with one twist ski rental lacks. Promoted items share a single classifier head, and each new item *crowds* its neighbors: it makes the head slightly worse at telling them apart. Purchases are therefore no longer independent, since each one changes the cost of the others.

Crowding can be charged in two ways, and the choice decides what can be proven. Under *lump accounting*, crowding is a one-time charge added to the purchase price. Under *flow accounting*, it is paid as extra errors, a small amount on every query, as long as the crowded classes share the head. The two charges coincide in total cost under four conditions: promotions are never undone, demand

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.

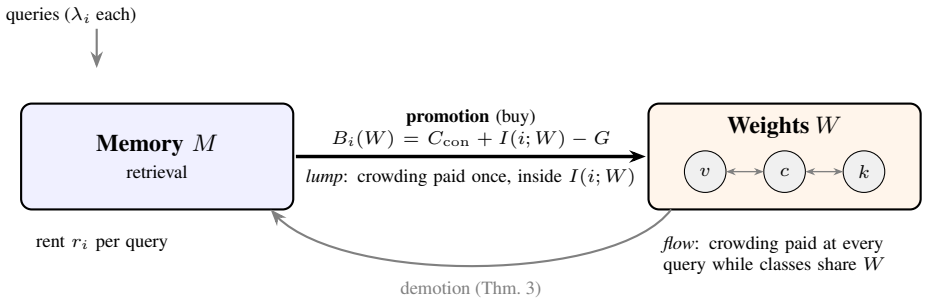

(a) The two-store model and its one priced decision.

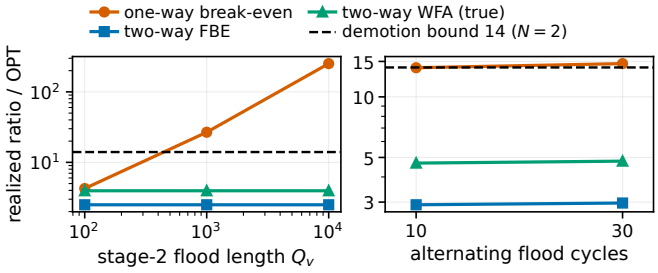

(b) The split as data (E-L1-8).

Figure 1: The whole paper in one figure. (a) An item pays rent on every query while it sits in memory. Promoting it into the weights costs once, $B_i(W) = C_{\text{con}} + I(i;W) - G$: the write cost, plus the harm $I(i;W)$ to neighbors already in the weights, minus a small credit G. So every purchase changes later prices. (b) The v, k example as data: cost divided by the best possible cost, as the flood grows (exact computation, no randomness). The rule that never demotes gets worse without limit (here up to 251.8 times the best), while the two rules that may demote stay flat. Theorem 2 proves the first behavior for every policy that never demotes; Theorem 3 bounds a classical policy from memory management (the work-function algorithm), while the plotted runs use the registered tie rule (Appendix E).

is stationary, future costs are discounted, and the crowding an item suffers on entry is billed to that item (Proposition 8).

Under lump accounting, the ski-rental answer survives: break-even stays within a factor $1 + \chi$ of the optimal schedule, where χ bounds how much crowding can change a price (Theorem 1). Under flow accounting, no policy that never demotes, deterministic or randomized, stays within any fixed factor of the optimal schedule (Theorem 2). The smallest counterexample breaks the stationarity condition. Two classes, k and v, share a head. First, queries ask about k until the policy promotes it; a policy that never promotes k already pays unboundedly more than one that does. Then the queries switch to v. From then on, every query about v pays either rent (if v stays in memory) or an extra crowding error (if v is promoted beside k). A planner who knew the future would rent k through the first phase, then promote v alone, paying a bounded amount after the switch. A policy that cannot undo a promotion keeps paying on every v query, and the adversary, which picks worst-case queries after seeing the policy's commitment, can extend those queries indefinitely, so the competitive ratio grows without bound. The fix is not a smarter promotion rule but the freedom to reverse one: moving an item back out of the weights (*demotion*) restores a bounded ratio (Theorem 3).

What is old and what is new? Without crowding, the problem splits into independent ski-rental instances, one per item. Lump accounting preserves this split, because each item's crowding is folded into its own price, so the $(1 + \chi)$ bound of Theorem 1 follows from the textbook argument applied item by item. Four results are new: the capitalization result (Proposition 8), which converts a flow of future costs into a single price; the impossibility under flow accounting (Theorem 2); the demotion guarantee for any fixed number of items N (Theorem 3); and exact price formulas in a Gaussian model, where each class's examples are drawn from a Gaussian around its class center. The impossibility is the most surprising result we prove. It holds with only two items and fixed

rent and purchase prices, and we know of no analogue in the classical ski-rental literature. The contributions are as follows:

1. *The split* (Section 4). Under lump accounting, break-even keeps a $1 + \chi$ bound, and no deterministic policy without demotion does better than break-even. Under flow accounting, no policy without demotion has a bounded ratio. Demotion restores a bounded ratio for any fixed number of items N, but simple greedy demotion rules still fail against two adversaries: a single long flood, a burst of queries for one class, and alternating floods, which switch between classes. We pre-registered both failures and observed them.

2. *The prices* (Sections 5 to 7). For frozen features under a Gaussian model, prices have closed forms, and five consequences follow. Promoting a class changes only the pairwise terms that involve it. A class's price can drop if a neighbor is promoted first. Without crowding and with steady demand, waiting never helps. The optimal memory size grows as the square root of demand, corrected by a neighbor tax, the accuracy a class's extra examples subtract from its neighbors. Estimating a class from few examples adds a price equal to a classical constant due to Okamoto. Theorem 15 gives the optimal order of promotions.

3. *The bets* (Section 8). We run six pre-registered experimental programs on synthetic data and thirteen real datasets. The split holds when prices are measured rather than assumed. The neighbor-tax threshold and the memory-sizing rule derived from it fail their tests. In all, six tests came out against our hypotheses, and we report every one.

2 Related work

Online algorithms. In standard online problems, costs are dictated by an adversary or a fixed list, whereas here they emerge dynamically from the learner itself, specifically from how previously stored classes are configured relative to one another. We know of no prior rent-or-buy analysis of model promotion. Ski rental and its extensions are the foundational ancestors without crowding, with fixed prices (Karlin et al., 1988; 1994; Lotker et al., 2012; 2008). File migration provides the closest two-store relative (Black & Sleator, 1989), while metrical task systems (a generalized framework for paying transition costs between states) provide the work-function algorithm underlying Theorem 3 (Borodin et al., 1992; Chrobak & Larmore, 1998). Other relevant relatives include leasing models with discounting (El-Yaniv et al., 1999), algorithms leveraging predictive advice (Purohit et al., 2018; Antoniadis et al., 2020), and offline formulations connecting to unconstrained submodular maximization (Feige et al., 2011; Buchbinder et al., 2015). Furthermore, online frameworks that revise past choices highlight a similar performance gap between irrevocable and revisable decisions (Imase & Waxman, 1991; Megow et al., 2016; Gu et al., 2016; Boyar et al., 2022), echoing dynamic right-sizing models that manage server activation costs subject to reversal penalties (Lin et al., 2013; Albers & Quedenfeld, 2022).

Memory systems. Existing memory-management frameworks rely entirely on heuristics to determine what to store, transfer, or discard. While some report empirical costs, none provide theoretical performance bounds relative to an optimal policy.

This separation between a temporary memory store and consolidated long-term weights is broadly motivated by complementary learning systems (McClelland et al., 1995), alongside standard machine learning paradigms such as regularization and rehearsal (Kirkpatrick et al., 2017; Rebuffi et al., 2017), retrieval-augmented prediction, episodic control, and exemplar buffers (Khandelwal et al., 2020; Borgeaud et al., 2022; Pritzel et al., 2017; Prabhu et al., 2020). Building on these foundations, a recent wave of research designs explicit memory systems for large language models, encompassing memory streams, operating-system-style paging, offline consolidation, self-reorganizing architectures, insight extraction, production memory layers, comprehensive surveys, and knowledge-editing techniques that write facts directly into weights (Park et al., 2023; Packer et al., 2023; Lin et al., 2025; Xu et al., 2025; Zhao et al., 2024; Chhikara et al., 2025; Hu et al., 2025; Wang et al., 2024).

Our work bridges these operational approaches with formal guarantees by analyzing two specific practices from this literature. First, we compare nearest-class-mean classifiers (Goswami et al., 2023; McDonnell et al., 2023; Hayes & Kanan, 2020) against trained heads, demonstrating via Proposition 10 that under infinite data, shared covariance, and balanced priors, the two coincide.

Second, iCaRL's herding heuristic (Welling, 2009) for exemplar selection is explicitly priced by Theorem 15.

3 THE MODEL: TWO STORES, TWO WAYS TO CHARGE

Time proceeds in rounds, where item i is queried with probability λ_i. The memory store M answers queries by looking up stored exemplars and serves every not-yet-promoted item. The weight store W consists of a trained model head — the classifier trained on top of the frozen features — that serves promoted items. Promoting item i incurs a one-time price, $B_i(W) = C_{\mathrm{con}}(i) + I(i; W) - G(i)$. In words: the write costs compute (C_{con}); it also makes life slightly harder for the classes already in the weights, the interference $I(i; W)$, and the newcomer is charged for that harm; and its small own-class accuracy gain is credited back ($G(i)$, vanishing under infinite data). While item i remains in M, every query to it incurs a maintenance rent ρ_i, covering lookup costs and retrieval-failure risk.

Because the promotion price depends on the contents of W, prices fluctuate. The *coupling ratio* $\chi = \max_i B_i^{\mathrm{pes}}/B_i^{\min}$ bounds how far any configuration of neighbors can move the price, from the lowest $B_i^{\min}$ to the highest B_i^{pes}; $\chi = 1$ denotes constant prices (Definition 1 and Assumption A1, Appendix B).

Rent represents the marginal extra cost of answering from memory rather than from weights; its lookup-success model and its exactness under infinite data (Proposition 10) are in Appendix B. Promotion involves two distinct choices: the *entry* decision of whether to introduce a class's first example (adding a new label to the head), and the *accumulation* decision of whether to consolidate subsequent exemplars for a class the head already recognizes.

The two accountings. Under *flow accounting*, every query to an item q currently in the head pays $\ell\,\Delta E_q(W)$: the continuous extra error q suffers due to neighboring classes sharing the head (zero when q is isolated, and monotonically non-decreasing as more neighbors arrive), converted to compute units at rate ℓ. In this regime, the buy price reduces to the fixed write cost C_i. Under *lump accounting*, queries are free post-purchase, and crowding is charged entirely upfront inside the dynamic promotion price $B_i(W)$. The two choices differ only in when the crowding is paid: once, at the buy, or a little on every query the crowded class ever serves.

When do the two accountings agree? The *capitalization lemma* (Proposition 8, Appendix B) answers: with one-way decisions (promotions never undone), stationary demand, and discounting ($\beta < 1$), the lump charged at the buy equals the expected discounted flow; demotion, shifting demand, or $\beta = 1$ breaks the equality; the theorems of Section 4 live in that gap. The v, k example of Section 1, worked through both accountings, is in Appendix B.

What the model covers. The derived price formulas hold exactly under infinite data for Gaussian classes sharing a covariance matrix over frozen features, with multi-class formulations expressed as sums over pairs. Estimation results require exemplars per class to substantially exceed the feature dimension ($m \gg d\max(1, 1/\gamma)$; Appendix A defines γ). While prices assume stationary demand, the core divergence between the two accountings persists under changing demand. Because large language model memory entries are not structured as labeled feature vectors, these parametric formulas do not directly transfer to LLM agent architectures without adaptation.

4 THE SPLIT: WHEN A GUARANTEE IS POSSIBLE

Can a non-demoting policy keep up with a clairvoyant scheduler (one that knows the future)? The answer depends entirely on how crowding is charged. In this section, an adversary selects the query sequence, and costs are plain totals without discounting ($\beta = 1$). Lump prices are fixed budgets established in advance (Assumption A1); they align with cumulative flow costs only when $\beta < 1$ under stationary demand, and even then only up to the entry term of Proposition 8. A policy is *one-way* if it never removes an item from W, and *two-way* if it permits removal. The agent inspects each query and may promote or demote before answering it, following the framework of metrical task systems (Theorem 2 holds under either order). Table 3 (Appendix E) contrasts the two models, both leveraging the Gaussian geometry of Section 5; at $\chi = 1$, the lump model reduces precisely to classical ski rental.

The break-even rule. The break-even rule rents item i until cumulative rent reaches the minimum possible price, $B_i^{\min} = \inf_W B_i(W)$ (the *uncrowded price*), and then buys at the prevailing market price, which may be up to a factor χ higher; Theorem 1 proves waiting for the lowest price is superior.

Theorem 1 (Ski rental survives lump pricing: $(1+\chi)$, optimal)**.** *In the lump model under Assumption A1, for every stream and price configuration,* $\mathrm{cost}_{\mathrm{BE}} \le (1+\chi)\cdot \mathrm{OPT} + \sum_i r_i$*, where* OPT *denotes the clairvoyant one-way optimum and the additive term is a single-query discretization slack. No deterministic one-way policy can achieve* $(1+\chi)-\varepsilon$*: an explicit family of sequences drives every such policy's competitive ratio to* $(1+\chi)-o(1)$*. At* $\chi = 1$*, this bound reduces to the classical factor of* 2 *over* OPT*. Furthermore, the pessimistic variant (triggering at* B_i^{pes}*) is* 2χ*-competitive (Appendix E).*

In essence: when crowding is charged as a lump sum, the classic break-even rule remains effective. It never exceeds $(1+\chi)$ times the optimal schedule's cost plus at most one query's rent per item, and no deterministic non-demoting rule can guarantee better. Without crowding ($\chi = 1$), this yields the familiar factor of 2 (proof sketch and full proof in Appendix E).

Theorem 2 (The flood: no bounded ratio under flow accounting)**.** *In the flow model, assume a single ordered pair* (v,k) *satisfies* $\varphi := \ell\,\Delta E_v(\{v,k\}) > 0$*. Then for any threshold* Γ*: (a) every deterministic one-way policy admits a finite stream with cost* $> \Gamma\cdot\mathrm{OPT}$*; (b) every randomized one-way policy admits a fixed oblivious finite stream with expected cost* $> \Gamma\cdot\mathrm{OPT}$*. This holds even for* $N = 2$ *with equal rents and constant buy prices; the parameters are realizable within the Gaussian model composed of the shared covariance and per-class means (Appendix E).*

In essence: this formalizes the v and k scenario from Section 1. When crowding is assessed per query, as deployed systems do, an adversary can force any non-demoting policy to perform arbitrarily poorly; two items suffice, and prices stay constant. The underlying cause: an irreversible promotion creates an unbounded debt paid through another item's future queries, and no upfront purchase price can cover a liability of unknown magnitude (proof sketch and full proof in Appendix E).

Theorem 3 (Demotion restores a bound at fixed N)**.** *Under flow accounting with demotion cost* $c_{\mathrm{dem},i} \ge 0$ *and buy prices* $C_i > 0$*, the work-function algorithm (classical tie-breaking rule in Appendix E) operates on the* 2^N*-state configuration space, equipped with the symmetrized movement metric* $\hat{d}(s,s') = \sum_{i\in s\triangle s'} \max(C_i, c_{\mathrm{dem},i})$*, where the symmetric difference counts the items present in one state but not the other. It achieves a realized cost of at most* $2(2^{N+1}-1)\max_i \max(1, c_{\mathrm{dem},i}/C_i)\cdot\mathrm{OPT} + c_0$ *relative to the two-way clairvoyant optimum, for some stream-independent constant* c_0*. Combined with Theorem 2, for any fixed* $N \ge 2$ *where* $\varphi > 0$*, one-way policies are unbounded whereas two-way policies remain* $O(1)$*-competitive.*

In essence: allowing the agent to demote items enables a classical online algorithm to maintain a bounded competitive ratio for any fixed number of items. Because the competitive factor scales exponentially with N, this result establishes the existence of a robust policy rather than a directly deployable algorithm. Crucially, it informs system designers that what saves memory promotion is the freedom to reverse decisions, rather than a refined accounting scheme (proof sketch and full proof in Appendix E).

A natural heuristic rescue is a meter, *flow break-even* (FBE): demote an item once its accumulated crowding reaches its price. It fails in the two ways the pre-registration predicted (Amendment A7, *overshoot* and *ping-pong*). On the registered adversarial streams of experiment E-L1-8, standard FBE incurs 60 to 180 times the optimal cost (the early-demotion variant up to 80), whereas the registered work-function execution stays at or below 1.48, well within the theoretical bound of 14 for $N = 2$ (full analysis in Appendix E). Whether a memory-efficient two-way policy with strong guarantees exists is open, as is whether deterministic policies can avoid the exponential dependence on N; randomization against oblivious streams lowers the ratio to polynomial in N (Remark 4). Figure 1b illustrates this empirical split.

Randomized break-even, drawing each item's trigger independently at the uncrowded price, guarantees $\mathbb{E}[\mathrm{cost}] \le (\chi + \frac{1}{e-1})\mathrm{OPT} + \sum_i r_i$ against predetermined query sequences (Theorem 17). Whether this factor is tight remains open, as do the hardness of flow-priced selection beyond the general result (Proposition 18) and the discounted-flow regime left open by the current split (Conjecture 19). The statements and open doors are collected in Appendix E.

5 PUTTING NUMBERS ON THE PRICES: A GAUSSIAN MODEL

To put numbers on the prices we need a model of the head. We take the simplest one in which every price has an exact formula. There are C classes in the head. Class r produces feature vectors from a Gaussian $\mathcal{N}(\mu_r, \Sigma)$, and all classes share the covariance $\Sigma \succ 0$. Class r is trained on $k_r > 0$ stored examples, so its training weight π_r is proportional to k_r, and the head $h_r(x) = w_r^\top x + b_r$ is a linear classifier trained with a ridge penalty that tends to zero. At test time the effect of the example counts is subtracted, so the head treats all classes as equally likely (*balanced priors*). In this model the prices can be written down: the error between any two classes, the bounds on crowding and the estimation constant all have formulas, and prices at finite counts follow by numerical integration (Lemma 5, Corollary 6, Lemma 7, Appendix B). Three facts from this model drive the rest of the paper.

First, promotion is *local*: adding one example of class c changes only the pairs that involve c. The harm it does to each neighbor has a formula, is never negative, and has explicit bounds (Proposition 9, Appendix B). When a *new* class joins the head, every class already there becomes harder to recognize: its best possible error rises, by a bounded amount. In the v, k example of Section 1, adding k beside v raises v's error against k from zero to $\Phi(-\sqrt{\gamma}/2)$ (Φ, the standard Gaussian cumulative distribution function), where the margin γ measures how far apart the two classes sit, in units of their spread (formalized in Appendix B).

Second, prices do not only rise: a class sitting in between can *shield* a neighbor from a newcomer's crowding, so the price of promoting c can *fall* when a neighbor arrives first (Remark 1, Appendix B), which is why Section 4 uses only the lowest and highest prices and never assumes that prices only rise.

Third, with infinite data and balanced priors, the prototype rule (assign each input to the nearest class mean) and the trained head are the *same classifier* (Proposition 10, Appendix B); what consolidation actually buys, and the benchmark puzzle this settles, are in Appendix B. With prices in hand, the first question a designer asks is when to promote.

6 WHEN TO PROMOTE, AND HOW MUCH TO KEEP

When should a new class be promoted, and how many examples of a class already in the head should the store keep? If demand were known, the answers would be two warehouse rules: promote above a threshold, and stock in proportion to the square root of demand. The formal statements (Lemma 11, Theorem 12, Corollary 13) are in Appendix C.

When items do not crowd each other and demand is steady, waiting never helps: the best time to promote is on arrival or never, so the only question is which (Lemma 11). The answer is a straight line: promote when the query rate is above a line that rises with the harm to neighbors, one line per discount factor (Theorem 12, Figure 2). Buffer sizes follow the warehouse manager's oldest rule: a warehouse stocks in proportion to the square root of demand, and a memory should keep examples in proportion to the square root of how often the class is queried (Corollary 13). Four times the traffic, twice the buffer.

But there is a side effect. When a class's estimated mean gets sharper, each neighbor loses a small boost it got from that class's noise, weighted by how often the neighbor is queried. With pretrained features, one more example therefore helps the class and costs each neighbor a little. The full law counts both; we call the neighbor side the *neighbor tax*:

Corollary 4 (The neighbor tax: completed stocking law and the 3Λ rule)**.** *Under the shared covariance and balanced class priors of Section 5, with both means estimated, the completed law keeps the square-root form with a corrected constant from the pairwise coefficients of Proposition 16 (the assembly is Appendix C), and promotes iff the demand-weighted marginal value of the next exemplar exceeds C_{con}. For a single dominant neighbor with $\varrho := 4(d-1)/\gamma_{cj} > 1$, marginal promotion is net negative for the demand-weighted objective iff*

$$\frac{\Lambda_j}{\Lambda_c} > \frac{3A_{cj} + B_{cj}}{A_{cj} - B_{cj}} = 3 + \frac{4}{\varrho - 1} \;\geq\; 3, \tag{1}$$

where Λ_c, Λ_j are the class's and the neighbor's total query demand, and A_{cj}, B_{cj} the two coefficients of Proposition 16. The threshold is 3.286 at $(\gamma, d) = (4, 16)$ and 3.0013 at the ViT-B/16

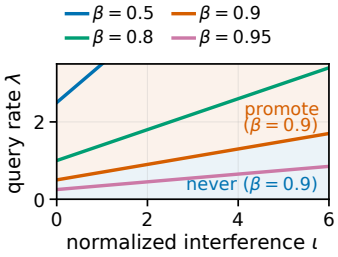

(a) The phase diagram of Theorem 12 (experiment E-L1-2): the promote-or-rent frontier is linear in the normalized interference for every discount factor; patient agents push it down and promote more.

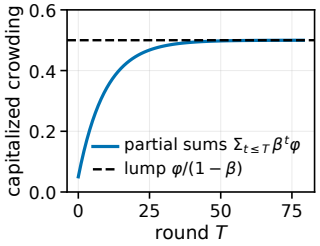

(b) Proposition 8: charging crowding once at promotion (lump) and a little on every query (flow) give the same total cost; the discounted partial sums of a per-round crowding flow converge geometrically to the lump.

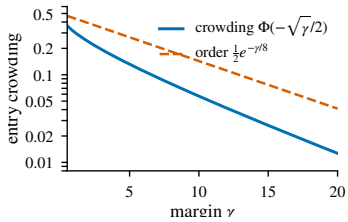

(c) Proposition 9(c): the entry-crowding envelope. Well-separated classes crowd each other exponentially weakly in the squared margin: the exact crowding (solid) stays under its order bound $\frac{1}{2}e^{-\gamma/8}$ (dashed).

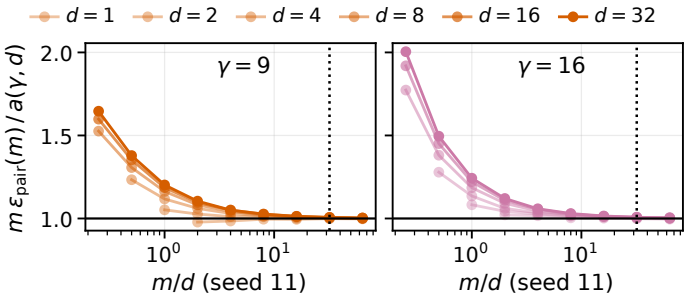

(d) E-L1-4: the measured error of the estimation constant, divided by its predicted value, approaches one in the region where the pre-registered examples-per-dimension criterion says it should.

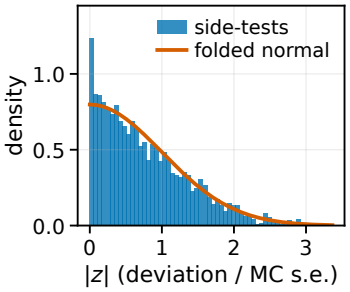

(e) E-L1-3: the Monte Carlo deviation distribution matches the folded-normal density that exactness predicts, wrong-side cells included.

Figure 2: The theory's checks at a glance (deterministic closed forms and canonical seeds 11, 13, 17): (a) the frontier family of Section 6; (b) the capitalization identity; (c) the entry-crowding envelope and its exponential order; (d) the validity map of the stocking constant; (e) the exactness check of the state-conditional formulas.

dimension $(1, 768)$. *The finite-m threshold can sit on either side of this asymptote (above it at* $(4, 16)$*, Appendix D; below it on the real-feature anchor, Appendix Q), so the field should use the finite-m marginal value.*

In words: with many examples, once one neighbor's demand passes three times the class's own, one more example of the class costs the store more than it earns. The boost it takes away from the neighbor outweighs the accuracy it adds for the class itself. The threshold is three plus a correction that shrinks as $4(d-1)/\gamma$ grows (3.0013 for ViT-B/16). With few examples the threshold can land on either side of three; on real features the measured flip (the demand ratio at which one more example stops helping the class) came nearer 1.5 (Section 8). So read the rule as too generous, and check the actual value of one more example before promoting into a crowded neighborhood (Figure 3).

The full entry and exemplar algorithms are Algorithms 3 and 4 (Appendix G). One question is left: how much does it cost to know a class only from examples?

7 The Cost of Not Knowing

A class known from m examples is known only roughly, and every pair of classes pays for that in extra errors. The extra error falls like $1/m$, and its constant is classical: for two classes with separation γ in dimension d, $a(\gamma, d) = \phi(\sqrt{\gamma}/2)\big((d-1)/(4\sqrt{\gamma}) + \sqrt{\gamma}/16\big)$ (Figure 3b). The sum

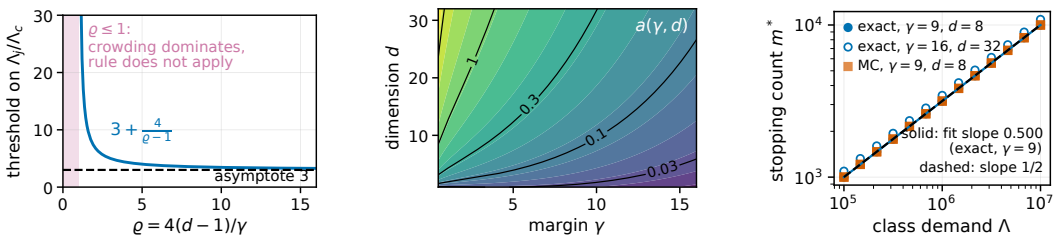

(a) The 3Λ rule of Corollary 4.

(b) The estimation constant $a(\gamma, d)$.

(c) E-L1-5: stopping counts ride the square root.

Figure 3: The stocking law at work: the rule, the constant, and the data (canonical seeds 11, 13, 17).

$a/m_c + a/m_j$ is the first-order extra error of the classifier built from estimated means when the covariance is known, and its split between the two sides is the known-covariance part of Okamoto's expansion (Okamoto, 1963; Kubokawa et al., 2013). What is new is using it as a price. Its two terms split the bill between errors in the direction of the estimated difference between the class means and errors in its length, and that split is what explains the rule of Section 6: through the direction term the noise in a class's own mean helps its neighbor by one quarter, the boost the neighbor tax takes away (Proposition 16, Appendix D).

Which example should be promoted next also has a formula: the value of promoting example z is, to first order, how far z lies from the current mean estimate $\hat{\mu}_c$ along the line between the two estimated class means, at right angles to the decision boundary, with a weight for each pair. The classical herding rule (Welling, 2009) has no such weights. The neighbor's side and the class's own side of the value, the Gaussian-tail error behind them, and how noise that pushes the estimated means apart *helps* the neighbor are worked out in Theorems 15 and 14 (Appendix D).

All of this needs many more examples than dimensions, $m \gg d \max(1, 1/\gamma)$, and instead of assuming that, the registered runs measure where it holds (Remark 3, Figure 2); the practical consequences, far-apart classes need few examples and high-dimensional heads need more examples per class, not fewer classes, are quantified in Appendix D. Everything so far is proved inside the model, the estimation results to first order in $1/m$. Section 8 asks what survives outside it.

8 Experiments: the bets and the verdicts

Before any official run, each program pre-registered the tests that would count as a loss for the theory; we call such a test a *falsifier*, and say it *fired* when the theory lost it. The theory made four bets: that its formulas are exact on data drawn from the model itself (Tier 1); that the $1/m$ rate survives when real feature spaces break the model's assumptions (Tier 1b); that the split and the prices hold on real features (Tiers 3 and 4); and that a wide real-feature study, run earlier on an older pipeline, agrees once re-analyzed the official way (E-R2 and its re-analysis E-R2C; Appendices S and T). The official runs used seeds 11, 13, 17 (trial runs: 0, 1, 2, Appendix L). Every plan change was logged before the runs it affects, with two exceptions, both stated where they occur (Appendices N and S). Every measured number below is written in by scripts from the logged results, every outcome is reported, and a falsifier that fired is reported as a loss.

Exactness. Eight experiments, E-L1-1 to E-L1-8, check the formulas where they claim to be exact (Table 11, Appendix L). A brute-force dynamic program agrees with the arrival-or-never rule and the threshold up to computer rounding, and finds the best-action boundary within 0.980 grid cells of the formula's line; simulated errors scatter around the formulas within sampling noise. The constant is recovered to 0.023% at every (γ, d), and the stocking exponents come out at 0.4997 and 0.4988, against the law's $1/2$ (Figure 3c). In the flood experiment (E-L1-8), the one-way rule's ratio grows with the flood to 251.8, while FBE and the registered work-function implementation stay flat (Figure 1b).

What survives broken assumptions. The second program breaks, one at a time, three assumptions that real feature spaces violate (Appendix O). When classes have different spreads, the $1/m$ rate survives, except in the one setting where the difference collapses the constant; there the rate falsifier

Table 1: The real-feature scoreboard: all seven Tier 3 falsifiers, on the registered *anchor* cell, our main setting (ViT-B/16 features of CIFAR-100), official seeds 11, 13, 17. As registered, F-1, F-3 and F-7 fire only if the whole interval falls outside the band; F-4 fires if its interval leaves the band or sign agreement drops below 0.70, and it fired on the second; F-5 fires unless the ratio exceeds 20; F-2 and F-6 check the selection and allocation code; $\hat{\rho}$ is the measured own-to-neighbor marginal ratio; ridge and CE are the refit head trained by ridge regression and cross-entropy. Six of seven did not fire; F-4 is reported as a loss.

Falsifier	Criterion (anchor cell)	Verdict	Margin
F-1	adequacy: origin slope, band $[0.5, 2]$	not fired	[0.904, 0.911]
F-2	selection machinery	not fired	random never dominates
F-3	stocking exponent, band $[0.40, 0.60]$	not fired	[0.308, 1.415] (underpowered)
F-4	externality: $\hat{\rho}$ band and sign test	**FIRED**	$\hat{\rho}$ [1.302, 2.135], sign 0.682
F-5	one-way ratio > 20 at the longest flood	not fired	ratio 107,496
F-6	allocation machinery	not fired	policies separate
F-7	refit-head rate, band $[0.7, 1.3]$	not fired	ridge [0.913, 0.957], CE [0.684, 0.716]

fired, and the rate comes back at larger m. With tails as heavy as t_3 (a Student-t distribution with 3 degrees of freedom), the rate survives and the constant falls to 0.46–0.85 of its Gaussian value, so the formulas over-charge, which is the safe side. And the program paid for itself: when features are stretched unevenly, both of its pre-declared checks turned against our analysis pipeline, exposing two defects that a logged change fixed before any official real-data run.

The split on real features. Tiers 3 and 4 ran on frozen features (Appendices Q and R): seven registered experiments on the main setting, ViT-B/16 (Dosovitskiy et al., 2021; Steiner et al., 2022) features of CIFAR-100 (Krizhevsky, 2009), plus descriptive runs on other models and datasets, and hard scenarios. Of the seven falsifiers, one fired (Table 1). On the main setting, the measured extra error is 0.907 ([0.904, 0.911]) times what the model predicts at the actual counts. The fit gives half weight to 87 of 59,400 cells. On the other vision settings the factor is 0.62 to 0.91, all below one (Tables 28 and 33).

Where the theory lost. F-4 fired on sign agreement (0.682, against a registered floor of 0.70). The harm to neighbors is real, seen on 94 percent of easily confused pairs, but the measured flip point, [1.302, 2.135], sits below the large-sample threshold of three; at $m = 256$ the reading is flat (Appendix R.3). The per-class budget allocator, greedy for demand-weighted value, lost to a uniform split, as registered as a possible loss. Its three repairs (a per-class floor, a positive own-side price, both together) did not save it (Appendices Q.6, R.1). The tenfold stocking rerun ended, by its pre-declared rule, without the power to detect a difference of the registered size (R.2).

What the real data taught us. Three lessons. First, on real vision features the formulas over-charge on average, although single-count settings can under-charge, and the two Banking77 text settings sit at the limit of measurement and do not test it. Second, the split is not an artifact of synthetic data: at measured prices the one-way rule's ratio reached 107,496 against a registered alarm level of 20; the gap survived 45 of 45 realistic query streams with Zipf-shaped demand and, under drift or bursts, 9 of 9 settings against the exact best schedule for $N = 10$ items. Third, the losses teach something: the neighbor tax bites earlier than the large-sample threshold, so the rule is too generous; and the allocator starves the classes whose value it cannot see.

Conclusion. Moving memory into weights has a price, and how that price is charged decides what can be guaranteed. Charged once, crowding leaves ski rental intact at $1 + \chi$; charged per query, as real systems pay, it defeats every policy that cannot undo a promotion, already with two items; the freedom to undo, not a cheaper charging, brings the bound back, though the cheap way to undo is not enough. For designers, five rules: charge crowding per query; keep promotions undoable; budget at the lowest price; stock by the square root of demand; and before promoting into a crowded neighborhood, check one more example's value: on real features the tax arrived nearer 1.5 than 3. Open: a two-way policy that remembers little yet keeps a guarantee (the supplement's attack suite is its test bed); and predictions; and the discounted case.

AI USE STATEMENT

The authors conceived, designed, and directed this research; every scientific decision, from the hypotheses to the pre-registered design, was theirs. A generative AI assistant, working under their direction, assisted with proofs, the pre-registered experiments and their verification, bibliography checking, and manuscript drafting. The verification pipeline released with the supplement independently re-derives every computational claim and regenerates every reported number from canonical artifacts. The authors reviewed and approved all AI-assisted content and take full responsibility. The submission form carries the same disclosure.

ETHICS STATEMENT

Theoretical work with synthetic data and, in the registered real-data programs, standard public datasets used only through frozen embeddings, whose contents we neither inspect nor release; no human subjects.

REPRODUCIBILITY STATEMENT

Every experimental program is pre-registered with frozen tolerances and a numbered amendment log (Appendix N). Canonical CSVs, manifests (code hashes, seeds, grids, tolerances, toolchain, timestamps), the runner libraries, the fill scripts, the figure and table generators, and the registration documents are released as supplementary material, and re-running the fill scripts and generators reproduces every number, table, and figure in this paper (Appendix V). Re-running the released runners with the recorded seeds reproduces every CSV on the pinned toolchain, with two exceptions stated in Appendix V: E-R2's first runner and toolchain were not archived (Appendix S), and most Tier 3 and 4 manifests record earlier versions of three feature and run scripts than the ones shipped.

REFERENCES

Susanne Albers and Jens Quedenfeld. Optimal algorithms for right-sizing data centers. *ACM Transactions on Parallel Computing*, 9(4):15:1–15:40, 2022. doi: 10.1145/3565513.

Antonios Antoniadis, Christian Coester, Marek Elias, Adam Polak, and Bertrand Simon. Online metric algorithms with untrusted predictions. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pp. 345–355. PMLR, 2020.

William J. Baumol. The transactions demand for cash: An inventory theoretic approach. *The Quarterly Journal of Economics*, 66(4):545–556, 1952.

David L. Black and Daniel D. Sleator. Competitive algorithms for replication and migration problems. Technical Report CMU-CS-89-201, Carnegie Mellon University, 1989.

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack W. Rae, Erich Elsen, and Laurent Sifre. Improving language models by retrieving from trillions of tokens. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, volume 162 of *Proceedings of Machine Learning Research*, pp. 2206–2240. PMLR, 2022.

Allan Borodin, Nathan Linial, and Michael E. Saks. An optimal on-line algorithm for metrical task system. *Journal of the ACM*, 39(4):745–763, 1992.

Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision (ECCV)*, pp. 446–461. Springer, 2014. doi: 10.1007/978-3-319-10599-4_29.

Joan Boyar, Lene M. Favrholdt, Michal Kotrbčík, and Kim S. Larsen. Relaxing the irrevocability requirement for online graph algorithms. *Algorithmica*, 84(7):1916–1951, 2022. doi: 10.1007/s00453-022-00944-w.

Sébastien Bubeck, Michael B. Cohen, James R. Lee, and Yin Tat Lee. Metrical task systems on trees via mirror descent and unfair gluing. *SIAM Journal on Computing*, 50(3): 909–923, 2021. doi: 10.1137/19M1237879. Preliminary version in Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2019, pages 89–97, doi:10.1137/1.9781611975482.6.

Niv Buchbinder, Moran Feldman, Joseph (Seffi) Naor, and Roy Schwartz. A tight linear time (1/2)-approximation for unconstrained submodular maximization. *SIAM Journal on Computing*, 44(5): 1384–1402, 2015. doi: 10.1137/130929205.

Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. Efficient intent detection with dual sentence encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pp. 38–45. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.nlp4convai-1.5.

Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. Mem0: Building production-ready AI agents with scalable long-term memory. In *ECAI 2025*, volume 413 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2025. doi: 10.3233/FAIA251160.

Marek Chrobak and Lawrence L. Larmore. Metrical task systems, the server problem and the work function algorithm. In Amos Fiat and Gerhard J. Woeginger (eds.), *Online Algorithms: The State of the Art*, volume 1442 of *Lecture Notes in Computer Science*, pp. 74–96. Springer, 1998.

Adam Coates, Honglak Lee, and Andrew Y. Ng. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 15 of *Proceedings of Machine Learning Research*, pp. 215–223. PMLR, 2011.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 248–255. IEEE, 2009. doi: 10.1109/CVPR.2009.5206848.

Vaggelis Dorovatas, Malte Schwerin, Andrew D. Bagdanov, Lucas Caccia, Antonio Carta, Laurent Charlin, Barbara Hammer, Tyler L. Hayes, Timm Hess, Christopher Kanan, Dhireesha Kudithipudi, Xialei Liu, Vincenzo Lomonaco, Jorge Mendez-Mendez, Darshan Patil, Ameya Prabhu, Elisa Ricci, Tinne Tuytelaars, Gido M. van de Ven, Liyuan Wang, Joost van de Weijer, Jonghyun Choi, Martin Mundt, and Rahaf Aljundi. Position: Modular memory is the key to continual learning agents. In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, Proceedings of Machine Learning Research. PMLR, 2026. Position paper. Preprint: arXiv:2603.01761.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021.

Ran El-Yaniv, Ron Kaniel, and Nathan Linial. Competitive optimal on-line leasing. *Algorithmica*, 25(1):116–140, 1999.

Uriel Feige, Vahab S. Mirrokni, and Jan Vondrák. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011. doi: 10.1137/090779346.

Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha Britan, Wouter Leeuwis, Gokhan Tur, and Prem Natarajan. MASSIVE: A 1M-example multilingual natural language understanding dataset with 51 typologically-diverse languages. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4277–4302. Association for Computational Linguistics, 2023. doi: 10.18653/v1/2023.acl-long.235.

Dipam Goswami, Yuyang Liu, Bartłomiej Twardowski, and Joost van de Weijer. FeCAM: Exploiting the heterogeneity of class distributions in exemplar-free continual learning. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pp. 6582–6595, 2023.

Albert Gu, Anupam Gupta, and Amit Kumar. The power of deferral: Maintaining a constant-competitive Steiner tree online. *SIAM Journal on Computing*, 45(1):1–28, 2016. doi: 10.1137/140955276.

Ford W. Harris. How many parts to make at once. *Factory, The Magazine of Management*, 10(2):135–136, 152, 1913. Reprinted in Operations Research, 38(6):947–950, 1990, doi:10.1287/opre.38.6.947.

Tyler L. Hayes and Christopher Kanan. Lifelong machine learning with deep streaming linear discriminant analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pp. 887–896, 2020.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778. IEEE, 2016. doi: 10.1109/CVPR.2016.90.

Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15979–15988. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01553.

Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. doi: 10.1109/JSTARS.2019.2918242.

Yuyang Hu, Shichun Liu, Yanwei Yue, Guibin Zhang, et al. Memory in the age of AI agents. *arXiv preprint arXiv:2512.13564*, 2025.

Makoto Imase and Bernard M. Waxman. Dynamic Steiner tree problem. *SIAM Journal on Discrete Mathematics*, 4(3):369–384, 1991. doi: 10.1137/0404033.

Anna R. Karlin, Mark S. Manasse, Larry Rudolph, and Daniel D. Sleator. Competitive snoopy caching. *Algorithmica*, 3(1):79–119, 1988.

Anna R. Karlin, Mark S. Manasse, Lyle A. McGeoch, and Susan Owicki. Competitive randomized algorithms for nonuniform problems. *Algorithmica*, 11(6):542–571, 1994.

Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Generalization through memorization: Nearest neighbor language models. In *International Conference on Learning Representations (ICLR)*, 2020.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

Tatsuya Kubokawa, Masashi Hyodo, and Muni S. Srivastava. Asymptotic expansion and estimation of EPMC for linear classification rules in high dimension. *Journal of Multivariate Analysis*, 115:496–515, 2013. doi: 10.1016/j.jmva.2012.11.001.

Ken Lang. NewsWeeder: Learning to filter netnews. In *Proceedings of the Twelfth International Conference on Machine Learning (ICML)*, pp. 331–339. Morgan Kaufmann, 1995. doi: 10.1016/B978-1-55860-377-6.50048-7.

Stefan Larson, Anish Mahendran, Joseph J. Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael A. Laurenzano, Lingjia Tang, and Jason Mars. An evaluation dataset for intent classification and out-of-scope prediction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 1311–1316. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1131.

Ya Le and Xuan Yang. Tiny ImageNet visual recognition challenge, 2015. CS231N course project report, Stanford University.

Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004. doi: 10.1016/S0047-259X(03)00096-4.

Kevin Lin, Charlie Snell, Yu Wang, Charles Packer, Sarah Wooders, Ion Stoica, and Joseph E. Gonzalez. Sleep-time compute: Beyond inference scaling at test-time. *arXiv preprint arXiv:2504.13171*, 2025.

Minghong Lin, Adam Wierman, Lachlan L. H. Andrew, and Eno Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Transactions on Networking*, 21(5):1378–1391, 2013. doi: 10.1109/TNET.2012.2226216.

Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9992–10002. IEEE, 2021. doi: 10.1109/ICCV48922.2021.00986.

Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A ConvNet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11966–11976. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01167.

Zvi Lotker, Boaz Patt-Shamir, and Dror Rawitz. Ski rental with two general options. *Information Processing Letters*, 108(6):365–368, 2008.

Zvi Lotker, Boaz Patt-Shamir, and Dror Rawitz. Rent, lease, or buy: Randomized algorithms for multislope ski rental. *SIAM Journal on Discrete Mathematics*, 26(2):718–736, 2012.

Hiroshi Maehara. Space graphs and sphericity. *Discrete Applied Mathematics*, 7(1):55–64, 1984.

James L. McClelland, Bruce L. McNaughton, and Randall C. O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457, 1995.

Mark D. McDonnell, Dong Gong, Amin Parvaneh, Ehsan Abbasnejad, and Anton van den Hengel. RanPAC: Random projections and pre-trained models for continual learning. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pp. 12022–12053, 2023.

Nicole Megow, Martin Skutella, José Verschae, and Andreas Wiese. The power of recourse for online MST and TSP. *SIAM Journal on Computing*, 45(3):859–880, 2016. doi: 10.1137/130917703.

Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Proceedings of the Sixth Indian Conference on Computer Vision, Graphics & Image Processing (ICVGIP)*, pp. 722–729. IEEE, 2008. doi: 10.1109/ICVGIP.2008.47.

Masashi Okamoto. An asymptotic expansion for the distribution of the linear discriminant function. *The Annals of Mathematical Statistics*, 34(4):1286–1301, 1963. doi: 10.1214/aoms/1177703864.

Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. URL `https://openreview.net/forum?id=a68SUt6zFt`.

Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. MemGPT: Towards LLMs as operating systems. *arXiv preprint arXiv:2310.08560*, 2023.

Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST)*, pp. 1–22, 2023.

Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3498–3505. IEEE, 2012. doi: 10.1109/CVPR.2012.6248092.

Ameya Prabhu, Philip H. S. Torr, and Puneet K. Dokania. GDumb: A simple approach that questions our progress in continual learning. In *European Conference on Computer Vision (ECCV)*, volume 12347 of *Lecture Notes in Computer Science*, pp. 524–540. Springer, 2020.

Alexander Pritzel, Benigno Uria, Sriram Srinivasan, Adrià Puigdomènech Badia, Oriol Vinyals, Demis Hassabis, Daan Wierstra, and Charles Blundell. Neural episodic control. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2827–2836. PMLR, 2017.

Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ML predictions. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pp. 9661–9670, 2018.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139, pp. 8748–8763. PMLR, 2021.

Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. iCaRL: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5533–5542, 2017. doi: 10.1109/CVPR.2017.587.

Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 3982–3992. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1410.

Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MPNet: Masked and permuted pre-training for language understanding. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, volume 33, pp. 16857–16867. Curran Associates, Inc., 2020.

Andreas Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your ViT? Data, augmentation, and regularization in vision transformers. *Transactions on Machine Learning Research*, 2022. URL https://openreview.net/forum?id=4nPswr1KcP.

James Tobin. The interest-elasticity of transactions demand for cash. *The Review of Economics and Statistics*, 38(3):241–247, 1956.

Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*, 2022. doi: 10.48550/arXiv.2212.03533.

Song Wang, Yaochen Zhu, Haochen Liu, Zaiyi Zheng, Chen Chen, and Jundong Li. Knowledge editing for large language models: A survey. *ACM Computing Surveys*, 57(3), 2024. doi: 10.1145/3698590.

Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, volume 33, pp. 5776–5788. Curran Associates, Inc., 2020.

Max Welling. Herding dynamical weights to learn. In *Proceedings of the 26th International Conference on Machine Learning (ICML)*, pp. 1121–1128. ACM, 2009.

Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muennighoff, Defu Lian, and Jian-Yun Nie. C-Pack: Packed resources for general Chinese embeddings. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pp. 641–649. ACM, 2024. doi: 10.1145/3626772.3657878.

Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. A-MEM: Agentic memory for LLM agents. In *Advances in Neural Information Processing Systems*, 2025. arXiv:2502.12110.

Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pp. 649–657, 2015.

Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. ExpeL: LLM agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 19632–19642, 2024.

A Notation and glossary

Table 2: Notation. Section 3 fixes the model objects, Section 4 the online objects, Section 7 the estimation objects.

Symbol	Meaning
λ_i	per-round query probability of item i
$\beta \in (0,1)$	discount factor
r_i	per-query rent of item i while served from memory
$B_i(W)$	buy price of item i given incumbent set W
C_{con}, C_i	write compute of one promotion (the constant buy price under flow accounting)
$I(i;W)$	interference increment on neighbors (lump accounting)
$G(i)$	spillover credit
χ	coupling ratio on buy prices, $\max_i \sup_W B_i(W)/\inf_W B_i(W)$ (Assumption A1)
φ, φ_{qk}	per-query crowding flow on a co-held pair (flow accounting)
$c_{\mathrm{dem},i}$	demotion cost of item i (two-way policies)
$\hat{d}$	symmetrized movement metric of Theorem 3
$\Delta\nu_i$	offset-mediated benefit of a promotion (Remark 2)
e_{hit}	per-query error of a successful retrieval, idealized to e_{W}
$d_i^{\max}$	per-round attributable crowding cap (Conjecture 19)
μ_r, Σ	class mean and shared covariance
π_r	training mixture weight of class r
k_r, m_r	buffer counts
$\tau_{rs} = \log(k_r/k_s)$	training log-odds offset
γ_{rs}	squared Mahalanobis margin of pair (r,s)
ϕ, Φ	standard normal density and distribution function
w	one-sided displacement of an estimated class mean
$v(z)$	mean displacement caused by promoting exemplar z
$\epsilon_c(m)$	excess per-query error of the head on class c at count m
$a(\gamma, d)$	expected-excess constant (Proposition 16)
A_{cj}, B_{cj}	the two channels of the constant (Corollary 4); $\varrho = A_{cj}/B_{cj}$
Λ_c	total class demand served by the head for class c
ℓ	error-to-compute currency multiplier
$p(m) = p_0 + \rho m$	retrieval failure probability at store size m
OPT	clairvoyant optimum (one-way in Theorems 1 and 2, two-way in Theorem 3)
BE, FBE, WFA	break-even at the uncrowded price; flow break-even (two-way, fixed attribution); work-function algorithm

Buy price $B_i(W)$**.** The one-time cost of promoting item i given the incumbent set W: write compute plus interference minus spillover.

Capitalization. Proposition 8's identity: the discounted sum of a stationary per-round flow φ is $\varphi/(1-\beta)$, so for $\beta < 1$ a lump capitalizes the crowding an entry imposes on incumbents; with each entrant's own crowding at its entry added, lumps equal the expected discounted flow for one-way stationary decisions.

Composed protocol. Evaluating the head with count offsets subtracted, so pairwise training biases τ_{rs} are corrected to the evaluation prior; the paper's default.

Coupling ratio χ. The worst ratio of a buy price across incumbent sets, on prices, not on the interference increment; $\chi = 1$ is the decoupled classical regime.

Entry crowding. The irreducible Bayes error increase an existing class suffers when a new class enters the label set; pairwise envelope $\Phi(-\sqrt{\gamma}/2)$, Proposition 9(c).

Exemplar tier. The economy of $k \to k+1$ promotions of an already-consolidated class; under composition a pure estimation public good, Corollary 13.

Flow accounting. Charging interference as the per-query serving error paid while classes coexist in the head; the operational reading, and the accounting under which one-way online promotion has no bounded competitive ratio (Theorem 2).

Lump accounting. Charging interference as a provisioning-time budget fixed in advance (for $\beta < 1$, the capitalized crowding the entry imposes on incumbents); the accounting under which break-even is $(1+\chi)$-competitive (Theorem 1).

Now-or-never. Lemma 11: in the decoupled regime, with stationary demand and discounting, the optimal promotion time is an endpoint, arrival or never.

Plug-in rule. The discriminant obtained by substituting empirical class means into the closed-form Bayes rule; the estimator all constants are stated for.

Population ledger. The bookkeeping in which population quantities (means, Bayes errors) are exact and estimation noise is absent; opposed to the estimation ledger of Section 7.

Public good. A promotion whose benefit accrues to every query of the class at total demand Λ_c, not to the promoted item alone.

Rent r_i. The differential per-query cost of serving item i from memory rather than from weights: retrieval compute plus failure-weighted coverage loss.

Shielding. Remark 1: an intermediate class absorbing an entrant's crowding, so that adding a neighbor lowers a buy price; the reason prices are non-monotone in the instantiation.

Spillover G. The within-class margin improvement an added exemplar would buy; identically zero at population level, an estimation-ledger object (Remark 2).

Stocking constant $a(\gamma, d)$. The leading coefficient of the expected excess error; the square-root law's engine.

Two-phase adversary. Theorem 2's stream: flood k to force a buy, then flood v to collect crowding forever.

Wrong side. Cells with $|w|^2 > \gamma$, where the estimated mean has wandered past the neighbor and the rule is worse than chance on its own class.

B PROOFS OF THE SUBSTRATE LEMMAS

Lemma 5 (Population optimum and gauge selection)**.** *The population CE risk is uniquely minimized in function space by the Bayes posterior, which the head realizes exactly on the family $w_r = \Sigma^{-1}\mu_r + a$, $b_r = -\frac{1}{2}\mu_r^\top \Sigma^{-1}\mu_r + \log \pi_r + c_0$, with (a, c_0) the gauge freedom; as $\lambda \to 0^+$ the ridge selects $a = -\frac{1}{C}\sum_r \Sigma^{-1}\mu_r$ (the bias gauge c_0 is softmax-null and stays free).*

Corollary 6 (Pairwise ledger)**.** *For pair (r, s) the decision statistic is Gaussian with mean $\gamma_{rs}/2 + \tau_{rs}$ and variance γ_{rs}, where $\gamma_{rs} = (\mu_r - \mu_s)^\top \Sigma^{-1}(\mu_r - \mu_s)$ and $\tau_{rs} = \log(k_r/k_s)$; the pairwise error is $\mathrm{err}_{r\to s} = \Phi\big(-\sqrt{\gamma_{rs}}/2 - \tau_{rs}/\sqrt{\gamma_{rs}}\big)$, and composition corrects every τ_{rs} to the evaluation prior at every count vector.*

Lemma 7 (Count continuity)**.** *Promoting one exemplar of class c ($m_c \to m_c + 1$) changes the pairwise training log-odds by $\Delta\tau_{cj} = \log \frac{m_c+1}{m_c}$ on pairs at c and by zero elsewhere, independent of which exemplar moved and of the feature geometry.*

This appendix is self-contained: it derives everything used about the Gaussian instantiation, in the order Lemma 5, Lemma 7, Proposition 9, composition neutrality, and store neutrality. Figures 4 to 2c visualize the exact objects; all are deterministic closed forms.

B.1 Lemma 5: population optimum and gauge selection

Proof. Write $q(y \mid x)$ for the head's predicted posterior. The population weighted cross-entropy risk is

$$R(q) = \mathbb{E}\big[-\log q(Y \mid X)\big],$$

an expectation over the true joint $p(x, y)$. For any competitor q,

$$R(q) - R(p) = \mathbb{E}_x\big[\mathrm{KL}\big(p(\cdot \mid x) \,\|\, q(\cdot \mid x)\big)\big] \geq 0$$

by Gibbs' inequality, with equality iff $q(\cdot \mid x) = p(\cdot \mid x)$ for p_X-almost every x. This is strict properness of log loss: the function-space minimizer is unique and equals the Bayes posterior.

We now exhibit the posterior in the head's function class. For the shared-covariance Gaussian mixture with mixture weights π_r proportional to buffer counts k_r,

$$\log p(Y = r \mid x) = \mu_r^\top \Sigma^{-1} x - \tfrac{1}{2}\mu_r^\top \Sigma^{-1}\mu_r + \log \pi_r - \log Z(x), \tag{2}$$

where $Z(x) = \sum_s \pi_s \exp(\mu_s^\top \Sigma^{-1} x - \frac{1}{2}\mu_s^\top \Sigma^{-1}\mu_s)$ is the partition function common to all classes. The unnormalized logit $\mu_r^\top \Sigma^{-1} x - \frac{1}{2}\mu_r^\top \Sigma^{-1}\mu_r + \log \pi_r$ is affine in the parameters $(w_r, b_r) = (\Sigma^{-1}\mu_r + a,\ -\frac{1}{2}\mu_r^\top \Sigma^{-1}\mu_r + \log \pi_r + c_0)$ for any common offset (a, c_0), because softmax is invariant to adding the same affine function of x (through $a^\top x + c_0$) to every logit. Hence the Bayes posterior lies exactly in the head family, and by the uniqueness above the minimizer set of the population risk within the family is precisely the gauge orbit

$$\mathcal{G} = \big\{(w_r, b_r)_r : w_r = \Sigma^{-1}\mu_r + a,\ b_r = -\tfrac{1}{2}\mu_r^\top \Sigma^{-1}\mu_r + \log \pi_r + c_0\big\}. \tag{3}$$

Because the function-space minimizer lies in the head class, the in-class argmin is the global argmin.

Ridge selection. On $\mathcal{G}$ the weight norm is

$$\|W\|^2 = \sum_r \|\Sigma^{-1}\mu_r + a\|^2, \tag{4}$$

a strictly convex quadratic in a alone (c_0 does not enter). Setting the gradient to zero:

$$\nabla_a \|W\|^2 = 2\sum_r (\Sigma^{-1}\mu_r + a) = 0 \quad \Longleftrightarrow \quad a = -\frac{1}{C}\sum_r \Sigma^{-1}\mu_r = -\Sigma^{-1}\bar{\mu}, \tag{5}$$

so the minimum-norm representative has centered directions $w_r = \Sigma^{-1}(\mu_r - \bar{\mu})$. Standard epi-convergence of the λ-regularized objectives to the constrained problem (the ridge objective increases in λ and any limit point of minimizers is feasible for $\lambda \to 0^+$ with minimal norm) selects exactly this representative; the bias gauge c_0 is unconstrained by a weight ridge and fixed by the sum-zero convention. □

Corollary 6: the pairwise ledger.

Proof. Gauge invariance first: pairwise logit differences cancel both gauge freedoms,

$$h_r(x) - h_s(x) = (\mu_r - \mu_s)^\top \Sigma^{-1} x - \tfrac{1}{2}(\mu_r^\top \Sigma^{-1}\mu_r - \mu_s^\top \Sigma^{-1}\mu_s) + \tau_{rs}, \qquad \tau_{rs} = \log(k_r/k_s). \tag{6}$$

For $x \sim \mathcal{N}(\mu_r, \Sigma)$ the statistic $h_r(x) - h_s(x)$ is Gaussian with mean

$$(\mu_r - \mu_s)^\top \Sigma^{-1}\mu_r - \tfrac{1}{2}\gamma_{rs} + \tau_{rs} = \tfrac{1}{2}\gamma_{rs} + \tau_{rs} \tag{7}$$

and variance $(\mu_r - \mu_s)^\top \Sigma^{-1}(\mu_r - \mu_s) = \gamma_{rs}$. The error of classifying x from class r into s is the upper tail at zero:

$$\mathrm{err}_{r\to s}(\tau) = \Phi\left(-\frac{\frac{1}{2}\gamma_{rs} + \tau_{rs}}{\sqrt{\gamma_{rs}}}\right) = \Phi\left(-\frac{\sqrt{\gamma_{rs}}}{2} - \frac{\tau_{rs}}{\sqrt{\gamma_{rs}}}\right). \tag{8}$$

Under the composed protocol the evaluator subtracts $\log k_r$ from logit r; every pairwise offset τ_{rs} is thereby corrected to the evaluation prior, at every count vector, which is Remark 2's cancellation in the population ledger. □

B.2 LEMMA 7: COUNT CONTINUITY

Proof. Lemma 5 uses the counts only through the weights $\pi_r \propto k_r$; positivity of all k_r is the only requirement (a zero-mass class has no posterior). The closed forms therefore hold at every strictly positive count vector. The training log-odds is $\tau_{rs} = \log(k_r/k_s)$, a function of counts alone. Under $k_c \mapsto k_c + 1$ with all other counts fixed:

$$\Delta\tau_{cj} = \log\frac{k_c+1}{k_j} - \log\frac{k_c}{k_j} = \log\frac{k_c+1}{k_c} \quad (j \neq c), \qquad \Delta\tau_{rs} = \log\frac{k_r}{k_s} - \log\frac{k_r}{k_s} = 0 \quad (r, s \neq c). \tag{9}$$

Neither expression involves the feature geometry or the identity of the moved exemplar; the margin geometry enters only through the error map $\mathrm{err}(\tau)$ of Corollary 6, not through the increment. □

B.3 PROPOSITION 9: INTERFERENCE LOCALITY AND THE AXIOMS

Proof. (a) Locality. By Lemma 7 the pairwise offsets on pairs not involving c are unchanged, and by gauge invariance (Corollary 6) the pairwise statistics are functions of (γ_{rs}, τ_{rs}) only. Marginals γ_{rs} are properties of the class means, which a count increment does not touch in the population ledger. Hence the entire error change of the promotion is supported on pairs (c, j).

(b) Calibration interference, uncomposed. In the uncomposed protocol the neighbor j's error against c is evaluated at the shifted offset $\tau_{cj} + \Delta\tau_{cj}$; the interference is the increase

$$I_{\mathrm{cal}}(c \to j) = \Phi\left(-\frac{\sqrt{\gamma_{cj}}}{2} + \frac{\tau_{cj} + \Delta\tau_{cj}}{\sqrt{\gamma_{cj}}}\right) - \Phi\left(-\frac{\sqrt{\gamma_{cj}}}{2} + \frac{\tau_{cj}}{\sqrt{\gamma_{cj}}}\right), \tag{10}$$

using $\mathrm{err}_{j\to c}(\tau) = \Phi(-\sqrt{\gamma}/2 + \tau/\sqrt{\gamma})$ from Corollary 6 (the sign of τ reflects viewing the offset from j's side). Nonnegativity and monotonicity in $\Delta\tau$ hold because Φ is increasing and $\Delta\tau_{cj} = \log\frac{m_c+1}{m_c} > 0$. For the envelope: the error map's derivative in τ is $\phi(-\sqrt{\gamma}/2 + \tau/\sqrt{\gamma})/\sqrt{\gamma} \leq 1/\sqrt{2\pi\gamma}$ unconditionally, giving

$$I_{\mathrm{cal}}(c \to j) \leq \frac{\Delta\tau_{cj}}{\sqrt{2\pi\gamma_{cj}}} \tag{11}$$

for free; at the composed balance point $\tau_{cj} = 0$ the slope at offset t is $\phi(\sqrt{\gamma}/2)\, e^{t/2 - t^2/(2\gamma)}/\sqrt{\gamma}$, which is at least its value at $t = 0$ for $t \leq \gamma$ and at most $e^{t/2}$ times it everywhere, so for $\Delta\tau_{cj} \leq \gamma_{cj}$

$$\phi\Big(\tfrac{\sqrt{\gamma_{cj}}}{2}\Big)\,\frac{\Delta\tau_{cj}}{\sqrt{\gamma_{cj}}} \;\leq\; I_{\mathrm{cal}}(c \to j) \;\leq\; e^{\Delta\tau_{cj}/2}\,\phi\Big(\tfrac{\sqrt{\gamma_{cj}}}{2}\Big)\,\frac{\Delta\tau_{cj}}{\sqrt{\gamma_{cj}}}, \tag{12}$$

so the interference is $\Theta\big(e^{-\gamma_{cj}/8}\big)\,\Delta\tau_{cj}/\sqrt{\gamma_{cj}}$ with the two constants within the factor $e^{\Delta\tau_{cj}/2} \leq \sqrt{2}$ of each other: the exponentially weak crowding in the margin. The exponential rate is a *balance-point* statement, not an absolute bound off balance: at $\tau = \gamma/2$, $\Delta\tau = \log 2$ the exact interference exceeds $e^{-\gamma/8}$ by a factor of 1.25 at $\gamma = 25$ and 4.14 at $\gamma = 36$ (exact arithmetic, verification script), which is why Proposition 9(b) states the unconditional linear bound and scopes the exponential envelope to the composed balance point.

(c) Entry crowding. Let $S \subseteq S'$ be label sets and let $\mathrm{reg}_j(S) = \{x : h_j(x) \geq h_k(x)\,\forall k \in S\}$ be the correct-classification region of j. Then $\mathrm{reg}_j(S') \subseteq \mathrm{reg}_j(S)$: winning against more competitors is a smaller event. Under the shared-covariance Gaussian model the probability of $\mathrm{reg}_j(S)$ under class j is nonincreasing in S, so j's error is nondecreasing in the label set; restricting to the single added competitor c gives the pairwise upper bound

$$\Pr_{x\sim\mathcal{N}(\mu_j,\Sigma)}\big[h_c(x) > h_j(x)\big] = \Phi(-\sqrt{\gamma_{cj}}/2),$$

exactly Corollary 6 at the balanced offset. Region containment delivers *level* monotonicity (more incumbents, smaller winning regions, higher per-class error) and nothing more: the *increment* an entrant imposes on a given incumbent is not monotone in the incumbent set, as the shielding configuration of Remark 1 shows, which is why Assumption A1 is stated on price infima and suprema. The pairwise envelope makes those suprema, and hence the coupling ratio χ, computable geometric quantities (ratios of neighbor-distance envelopes over reachable sets) rather than assumed constants. □

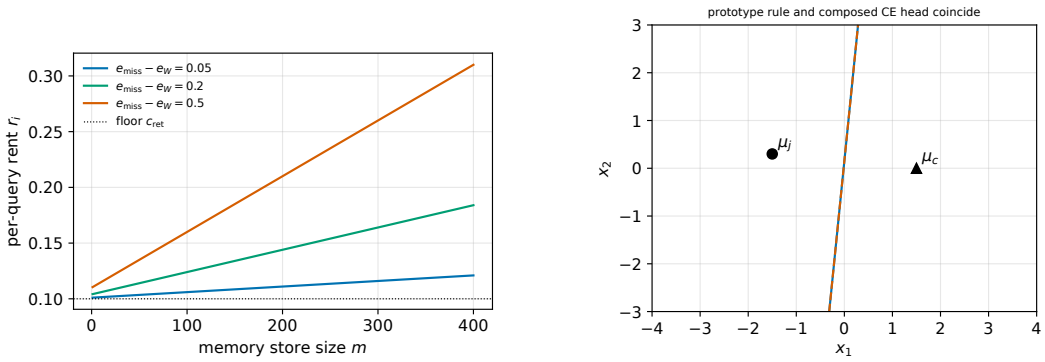

Figure 4: Left: the rent decomposition against memory store size m, a retrieval floor c_{ret} plus a failure channel $\ell\, p(m)(e_{\text{miss}} - e_{\text{W}})$ that grows with the store and with the coverage gap; large stores pay for their own retrieval unreliability. Right: store neutrality geometrically, the nearest-class-mean boundary (solid) and the composed CE head's boundary (dashed, coincident) are the same perpendicular bisector at balanced priors (Proposition 10). Deterministic closed forms.

B.4 Shielding (Remark 1), arithmetic

With unit variance and means $\mu_j = 0$, $\mu_k = 2$, $\mu_c = 4$ under the composed protocol, the Bayes regions are the Voronoi cells of the means. Against $W = \{j\}$ the entrant c moves the (j, c) boundary to 2, so j's region shrinks from $\mathbb{R}$ to $(-\infty, 2)$ and $\Delta E_j = \Phi(-2)$. Against $W = \{j, k\}$, j's region is already $(-\infty, 1)$ and the entrant at 4 does not touch it: $\Delta E_j = 0$ exactly, while k's region shrinks from $(1, \infty)$ to $(1, 3)$ and $\Delta E_k = \Phi(-1)$. Demand-weighting, $I(c; \{j\}) = \kappa\lambda_j\Phi(-2)$ and $I(c; \{j, k\}) = \kappa\lambda_k\Phi(-1)$, whose ratio is $(\lambda_k/\lambda_j)\,\Phi(-1)/\Phi(-2) \approx 6.97\,\lambda_k/\lambda_j$: the second is smaller whenever $\lambda_k < (\Phi(-2)/\Phi(-1))\,\lambda_j \approx 0.143\,\lambda_j$, at $\lambda_k = 0.05\,\lambda_j$ the interference term falls by 65 percent, and as $\lambda_k \to 0$ it vanishes and the entrant pays write compute alone. Error levels are monotone throughout (every class's error is nondecreasing in the label set); only the increments are not. The arithmetic is re-verified in the released verification script.

B.5 Composition neutrality (Remark 2)

Proof. Under composition the evaluator removes the count offsets: logit r loses $\log k_r$. Write the composed pairwise offset at evaluation prior π^{eval} as $\tau^{\text{comp}}_{rs} = \log(\pi^{\text{eval}}_r/\pi^{\text{eval}}_s)$, independent of the training counts. A $k \to k+1$ promotion of an already-consolidated class c changes τ_{cj} in training but not after composition: $\Delta\tau^{\text{comp}}_{cj} = 0$. Both terms of the promotion's population effect, the offset-mediated benefit $\Delta\nu_i$ and the offset-mediated interference I_{cal}, are functions of $\Delta\tau$; both vanish identically. The real effect of such a promotion is finite-sample: the empirical mean of class c improves at rate $1/\sqrt{m_c}$, which is estimation ledger (Section 7), not population ledger.

The entry promotion ($0 \to 1$) is different in kind. It changes the label set: c's queries stop paying the missing-class error $e_{\text{miss}}(c)$ (benefit, not removable by any protocol), and every existing class j pays the irreducible Bayes crowding of Proposition 9(c) (cost, not removable by any protocol). These two terms are the genuine protocol-independent buy price of the entry tier. □

Model definitions and assumptions (from Section 3)

A memory lookup succeeds with probability $1 - p(m)$ (failing with probability $p(m)$ when the store holds m items), and upon success, it is assumed to share the error rate of the head, an exact equivalence under infinite data (Proposition 10, above) and a clean approximation for finite data.

Definition 1 (Two-store consolidation)**.** Items arrive over discrete rounds; item i is queried with per-round probability λ_i, discount $\beta \in (0, 1)$. The memory store M (a nonparametric predictor over retrieved exemplars) serves unpromoted items; the weight store W (a parametric head trained on the consolidated buffer) serves promoted ones. Promoting i costs $B_i(W) = C_{\text{con}}(i) + I(i; W) - G(i)$: write compute, the interference increment on neighbors, and a spillover credit. While $i \in M$ each

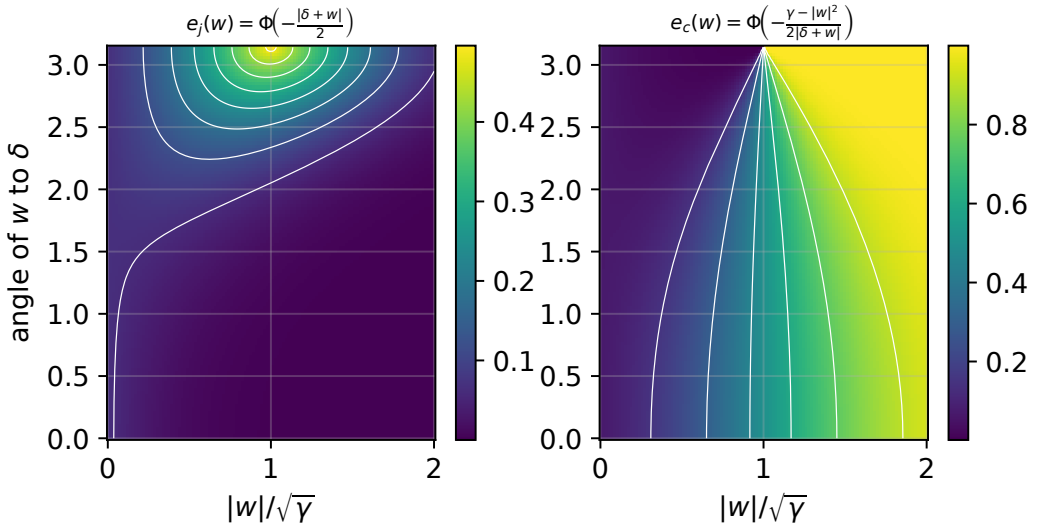

Figure 5: Exact state-conditional error maps of Theorem 14 at $\gamma = 9$ over the displacement magnitude and angle. Left: the neighbor side depends on the estimate only through $|\hat{\delta}|$. Right: the own side pays through both the length and the $|w|^2$ numerator; past $|w|^2 = \gamma$ the rule is worse than chance on c (the bright band beyond $|w|/\sqrt{\gamma} = 1$ at wide angles). Deterministic closed forms.

query pays rent $r_i = c_{\text{ret}} + \ell\, p(m_i)\,(e_{\text{miss}}(i) - e_{\text{W}}(i))$, with c_{ret} the per-retrieval context cost, $p(m) = p_0 + \rho m$ the retrieval-failure probability at store size m, ℓ the error-to-compute currency, and $e_{\text{miss}}, e_{\text{W}}$ the missing-class and head errors on i's queries.

Assumption A1 (Bounded positive buy prices)**.** For every item i, $0 < B_i^{\min} := \inf_W B_i(W) \le \sup_W B_i(W) =: B_i^{\text{pes}} < \infty$. The *coupling ratio* is $\chi := \max_i \chi_i$, $\chi_i := B_i^{\text{pes}}/B_i^{\min}$; $\chi = 1$ is the decoupled case. Monotonicity of $B_i(\cdot)$ in W is *not* assumed.

Proposition 8 (Capitalization: lump plus entry crowding equals expected flow)**.** *Define the interference lump as the capitalized demand-weighted crowding flow, $I(i; W) := \frac{\ell}{1-\beta} \sum_{j \in W} \lambda_j \, \Delta E_j(i; W)$, with ΔE_j the per-query error increase of neighbor j at i's entry. The total discounted cost of the promote-at-0 policy under flow-accounted interference equals $C_{\text{con}} + I - G$ plus the item's own serving costs, so the lump-price threshold of Theorem 12 is exact; and by telescoping of the entry-time increments, entry-priced lumps, plus each entrant's own crowding at its entry, equal the expected discounted flow for every one-way schedule. The identity needs one-way decisions, stationary demand and $\beta < 1$: demote an item, let its demand change after the buy, or set $\beta = 1$, and it breaks. Section 4 lives in that gap.*

The capitalization identity in full (from Section 3)

The *capitalization lemma* (Proposition 8) establishes when the two accountings of Section 3 align. Assuming moves are irreversible, demand is stationary, and future costs are discounted at rate $\beta < 1$ (where a cost incurred t rounds ahead is weighted by β^t), the lump sum plus each newcomer's initial crowding equals the expected discounted flow. Reverting to the v and k example from Section 1, lump accounting charges v's promotion once for the crowding it imposes on k (excluding v's own entry crowding), whereas flow accounting levies a continuous tax of approximately 0.16ℓ on every subsequent query to v for as long as both items occupy the head. This equivalence breaks down when items can be demoted or when demand shifts post-purchase, a gap explored formally by the theorems in Section 4.

Instantiation results (from Section 5)

Proposition 9 (Locality, envelopes, entry crowding)**.** *(a) A $k \to k{+}1$ promotion of class c leaves every pairwise statistic not involving c unchanged. (b) Calibration interference on a neighbor j is*

closed form, nonnegative, and monotone, with an explicit unconditional upper envelope and a two-sided one at the balance point; the formulas and their off-balance failure modes are in Appendix B. (c) When c first enters the label set, every incumbent j's Bayes error rises by a crowding term bounded pairwise by $\Phi(-\sqrt{\gamma_{cj}}/2)$, monotone in the label set by region containment.

Remark 1 (Shielding: prices are non-monotone in the instantiation). Increment monotonicity fails inside the Gaussian instantiation of Section 5 even though error levels are monotone: with unit variance and means $\mu_j = 0$, $\mu_k = 2$, $\mu_c = 4$, the entrant c against $W = \{j\}$ moves the (j, c) boundary to 2 and costs $\Delta E_j = \Phi(-2)$, while against $W = \{j, k\}$ the intermediate class k shields j ($\Delta E_j = 0$ exactly) and the crowding lands on k as $\Delta E_k = \Phi(-1)$. Demand-weighted, the price *falls* when a neighbor is added whenever $\lambda_k < (\Phi(-2)/\Phi(-1))\,\lambda_j \approx 0.143\,\lambda_j$, by 65 percent at $\lambda_k = 0.05\,\lambda_j$. The online theory therefore runs on the infimum and supremum of the price system, a weaker hypothesis and so a wider theorem than any monotone-price axiom.

Remark 2 (Composition neutrality). Under composition the offset-mediated benefit and the offset-mediated interference cancel identically for $k \to k{+}1$ promotions of consolidated classes. The entry promotion ($0 \to 1$) is categorically different: it changes the label set, removes e_{miss}, and imposes the irreducible Bayes crowding of (c). The exemplar tier is therefore estimation economics alone (Section 7); the uncomposed variant is Appendix F.

Proposition 10 (Store neutrality)**.** *Over matched label sets, the prototype rule and the composed CE head are the same classifier: the nearest class mean in the Mahalanobis metric, Bayes at balanced priors. So $e_{\mathrm{M}} = e_{\mathrm{W}}$, and $r_i = c_{\mathrm{ret}} + \ell\, p\,(e_{\mathrm{miss}} - e_{\mathrm{W}}) > 0$ strictly. Consolidation buys lower serving cost, retrieval reliability, and, through coverage, the removal of e_{miss}. It does not buy per-query accuracy on covered classes.*

A BENCHMARK PUZZLE, EXPLAINED (FROM SECTION 5)

So what does consolidation buy? Cheaper answers, answers that do not depend on a lookup succeeding, and answers for classes that memory would otherwise miss. It does not buy better accuracy on classes the head already serves. Under balanced priors, benchmarks that compare prototype-based continual learners with trained heads are measuring how well each estimates from finite data, not a difference between the classifiers themselves (Proposition 10, above).

C PROOFS OF THE THRESHOLD THEOREMS

The difference identity and the threshold algebra are verified symbolically in the released instruments; Figure 6 and panels (a) and (b) of Figure 2 visualize the exact objects.

THE COMPLETED CONSTANT ASSEMBLED

Per pair (c, j) let $A_{cj} = \phi(\sqrt{\gamma_{cj}}/2)\frac{d-1}{4\sqrt{\gamma_{cj}}}$ and $B_{cj} = \phi(\sqrt{\gamma_{cj}}/2)\frac{\sqrt{\gamma_{cj}}}{16}$, so that the own-side and neighbor-side coefficients of Proposition 16 are $3A_{cj} + B_{cj}$ and $B_{cj} - A_{cj}$. With both means estimated, the demand-weighted marginal value of $m_c \to m_c + 1$ is $\Delta V_c = \frac{\ell}{1-\beta} S_c\big(\frac{1}{m_c} - \frac{1}{m_c+1}\big)$, $S_c := \sum_j [\Lambda_c(3A_{cj} + B_{cj}) + \Lambda_j(B_{cj} - A_{cj})]$, and the completed law promotes iff $\Delta V_c > C_{\mathrm{con}}$, giving $m_c^* = \Theta\big(\sqrt{\ell S_c/((1-\beta)C_{\mathrm{con}})}\big)$ when $S_c > 0$: the square root survives with a corrected constant, and the own-error form is the special case $\Lambda_j = 0$.

The three comes from the parts of the two coefficients that grow with dimension, which stand at 3:1 under balanced priors; unequal priors move it.

C.1 LEMMA 11: NOW-OR-NEVER

Proof. Fix an item with parameters (λ, r, B, β) and write the discounted cost of the promote-at-t policy (rent until t, buy at t; per-query costs common to both stores cancel because r is differential):

$$V(t) = \lambda r \frac{1 - \beta^t}{1 - \beta} + \beta^t B. \tag{13}$$

The forward difference is

$$V(t+1) - V(t) = \lambda r \frac{\beta^t - \beta^{t+1}}{1-\beta} + (\beta^{t+1} - \beta^t)B = \beta^t \big[\lambda r - (1-\beta)B\big],$$

constant in sign. If $\lambda r > (1-\beta)B$ the difference is positive, V is increasing, and $t = 0$ is optimal. If the inequality is reversed, V decreases to its infimum at $t \to \infty$, which is the never-promote policy of value $\lambda r/(1-\beta)$; at equality all t are optimal. There is no interior optimum at any horizon: monotone functions on the integers attain their minima at endpoints. □

C.2 Theorem 12: threshold and frontier

Proof. Lemma 11 gives the per-item decision promote-at-0 iff $\lambda r > (1-\beta)B$; rearranging,

$$\lambda > \lambda^* = \frac{(1-\beta)B}{r} = \frac{(1-\beta)(C_{\text{con}} + I - G)}{r}. \tag{14}$$

For the frontier, write the interference as $I = \iota\, \kappa_{\text{scale}}$ by the normalization of susceptibility (demand-weighted crowding envelope, Propositions 9(c) and 8); since B is affine in I, the frontier

$$\lambda^*(\iota) = \frac{(1-\beta)(C_{\text{con}} + \iota\kappa_{\text{scale}} - G)}{r} \tag{15}$$

is linear in ι, decreasing in β (patient agents promote more) and increasing in write compute. The three regions follow: below the frontier never promote, above it always promote at arrival, on it indifferent. With $G = 0$ in the population ledger (Remark 2), the frontier's intercept is $(1-\beta)C_{\text{con}}/r$ and its slope is $(1-\beta)\kappa_{\text{scale}}/r$; Figure 2 plots the family. □

C.3 Proposition 8: capitalization

Proof. Consider the promote-at-0 policy for item i and account interference as a flow: from round 0 onward, each query to a neighbor $j \in W$ pays the per-query error increase $\Delta E_j(i; W)$ converted at rate ℓ, and such queries arrive at rate λ_j per round. The discounted sum of this flow is

$$\sum_{t=0}^{\infty} \beta^t\, \ell \sum_{j \in W} \lambda_j \Delta E_j = \frac{\ell}{1-\beta} \sum_{j \in W} \lambda_j \Delta E_j = I(i; W), \tag{16}$$

the geometric series collapsing because the flow is stationary once i is in place (one-way decision, stationary demand). Add the write compute C_{con} and subtract the spillover credit; the total discounted cost of the promote-at-0 policy is exactly $C_{\text{con}} + I - G$ plus the item's own serving costs, independent of the accounting convention. Hence the lump-price threshold of Theorem 12 is the exact one-way decision rule, not an approximation. The seam the lemma opens, and Section 4 drives a theorem through, is that the equivalence needs one-way and stationary: a demotable item, or demand that changes after the buy, breaks the geometric-series collapse. □

Sequences telescope. The lemma extends from one promotion to any one-way schedule: with $\Delta E_j(i; W)$ denoting the marginal error increment on j at i's entry-time incumbent set and W_{t_j} the store just after j's own entry, j's crowding telescopes,

$$E_j(W_t) - E_j(\{j\}) \;=\; \big[E_j(W_{t_j}) - E_j(\{j\})\big] + \sum_{\text{entrants } i \text{ after } j,\ i \le t} \Delta E_j(i; W \text{ at } i\text{'s entry}), \tag{17}$$

so entry-priced lumps, plus each entrant's own crowding at its entry (the bracket), equal the expected discounted flow for every one-way schedule under stationary demand and $\beta < 1$, not just the single-item promote-at-0 case. The lump alone omits the bracket, which is the crowding Theorem 2's adversary charges; at $\beta = 1$ the lump is a bounded budget fixed in advance and the flow a per-query liability, and the adversarial ordering separates them.

C.4 Proposition 10: store neutrality

Proof. The nearest-class-mean rule assigns

$$\arg\min_r (x - \mu_r)^\top \Sigma^{-1} (x - \mu_r) = \arg\max_r \big[\mu_r^\top \Sigma^{-1} x - \tfrac{1}{2} \mu_r^\top \Sigma^{-1} \mu_r\big],$$

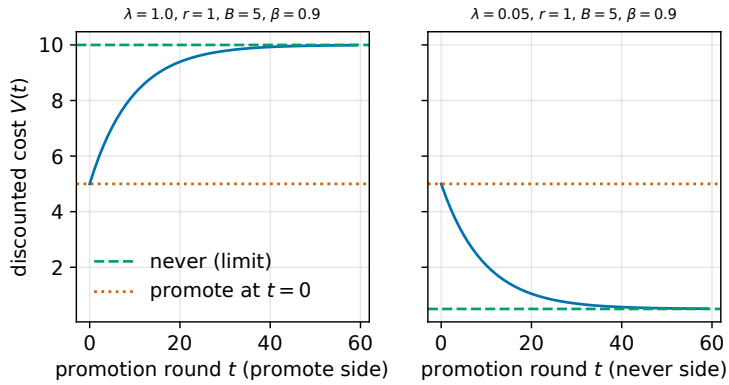

Figure 6: Lemma 11: the discounted promotion cost $V(t)$ is monotone in t, so the optimum is an endpoint. Left: on the promote side the cost climbs from $t = 0$. Right: on the never side it declines to the rent-forever limit. Exact closed forms; E-L1-1 verifies endpoint optimality against a finite-horizon dynamic program.

which is Bayes for the shared-covariance Gaussian model at balanced priors (no $\log \pi_r$ term). By Lemma 5 the composed CE head implements $\arg\max_r[\mu_r^\top \Sigma^{-1} x - \frac{1}{2}\mu_r^\top \Sigma^{-1}\mu_r + \log \pi_r^{\text{eval}}]$; at the balanced evaluation prior the offsets vanish and the two rules coincide classifier-by-classifier. Identical classifiers have identical errors: $e_{\text{M}} = e_{\text{W}}$ over matched label sets. The rent is then

$$r_i = c_{\text{ret}} + \ell\, p(m_i)(e_{\text{miss}}(i) - e_{\text{W}}(i)) > 0 \tag{18}$$

strictly, since $e_{\text{miss}} \geq e_{\text{W}}$ (an uncovered class errs at least as much as a covered one), $c_{\text{ret}} > 0$, and $p > 0$: consolidation buys lower serving cost and retrieval reliability, and through coverage the removal of e_{miss}; it does not buy per-query accuracy on covered classes. The matched-label-sets hypothesis is stated in the proposition because it is load-bearing: with unmatched coverage the two stores answer different questions, and the comparison of per-query errors is not well-posed. □

C.5 Corollary 13: the public-good tier and the conditional square-root law

Proof. Under composition (Remark 2) a $k \to k+1$ promotion of class c has zero population effect; its benefit is the estimation improvement of the fitted head on class c, and this benefit accrues to every query the head answers for class c, at total rate $\Lambda_c = \sum_{i \in c} \lambda_i$. The marginal saving of the $(m \to m+1)$ write is $\ell\Lambda_c[\epsilon_c(m) - \epsilon_c(m+1)]$ per round, capitalized at $1/(1-\beta)$; promoting while this exceeds C_{con} is exactly Lemma 11's comparison with the public-good rent in place of the private one, which proves the general threshold.

For the conditional law, suppose the marginal gains $\epsilon_c(m) - \epsilon_c(m+1)$ are nonincreasing in m and $m^2[\epsilon_c(m) - \epsilon_c(m+1)] \to a_c$ (which gives $\epsilon_c(m) = a_c/m + o(1/m)$). Then $\epsilon_c(m) - \epsilon_c(m+1) = a_c/(m(m+1)) + o(1/m^2)$, the first crossing and the last profitable count coincide, and the stopping count is the largest m with

$$\frac{\Lambda_c \ell\, a_c}{(1-\beta)C_{\text{con}}} > m(m+1) + o(m^2), \tag{19}$$

so $m^*(m^*+1) \approx \Lambda_c \ell a_c/((1-\beta)C_{\text{con}})$ and

$$m^* = \Theta\left(\sqrt{\Lambda_c \ell a_c/((1-\beta)C_{\text{con}})}\right). \tag{20}$$

Buffer counts scale as the square root of class demand: the economic-order-quantity law transplanted to consolidation. The law is conditional on the $1/m$ rate; Section 7 discharges the condition for the plug-in rule with the explicit a_c, and E-L1-5 measures the exponent through the exact benchmark rather than the asymptotic. □

THRESHOLD ECONOMICS STATEMENTS (FROM SECTION 6)

Lemma 11 (Now-or-never). *Decoupled regime, stationary λ_i, discount β. With $V(t) = \lambda r \frac{1-\beta^t}{1-\beta} + \beta^t B$ the cost of promoting at round t, the forward difference $V(t+1) - V(t) = \beta^t[\lambda r - (1-\beta)B]$ has constant sign, so the optimal promotion time is 0 or never.*

Theorem 12 (Threshold and frontier). *In the decoupled regime (Lemma 11; equivalently, a single item facing a fixed incumbent set), the optimal policy promotes item i at arrival iff $\lambda_i > \lambda_i^* = (1-\beta)B_i/r_i$; with coupled prices the per-item threshold is the greedy rule and the joint optimum is NP-hard (Proposition 18). Writing the interference as $\iota\,\kappa_{\text{scale}}$, the frontier $\lambda^*(\iota) = (1-\beta)(C_{\text{con}} + \iota\kappa_{\text{scale}} - G)/r$ is linear in the normalized interference, decreasing in β at fixed prices and increasing in write compute, and splits demands into never, always, and indifferent regions. The theorem is the full-information benchmark: rates are known to the policy, unlike the online setting of Section 4.*

Corollary 13 (Exemplar tier as public good; conditional square-root law). *Under composition a $k \to k{+}1$ promotion of class c has zero population effect; its benefit is the head's estimation improvement on class c, accruing to every item of the class at total demand Λ_c. Promote iff $\Lambda_c\,\ell\,[\epsilon_c(m_c) - \epsilon_c(m_c{+}1)]/(1-\beta) > C_{\text{con}}$. Conditional on nonincreasing marginal gains with $m^2[\epsilon_c(m) - \epsilon_c(m{+}1)] \to a_c$ the stopping count satisfies $m_c^* = \Theta\big(\sqrt{\Lambda_c \ell a_c/((1-\beta)C_{\text{con}})}\big)$: buffer counts scale as the square root of class demand, the economic-order-quantity law of inventory theory (Harris, 1913), which Baumol (1952) and Tobin (1956) carried to the demand for cash, transplanted to consolidation. Section 7 discharges the condition with the explicit constant. The condition is an $m \to \infty$ asymptotic whose domain is $m \gg d\max(1, 1/\gamma)$; the registered runs map that region and assert nothing beyond it.*

WHICH EXAMPLE NEXT: THE TWO SIDES OF THE VALUE (FROM SECTION 7)

Which example should be promoted next? To first order, the value of promoting example z is how far z lies from the current mean estimate $\hat{\mu}_c$ along the line between the two estimated class means, which is at right angles to the decision boundary, with a weight for each pair given by a formula (Theorem 15, below). That is the neighbor's side of the value, the half the agent can compute from its own estimates; the class's own side depends on the true mean. The classical herding rule (Welling, 2009) has no such weights. Behind this, the error of each pair, given the current estimates, is an explicit Gaussian tail probability (Theorem 14, below). Noise that pushes the two estimated means apart *helps* the neighbor, so estimation noise does not hurt both sides equally.

D PROOFS OF THE ESTIMATION THEOREMS

Theorem 14 (Exact state-conditional pairwise errors). *One-sided estimation with displacement w ($\hat{\mu}_c = \mu_c + w$, $\hat{\mu}_j = \mu_j$): exactly,*

$$e_j(w) = \Phi\Big(-\frac{|\delta + w|}{2}\Big), \qquad e_c(w) = \Phi\Big(-\frac{\gamma - |w|^2}{2|\delta + w|}\Big), \tag{21}$$

both $\Phi(-\sqrt{\gamma}/2)$ at $w = 0$; for $|w|^2 > \gamma$ the rule is worse than chance on c. The neighbor's error depends on the estimate only through $|\hat{\delta}|$: wander that lengthens the estimated difference relieves the neighbor.

Theorem 15 (Position-dependent promotion value). *Promoting exemplar z of class c at count m moves the empirical mean by $v(z) = (z - \hat{\mu}_c)/(m{+}1)$, so $\Delta e_j(z) = \Phi(-|\delta + w + v|/2) - \Phi(-|\delta + w|/2)$ exactly, and to first order*

$$\Delta e_j(z) = -\phi\Big(\frac{|\hat{\delta}|}{2}\Big)\frac{\hat{\delta}^\top v(z)}{2|\hat{\delta}|} + O(|v|^2), \tag{22}$$

a function of the agent's own state: the boundary-aligned component of $z - \hat{\mu}_c$ with the pair weights $\phi(|\hat{\delta}|/2)/(2|\hat{\delta}|)$ that the herding heuristic lacks. The e_c side involves the true displacement and is oracle-only: what-to-promote splits exactly into a computable half and an estimation target.

Proposition 16 (Expected excess and the stocking constant)**.** *With $w \sim \mathcal{N}(0, I/m)$, the balanced pairwise excess is $\epsilon_{\text{pair}}(m) = \mathbb{E}[\frac{1}{2}(e_c + e_j)] - \Phi(-\sqrt{\gamma}/2) = a(\gamma, d)/m + O(m^{-2})$ with*

$$a(\gamma, d) = \phi\Big(\frac{\sqrt{\gamma}}{2}\Big)\Big(\frac{d-1}{4\sqrt{\gamma}} + \frac{\sqrt{\gamma}}{16}\Big), \tag{23}$$

and with both means estimated the leading term is $a(\gamma, d)(1/m_c + 1/m_j)$. Per side, $\mathbb{E}[e_j] - \Phi(-\sqrt{\gamma}/2) = \frac{\phi(\sqrt{\gamma}/2)}{m}\big(-\frac{d-1}{4\sqrt{\gamma}} + \frac{\sqrt{\gamma}}{16}\big)$ and $\mathbb{E}[e_c] - \Phi(-\sqrt{\gamma}/2) = \frac{\phi(\sqrt{\gamma}/2)}{m}\big(\frac{3(d-1)}{4\sqrt{\gamma}} + \frac{\sqrt{\gamma}}{16}\big)$, up to $O(m^{-2})$: through the $(d-1)$ channel a class's own mean noise burdens itself at $3/4$ and relieves the neighbor at $1/4$. The $\xi_\parallel^2$ terms of the own-side numerator and inverse length cancel exactly; a derivation that skips the cancellation produces d in place of $d-1$, and the exact $d = 1$ law $\sqrt{\gamma}\,\phi(\sqrt{\gamma}/2)/16m$ excludes that variant.

The per-side coefficients are the known-covariance first-order terms of Okamoto's expansion of the plug-in discriminant's error (Okamoto, 1963). We quote that expansion in the form restated by Kubokawa et al. (2013, eq. (20)), who identify it with Okamoto's, with its $1/n$ covariance term dropped; substituting $\Delta = \sqrt{\gamma}$ and $p = d$ reproduces the coefficients above, and at $(\gamma, d) = (4, 16)$ their values 1.391332 and -0.423449. The constant and its split are therefore classical, and what the ledger adds is their use as prices.

Theorem 14's formulas are machine-verified to 6×10^{-16} over random configurations; the series coefficients of Proposition 16 are verified symbolically (identically zero residuals against the stated constant) and by Monte Carlo, and the two-sided law at general d at instrument level. Per-side constants are written in the multiplicative form $\phi(\sqrt{\gamma}/2) \cdot (\cdot)/m$: the pdf value at the balanced margin times the stated coefficients, not the pdf evaluated at them. Figures 3b and 7 visualize the constant and its split.

D.1 Theorem 14: exact state-conditional errors

Proof. Whiten so $\Sigma = I$ and place μ_j at the origin, $\mu_c = \delta$ with $|\delta|^2 = \gamma$. The plug-in rule assigns class c iff $\hat{\delta}^\top x > \frac{1}{2}(|\hat{\mu}_c|^2 - |\hat{\mu}_j|^2)$ with $\hat{\delta} = \delta + w$, $\hat{\mu}_c = \delta + w$, $\hat{\mu}_j = 0$; the boundary is the perpendicular bisector of the segment $[\hat{\mu}_j, \hat{\mu}_c]$, i.e. the hyperplane $\hat{\delta}^\top x = \frac{1}{2}|\hat{\delta}|^2$.

Neighbor side. For $x \sim \mathcal{N}(0, I)$ the statistic $\hat{\delta}^\top x$ is $\mathcal{N}(0, |\hat{\delta}|^2)$, so

$$e_j(w) = \Pr\big[\hat{\delta}^\top x > \tfrac{1}{2}|\hat{\delta}|^2\big] = \Phi\left(-\frac{|\hat{\delta}|}{2}\right) = \Phi\left(-\frac{|\delta + w|}{2}\right).$$

Geometrically: the true mean of j sits at one endpoint of the estimated segment, the distance to the bisector is half the segment length, and the Gaussian mass beyond a hyperplane at signed distance s is $\Phi(-s)$. The neighbor's error depends on the estimate only through the length $|\hat{\delta}|$: any wander of $\hat{\mu}_c$ that lengthens the estimated difference relieves the neighbor.

Own side. For $x \sim \mathcal{N}(\delta, I)$ the statistic has mean $\hat{\delta}^\top \delta = (\delta + w)^\top \delta = \gamma + w^\top \delta$ and variance $|\hat{\delta}|^2$, so

$$e_c(w) = \Pr\big[\hat{\delta}^\top x \le \tfrac{1}{2}|\hat{\delta}|^2\big] = \Phi\left(\frac{\frac{1}{2}|\delta + w|^2 - \gamma - w^\top \delta}{|\delta + w|}\right).$$

The numerator simplifies by the polarization identity $|\delta + w|^2 = \gamma + 2\delta^\top w + |w|^2$:

$$\tfrac{1}{2}|\delta + w|^2 - \gamma - \delta^\top w = \tfrac{1}{2}(|w|^2 - \gamma), \tag{24}$$

the linear terms cancelling, giving

$$e_c(w) = \Phi\left(-\frac{\gamma - |w|^2}{2|\delta + w|}\right). \tag{25}$$

The own side pays twice, through the length in the denominator and through $|w|^2$ in the numerator; for $|w|^2 > \gamma$ the argument flips sign and the rule errs on c more than half the time, as it must once the estimate has wandered past the neighbor. □

D.2 Theorem 15: position-dependent promotion value

Proof. Promoting exemplar z of class c at count m moves the empirical mean deterministically:

$$\hat{\mu}_c^{\text{new}} = \frac{m\,\hat{\mu}_c + z}{m+1}, \qquad v(z) = \hat{\mu}_c^{\text{new}} - \hat{\mu}_c = \frac{z - \hat{\mu}_c}{m+1}, \tag{26}$$

so the displacement changes $w \mapsto w + v(z)$ and Theorem 14 applies verbatim at the shifted displacement:

$$\Delta e_j(z) = \Phi\left(-\frac{|\delta + w + v|}{2}\right) - \Phi\left(-\frac{|\delta + w|}{2}\right), \tag{27}$$

with the analogous expression for Δe_c from the second formula. This is exact, and it answers the design question in the strong direction: the promotion increment depends on the exemplar's feature position, in closed form, and the count-only increment of Lemma 7 is its population shadow.

First order in v. Expand $e_j = \Phi(-s/2)$, $s = |\hat{\delta}|$, in the displacement of s; the differential of s along v is $\hat{\delta}^\top v/|\hat{\delta}|$, so

$$\Delta e_j = \Phi'(-s/2) \cdot \left(-\frac{1}{2} \cdot \frac{\hat{\delta}^\top v}{|\hat{\delta}|}\right) + O(|v|^2) = -\phi\left(\frac{|\hat{\delta}|}{2}\right) \frac{\hat{\delta}^\top v}{2|\hat{\delta}|} + O(|v|^2). \tag{28}$$

Every quantity on the right is a function of the agent's own state (empirical means, counts): the agent can evaluate its neighbor-side promotion value exactly. It is the boundary-aligned component of $(z - \hat{\mu}_c)$, weighted by the pair weights $\phi(|\hat{\delta}|/2)/(2|\hat{\delta}|)$; the connection to the γ-preserving coreset objective is a shared privileged direction, not a shared weighting, and the derivation supplies the weights the heuristic lacked. The own side Δe_c involves the true δ and the true displacement w through both numerator and denominator; its first-order coefficient is not a function of the agent's state. The what-to-promote question therefore splits exactly: a computable half and an oracle half. □

D.3 Proposition 16: expected excess and the constant

Proof. Let $w = \xi \sim \mathcal{N}(0, I/m)$ and decompose $\xi = \xi_{\text{par}} + \xi_\perp$ into the along-δ coordinate and the orthogonal block of dimension $d-1$, with $\mathbb{E}[\xi_{\text{par}}^2] = 1/m$ and $\mathbb{E}[\,|\xi_\perp|^2\,] = (d-1)/m$.

The neighbor side. From $e_j = \Phi(-s/2)$ at $s = |\delta + \xi|$: expand the length (the ξ_{par}^2 contributions of the two binomial orders cancel),

$$|\delta + \xi| = \sqrt{\gamma} + \xi_{\text{par}} + \frac{|\xi_\perp|^2}{2\sqrt{\gamma}} + O(m^{-3/2}), \tag{29}$$

so $\mathbb{E}[s] - \sqrt{\gamma} = \frac{d-1}{2\sqrt{\gamma}\,m} + O(m^{-2})$ and $\mathrm{Var}(s) = 1/m + O(m^{-2})$ (the linear term dominates the variance). A second-order Taylor expansion of $g(s) = \Phi(-s/2)$ at $s = \sqrt{\gamma}$ gives

$$\mathbb{E}[g(s)] - g(\sqrt{\gamma}) = g'(\sqrt{\gamma})\,\mathbb{E}[s - \sqrt{\gamma}] + \tfrac{1}{2} g''(\sqrt{\gamma})\,\mathrm{Var}(s) + O(m^{-2}),$$

with $g'(s) = -\phi(s/2)/2$ and $g''(s) = -\phi'(s/2)/4 = (s/8)\phi(s/2)$, so $g'(\sqrt{\gamma}) = -\phi(\sqrt{\gamma}/2)/2$ and $g''(\sqrt{\gamma}) = \frac{\sqrt{\gamma}}{8}\phi(\sqrt{\gamma}/2)$, hence (the first-order term contributing $-\frac{d-1}{4\sqrt{\gamma}}$ and the curvature term $\frac{1}{2} \cdot \frac{\sqrt{\gamma}}{8} \cdot \frac{1}{m} = \frac{\sqrt{\gamma}}{16m}$)

$$\mathbb{E}[e_j] - \Phi(-\sqrt{\gamma}/2) = \frac{\phi(\sqrt{\gamma}/2)}{m}\left(-\frac{d-1}{4\sqrt{\gamma}} + \frac{\sqrt{\gamma}}{16}\right) + O(m^{-2}).$$

The own side and the cancellation. The own-side argument is $u = (\gamma - |\xi|^2)/|\delta + \xi|$. Expand numerator and denominator:

$$\gamma - |\xi|^2 = \gamma - \xi_{\text{par}}^2 - |\xi_\perp|^2, \qquad \frac{1}{|\delta + \xi|} = \frac{1}{\sqrt{\gamma}}\left(1 - \frac{\xi_{\text{par}}}{\sqrt{\gamma}} - \frac{|\xi_\perp|^2}{2\gamma} + \frac{\xi_{\text{par}}^2}{\gamma}\right) + O(m^{-3/2}). \tag{30}$$

The ξ_{par}^2 contribution of the numerator, $-\xi_{\mathrm{par}}^2/\sqrt{\gamma}$, cancels *exactly* against the ξ_{par}^2 term of the inverse length, $+\xi_{\mathrm{par}}^2/\gamma \cdot \gamma/\sqrt{\gamma} = +\xi_{\mathrm{par}}^2/\sqrt{\gamma}$ (writing $|\delta+\xi|^{-1} = \gamma^{-1/2}(1+h)^{-1/2}$ with $h = 2\xi_{\mathrm{par}}/\sqrt{\gamma} + |\xi|^2/\gamma$, the $3h^2/8$ term contributes $+3\xi_{\mathrm{par}}^2/(2\gamma)$ on top of the $-\xi_{\mathrm{par}}^2/(2\gamma)$ of the linear term); this is the cancellation a first derivation missed (it attributed $-d/\sqrt{\gamma}$ to the numerator and produced the retracted d-version of the constant). Only the perpendicular block survives:

$$u = \sqrt{\gamma} - \xi_{\mathrm{par}} - \frac{3}{2}\frac{|\xi_\perp|^2}{\sqrt{\gamma}} + O(m^{-3/2}), \qquad \mathbb{E}[u] - \sqrt{\gamma} = -\frac{3(d-1)}{2\sqrt{\gamma}\,m}, \quad \mathrm{Var}(u) = \frac{1}{m} + O(m^{-2}).$$

The same Taylor step, with the sign of the mean shift reversed, gives

$$\mathbb{E}[e_c] - \Phi(-\sqrt{\gamma}/2) = \frac{\phi(\sqrt{\gamma}/2)}{m}\left(\frac{3(d-1)}{4\sqrt{\gamma}} + \frac{\sqrt{\gamma}}{16}\right) + O(m^{-2}).$$

Averaging. The balanced pairwise excess averages the two sides: the $(d-1)$-terms combine as $(-\frac{1}{4}+\frac{3}{4})\frac{d-1}{\sqrt{\gamma}} = \frac{d-1}{2}\cdot\frac{1}{\sqrt{\gamma}}\cdot\frac{1}{2} = \frac{d-1}{4\sqrt{\gamma}}$, the curvature terms as $\frac{2\sqrt{\gamma}}{16}/2 = \frac{\sqrt{\gamma}}{16}$, and

$$\epsilon_{\mathrm{pair}}(m) = \frac{a(\gamma,d)}{m} + O(m^{-2}), \qquad a(\gamma,d) = \phi\left(\frac{\sqrt{\gamma}}{2}\right)\left(\frac{d-1}{4\sqrt{\gamma}} + \frac{\sqrt{\gamma}}{16}\right). \tag{31}$$

The own/neighbor asymmetry is exactly 3:1 in the $(d-1)$ channel: own mean noise burdens the estimated class at $3/4$ of $(d-1)/\sqrt{\gamma}$ per unit of ϕ/m and relieves the neighbor at $1/4$.

The $d = 1$ certificate. At $d = 1$ the geometry collapses to pure threshold noise: $\xi_\perp$ vanishes and $e_{\mathrm{pair}} = \frac{1}{2}[\Phi(-(\sqrt{\gamma}+\xi)/2) + \Phi(-(\sqrt{\gamma}-\xi)/2)]$, whose exact expectation has excess $\sqrt{\gamma}\,\phi(\sqrt{\gamma}/2)/16m + O(m^{-2})$; the corrected constant reduces to exactly this at $d = 1$ (the $(d-1)$ term dies) and the retracted d-version does not. This exact special case is what first discriminated the two constants (Appendix J.1).

Both means estimated. With $\hat{\mu}_c$ and $\hat{\mu}_j$ carrying independent noises, the excess is quadratic in two independent perturbations whose parallel second-order effects cancel individually and whose perpendicular effects are additive at leading order; the joint leading term is the superposition

$$\epsilon_{\mathrm{pair}}(m_c, m_j) = a(\gamma,d)\left(\frac{1}{m_c} + \frac{1}{m_j}\right) + O\left(\frac{1}{m^2}\right), \tag{32}$$

the two-sided additivity law, with the per-side asymmetric constants as above (verified at general d by the released instruments, seed-stable; Figure 7). □

D.4 Corollary 4: the completed stocking law and the 3Λ rule

Proof. Instantiate Proposition 16's per-side coefficients per pair (c, j): with only c's mean noisy at count m_c, the expected own-side error shift per unit of ϕ/m_c is $3A_{cj} + B_{cj}$ and the neighbor-side shift is $-(A_{cj} - B_{cj})$, where

$$A_{cj} = \phi(\sqrt{\gamma_{cj}}/2)\frac{d-1}{4\sqrt{\gamma_{cj}}}, \qquad B_{cj} = \phi(\sqrt{\gamma_{cj}}/2)\frac{\sqrt{\gamma_{cj}}}{16}; \tag{33}$$

the two-sided law assigns c's noise the coefficients $(3A+B)/m_c$ on e_c and $(B-A)/m_c$ on e_j, and symmetrically for j's noise. The demand-weighted value of one promotion of class c (which changes the $1/m_c$ prefactor, not the coefficients, at leading order) is therefore

$$\begin{aligned}\Delta V_c &= \frac{\ell}{1-\beta}\sum_j \left[\Lambda_c(3A_{cj}+B_{cj}) + \Lambda_j(B_{cj}-A_{cj})\right]\left(\frac{1}{m_c} - \frac{1}{m_c+1}\right) + O(m^{-3}) \\ &= \frac{\ell}{1-\beta}\,S_c\,\frac{1}{m_c(m_c+1)} + O(m^{-3}),\end{aligned}$$

with S_c the pair sum as stated. Promoting while $\Delta V_c > C_{\mathrm{con}}$ stops at $m_c^* = \Theta(\sqrt{\ell S_c/((1-\beta)C_{\mathrm{con}})})$ whenever $S_c > 0$, by the same $m^*(m^*+1)$ arithmetic as Corollary 13.

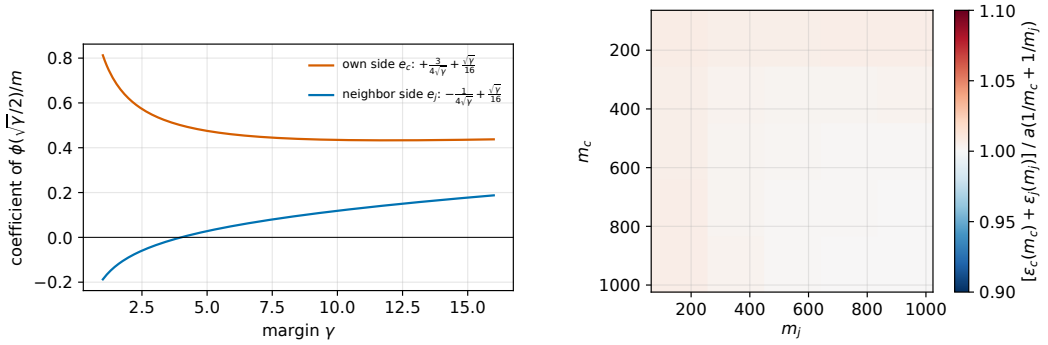

Figure 7: Left: the asymmetric per-side split behind Proposition 16; through the $(d-1)$ channel, own-side noise burdens the estimated class three times as much as it relieves the neighbor; the plotted coefficients of $\phi(\sqrt{\gamma}/2)/m$ are the $d = 2$ slice, on which the neighbor-side coefficient changes sign at $\gamma = 4$ (at general d, at $\gamma = 4(d-1)$). Right: the two-sided additivity law $a(\gamma, d)(1/m_c + 1/m_j)$ against the superposition of one-sided exact benchmark excesses at $\gamma = 6.25$, $d = 4$; the ratio is flat at one with the $O(1/m^2)$ deviation shrinking in the counts. Exact deterministic quadrature.

For the threshold: with a single dominant neighbor and $A_{cj} > B_{cj}$ (equivalently $d - 1 > \gamma_{cj}/4$), $S_c < 0$ iff $\Lambda_j(A_{cj} - B_{cj}) > \Lambda_c(3A_{cj} + B_{cj})$ iff

$$\frac{\Lambda_j}{\Lambda_c} > \frac{3A_{cj} + B_{cj}}{A_{cj} - B_{cj}} = 3 + \frac{4B_{cj}}{A_{cj} - B_{cj}} \; > \; 3, \tag{34}$$

decreasing to 3 as $A_{cj}/B_{cj} \to \infty$. Below the flip $d - 1 = \gamma_{cj}/4$, $A_{cj} < B_{cj}$, both bracket terms are positive, and every promotion is a public good to both sides. Numerically, at $(\gamma, d) = (4, 16)$: $A = 0.45370$, $B = 0.03025$, threshold $= 1.39133/0.42345 = 3.286$, against the independently measured per-side Monte Carlo ratios $1.3901/0.4183 = 3.323$ at $m = 512$ and $1.3898/0.4208 = 3.303$ at $m = 2048$ (independent verification runs, seed 11, non-canonical), the difference being the $O(1/m)$ residual in the per-side measurements. □

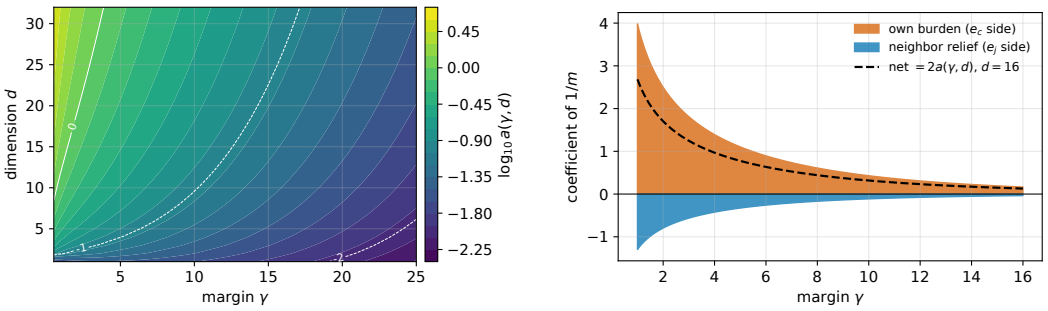

Figure 8: Left: the stocking constant as a heatmap, $\log_{10} a(\gamma, d)$ with white decade contours; the near-vertical contours at small γ are the exponential decay in the margin and their rightward drift with d is the linear $(d - 1)$ factor. Right: the per-side split as a ledger at $d = 16$, own-side burden above zero, neighbor-side relief below zero, dashed net exactly $2a(\gamma, d)$, the coefficient of the two-sided excess $2a/m$; the relief stays negative over the plotted range because the curvature term overtakes it only at $\gamma = 4(d - 1) = 60$. Deterministic closed forms.

Estimation remarks (from Section 7)

Remark 3 (Validity domain). The expansion is an $m \to \infty$ asymptotic at fixed (d, γ) whose relative remainder scales like $d/(\gamma m)$ at small margins and like d/m at large ones, so the domain is $m \gg d \max(1, 1/\gamma)$: at $\gamma = 0.25$, $d = 8$ the asymptote is 31 percent off at $m = 4d$, and at large

margins it is off by tens of percent at $m\gamma/d \approx 32$ with $m < d$ (Appendix J.4). The continual-incremental regime $m < d$ sits outside the domain at every margin; there the exact one-dimensional integral representation (Appendix H) replaces the asymptote, and the registered runs map the region (Figure 2). Comparative statics: for $d \geq 2$, a falls in γ through ϕ (separated classes need few exemplars; at $d = 1$ it rises up to $\gamma = 4$) and grows linearly in d (high-dimensional heads are estimation-hungry, so patient agents in high dimension consolidate more exemplars per class, not fewer classes).

The constant's two terms split the bill: the $(d-1)$ term comes from errors in the direction of the estimated difference between the class means, and the $\sqrt{\gamma}$ term from errors in its length (Proposition 16, above). Through the direction term, the noise in a class's own mean costs the class itself three quarters and helps its neighbor by one quarter; that quarter is the boost the neighbor tax takes away.

E PROOFS OF THE ONLINE THEOREMS

This appendix proves Theorems 1, 2, and 3, states and proves the randomized bound (Theorem 17) and the hardness result (Proposition 18), and states the discounted-flow conjecture. Where numbers from the released verification instruments are quoted they are labeled as such; they are non-canonical and never enter a claim. Figure 9 computes the tightness family (left) and the two-phase adversary (right).

Sketch of Theorem 1. Per item, $\mathrm{OPT}_i = \min(r_i Q_i, B_i^{\min})$, because every price the optimum can face lies in $[B_i^{\min}, B_i^{\mathrm{pes}}]$. If break-even never buys i, its rent stayed below $B_i^{\min}$, so $\mathrm{BE}_i = r_i Q_i = \mathrm{OPT}_i$. If it buys at query q_i, the trigger gives $r_i Q_i \geq r_i q_i \geq B_i^{\min}$, so $\mathrm{OPT}_i = B_i^{\min}$, and $\mathrm{BE}_i < B_i^{\min} + r_i + B_i^{\mathrm{pes}} \leq (1+\chi_i)\mathrm{OPT}_i + r_i$. Lower bound: a cheap forcer obliges any competitive policy to buy it first, inflating item 1's price to χB; stopping the stream right after the policy buys item 1 at accumulated rent T leaves it a ratio of at least $(T+\chi B)/(\min(T+r, B)+\varepsilon')$, since the optimum rents item 1 when T is small and bought it at B before the forcer when T is large; the infimum over T is $(1+\chi) - o(1)$. □

Sketch of Theorem 2. (a) Stage 1 floods k for $Q^* = \lceil (\Gamma+1)B/r \rceil$ queries: a policy that never buys already loses against an optimum that bought k at time zero. A policy that buys holds k irrevocably through stage 2, a flood of v, where each query costs at least $\min(r_v, \varphi) > 0$ (rent if v is outside, the crowding flow if inside), while OPT bought v at time zero and rented k: $\mathrm{OPT} \leq B + rQ^*$ is fixed and the stage-2 length is free. (b) On the common prefix the randomized policy either fails to buy k with probability above $1/2$, in which case the pure flood suffices, or buys with probability at least $1/2$, in which case the composed oblivious stream inherits the stage-2 crowding on that event; restricting a nonnegative cost to the event only lowers the expectation. Appendix E. □

Sketch of Theorem 3. Configurations $s \subseteq [N]$ with service cost $c_q(s) = r_q \mathbf{1}[q \notin s] + \sum_{\{q,k\} \subseteq s} \varphi_{qk} \geq 0$ form a metrical task system under the weighted-Hamming $\hat{d}$, on which the work-function algorithm is $(2n-1)$-competitive with $n = 2^N$ (Chrobak & Larmore, 1998); $(2n-1)$ is the optimal deterministic ratio for n-state systems (Borodin et al., 1992). Its realized cost never exceeds its $\hat{d}$-cost (true additions and removals cost at most $\max(C_i, c_{\mathrm{dem},i})$), and any schedule from $\emptyset$ has each item's removals preceded by and alternating with its additions, so its $\hat{d}$-movement is at most twice its true movement, which bounds $\mathrm{OPT}_{\hat{d}}$ and chains to the constant (a pairing lemma that needs no lower bound on the prices). Appendix E. □

E.1 THEOREM 1: $(1+\chi)$-COMPETITIVENESS UNDER LUMP PRICES

Proof. Fix a finite query stream and write Q_i for item i's total queries. Attribute to each item its rent payments plus its buy payment (if any); both break-even (BE) and the optimum decompose over items this way, with the coupling entering only through the price each buyer faces. Throughout, only Assumption A1 is used: every price $B_i(W)$ lies in $[B_i^{\min}, B_i^{\mathrm{pes}}]$; monotonicity of $B_i(\cdot)$ is nowhere assumed, and the shielding configuration of Remark 1 is a legitimate price system under this proof.

Table 3: Comparison of the two crowding accounting frameworks. Under discounting ($\beta < 1$) and stationary demand, they agree up to each newcomer's initial entry crowding (Proposition 8); against an undiscounted adversary ($\beta = 1$), they yield divergent guarantees.

	Lump: crowding charged once	Flow: crowding charged per query
Rent	r_i per query outside W	r_i per query outside W
Promotion Price	$B_i(W) > 0$, bounded, dynamic (Assumption A1)	Fixed $C_i > 0$
Post-Promotion Query	Free	$\ell\,\Delta E_q(W) \geq 0$; zero when isolated
Governing Parameter	$\chi = \max_i B_i^{\mathrm{pes}}/B_i^{\min}$	$\varphi = \ell\,\Delta E_v(\{v,k\})$ on a single pair
Best Bound (One-Way)	$1+\chi$ deterministic (Thm. 1); $\chi + \frac{1}{e-1}$ randomized (Thm. 17)	None, with or without randomization (Thm. 2)
With Demotion	Not studied here	Bounded for fixed N (Thm. 3)
Design Rule	Budget at the uncrowded price	Maintain demotion capability

The OPT lower bound (separable). If OPT never buys i it pays r_iQ_i on i; if it buys i, its price is at least the infimum over reachable sets, $B_i(W_{\mathrm{OPT}}) \geq B_i^{\min}$. Hence, writing $\mathrm{OPT}_i := \min(r_iQ_i, B_i^{\min})$,

$$\mathrm{OPT} \;\geq\; \sum_i \mathrm{OPT}_i. \tag{35}$$

The BE upper bound, three cases. BE buys i at the first query where accumulated rent reaches $B_i^{\min}$, paying the true current price $B_i(W) \leq B_i^{\mathrm{pes}} = \chi_i B_i^{\min}$.

First, a constraint that deletes a case: *if BE buys i, then $r_iQ_i \geq r_iq_i \geq B_i^{\min}$*, so $\mathrm{OPT}_i = B_i^{\min}$ *whether or not the optimum itself buys*. The situation "BE buys, OPT rents" is therefore vacuous, and the realizable cases are:

(i) *BE never buys, OPT rents:* identical costs.

(ii) *BE never buys, OPT buys:* BE's accumulated rent stayed below the trigger, so $\mathrm{BE}_i = r_iQ_i < B_i^{\min} \leq \mathrm{OPT}_i$.

(iii) *BE buys:* $\mathrm{OPT}_i = B_i^{\min}$, and $\mathrm{BE}_i < (B_i^{\min} + r_i) + B_i^{\mathrm{pes}} = (1+\chi_i)B_i^{\min} + r_i \leq (1+\chi_i)\mathrm{OPT}_i + r_i$.

Summing: $\mathrm{cost}_{\mathrm{BE}} \leq (1+\chi)\cdot \mathrm{OPT} + \sum_i r_i$, the additive term being the standard one-query discretization slack. At $\chi = 1$ this is the classical deterministic bound. One could instead wait for the highest possible price; in the Gaussian instantiation both bounds can be computed beforehand using the pairwise crowding limits of Proposition 9(c), and the theorem (and the argument above) proves that waiting for the lowest price is superior.

The pessimistic variant. Triggering at B_i^{pes} instead runs the same three cases with case (iii) giving $< 2B_i^{\mathrm{pes}} + r_i \leq 2\chi_i\mathrm{OPT}_i + r_i$, hence 2χ; the family below drives it to 2χ exactly, so budgeting at the worst price costs the factor $2\chi/(1+\chi)$ and buys no guarantee in return. □

Universal deterministic lower bound (optimality of $1+\chi$). Two items: item 2 is a forcer with $B_2 = \varepsilon' \ll B$ and rent $r_2 = \varepsilon'$ (any $r_2 = o(B)$ keeps the additive slack $\sum_i r_i$ at $o(\mathrm{OPT})$); item 1 has $B_1(\emptyset) = B$ and $B_1(\{2\}) = \chi B$. Let A be a deterministic one-way policy with ratio at most Γ on every stream. (i) Flood item 2; A must buy it (else its cost grows as r_2Q against ε'), and cannot buy item 1 during the flood (else the stream stops and the ratio exceeds Γ). (ii) With item 2 held, item 1's price is χB. Query item 1 repeatedly and stop after A buys it. If A buys after accumulating rent T on item 1, it pays at least $T + \chi B$ on this segment, while the clairvoyant optimum pays at most $\min(T+r, B) + \varepsilon'$: it rents item 1 throughout when T is small, or buys it at B before the forcer arrives when T is large. The ratio is at least $(T+\chi B)/(\min(T+r, B)+\varepsilon')$, which decreases in T while $T + r \leq B$ and increases beyond, so its infimum is attained at $T = B - r$ and equals $(1+\chi) - o(1)$. Hence no deterministic one-way policy achieves $(1+\chi) - \varepsilon$.

Tightness for BE. Two items; item 2 is a cheap forcer; item 1 has $B_1(\emptyset) = B$ and $B_1(\{2\}) = \chi B$. The adversary queries item 2 once (BE buys at ε), then queries item 1 exactly $\lceil B/r \rceil$ times. BE rents

item 1 up to the trigger B and then buys at the incumbent price χB, paying $(1+\chi)B + O(\varepsilon + r)$, while OPT buys 1 first at B and 2 at ε, paying $B + \varepsilon$ with no rent. The ratio tends to $1 + \chi$; the pessimistic variant, queried $\lceil \chi B/r \rceil$ times, pays $2\chi B + O(\varepsilon + r)$ on the same family. Figure 9 (left, solid and dashed) computes both.

E.2 Theorem 2: no bounded ratio under flow accounting

Proof. Fix the flow model at $N = 2$ with items v, k, rents $r > 0$, buys $B > 0$, and $\varphi = \ell\, \Delta E_v(\{v, k\}) > 0$; crowding is nondecreasing in the consolidated set and $\Delta E_i(\{i\}) = 0$ (a class alone in the head is uncrowded).

(a) Deterministic policies. Let $\Gamma \geq 1$ and let A be any deterministic one-way policy. Stage 1 is the stream σ_1 of

$$Q^* = \lceil (\Gamma + 1)B/r \rceil \tag{36}$$

queries to k. If A has not bought k by the end of σ_1, then $\mathrm{cost}_A(\sigma_1) \geq rQ^* > \Gamma B \geq \Gamma \cdot \mathrm{OPT}(\sigma_1)$, since OPT buys k at time zero and pays B with all subsequent k-queries free ($\Delta E_k(\{k\}) = 0$). Done.

Otherwise A first holds k after some query $q \leq Q^*$, and since A is deterministic its behavior on the common prefix of length q is fixed; let σ be q queries to k followed by Q_v queries to v. Irrevocability keeps k in A's set throughout stage 2. Each stage-2 query costs A at least $\min(r, \varphi) > 0$: either $v \notin A$'s set and the query pays rent r, or $v \in A$'s set, which then contains $\{v, k\}$, and the query costs $\ell\, \Delta E_v(W) \geq \ell\, \Delta E_v(\{v, k\}) = \varphi$ by monotonicity. Hence $\mathrm{cost}_A(\sigma) \geq \min(r, \varphi)\, Q_v$. Meanwhile

$$\mathrm{OPT}(\sigma) \leq B + r\, Q^* : \tag{37}$$

buy v at time zero (its queries are then uncrowded and free) and rent k's q queries, never buying k. Choosing

$$Q_v > \frac{\Gamma(B + r\, Q^*)}{\min(r, \varphi)} \tag{38}$$

gives $\mathrm{cost}_A(\sigma) > \Gamma \cdot \mathrm{OPT}(\sigma)$.

(b) Randomized policies. Let A be randomized and suppose, for contradiction, that it is Γ-competitive in expectation against oblivious streams. Let $Q_1 = \lceil 2\Gamma B/r \rceil$ and consider first the pure stream of Q_1 queries to k. There, $\mathrm{OPT} = B$, so $\mathbb{E}[\mathrm{cost}_A] \leq \Gamma B$. Dichotomy on the buy probability: if $\Pr[\text{no } k\text{-buy within } Q_1] > 1/2$, this pure stream alone gives $\mathbb{E}[\mathrm{cost}_A] \geq \frac{1}{2} r Q_1 \geq \Gamma B = \Gamma \cdot \mathrm{OPT}$ (up to the ceiling), and we are done; so assume $\Pr[\text{no } k\text{-buy}] \leq \frac{1}{2}$ and pass to the composed stream: the same Q_1 queries to k, then Q_v queries to v. The stage-1 prefix is identical, so A's stage-1 behavior has the same law, and with probability at least $1/2$ the event $E =$ "k bought by the end of stage 1" holds. On E, irrevocability keeps k in A's set through stage 2 and the pointwise argument of part (a) applies: stage-2 cost at least $\min(r, \varphi) Q_v$. Costs are nonnegative, so restricting the expectation to E only decreases it:

$$\mathbb{E}[\mathrm{cost}_A] \geq \Pr(E) \cdot \min(r, \varphi)\, Q_v \geq \tfrac{1}{2} \min(r, \varphi)\, Q_v,$$

while $\mathrm{OPT} \leq B + rQ_1$ is fixed. Taking $Q_v > 2\Gamma(B + rQ_1)/\min(r, \varphi)$ yields $\mathbb{E}[\mathrm{cost}_A] > \Gamma \mathrm{OPT}$: contradiction. The stream is oblivious (fixed in advance; no adaptation to A's coins).

Realizability. Two classes at whitened margin γ under the composed protocol give the instance concretely: the pairwise crowding at the balanced offset is the envelope of Proposition 9(c), $\varphi = \ell\, \Phi(-\sqrt{\gamma}/2) > 0$; the singleton value is zero ($\Delta E_i(\{i\}) = 0$); monotonicity is region containment; and the rent $r = c_{\mathrm{ret}} + \ell\, p(e_{\mathrm{miss}} - e_{\mathrm{W}})$ is strictly positive by Proposition 10. Numbers: at $\gamma = 4$, $\ell = 1$, the crowding flow is $\varphi = \Phi(-1) = 0.1587$ per v-query once both classes are held; with $c_{\mathrm{ret}} = 0.05$, $p = 0.1$, $e_{\mathrm{miss}} = 0.5$ the rents are equal and positive, $r \approx 0.10$. And under flow accounting the buy prices are the constant C_{con}, so the price coupling ratio is *exactly* 1: the impossibility holds with no price coupling at all, the liability being created purely by the flow. The impossibility is therefore not an artifact of strong coupling or of adversarial set functions, and the discontinuity at $\varphi = 0$ (where the classical guarantees hold exactly) is sharp. □

Corollaries. (a) *Discontinuity at $\varphi = 0$:* the classical guarantee holds at $\varphi = 0$ and nothing holds for any $\varphi > 0$, even at coupling ratio $\chi = 1$; under flow accounting the coupling ratio

does not parameterize difficulty at all. (b) *No repair:* the theorem quantifies over all one-way policies, so no modified trigger, externality surcharge at buy time, or randomization rescues one-way promotion; structurally, an irrevocable buy creates a liability denominated in another item's unbounded future demand, and no finite charge covers an adversarially long stream. (c) *Demotion necessity:* jointly with the capitalization lemma, a coupled online theory with worst-case guarantees must price interference as a flow and allow demotion.

Instrument note (non-canonical). The released instrument for Theorem 2 ran 269 one-way policies (a 13×13 threshold grid including the pessimistic-flow transplant, plus 100 random policies) against the two-phase adversary at $r = 1, B = 5, \varphi = 0.5$: the minimum over policies of the worst-case ratio grew 2.01, 5.81, 48.7, 400 as the stage-2 length swept 10^2 to 10^5; the naive pessimistic-flow transplant alone measured 5.0, 35.0, 335.0. Figure 9 recomputes the threshold-policy family deterministically.

E.3 Theorem 3: demotion suffices at fixed N

Proof. (i) The configuration system is a metrical task system after symmetrization. Timing: at each step the query is revealed, the policy may move, and the query is served in the new configuration, as in metrical task systems. Configurations are the subsets $s \subseteq [N]$; a query to item q presents the task vector

$$c_q(s) = r_q\,\mathbf{1}[q \notin s] + \sum_{\{q,k\}\subseteq s} \varphi_{qk} \;\geq\; 0,$$

rent if the queried item is unowned, plus the pairwise crowding flows on every owned conflicting pair; nonnegativity makes these legitimate MTS tasks. The true movement cost $d(s \to s') = \sum_{i\in s'\setminus s} C_i$ (additions priced, removals free, or priced $c_{\mathrm{dem},i}$ in the costly variant) is asymmetric; the weighted Hamming

$$\hat{d}(s, s') = \sum_{i \in s \triangle s'} w_i, \qquad w_i = \max(C_i, c_{\mathrm{dem},i}), \tag{39}$$

is a genuine metric (symmetric, and the triangle inequality is inherited from the ℓ_1 structure on the hypercube).

(ii) The work-function bound. On an n-state metric task system the work-function algorithm, which moves at each step to a minimizer s of $w_t(s) + \hat{d}(s_{t-1}, s)$ whose work function paid the current task ($w_t(s) = w_{t-1}(s) + c_t(s)$), is $(2n-1)$-competitive (Chrobak & Larmore, 1998), which meets the optimal deterministic ratio for n-state systems (Borodin et al., 1992): its $\hat{d}$-cost satisfies $\mathrm{cost}_{\hat{d}}(\mathrm{WFA}) \leq (2n-1)\mathrm{OPT}_{\hat{d}} + c_0$ with c_0 stream-independent (of the order of the metric's diameter); here $n = 2^N$ and $2n - 1 = 2^{N+1} - 1$.

(iii) Realized cost at most $\hat{d}$-cost. WFA is executed on $\hat{d}$, and its realized charges are the true movement plus service: every addition costs $C_i \leq w_i$ and every removal costs at most $c_{\mathrm{dem},i} \leq w_i$, so step by step the true bill never exceeds the $\hat{d}$ bill; service costs are identical in the two accountings.

(iv) Pairing lemma: $\mathrm{OPT}_{\hat{d}} \leq 2\max_i \max(1, c_{\mathrm{dem},i}/C_i) \cdot \mathrm{OPT}$. Take the true optimum's schedule, which starts at $\emptyset$. For each item i, removals of i can occur only after an addition of i that has not yet been removed: along the schedule the removals of each item alternate with, and are preceded by, its additions, so $\#\mathrm{rem}(i) \leq \#\mathrm{add}(i)$. The schedule's $\hat{d}$-movement is therefore

$$\begin{aligned}\sum_i \big(\#\mathrm{add}(i) + \#\mathrm{rem}(i)\big) w_i \;&\leq\; 2\sum_i \#\mathrm{add}(i)\, w_i \;\leq\; 2\max_i \frac{w_i}{C_i} \sum_i \#\mathrm{add}(i)\, C_i \\ &\leq\; 2\max_i \max(1, c_{\mathrm{dem},i}/C_i) \cdot (\text{true movement}),\end{aligned}$$

and its true movement also dominates its true demotion charges, while service costs coincide under both accountings. Running the same schedule in the $\hat{d}$ system therefore certifies $\mathrm{OPT}_{\hat{d}} \leq 2\max_i \max(1, c_{\mathrm{dem},i}/C_i) \cdot \mathrm{OPT}$. Note the lemma needs no lower bound on the buy prices: the alternation counting is pure combinatorics of schedules that start at $\emptyset$.

(v) Chain.

$$\begin{aligned}\mathrm{cost}_{\mathrm{true}}(\mathrm{WFA}) \;&\le\; \mathrm{cost}_{\hat{d}}(\mathrm{WFA}) \;\le\; (2^{N+1}-1)\mathrm{OPT}_{\hat{d}} + c_0 \\ &\le\; 2(2^{N+1}-1)\max_i \max(1, c_{\mathrm{dem},i}/C_i)\,\mathrm{OPT} + c_0,\end{aligned}$$

as claimed. At $N = 2$ the multiplicative constant is $2(2^3 - 1) = 14$. □

Notes. (a) No positive lower bound on the buy prices is needed for the symmetrization: the pairing lemma uses only start-at-$\emptyset$ alternation. (b) The free-demotion case $c_{\mathrm{dem}} = 0$ recovers the factor exactly as stated. (c) Combined with Theorem 2: at every fixed N, one-way policies are unbounded and two-way policies are $O(1)$-competitive, so reversibility, not the accounting convention, is the binding constraint; what the accounting controlled was the one-way world, where lump restores $1 + \chi$ and flow destroys everything. (d) The simulation evidence is E-L1-8 (Section L.8): realized WFA ratios 3.95–4.79 and FBE ratios at most 2.97 on the designed streams against the bound 14; the $\hat{d}$-cost ratios that the theorem directly bounds reach at most 5.78 across the full registered grid, comfortably inside $14+$ the additive allowance; and on the adversarial attack rows (Amendment A7) WFA stays at or below 1.48 while the simple policies blow up, the empirical content of the necessity note in Section 4.

Remark 4 (Randomization: polynomial in N against oblivious streams). Keep the configuration system, the metric $\hat{d}$, the start state $\emptyset$ and the metrical-task-system timing of the proof of Theorem 3 (a query is revealed, the policy may move, and the query is served in the new configuration), and replace the work-function algorithm by the randomized metrical-task-system algorithm of Bubeck et al. (2021). Against every oblivious stream (fixed in advance, independently of the policy's coins),

$$\mathbb{E}[\mathrm{cost}_{\mathrm{true}}] \;\le\; O(N^2)\max_i \max(1, c_{\mathrm{dem},i}/C_i)\cdot \mathrm{OPT} + c_1,$$

with an absolute constant in $O(\cdot)$, OPT the two-way clairvoyant optimum, and c_1 independent of the stream. The exponential factor of Theorem 3 therefore comes from the deterministic $(2n-1)$ bound: against oblivious streams, randomization makes the ratio at most quadratic in N, although the policy still computes over 2^N states. The constant in $O(\cdot)$ is not explicit, so nothing here says the randomized bound beats Theorem 3's at small N. Streams that react to the policy's configuration are not covered, and whether a deterministic policy can avoid the exponential is open.

Proof. On every n-point metric and from every start state, Bubeck et al. (2021) give a randomized algorithm with $\mathbb{E}[\mathrm{cost}] \le O(\log^2 n)\,\mathrm{OPT} + \beta$ on every fixed sequence of nonnegative task vectors, OPT being the offline optimum from the same start and β independent of the sequence. An oblivious stream fixes the tasks $c_{q_t} \ge 0$ of step (i); run from $\emptyset$, with $\log^2 2^N = O(N^2)$, the algorithm has $\mathbb{E}[\mathrm{cost}_{\hat{d}}] \le O(N^2)\,\mathrm{OPT}_{\hat{d}} + \beta$. Step (iii) compares the true and $\hat{d}$ charges of an arbitrary trajectory, so it holds on every realization and hence in expectation; step (iv) involves only the optimum. Chaining, $\mathbb{E}[\mathrm{cost}_{\mathrm{true}}] \le \mathbb{E}[\mathrm{cost}_{\hat{d}}] \le 2\max_i \max(1, c_{\mathrm{dem},i}/C_i)\,O(N^2)\,\mathrm{OPT} + \beta$; take $c_1 = \beta$. □

The registered implementation. The E-L1-8 runner breaks ties toward moving, then by state index (Amendment A6), which is not the classical rule of step (ii), so the registered runs measure that implementation, not the theorem's algorithm. Outside the registered grid the A6 rule can exceed the bound, which the classical rule cannot; a non-canonical check ships as `runners/audit/wfa_tie_rule_check.py`.

E.4 The flow break-even heuristic fails, and the open directions (from Section 4)

A natural heuristic rescue is a meter: accumulate the crowding incurred by each item in the head and demote it once the accumulated penalty reaches its price (*flow break-even*, FBE, paired with a fixed blame allocation rule). However, this heuristic fails. Prior analysis identified two distinct failure modes, which were pre-registered along with their predicted defeat before empirical execution (Amendment A7). Under *overshoot*, a single highly crowded query incurs a cost far exceeding the item price ($\varphi \gg B$), inflicting irreversible damage before the meter can react. Under *ping-pong*, low refill costs trigger repeated re-purchases of the same item while an optimal schedule retains it

indefinitely. Attempting to repair overshoot via early demotion simply exacerbates ping-pong failures. On the registered adversarial test streams of experiment E-L1-8, standard FBE incurs 60 to 180 times the optimal cost (with the early-demotion variant reaching up to 80), whereas the registered work-function execution remains at or below 1.48, well within the theoretical bound of 14 for $N = 2$. This execution utilizes a tie-breaking rule differing from the classical algorithm analyzed in the theorem (Amendment A6, above).

This performance guarantee comes at the cost of memory overhead: the work-function algorithm must maintain a history of past conflicts. Two open questions remain. First, does there exist a memory-efficient two-way policy that retains strong theoretical guarantees? The proposed attack suite serves as a natural testbed for this question. Second, do deterministic policies inherently require an exponential dependence on N? With randomization against oblivious query sequences, the algorithm of Bubeck et al. (2021) over the same 2^N states achieves an expected cost bounded by $O(N^2) \max_i \max(1, c_{\mathrm{dem},i}/C_i) \cdot \mathrm{OPT}$ plus a constant (Remark 4).

Selecting break-even trigger thresholds independently and randomly per item according to the classical distribution (Karlin et al., 1994) at price $B_i^{\min}$ guarantees $\mathbb{E}[\mathrm{cost}] \le (\chi + \frac{1}{e-1})\mathrm{OPT} + \sum_i r_i$ against predetermined query sequences (Theorem 17, below); whether this factor is tight remains open. Computing the optimal item set with complete future foresight is NP-hard under quadratic programming formulations (Proposition 18, below); determining whether this hardness persists for instances generated by the Gaussian model is an open direction. Flow accounting with discounting ($\beta < 1$), where discounting bounds debt accumulation, represents the sole regime left unresolved by the current dichotomy (Conjecture 19). While the structural split holds universally across price functions, determining the exact pricing dynamics of actual network heads represents the immediate next step.

E.5 Randomized break-even at the uncrowded price

Theorem 17 (Randomized break-even)**.** *Draw, independently per item, a threshold $T_i = z_i B_i^{\min}$ with z_i of density $e^z/(e-1)$ on $[0,1]$, and buy item i at the current price when its accumulated rent reaches T_i. Against any oblivious stream, $\mathbb{E}[\mathrm{cost}] \le \left(\chi + \frac{1}{e-1}\right) \cdot \mathrm{OPT} + \sum_i r_i$. At $\chi = 1$ this is the classical $\frac{e}{e-1}$; the coupling enters additively. The matching lower bound is open.*

Proof. Fix an item, write $p = B_i^{\min}$ and $x = r_i Q_i / p$. If $x \ge 1$ the policy buys with probability one, paying expected rent $\frac{p}{e-1}$ and at most $\chi_i p$ for the purchase, while $\mathrm{OPT}_i = p$; the ratio is $\frac{1}{e-1} + \chi_i$. If $x < 1$ the policy buys iff $z \le x$ (probability $(e^x - 1)/(e-1)$), and the expected cost is at most $p\,[xe + (\chi_i - 1)(e^x - 1)]/(e-1)$ against $\mathrm{OPT}_i = xp$; the ratio is bounded by $\chi_i + \frac{1}{e-1}$ because $(e^x - 1)/x \le e - 1$ on $(0, 1]$. Summing over items gives the theorem. □

The open lower bound. The deterministic construction does not lift verbatim to randomized policies. Whether the optimal randomized ratio equals $\chi + \frac{1}{e-1}$ or is strictly smaller is open.

E.6 NP-hardness of flow-priced selection

Proposition 18 (NP-hardness of flow-priced selection)**.** *Minimizing the capitalized flow objective $\sum_i (C_i - R_i) x_i + \sum_{i<j} M_{ij} x_i x_j$ over $x \in \{0,1\}^N$ with $M_{ij} \ge 0$ is NP-hard. The hardness is of the abstract quadratic objective; hardness restricted to instances realizable inside the Gaussian ledger is open. The objective is supermodular, so nonmonotone submodular maximization approximations transfer for nonnegative objectives (Feige et al., 2011; Buchbinder et al., 2015).*

Proof. Reduction from maximum independent set. Given a graph $G = (V, E)$ on N vertices, set $C_i - R_i = -1$ for all i and $M_{ij} = 2$ on edges, 0 otherwise. The objective becomes

$$f(x) = -|x| + 2\,|\{(i,j) \in E : x_i = x_j = 1\}|. \tag{40}$$

If a solution x contains both endpoints of some edge, dropping one endpoint u changes the cost by

$$f(x - e_u) - f(x) = -(|x| - 1) - 2\,e(u) + |x| = 1 - 2\,e(u) \le -1 < 0, \tag{41}$$

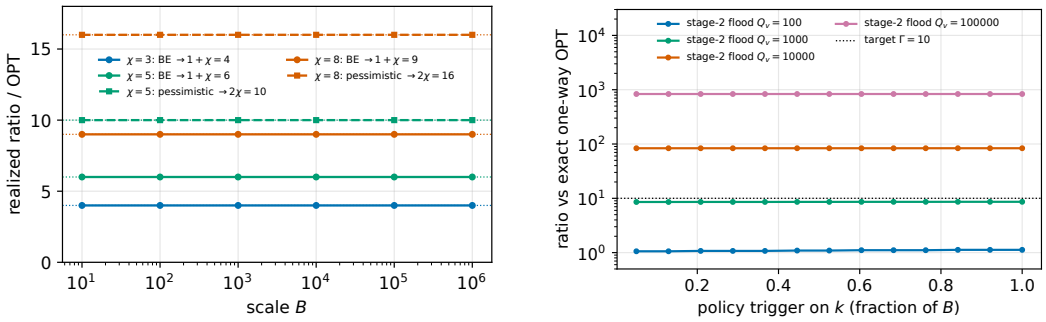

Figure 9: Left: the tightness family of Theorem 1, a cheap forcer then a flood of item 1 up to the uncrowded-price trigger, drives the realized BE/OPT ratio to $1 + \chi$ at three coupling levels (at scale $B = 10^6$ the exact family values are 4.0000 of 4, 6.0000 of 6, 9.0000 of 9; Algorithm 9); the pessimistic variant on the same family (dashed, drawn at $\chi = 5$ and 8) reaches 2χ. Right: Theorem 2 on a policy grid: Algorithm 10 at $\Gamma = 10$ ($r = 1$, $B = 5$, $\varphi = 0.5$, stage 1 of $Q^* = 55$ queries to k), exact ratios against the one-way optimum for threshold policies on item k, one curve per stage-2 flood length (log scale); every such policy buys k inside stage 1, inherits the crowding flow, and its ratio grows linearly in the flood, crossing Γ (dotted) between $Q_v = 10^3$ and 10^4; the escape of not buying is closed by stage 1 itself. Deterministic recomputation of the constructions.

where $e(u) \geq 1$ counts selected edges at u: every non-independent solution is strictly dominated by a smaller one, so some optimum is an independent set. On independent sets the edge term vanishes and $f(x) = -|x|$, so the minimum equals $-\alpha(G)$, the negative of the independence number; computing it is NP-hard.

The objective is supermodular ($M_{ij} \geq 0$ gives the pair term increasing differences), so maximizing its negative is nonmonotone submodular maximization in disguise, with the known constant-factor approximations for the unconstrained problem transferable (Feige et al., 2011; Buchbinder et al., 2015). Geometric realizability of arbitrary instances inside the Gaussian ledger needs an embedding in which designated pairs are near and all others far; the sphericity literature (Maehara, 1984) supplies the existence direction (graphs realized as threshold-distance graphs in some dimension), which is the only direction used, and the remark is stated existence-only. □

E.7 The open regime: discounted flow

The one regime the split leaves open is discounted flow, where $\beta < 1$ caps the externality liability. Define the per-round attributable crowding cap $d_i^{\max} := \sup_W \sup_j \Delta E_j(W \cup \{i\}) - \Delta E_j(W)$ (the largest pairwise marginal crowding increment item i can impose, closed form in the instantiation by Proposition 9), the capitalized pessimistic price $B_i^{\text{cap}} = C_i + d_i^{\max}/(1 - \beta)$, and $\chi_\beta = \max_i B_i^{\text{cap}}/C_i$.

Conjecture 19 (Discounted flow)**.** *The break-even rule on the capitalized pessimistic ledger is $c(\chi_\beta)$-competitive in discounted cost, with c linear (target $2\chi_\beta + O(1)$). The modular attribution bound (the true per-round crowding is at most the sum of pairwise streams) is expected to be provable with present techniques. Single-item discounted ski rental is solved (El-Yaniv et al., 1999); the step not yet closed is carrying that solution to buy prices that the capitalized externality makes state-dependent, where delaying a buy both deflates its discounted price and changes it.*

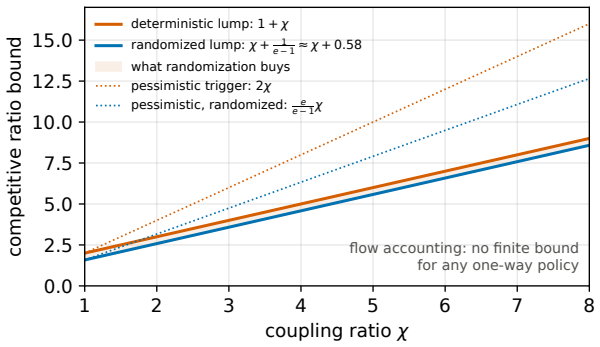

Figure 10: The two lump-accounting constants of the online split against the coupling ratio: the deterministic $1+\chi$ of Theorem 1 and the randomized $\chi+\frac{1}{e-1}$ of Theorem 17; the shaded band, of constant width $1-\frac{1}{e-1}\approx 0.42$, is what randomization buys; the dotted curves are the pessimistic-trigger variants, 2χ and $\frac{e}{e-1}\chi$. Flow accounting has no curve to draw: by Theorem 2 no one-way policy, deterministic or randomized, has a finite bound there. Deterministic closed forms.

F THE UNCOMPOSED PROTOCOL

In the uncomposed protocol, the offset-mediated calibration terms of Proposition 9(b) do not cancel. The $k \to k{+}1$ promotion of an already-consolidated class has a nonzero population effect through the closed-form $\Delta\tau_{cj} = \log((m_c{+}1)/m_c)$: the promoted class trains relatively stronger and its neighbors pay the calibration interference I_{cal}. Three things change and one does not. Changed: (i) the exemplar tier is no longer a pure public good; its buy price carries the (exponentially small in the margin, by the envelope of Proposition 9(b)) calibration interference; (ii) the optimal policy is still now-or-never (Lemma 11 is protocol free), but the threshold uses the uncomposed price; (iii) the composed protocol is not available to an evaluator who cannot subtract count offsets, e.g. a black-box head. Unchanged: the entry-tier economics, whose e_{miss} removal and Bayes crowding are protocol independent (Remark 2). The manuscript's default is the composed protocol because it is what an evaluator should do; the uncomposed variant is exactly the calibration-interference ledger above.

G ALGORITHMS

This appendix collects the algorithmic content of the theory: optimal policies, exact price evaluations, verification harnesses, and the adversarial constructions. All pseudocode mirrors the registered runner implementations (Appendix N); the runner library (`l1_lib.py`, hash-recorded in every manifest) is the executable form of this appendix. Algorithm 1 is the system-level loop the policies live in.

Algorithm 1 The consolidation loop (system-level view; the online analysis of Section 4 lets the policy act on a query before serving it)

1: each item i starts in the memory store M; the consolidated set $W \leftarrow \emptyset$; buffer counts k_r per consolidated class
2: **for** each round **do**
3: a query arrives for item i with probability λ_i
4: **if** $i \in M$ **then** pay rent r_i
5: **else** serve from the head W
6: **end if**
7: the promotion policy decides whether to promote some item z of class c from M to W (paying $B(W)$, updating $k_c \leftarrow k_c + 1$ and refitting); a two-way policy may also demote
8: **end for**

Algorithm 2 Closed-form threshold policy (Theorem 12; optimal in the decoupled regime, the greedy rule under coupled prices)

Require: per-item λ_i, r_i, B_i; discount β

for each arriving item i **do**
  if $\lambda_i r_i > (1-\beta)B_i$ **then**
    promote i at arrival (pay B_i)
  else
    keep i in memory (pay r_i per query, forever)
  end if
end for

Algorithm 3 Entry-tier promotion with exact prices (Remark 2, Proposition 8)

Require: class c arriving with demand Λ_c; incumbents W with margins γ_{cj} and demands λ_j; $\beta, \ell, C_{\mathrm{con}}$

$\Delta E_j \leftarrow \Phi(-\sqrt{\gamma_{cj}}/2)$ for each $j \in W$ ▷ Proposition 9(c), pairwise envelope

$I(c;W) \leftarrow \frac{\ell}{1-\beta}\sum_{j\in W}\lambda_j \Delta E_j$ ▷ capitalized, demand-weighted

$B_c(W) \leftarrow C_{\mathrm{con}} + I(c;W)$

if $\Lambda_c\,[\,c_{\mathrm{ret}} + \ell\, p\,(e_{\mathrm{miss}}(c) - e_{\mathrm{W}}(c))\,]/(1-\beta) > B_c(W)$ **then** promote the first exemplar of c ▷ $e_{\mathrm{W}}(c)$ is the post-entry balanced Bayes error at the incumbent margins; at the population level over matched coverage Proposition 10 evaluates it in closed form

end if

Algorithm 4 Exemplar-tier marginal rule (Corollary 13)

Require: class c at buffer count m_c, class demand Λ_c, excess function $\epsilon_c(\cdot)$ (asymptotic a_c/m or exact benchmark), $\beta, \ell, C_{\mathrm{con}}$

while $\Lambda_c \ell\,[\epsilon_c(m_c) - \epsilon_c(m_c+1)]/(1-\beta) > C_{\mathrm{con}}$ **do**
  promote one more exemplar of c; $m_c \leftarrow m_c + 1$
end while ▷ stops at $m_c^* = \Theta(\sqrt{\Lambda_c \ell a_c/((1-\beta)C_{\mathrm{con}})})$

Algorithm 5 Agent-computable promotion value (Theorem 15)

Require: candidate exemplar z of class c; empirical means $\hat{\mu}_c$; neighbors j with $\hat{\delta}_{cj} = \hat{\mu}_c - \hat{\mu}_j$

for each neighbor j **do**
  $v \leftarrow (z - \hat{\mu}_c)/(m_c+1)$
  $\Delta e_j^{\mathrm{comp}} \leftarrow -\phi(|\hat{\delta}_{cj}|/2)\,\hat{\delta}_{cj}^{\top} v/(2|\hat{\delta}_{cj}|)$ ▷ first order; agent-computable
end for
return $\sum_j \Delta e_j^{\mathrm{comp}}$ ▷ the e_c side is oracle-only: it involves the true displacement

Algorithm 6 Exact integral benchmark for $\epsilon_{\mathrm{pair}}(m)$ (Remark 3; derivation in Appendix H)

Require: pair (c, j) whitened to γ, d; count m

$E[e_j] \leftarrow \int \Phi(-\frac{1}{2}\sqrt{q/m})\, f_{\chi'^2_d(q;\,m\gamma)}(q)\, dq$ over the mean $\pm$ 12 sd window ▷ noncentral chi-square

$E[e_c] \leftarrow \iint \Phi\left(-\frac{\gamma - s^2}{2\sqrt{\gamma + s^2 + 2\sqrt{\gamma}\, s\, u}}\right) f_{\chi_d}(s\sqrt{m})\sqrt{m}\; f_U(u)\; du\, ds$ ▷ $u = \cos\theta$; f_U is the $(1-u^2)^{(d-3)/2}$ angular law, two point masses at ± 1 when $d = 1$; Gauss–Jacobi nodes at $d \geq 2$

return $\frac{1}{2}(E[e_j] + E[e_c]) - \Phi(-\sqrt{\gamma}/2)$

Algorithm 7 Finite-horizon DP verifier (E-L1-1, E-L1-2)

Require: cell (λ, r, B, β); tail tolerance
1: $H \leftarrow$ smallest doubling horizon with $\beta^H(\lambda r/(1-\beta) + B) \le 10^{-13}$
2: DP $\leftarrow \min_{0 \le t \le H} V(t)$ with $V(t) = \lambda r \frac{1-\beta^t}{1-\beta} + \beta^t B$
3: threshold cost $\leftarrow B$ if $\lambda r > (1-\beta)B$ else $\lambda r/(1-\beta)$
4: **return** gap = threshold cost − DP ▷ negative beyond tolerance falsifies Theorem 12

Algorithm 8 Break-even rule BE (Theorem 1)

Require: per-item r_i, price system $B_i(\cdot)$, uncrowded prices $B_i^{\min} = \inf_W B_i(W)$
1: **for** each query to item i outside W **do**
2: accumulated rent $\leftarrow$ accumulated rent $+r_i$; pay r_i
3: **if** accumulated rent $\ge B_i^{\min}$ **then**
4: buy i at the true current price $B_i(W)$; $W \leftarrow W \cup \{i\}$
5: **end if**
6: **end for**

Algorithm 9 Tightness adversary for $1 + \chi$ (Theorem 1)

Require: $\chi > 1$, scale B, rent r, forcer price $\varepsilon \le r$, $\varepsilon \ll B$
1: query item 2 once ▷ BE buys at ε
2: query item 1 for $\lceil B/r \rceil$ rounds ▷ BE rents to the trigger B, then buys at the true incumbent price χB
3: BE $\approx (1+\chi)B$, OPT $= B + \varepsilon$ ▷ buy 1 first at B, then 2 at ε; ratio $\to 1 + \chi$; the pessimistic variant needs $\lceil \chi B/r \rceil$ rounds and pays $\approx 2\chi B$

Algorithm 10 Two-phase adversary for flow accounting (Theorem 2)

Require: target ratio Γ; pair (v, k) with $\varphi = \ell\, \Delta E_v(\{v, k\}) > 0$; rents r, buys B
1: Stage 1: $Q^* = \lceil (\Gamma + 1)B/r \rceil$ queries to k ▷ forces any Γ-competitive policy to buy k
2: **if** k is not held after stage 1 **then**
3: stop ▷ the policy already paid $rQ^* > \Gamma B \ge \Gamma\,\mathrm{OPT}$
4: **end if**
5: Stage 2: $Q_v > \Gamma(B_v + rQ^*)/\min(r_v, \varphi)$ queries to v ▷ each costs $\ge \min(r_v, \varphi)$ once k is irrevocably held
6: OPT $\le B_v + rQ^*$ ▷ buy v at time 0, rent k

Algorithm 11 Allocation-consistency harness (E-L1-7)

Require: class demands Λ_i, excess functions $\epsilon_i(m) = a_i/m$, budget M, baseline $m_0 = 1$
1: greedy: start $m_i = m_0$; repeat $M - m_0 N$ times, add one count to the class with the largest $\Lambda_i(\epsilon_i(m_i) - \epsilon_i(m_i + 1))$
2: DP: solve $\max \sum_i \Lambda_i(\epsilon_i(m_0) - \epsilon_i(m_i))$ subject to $\sum_i m_i = M$, $m_i \ge m_0$, exactly over the count simplex
3: **return** whether objectives and allocations agree beyond 10^{-9} ▷ promotion theory equals allocation theory at the intersection

H THE EXACT INTEGRAL BENCHMARK

This appendix derives the deterministic quadrature used as the exact benchmark in E-L1-4 and E-L1-5 (Algorithm 6) and states its numerical validation. The benchmark computes $\epsilon_{\mathrm{pair}}(m) = \mathbb{E}[\frac{1}{2}(e_c(\xi) + e_j(\xi))] - \Phi(-\sqrt{\gamma}/2)$ for $\xi \sim \mathcal{N}(0, I/m)$ without sampling.

H.1 THE NEIGHBOR SIDE: A NONCENTRAL CHI INTEGRAL

Since $m|\delta+\xi|^2$ is noncentral chi-square with d degrees of freedom and noncentrality $m\gamma$,

$$\mathbb{E}[e_j] = \int_0^\infty \Phi\left(-\frac{1}{2}\sqrt{q/m}\right) f_{\chi'^2}(q;\, d,\, m\gamma)\, dq.$$

The integrand is smooth and the density is concentrated near its mean $d+m\gamma$ with standard deviation $\sqrt{2(d+2m\gamma)}$; integrating over the window mean ± 10 sd with adaptive quadrature (relative tolerance 10^{-11}) is exact to machine precision. The batch evaluator uses a fixed-order Gauss–Legendre rule on the window mean ± 12 sd; it is certified against the adaptive value only at the two E-L1-5 regimes, where it is used (validation (i) below).

H.2 THE OWN SIDE: A CHI LAW CROSSED WITH AN ANGULAR LAW

Condition on $s = |\xi|$ and $u = \cos\theta$, the cosine of the angle between ξ and δ. The error depends on the pair (s, u) only:

$$e_c = \Phi\left(-\frac{\gamma - s^2}{2\sqrt{\gamma + s^2 + 2\sqrt{\gamma}\, s\, u}}\right), \qquad s\sqrt{m} \sim \chi_d, \qquad u \sim f_U(u) = c_d\,(1-u^2)^{(d-3)/2}, \tag{42}$$

the angular law of a uniform direction in d dimensions, with the $d = 1$ degenerate case $u \in \{-1, +1\}$ with probability $1/2$ each. Hence

$$\mathbb{E}[e_c] = \int_0^\infty \int_{-1}^1 \Phi\left(-\frac{\gamma - s^2}{2\sqrt{\gamma + s^2 + 2\sqrt{\gamma}\, s\, u}}\right) f_{\chi_d}(s\sqrt{m})\sqrt{m}\; f_U(u)\; du\; ds.$$

For $d \geq 2$ the angular law carries the weight $(1-u^2)^{(d-3)/2}$ exactly in the Gauss–Jacobi family (Chebyshev at $d = 2$, Legendre at $d = 3$), so a 48-node Gauss–Jacobi product rule integrates u with exponential accuracy including the integrable endpoint singularity at $d = 2$; a fixed-order Gauss–Legendre rule over the chi window handles s. At $d = 1$ the angular integral is the two-point average. Two implementation notes recorded for reproducibility: a plain Gauss–Legendre rule in u is *not* adequate at $d = 2$ (slow convergence through the endpoint singularity) and a continuous angular law is *wrong* at $d = 1$; both were caught and fixed during smoke validation, before any canonical run.

H.3 VALIDATION

Three certificates, all pre-canonical: (i) the two quadrature routes (adaptive and batch fixed-order) agree to within 10^{-11} at the two E-L1-5 regimes $(9, 8)$ and $(16, 32)$, at six counts from $m = 2$ to $m = 2048$ (`validate_batch.py`); E-L1-4, whose grid includes $d = 1, 2$, uses the adaptive route only; (ii) against Monte Carlo with exact conditional errors, per-cell deviations sit within three standard errors at 2×10^5 draws (pooled $p = 0.700$ across the whole canonical E-L1-4 grid; the standard errors are evaluated at the exact rates rather than the empirical ones, the correct Wald form for testing against a hypothesized exact value, avoiding the degeneracy of a zero-or-all empirical rate); (iii) the benchmark's own large-m behavior reproduces the corrected constant: OLS of $m\,\epsilon$ on $1/m$ over asymptotic cells recovers $a(\gamma, d)$ to at most 0.023% relative deviation (Table 12). The benchmark is the reference E-L1-5 uses for its exact-excess variant, evaluated at consecutive counts in blocks by the vectorized rule.

I RUNNER LIBRARY REFERENCE

The executable form of Appendix G and of the experimental program is the runner library shipped with the supplement:

- `ll_lib.py`: exact formulas ($e_j, e_c, a(\gamma, d)$), the integral benchmark (adaptive and vectorized batch forms), the DP verifier, the stream generator and policies for the autopsies, the greedy and DP allocation solvers, the frozen tolerance constants, and the manifest writer (SHA-256 code hashes, seeds, grids, library versions, timestamps).

- `l1_e1.py` ... `l1_e8.py`: one runner per registered Tier 1 experiment, each writing its CSV plus manifest and printing its falsifier verdict; smoke seeds are passed explicitly and recorded as non-canonical in the manifest.
- `l1b_lib.py`, `l1b_e1.py` ... `l1b_e3.py`, `fill_l1b.py`, `fig_results_l1b.py`: the Tier 1b library, runners, fill script, and figure generator (Appendix O).
- `er1_lib.py`, `er1_run.py`, `er1_colab.py`: the Tier 2 package (Appendix P).
- `fill_l1_extended.py`, `fill_conference.py`: read the canonical CSVs and manifests and emit the macro files this paper inputs; the only path by which Tier 1 numbers enter the tex.
- `fig_gen.py`: deterministic closed-form figures (no stochastic content).
- `fig_results.py`: canonical-data figures, reading only manifest-backed CSVs.

Every figure caption in this manuscript states whether its content is an exact closed form (deterministic) or canonical data (registered run); no figure mixes the two provenances in one panel.

J Extended derivations

Five derivations that support the main text but are not proofs of named theorems: the exact $d = 1$ law that certifies the constant, the angular quadrature's justification, the second-order term behind the amended falsifiers, the validity domain, and the crowding envelope's tail bound.

J.1 The exact $d = 1$ threshold-noise law

At $d = 1$ there is no perpendicular block and the balanced pair excess has an exact one-dimensional form. With $\xi \sim \mathcal{N}(0, 1/m)$ scalar,

$$\epsilon_{\text{pair}}(m) = \tfrac{1}{2}\,\mathbb{E}\big[\Phi\Big(-\tfrac{\sqrt{\gamma}+\xi}{2}\Big) + \Phi\Big(-\tfrac{\sqrt{\gamma}-\xi}{2}\Big)\,\big] - \Phi\Big(-\tfrac{\sqrt{\gamma}}{2}\Big)\,.$$

Write $g(\xi)$ for the bracketed average; g is even in ξ, so its first derivative vanishes at zero and the Taylor expansion in ξ starts quadratically:

$$\mathbb{E}[g(\xi)] - g(0) = \tfrac{1}{2} g''(0)\,\mathbb{E}[\xi^2] + O(\mathbb{E}[\xi^4]), \qquad \mathbb{E}[\xi^2] = 1/m.$$

Computing $g''(0)$: each side contributes $\frac{1}{4}\Phi''(-\sqrt{\gamma}/2) = \frac{\sqrt{\gamma}}{8}\phi(\sqrt{\gamma}/2)$, and averaging the two equal sides gives

$$\epsilon_{\text{pair}}(m) = \frac{\sqrt{\gamma}\;\phi(\sqrt{\gamma}/2)}{16}\,\frac{1}{m} + O(m^{-2}), \tag{43}$$

which is exactly the corrected $a(\gamma, 1)$: the $(d-1)$ channel dies and the curvature channel survives. The retracted constant's $d/4\sqrt{\gamma}$ term would leave a spurious $\phi/(4\sqrt{\gamma}\,m)$ piece at $d = 1$, which this exact law excludes; this is the special-case certificate that first discriminated the two constants, and E-L1-4's extrapolation re-verifies it through the canonical pipeline (Table 12, γ rows at $d = 1$).

J.2 Why the angular rule is Gauss–Jacobi

The angle between an isotropic Gaussian vector and a fixed direction has density $f_U(u) = c_d(1 - u^2)^{(d-3)/2}$ on $[-1, 1]$: for $d = 3$ uniform, for $d = 2$ the arcsine density $(1 - u^2)^{-1/2}/\pi$ with integrable endpoint singularities, and for $d = 1$ the degenerate two-point law $u \in \{-1, +1\}$ half each. Expectations against this density are exactly the Gauss–Jacobi family with parameters $\alpha = \beta = (d-3)/2$:

$$\int_{-1}^{1} h(u)(1-u)^{\alpha}(1+u)^{\alpha} du \approx \sum_{i=1}^{n} w_i\, h(u_i), \tag{44}$$

with exponential convergence for analytic h and, importantly, *with the endpoint behavior absorbed into the weight*, which is why the $d = 2$ singularity costs nothing. A plain Gauss–Legendre rule in u does not absorb the weight (slow algebraic convergence at $d = 2$), and a continuous law is simply the wrong measure at $d = 1$; both errors were caught in smoke validation (Appendix H records the notes) and fixed before any canonical run.

J.3 The second-order term behind Amendments A4 and A5

Two appearances of the $O(1/m^2)$ correction motivated the amended falsifiers, and its sign and rough size are worth recording. From the benchmark one reads off empirically that

$$m\,\epsilon_{\mathrm{pair}}(m) = a(\gamma,d) + \frac{c_2(\gamma,d)}{m} + O(m^{-2}), \qquad c_2 = c_2(\gamma,d) \text{ of either sign on the registered grids,} \tag{45}$$

so the measured $m\epsilon$ approaches a with a deviation that shrinks as $m\gamma/d$ grows and whose sign depends on (γ, d) (Figure 2): at $m/d = 32$ and large d the c_2/m term still exceeds Monte Carlo error at 2×10^5 draws, which is why the raw registered comparison conflates second order with leading order (A4). For the stocking law, inserting $\epsilon = a/m + c_2/m^2$ into the marginal rule gives

$$m^* = \sqrt{\frac{\Lambda \ell a}{(1-\beta)C_{\mathrm{con}}}}\left(1 + O\big(\tfrac{1}{m^*}\big)\right), \tag{46}$$

so the log-log slope at finite Λ is $1/2$ minus a term of order $1/\sqrt{\Lambda}$: exactly the small negative offset the E-L1-5 fits measure (exponents 0.4997 and 0.4988 with the $(16, 32)$ regime, whose offsets are larger, sitting farther from one half), and the reason A5 tests the band $[0.48, 0.52]$ rather than the point $1/2$.

J.4 The validity domain is $m \gg d\max(1, 1/\gamma)$

The relative $O(1/m)$ remainder of Proposition 16 scales like $d/(\gamma m)$ at small margins and like d/m at large ones, so the asymptotic domain is $m \gg d\max(1, 1/\gamma)$: at small margins $m\gamma/d$ is the binding coordinate, at large margins m/d. Table 4 shows both with the exact benchmark: at equal $m\gamma/d$ the error runs from about -1 to $+31$ percent, while at $m/d = 32$ it stays within about one percent.

Table 4: The exact benchmark's relative deviation from the asymptote, $m\,\epsilon_{\mathrm{pair}}(m)/a(\gamma,d) - 1$, at cells of equal $m\gamma/d$ and of equal m/d (adaptive route). Exact deterministic quadrature.

γ	d	m	$m\gamma/d$	m/d	$m\epsilon/a - 1$
0.25	8	1024	32	128	-1.3%
1	32	1024	32	32	-1.2%
9	8	28	31.5	3.5	$+4.8\%$
16	32	64	32	2	$+12.0\%$
36	32	28	31.5	0.875	$+30.7\%$
16	32	16	8	0.5	$+49.6\%$
9	8	256	288	32	$+0.5\%$
16	32	1024	512	32	$+0.7\%$

At small margins a non-canonical Monte Carlo check (seed 11, 4×10^5 draws, exact conditional errors) makes the distinction concrete at $\gamma = 0.25$, $d = 8$, where $a(\gamma, d) = 1.3654$: the measured m times excess is 0.937 at $m = 32$ (31 percent below the constant despite $m/d = 4$), 1.121 at $m = 64$, 1.300 at $m = 256$, and 1.358 at $m = 2048$: convergence tracks $m\gamma/d$ (respectively 1, 2, 8, 64), in line with the d/γ-scaled remainder. Two consequences. For the registered map: E-L1-4's frozen criterion ($m/d \geq 32$) is the right one on the registered grid ($\gamma \geq 1$), and Figure 2 reads the map against m/d. For the stocking law: hard pairs (small γ) need more exemplars than the m/d framing suggests, by the factor $1/\gamma$, sharpening the square-root law's practical reading; the unregistered small-γ cells above are non-canonical and would require a numbered amendment before entering any canonical claim.

J.5 The crowding envelope's tail bound

Proposition 9(b) claims two envelopes. The unconditional one: the error map's derivative in the offset is

$$\frac{\partial}{\partial\tau}\,\Phi\!\left(-\frac{\sqrt{\gamma}}{2} + \frac{\tau}{\sqrt{\gamma}}\right) = \frac{1}{\sqrt{\gamma}}\,\phi\!\left(-\frac{\sqrt{\gamma}}{2} + \frac{\tau}{\sqrt{\gamma}}\right) \leq \frac{1}{\sqrt{2\pi\gamma}}, \tag{47}$$

so the calibration interference is at most $\Delta\tau/\sqrt{2\pi\gamma}$ for every τ. The exponential one is a balance-point statement: at $\tau = 0$ the slope is $\phi(\sqrt{\gamma}/2)/\sqrt{\gamma} = \Theta(e^{-\gamma/8}/\sqrt{\gamma})$, and with $\Delta\tau \le \log 2$ the interference there is $\Theta(e^{-\gamma/8}\Delta\tau/\sqrt{\gamma})$; the same Mills-ratio inequality $\Phi(-t) \le \frac{1}{2}e^{-t^2/2}$ at $t = \sqrt{\gamma}/2$ controls the entry crowding envelope $\Phi(-\sqrt{\gamma}/2)$. Figure 2c plots both against the exact quantities. The scoping matters: off balance the bare $e^{-\gamma/8}$ is false, and the verification script exhibits $\tau = \gamma/2$, $\Delta\tau = \log 2$ giving ratios 1.25 at $\gamma = 25$ and 4.14 at $\gamma = 36$ (exact Φ arithmetic, verification script).

K WORKED EXAMPLES WITH EXACT NUMBERS

The tables in this appendix are exact closed-form computations, the tabular analogue of the deterministic figures: every value is computed by the table generator (`gen_exact_tables.py`, released with the supplement) from the proved formulas, with no sampling and no registered-run content. They exist so that a reader can check the theory's arithmetic by hand.

K.1 THE ENTRY-TIER DECISION, END TO END

A candidate class arrives at pairwise margin γ to each of two incumbents queried at rate 0.5; $\beta = 0.9$, $\ell = 1$, $C_{\mathrm{con}} = 5$, missing-class error $e_{\mathrm{miss}} = 1$ (a query to an uncovered class is always wrong), $c_{\mathrm{ret}} = 0.1$, and the retrieval-failure probability at this store size is read at its certain-failure ceiling $p = 1$ (the worst case the rent formula permits; any smaller p scales the coverage-gap term proportionally, lowers the rent, and so raises the frontier). The post-entry head error $e_{\mathrm{W}}(c)$ that enters the rent is the candidate's balanced Bayes error against the two incumbents, the pairwise sum $2\,\Phi(-\sqrt{\gamma}/2)$ of Corollary 6 (Algorithm 3); it is not negligible at small margins and the table carries it. Table 5 walks the pricing: the pairwise crowding envelope $\Phi(-\sqrt{\gamma}/2)$ of Proposition 9(c), the capitalized demand-weighted interference of Proposition 8, the buy price, the rent, the frontier query rate, and the decision at each candidate demand. The economics is visible in the rows: a low-margin candidate ($\gamma = 2$) crowds the incumbents hard, capitalizes to a buy price above 7, and at the same time earns less rent relief because the head it would join misclassifies it often, so its frontier sits highest; a high-margin candidate ($\gamma = 25$) crowds nobody, pays essentially C_{con}, and has the lowest frontier query rate of the five, so it is kept only at the two smallest demands in the grid and promoted from $\Lambda = 0.5$ upward.

Table 5: Entry-tier walk-through (Section K.1): exact pricing across candidate margins and demands, with the post-entry head error $e_{\mathrm{W}} = 2\Phi(-\sqrt{\gamma}/2)$ inside the rent. Crowding is exponentially weak in the margin, so the frontier collapses toward $(1-\beta)C_{\mathrm{con}}/r$ as separation grows. Exact closed forms.

γ	Λ	$\Phi(-\sqrt{\gamma}/2)$	$I(c;W)$	$B_c(W)$	r	λ^*	decision
$\gamma = 2$	0.05	0.2398	2.398	7.398	0.620	1.192	keep
$\gamma = 2$	0.2	0.2398	2.398	7.398	0.620	1.192	keep
$\gamma = 2$	0.5	0.2398	2.398	7.398	0.620	1.192	keep
$\gamma = 2$	1.0	0.2398	2.398	7.398	0.620	1.192	keep
$\gamma = 2$	2.0	0.2398	2.398	7.398	0.620	1.192	promote
$\gamma = 4$	0.05	0.1587	1.587	6.587	0.783	0.842	keep
$\gamma = 4$	0.2	0.1587	1.587	6.587	0.783	0.842	keep
$\gamma = 4$	0.5	0.1587	1.587	6.587	0.783	0.842	keep
$\gamma = 4$	1.0	0.1587	1.587	6.587	0.783	0.842	promote
$\gamma = 4$	2.0	0.1587	1.587	6.587	0.783	0.842	promote
$\gamma = 9$	0.05	0.0668	0.668	5.668	0.966	0.587	keep
$\gamma = 9$	0.2	0.0668	0.668	5.668	0.966	0.587	keep
$\gamma = 9$	0.5	0.0668	0.668	5.668	0.966	0.587	keep
$\gamma = 9$	1.0	0.0668	0.668	5.668	0.966	0.587	promote
$\gamma = 9$	2.0	0.0668	0.668	5.668	0.966	0.587	promote
$\gamma = 16$	0.05	0.0228	0.228	5.228	1.054	0.496	keep
$\gamma = 16$	0.2	0.0228	0.228	5.228	1.054	0.496	keep
$\gamma = 16$	0.5	0.0228	0.228	5.228	1.054	0.496	promote
$\gamma = 16$	1.0	0.0228	0.228	5.228	1.054	0.496	promote
$\gamma = 16$	2.0	0.0228	0.228	5.228	1.054	0.496	promote
$\gamma = 25$	0.05	0.0062	0.062	5.062	1.088	0.465	keep
$\gamma = 25$	0.2	0.0062	0.062	5.062	1.088	0.465	keep
$\gamma = 25$	0.5	0.0062	0.062	5.062	1.088	0.465	promote
$\gamma = 25$	1.0	0.0062	0.062	5.062	1.088	0.465	promote
$\gamma = 25$	2.0	0.0062	0.062	5.062	1.088	0.465	promote

K.2 The exemplar-tier decision and the square-root law

Table 6 runs the marginal rule of Corollary 13 under the exact benchmark at the two E-L1-5 regimes and the canonical demand grid, against the square-root prediction $\sqrt{\Lambda a/((1-\beta)C_{\mathrm{con}})}$. Every exact stopping count in the table equals the canonical E-L1-5 value at the same parameters (the generator asserts equality against the registered CSV and fails the build otherwise); the deviations from the prediction are the genuine positive $O(1/m^2)$ offsets, shrinking in Λ and larger in the $(16, 32)$ regime, exactly the structure E-L1-5 measures at scale. The generator asserts equality of every exact stopping count against the registered E-L1-5 CSV and fails the build otherwise, so the worked table cannot drift from the canonical run.

Table 6: Exemplar-tier walk-through: exact-benchmark stopping counts versus the square-root law at $C_{\mathrm{con}} = 0.1$, $\beta = 0.9$, $\ell = 1$. Exact deterministic quadrature.

(γ, d)	Λ	exact m^*	$\sqrt{\Lambda a/((1-\beta)C)}$	rel. dev.
$(9, 8)$	$1.00e+05$	1001	999	+0.2%
$(9, 8)$	$3.16e+05$	1778	1777	+0.1%
$(9, 8)$	$1.00e+06$	3161	3160	+0.0%
$(9, 8)$	$3.16e+06$	5620	5619	+0.0%
$(9, 8)$	$1.00e+07$	9993	9992	+0.0%
$(16, 32)$	$1.00e+05$	1094	1087	+0.7%
$(16, 32)$	$3.16e+05$	1940	1933	+0.4%
$(16, 32)$	$1.00e+06$	3444	3437	+0.2%
$(16, 32)$	$3.16e+06$	6119	6111	+0.1%
$(16, 32)$	$1.00e+07$	10875	10868	+0.1%

K.3 Flow accounting's unbounded liability, numerically

Table 7 is Theorem 2(a) at $\Gamma = 10$: stage 1 floods k for $Q^* = \lceil 11 \cdot 5 \rceil = 55$ queries, forcing any 10-competitive policy to buy k; stage 2 floods v. The billed policy buys v at its first query and then pays the crowding flow $\varphi = 0.5$ on every v query; OPT is the exact one-way optimum over all buy patterns (for $Q_v \geq 100$ it buys v at time zero and rents k; at $Q_v = 10$ it buys k and rents v). The ratio crosses the target 10 between $Q_v = 10^3$ and 10^4 and grows without bound thereafter; no one-way policy escapes, which is the theorem.

Table 7: The two-phase adversary at $\Gamma = 10$, $r = 1$, $B = 5$, $\varphi = 0.5$: once k is irrevocably bought, the stage-2 flood converts another item's future demand into unbounded liability. The billed policy is the stage-1-forced one-way buyer of k (rent Q_1, buy k), which then buys v at its first query and serves the stage-2 v-flood at the crowding flow φ per query because irrevocability keeps k in the head; OPT is the exact one-way optimum over buy patterns. Exact closed-form construction (Algorithm 10).

Q_1	Q_v	policy cost	OPT	ratio	verdict
55	10	70	15	4.7	≤ 10
55	100	115	60	1.9	≤ 10
55	1000	565	60	9.4	≤ 10
55	10000	5065	60	84.4	> 10
55	100000	50065	60	834.4	> 10

K.4 Calibration interference in the uncomposed protocol

Table 8 tabulates the closed form of Proposition 9(b) at the balanced offset: the interference a $k \to k + 1$ promotion imposes on a neighbor in the uncomposed protocol, across margins and counts. The values are tiny and shrink in both the margin and the count, which is F1's content quantified: composition removes an effect that was already exponentially weak in the margin.

Table 8: Uncomposed calibration interference (Proposition 9(b)) at $\tau = 0$: exact values across margins and buffer counts. Exact closed forms.

γ	m	$\Delta\tau$	$\Phi(-\sqrt{\gamma}/2)$	$\mathrm{err}(\Delta\tau)$	I_{cal}
$\gamma = 4$	1	0.6931	0.15866	0.25674	$9.81e-02$
$\gamma = 4$	4	0.2231	0.15866	0.18716	$2.85e-02$
$\gamma = 4$	16	0.0606	0.15866	0.16610	$7.45e-03$
$\gamma = 9$	1	0.6931	0.06681	0.10223	$3.54e-02$
$\gamma = 9$	4	0.2231	0.06681	0.07699	$1.02e-02$
$\gamma = 9$	16	0.0606	0.06681	0.06946	$2.66e-03$
$\gamma = 16$	1	0.6931	0.02275	0.03387	$1.11e-02$
$\gamma = 16$	4	0.2231	0.02275	0.02593	$3.18e-03$
$\gamma = 16$	16	0.0606	0.02275	0.02358	$8.31e-04$
$\gamma = 25$	1	0.6931	0.00621	0.00910	$2.89e-03$
$\gamma = 25$	4	0.2231	0.00621	0.00704	$8.27e-04$
$\gamma = 25$	16	0.0606	0.00621	0.00643	$2.16e-04$

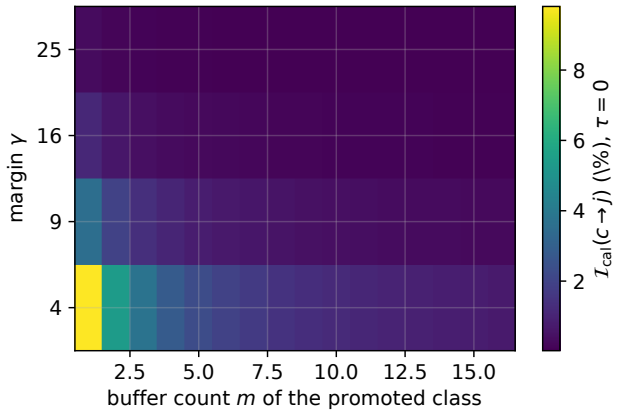

Figure 11: The calibration interference surface of Table 8 over the full count range: the uncomposed protocol's promotion cost dies in both the margin and the count, and the composed protocol removes it exactly. Deterministic closed forms.

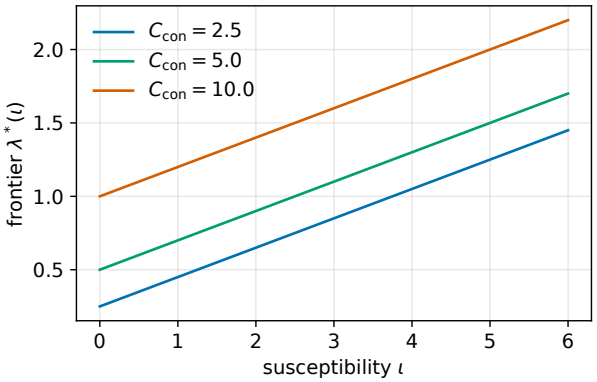

Figure 12: The frontier family over write compute C_{con} at fixed $\beta = 0.9$: the phase diagram of Theorem 12 translates compute prices directly into promotion thresholds. Deterministic closed forms.

K.5 Algorithm costs

Table 9: Per-operation costs of Appendix G's algorithms as implemented by the released runners.

Algorithm	Cost
Threshold policy (Alg. 2)	$O(1)$ per item
Entry-tier pricing (Alg. 3)	$O(\lvert W\rvert)$ per candidate
Exemplar marginal rule (Alg. 4)	$O(m^*)$ with ε oracle; closed form with a/m
Promotion value, computable half (Alg. 5)	$O(\lvert W\rvert d)$ per candidate
Integral benchmark (Alg. 6)	$O(n_q n_u)$ per count; batched over counts
DP verifier (Alg. 7)	$O(H)$ per cell, $H = O(\log(\mathrm{tol})/\log\beta)$
Break-even (Alg. 8)	$O(1)$ per query
Tightness adversary (Alg. 9)	$O(B/r)$ queries
Two-phase adversary (Alg. 10)	$O(\Gamma B/r + \Gamma^2 B/\min(r,\varphi))$ queries
Allocation harness (Alg. 11)	greedy $O(MN)$; DP $O(N\tilde{m}^2)$

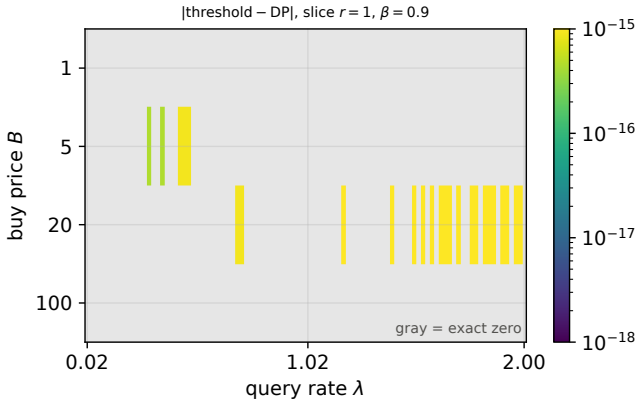

Figure 13: E-L1-1: the absolute gap surface on the (λ, B) slice at $r = 1$, $\beta = 0.9$ (log color scale, 10^{-18} to 10^{-15}; gray cells agree to exactly zero, 375 of the 400 on the slice): the agreement is uniform across the plane, on both sides of the frontier and through the frontier crossing. Canonical data.

L TIER 1 IN FULL: THE PRE-REGISTERED PROGRAM, EXPERIMENT BY EXPERIMENT

Governance. The program was registered in `experiments_design_L1.md` and frozen on 2026-08-16 with numbered pre-execution amendments (A1: tolerance values; A2–A4: multiplicity and second-order corrections to the statistical falsifiers, each preserving the falsifier's intent with correct error control; A5: the finite-demand band for the exponent test; A6–A8: the added flow-dichotomy experiment, its attack rows, and a relabeling; A9–A10: a single rent charge per query on the attack rows and one division in the exponent standard error), all summarized with their motivations in Appendix N. Smoke seeds 0, 1, 2 validated the runners and are excluded from canon; all results below are from canonical seeds 11, 13, 17. Every canonical CSV carries a manifest recording code hash, seeds, grid, tolerances, library versions, and timestamp (Appendix M); every number quoted below enters this paper through the fill script from those CSVs. Two amendments warrant emphasis: the raw three-sigma any-cell rules as originally drafted would have fired on perfect formulas (expected false-fire counts of about 5.4 and 1.1 cells respectively), and the raw MC-versus-constant comparison in asymptotic cells conflates the $O(1/m^2)$ correction with the leading term; the amendments fix the statistics, not the data.

L.1 E-L1-1: NOW-OR-NEVER AND THRESHOLD EQUALITY

Design. Finite-horizon dynamic programming over promotion times (Algorithm 7, Appendix G) on a grid of 6000 cells spanning (λ, r, B, β), horizons long enough that the discounted tail is below 10^{-13}, against the closed-form threshold policy of Theorem 12. Deterministic computation.

Result. K-L1-1 not fired. The maximum DP advantage over the threshold policy (one-sided) was 1.42×10^{-14}, the maximum absolute two-sided agreement gap was 1.4×10^{-14} (reached at the highest β; at machine precision the one-sided advantage and the two-sided gap coincide), and the maximum violation of now-or-never endpoint-optimality was 1.42×10^{-14} (Figure 21). The residuals are roundoff, not disagreement: 5738 of the 6000 cells agree to exactly zero, and every nonzero residual lies within 1.6 units of double-precision roundoff of the cost it sits on (2^{-52} times the cost), far below the 1×10^{-8} tolerance. The DP optimum coincides with the closed form at every grid cell on both sides of the frontier, as Lemma 11 demands.

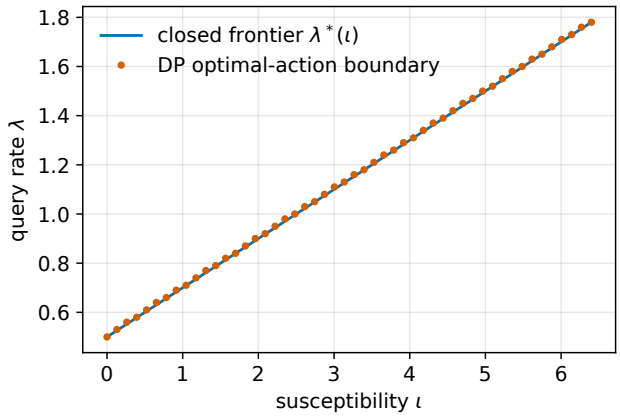

Figure 14: E-L1-2: the dynamic program's optimal-action boundary against the closed frontier of Theorem 12, within 0.980 grid cells on canonical seeds 11, 13, 17.

L.2 E-L1-2: FRONTIER LOCATION

Design. Cell-wise DP optimal actions on a (λ, ι) grid of 10000 cells at fixed $(C_{\text{con}}, \kappa_{\text{scale}}, G = 0, r, \beta)$, versus the closed frontier $\lambda^*(\iota) = (1-\beta)(C_{\text{con}} + \iota\kappa_{\text{scale}})/r$ of Theorem 12. Falsifier K-L1-2: boundary displacement exceeding one grid cell anywhere.

Result. K-L1-2 not fired. Maximum boundary displacement 0.980 grid cells (Figure 14), i.e. the empirical optimal-action boundary sits on the closed frontier up to discretization. The displacement sits just under the falsifier's one-cell ceiling, and near-ceiling values are typical across cells (detail table), so the frontier is confirmed at the grid's resolution and no finer; the three regions of the phase diagram (Figure 2) are exactly the DP's per-cell actions.

L.3 E-L1-3: EXACTNESS OF THE STATE-CONDITIONAL FORMULAS

Design. Monte Carlo conditional errors at fixed displacement w against the exact expressions of Theorem 14, on a grid of 336 cells over $(\gamma, |w|, \angle w, d)$ per seed (1008 rows over the three canonical seeds; the 1008 macro reports rows) (20000 draws per side-test), including wrong-side cells $|w|^2 > \gamma$. Falsifier K-L1-3 as amended (A2): any cell beyond 5 standard errors, or pooled chi-square rejection at 10^{-3}.

Result. K-L1-3 not fired. Maximum $|z| = 3.37$; 4 side-tests beyond three sigma (descriptive; about 5.4 are expected under exactness at this grid size); pooled goodness-of-fit $p = 0.596$ (Figure 2e). The z distribution is consistent with the folded normal that exactness predicts, wrong-side cells included.

L.4 E-L1-4: THE CONSTANT AND ITS VALIDITY MAP

Design. 192 cells of a (γ, d, m) grid per seed (576 rows over three seeds; 576 reports rows), with m spanning both $m \gg d$ and the continual-incremental regime $m < d$, measuring m times the balanced-pair excess with exact conditional errors (200000 draws), against $a(\gamma, d)$ and the exact one-dimensional integral benchmark (Algorithm 6) computed per cell. Falsifier K-L1-4 as amended (A3/A4): machinery failure (MC versus benchmark beyond 5σ anywhere or pooled rejection), or constant failure (the benchmark's OLS intercept $\hat{a}$ in $\hat{a} + c/m$ deviating from the stated $a(\gamma, d)$ by more than 0.5% relative).

Result. K-L1-4 not fired. The machinery test passes with pooled $p = 0.700$; the recovered constant matches the stated $a(\gamma, d)$ at every (γ, d) with maximum relative deviation 0.023% (Table 12); the superseded d-version of the constant is excluded by wide margins (it deviates by 2.9% and more, far beyond the 0.5% tolerance). The validity map (Figure 2) shows exactly the regime structure

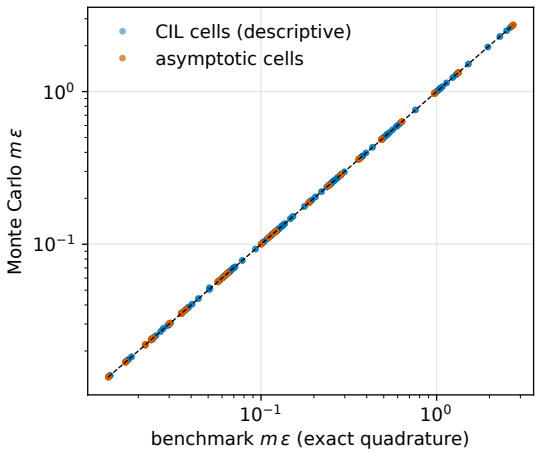

Figure 15: E-L1-4: Monte Carlo $m\epsilon$ against the exact benchmark $m\epsilon$ on every cell (log-log), CIL cells in blue (descriptive), asymptotic cells in orange. The points sit on the identity line across four decades: the stochastic arm measures what the deterministic benchmark computes.

Remark 3 predicts: the ratio $m\epsilon/a$ approaches one from above as m/d grows, with the small-m cells mapping the continual-incremental departure rather than asserting the asymptotic; on the registered grid ($\gamma \geq 1$) the m/d and $m\gamma/d$ readings coincide up to constants, and the honest domain $m\gamma/d \to \infty$ of Appendix J.4 is the one to use off it. The direct MC-versus-benchmark agreement across all cells (including the descriptive small-m cells) is shown in Figure 15.

L.5 E-L1-5: THE SQUARE-ROOT STOCKING EXPONENT

Design. The marginal promotion rule of Corollary 13 (Algorithm 4) simulated under the exact excess (the integral benchmark, not the asymptotic), across class demand Λ spanning two orders of magnitude (10^5 to 10^7), in two regimes flagged asymptotically valid by E-L1-4, plus a Monte Carlo-excess variant per canonical seed (50000 draws). Falsifier K-L1-5 as amended (A5): a fit's 95% confidence interval entirely outside $[0.48, 0.52]$.

Result. K-L1-5 not fired. Exact-excess exponents: 0.4997 at $(\gamma, d) = (9, 8)$ with CI $(0.49961, 0.49981)$, and 0.4988 at $(16, 32)$ with CI $(0.49854, 0.49898)$; Monte Carlo exponents across canonical seeds: 0.49968–0.49975 (Figure 3c). All fits sit inside the band with the expected small negative offset from one half, which is the finite-demand second-order correction $1+O(1/m^*)$ of the conditional law; the data decorate the square-root law rather than merely tolerating it.

L.6 E-L1-6: HEURISTIC AUTOPSIES

Design. Adversarial per-item query streams (stop-just-before-break-even families and randomized smoothings) on decoupled instances; realized competitive ratios of break-even, promote-after-N ($N \in \{2, 5, 10\}$), and a recency-frequency score policy. Falsifier K-L1-6: break-even's realized ratio exceeding 2.02 on any decoupled instance. The unboundedness demonstrations for the heuristics are constructive outputs, not falsifiers.

Result. K-L1-6 not fired. Over the 52 adversarial streams and three seeds (156 rows per policy), break-even's maximum realized ratio was 2.0000, the classical bound touched from below on 20 rows and exceeded on 0, under the 2.02 ceiling (Figure 16, left); every heuristic exceeds 2 on some stream, with worst cases from 2.05 to 4.75 and between 4 and 28 rows above the bound. The medians say nothing here (competitive analysis is a worst-case statement, and the heuristics' typical ratios are unremarkable); the maxima are the reading. The autopsies are stark: promote-after-N is driven to ratio 50.0 on the constructive family (right panel), and the recency-frequency score inherits the same exposure, which is the paper's empirical case for pricing promotion decisions rather than scoring them.

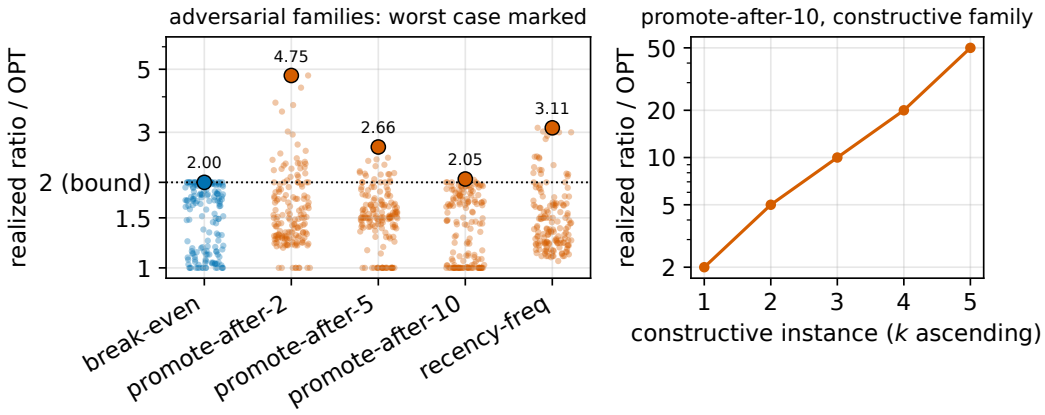

Figure 16: E-L1-6: heuristic autopsies. Left: every realized ratio per policy over the 52 adversarial streams and three seeds (small markers), the worst case marked with its value, the classical bound of 2 dotted; break-even's worst case is exactly 2.0000, touching the bound from below, and every heuristic's worst case lies above it. Right: the constructive unboundedness family drives promote-after-10 to ratio 50.0. Canonical seeds 11, 13, 17.

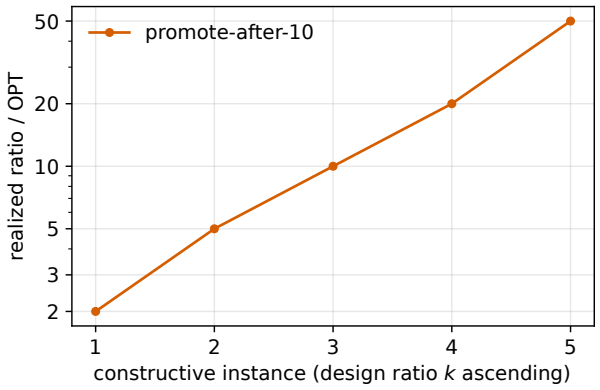

Figure 17: E-L1-6: the constructive family across promote-after-N policies: the realized ratio equals the design parameter k identically at every N (the policy pays $(N-1)$ rents against an optimum of $B = (N-1)r/k$, by construction), so the ascent in k is designed, not measured luck; the quantitative form of the autopsies' conclusion that fixed-count heuristics carry unbounded worst-case exposure. Canonical constructive outputs.

L.7 E-L1-7: ALLOCATION CONSISTENCY

Design. Capped-budget promotion with known demand (Algorithm 11) versus the exact allocation DP on matched instances at the intersection with allocation theory. Falsifier K-L1-7: any instance where objectives or allocations differ beyond 10^{-9}.

Result. K-L1-7 not fired. Across 72 matched instances (fixed and randomized per canonical seed), the greedy marginal rule and the allocation DP agreed to 0 (exactly) (Figure 18): promotion theory with known demand and a capacity constraint is exactly allocation theory at the intersection.

L.8 E-L1-8: THE FLOW DICHOTOMY, TWO-WAY (AMENDMENT A6)

Design. Two items $\{v, k\}$ under flow accounting, registered by Amendment A6: base instance $r = 1$, $B = 10$, $\varphi = 0.5$ plus the full 3×3 grid over $(B/r, \varphi/r)$; two-stage floods ($Q_1 = 20$, $Q_v \in \{10^2, 10^3, 10^4\}$) and alternating flood cycles; policies one-way break-even, two-way flow

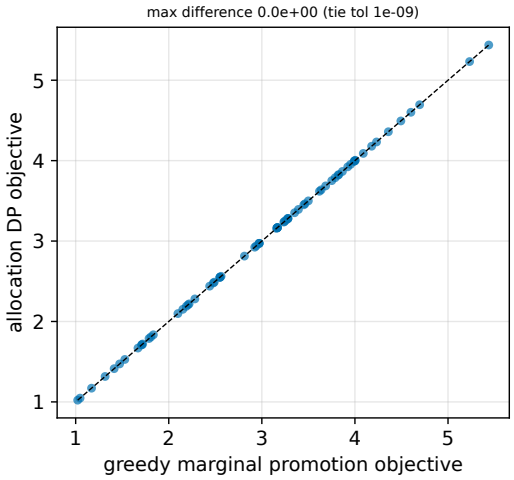

Figure 18: E-L1-7: greedy marginal promotion objectives against allocation DP objectives on all matched instances. The maximum difference is 0 (exactly); tie tolerance 10^{-9}. Canonical seeds 11, 13, 17.

break-even (FBE), and the work-function algorithm (WFA, with the amendment's fixed tie-break convention); comparator: the exact DP two-way optimum with free demotion. Falsifiers K-L1-8a (one-way ratio fails to exceed 50 at $Q_v = 10^4$), K-L1-8b (FBE ratio exceeds 5 on the non-attack streams), and K-L1-8c (WFA $\hat{d}$-cost ratio exceeds the Theorem 3 bound at $N = 2$ plus an additive allowance of 80, the scale of the $\hat{d}$-diameter constant c_0 written in ratio form on the short registered streams; the frozen number is the bound plus that allowance, not a tuned ceiling); the falsifiers were relabeled from the runner's internal F1–F4 tags by Amendment A8.

Result. K-L1-8 not fired. On the base instance (Figure 1b), the one-way rule's realized ratio climbs with the flood to 251.8 (K-L1-8a stands: the blow-up is the theorem's prediction), while FBE stays at or below 2.97 and WFA's realized true-cost ratios sit in 3.95–4.79 on the base instance, its $\hat{d}$-cost ratio (the quantity Theorem 3 bounds) reaching at most 5.78 over the entire 45-cell grid against the bound 14 (K-L1-8c stands). The dichotomy is now a measured object: identical instances, one-way liability growing linearly in the adversary's flood, two-way policies flat. The WFA convention recorded in Amendment A6 is itself a small finding: an argmin that prefers the current state on ties can rent indefinitely on this system while the work functions cap, so WFA's behavior is tie-break-sensitive, and the registration pins the convention down (post-task work functions, argmin of $w_t(s) + \hat{d}(\mathrm{cur}, s)$, ties broken toward moving, true costs charged).

Amendment A7 adds the adversarial attack rows, which invert the FBE reading of the designed streams: on the overshoot family ($\varphi = 100 \gg B$) and the ping-pong family (cheap refills at $r_v = 0.2$), FBE's realized ratio reaches 180 and even the preemptive repair reaches 80 (falsifier K-L1-8d, the counterexample prediction: the refutation must reproduce, and it does, with FBE at least 60 on every attack row), while WFA stays at or below 1.48 on the same instances. The lesson is now part of the paper's claims rather than a caveat: simple fixed-attribution two-way policies are not safe, and the work function's memory is what the guarantee costs; the designed streams measure the dichotomy, the attack rows measure its price. Figure 19 lays the whole grid out as trajectories and Figure 20 isolates the attack rows.

L.9 Summary

No fired falsifiers are hidden in Tier 1: none fired. The amendments (A2–A5) are logged before execution in Appendix N with their rationales; the raw drafted rules would have produced spurious fires on perfect formulas, and reporting those fires as falsifications would have been a statistical error in the opposite direction.

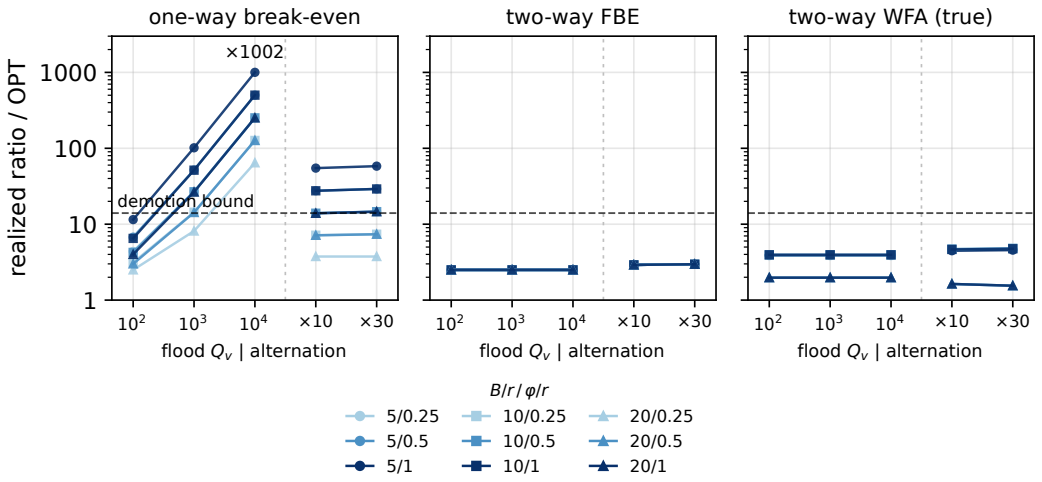

Figure 19: E-L1-8 anatomy: the full registered grid as trajectories, one line per $(B/r, \varphi/r)$ grid point (color: φ/r; marker: B/r), one panel per policy family; the flood family ($Q_v = 10^2$ to 10^4) is separated from the alternation family by the faint divider. The one-way lines climb with the flood as the break-even rule chases the adversary, to 1002 at the most expensive grid point, piercing the Theorem 3 bound (dashed); the two-way lines stay flat far below it, the visual form of the dichotomy. Canonical data (deterministic computation).

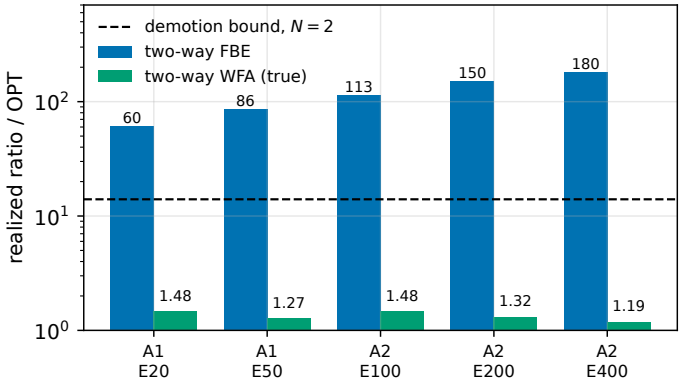

Figure 20: E-L1-8 attacks (Amendment A7): the adversarial E-scaling rows of the canonical CSV. FBE's realized ratio climbs from 60 to 180 while WFA stays at or below 1.48, an order of magnitude inside the Theorem 3 bound (dashed). Canonical data (deterministic computation).

Table 10: The registered Tier 1 program on canonical seeds 11, 13, 17: all falsifiers pass, including Amendment A6's E-L1-8; no stated negative results fired in Tier 1 (Tier 1b, Appendix O, is reported separately). Every number in this table is generated from the canonical CSVs by the fill script.

Experiment	Falsifier	Outcome	Headline number
E-L1-1 threshold equality	K-L1-1	not fired	max DP advantage 1.42×10^{-14}
E-L1-2 frontier location	K-L1-2	not fired	displacement 0.980 cells
E-L1-3 exact formulas	K-L1-3 (A2)	not fired	pooled $p = 0.596$
E-L1-4 constant + map	K-L1-4 (A3/A4)	not fired	max deviation 0.023%
E-L1-5 stocking exponent	K-L1-5 (A5)	not fired	exponents 0.4997, 0.4988
E-L1-6 heuristic autopsies	K-L1-6	not fired	break-even max 2.0000
E-L1-7 allocation consistency	K-L1-7	not fired	max diff 0 (exactly)
E-L1-8 flow dichotomy (A6, A7)	K-L1-8a–d	not fired	one-way max 251.8, FBE 2.97

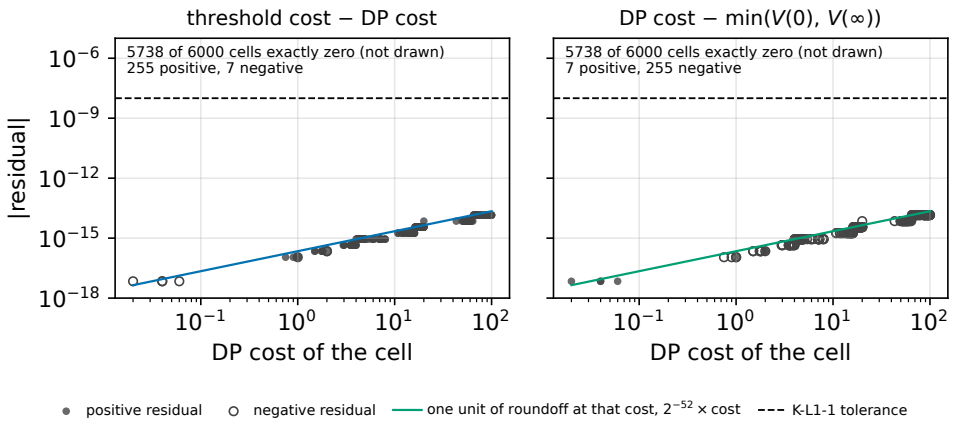

Figure 21: E-L1-1, all 6000 cells: the absolute threshold-minus-DP residual (left) and the now-or-never violation (right) against the DP cost of the cell, log-log, nonzero cells only. The dynamic program confirms now-or-never and the threshold to machine precision.

Table 11: Falsifier margins on canonical seeds 11, 13, 17. Tier 1 (top): amendments A2–A4 repair any-cell rules that would fire spuriously on exact formulas (about 5.4 expected false fires at three sigma); A5's band carries the law's own $O(1/m^*)$ offset. Tier 1b (bottom, Appendix O): machinery falsifiers K-L1b-1m and K-L1b-2m did not fire; the rate falsifier of E-L1b-1 fired at 1 of 16 configurations on every canonical seed, a stated negative result.

Experiment	Falsifier	Measured	Threshold	Margin	Status
E-L1-1 threshold equality	K-L1-1	DP adv 1.4×10^{-14}	1×10^{-8}	6 orders	pass
E-L1-2 frontier location	K-L1-2	0.98 cells	$\leq$ 1 cell	0.02	pass
E-L1-3 exact formulas	K-L1-3 (A2)	max z 3.37, pooled p 0.60	$z \leq 5$; pooled $> 10^{-3}$	holds	pass
E-L1-4 constant + map	K-L1-4 (A3/A4)	dev 0.023%; machinery p 0.70	$\leq 0.5\%$	20×	pass
E-L1-5 stocking exponent	K-L1-5 (A5)	0.4988–0.4997	CI in $[0.48, 0.52]$	inside	pass
E-L1-6 heuristic autopsies	K-L1-6	break-even 2.00	≤ 2.02	0.02	pass
E-L1-7 allocation consistency	K-L1-7	diff 0	$\leq 10^{-9}$	exact	pass
E-L1-8 flow dichotomy	K-L1-8a–d	one-way 252; FBE 2.97; WFA $\hat{d}$ 5.8	see A6/A7	holds	pass
E-L1b-1 covariance heterogeneity	K-L1b-1 (rate)	p 0.985–1.160	$[0.90, 1.10]$	out at η=1	FIRED 1/16
E-L1b-2 heavy tails	K-L1b-2 (rate)	p 0.987–1.090	$[0.90, 1.10]$	inside	pass
E-L1b-3 pipeline replica	K-L1b-3 (machinery)	slope 0.972–1.034	$[0.8, 1.25]$	inside	pass

M RESULT TABLES AND MANIFESTS

Table 12: E-L1-4 constant recovery: the OLS intercept $\hat{a}$ of the exact benchmark over asymptotic cells ($m/d \geq 32$) against the stated constant, per (γ, d). Maximum relative deviation 0.023%; the 0.5% tolerance of Amendment A4 holds everywhere, and the retracted d-version would fail it by an order of magnitude. Generated from the canonical manifest.

Regime	d	$\hat{a}$ (benchmark OLS)	stated $a(\gamma, d)$	rel. dev.
$\gamma = 1.0$	1	0.02200	0.02200	0.002%
$\gamma = 1.0$	2	0.11000	0.11002	0.023%
$\gamma = 1.0$	4	0.28607	0.28605	0.005%
$\gamma = 1.0$	8	0.63814	0.63812	0.004%
$\gamma = 1.0$	16	1.34223	1.34225	0.001%
$\gamma = 1.0$	32	2.75039	2.75051	0.005%
$\gamma = 4.0$	1	0.03025	0.03025	0.001%
$\gamma = 4.0$	2	0.06049	0.06049	0.001%
$\gamma = 4.0$	4	0.12099	0.12099	0.001%
$\gamma = 4.0$	8	0.24197	0.24197	0.002%
$\gamma = 4.0$	16	0.48395	0.48394	0.002%
$\gamma = 4.0$	32	0.96790	0.96788	0.002%
$\gamma = 9.0$	1	0.02428	0.02428	0.000%
$\gamma = 9.0$	2	0.03508	0.03508	0.000%
$\gamma = 9.0$	4	0.05666	0.05666	0.000%
$\gamma = 9.0$	8	0.09984	0.09984	0.000%
$\gamma = 9.0$	16	0.18618	0.18618	0.000%
$\gamma = 9.0$	32	0.35887	0.35887	0.000%
$\gamma = 16.0$	1	0.01350	0.01350	0.001%
$\gamma = 16.0$	2	0.01687	0.01687	0.000%
$\gamma = 16.0$	4	0.02362	0.02362	0.000%
$\gamma = 16.0$	8	0.03712	0.03712	0.000%
$\gamma = 16.0$	16	0.06411	0.06411	0.000%
$\gamma = 16.0$	32	0.11810	0.11811	0.000%

Canonical artifacts. The released bundle contains, per experiment: the CSV (`l1_e1.csv` ... `l1_e8.csv`, `l1b_e1.csv` ... `l1b_e3.csv`) and its manifest (`.json`) recording the frozen design pointer, runner name and SHA-256, library SHA-256, seed list and class (canonical vs. smoke), the full grid, the tolerance constants, Python/NumPy/SciPy versions, the library version string, and the timestamp; the runner sources; the fill script; and both figure generators. Smoke-seed outputs (`data_smoke/`) are shipped alongside, marked non-canonical in their manifests, so that every smoke claim in this appendix can be re-derived. The verification instruments (`runners/verify_*.py`) and the audit battery (`runners/audit/`) ship with the runners; their runs use fixed seeds of their own and are non-canonical by standing rule.

Determinism statement. E-L1-1, E-L1-2, and E-L1-8 are deterministic computations (executed once; manifests carry the canonical seed set with a determinism flag). E-L1-3, E-L1-4 (Monte Carlo arm), E-L1-5 (Monte Carlo arm), E-L1-6 (smoothing families), E-L1-7 (random instances), and the three Tier 1b runners are seeded per canonical seed with NumPy PCG64 generators; re-running the released runners with the recorded seeds reproduces every CSV bit-for-bit on the recorded library versions.

M.1 MANIFEST SUMMARY

Every canonical CSV is accompanied by a JSON manifest recording the frozen design pointer, runner name and SHA-256, library SHA-256, seed list and class, the full parameter grid, the frozen tolerance constants, the toolchain versions, and the execution timestamp. Table 13 summarizes the canonical Tier 1 runs; the full JSON files ship verbatim in the supplement, as do the Tier 1b manifests (Appendix O).

Table 13: The canonical Tier 1 runs at a glance. Grid entries name the per-axis grid sizes (e.g. lambda grid 100 is the λ axis of the E-L1-1 four-axis grid); cells is the CSV row count for stochastic runners and the grid size for deterministic ones. Generated from the manifests.

Runner	seeds	cells	grid summary	falsifier	timestamp
l1_e1	11,13,17	6000	lambda grid 100	pass	2026-08-16 16:57:05 -0400
l1_e2	11,13,17	10000	lambda grid 200; iota grid 50	pass	2026-08-16 16:57:07 -0400
l1_e3	11,13,17	336	gamma 4; d 3	pass	2026-08-16 16:57:10 -0400
l1_e4	11,13,17	192	gamma 4; d 6; m mults 9	pass	2026-08-16 16:58:04 -0400
l1_e5	11,13,17		regimes 2; Lambda grid 13	pass	2026-09-23 23:40:40 -0400
l1_e6	11,13,17		policy N grid	pass	2026-08-16 17:02:56 -0400
l1_e7	11,13,17		budget grid	pass	2026-08-16 17:02:57 -0400
l1_e8	11,13,17	50	B/r, phi/r grid; streams; attack rows	pass	2026-09-23 23:27:38 -0400

Reproduction protocol. With the released bundle: (1) verify the runner and library SHA-256 hashes against the manifests; (2) re-run each runner with its recorded seed list (Python 3.12, NumPy and SciPy versions pinned in the manifest); (3) re-run `fill_l1_extended.py` and the two figure generators; (4) recompile the manuscript. Steps (2) and (3) reproduce every number and figure in this document bit-for-bit on the pinned toolchain; step (1) is the trust anchor that the executed code is the released code.

On the runtime. The original seven Tier 1 experiments first ran within about six minutes on 2026-08-16, and E-L1-8 the same day under Amendment A6; E-L1-5 and E-L1-8 were re-executed on 2026-09-23 under Amendments A10 and A9, and the table shows the current manifests. That is the workload's true cost on the pinned toolchain: the deterministic experiments are closed-form arithmetic, the Monte Carlo arms are vectorized over 2×10^5 draws per cell, and the E-L1-5 ladder runs on the batched fixed-order quadrature (Appendix H); every run regenerates bit-for-bit from the recorded seeds.

Smoke artifacts. The `data_smoke/` directory ships the seeds-0, 1, 2 outputs with manifests marked non-canonical, so that every smoke-level claim in Appendix N's amendment log can be re-derived too. The verification instruments and the audit battery ship with the runners; their runs are non-canonical by standing rule and are quoted in this paper only as non-canonical evidence, always labeled as such.

M.2 Detailed result tables

All tables in this appendix are generated by `gen_results_tables.py` from the canonical CSVs; rows are selected examples or summaries, and the full CSVs ship in the supplement.

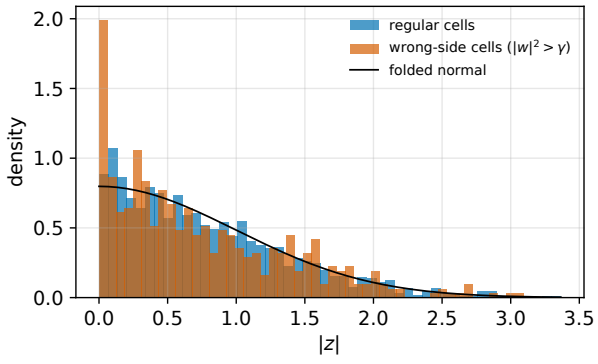

Figure 22: E-L1-3: the deviation distribution split by wrong-side cells ($|w|^2 > \gamma$). The wrong-side cells follow the same folded normal as the rest: the exact formulas hold precisely where the rule is worse than chance on its own class. Canonical data.

Table 14: E-L1-2: the DP optimal-action flip point against the closed frontier for every ι column of the registered grid, with the displacement in grid cells. The maximum 0.980 is the registered headline; the table shows it is typical, not a corner. Canonical data (deterministic).

ι	DP flip λ	closed $\lambda^*(\iota)$	displacement (cells)
0.000	0.5000	0.5000	0.000
0.131	0.5300	0.5261	0.388
0.261	0.5600	0.5522	0.776
0.392	0.5800	0.5784	0.163
0.522	0.6100	0.6045	0.551
0.653	0.6400	0.6306	0.939
0.784	0.6600	0.6567	0.327
0.914	0.6900	0.6829	0.714
1.045	0.7100	0.7090	0.102
1.176	0.7400	0.7351	0.490
1.306	0.7700	0.7612	0.878
1.437	0.7900	0.7873	0.265
1.567	0.8200	0.8135	0.653
1.698	0.8400	0.8396	0.041
1.829	0.8700	0.8657	0.429
1.959	0.9000	0.8918	0.816
2.090	0.9200	0.9180	0.204
2.220	0.9500	0.9441	0.592
2.351	0.9800	0.9702	0.980
2.482	1.0000	0.9963	0.367
2.612	1.0300	1.0224	0.755
2.743	1.0500	1.0486	0.143
2.873	1.0800	1.0747	0.531
3.004	1.1100	1.1008	0.918
3.135	1.1300	1.1269	0.306
3.265	1.1600	1.1531	0.694
3.396	1.1800	1.1792	0.082
3.527	1.2100	1.2053	0.469
3.657	1.2400	1.2314	0.857
3.788	1.2600	1.2576	0.245
3.918	1.2900	1.2837	0.633
4.049	1.3100	1.3098	0.020
4.180	1.3400	1.3359	0.408
4.310	1.3700	1.3620	0.796
4.441	1.3900	1.3882	0.184
4.571	1.4200	1.4143	0.571
4.702	1.4500	1.4404	0.959
4.833	1.4700	1.4665	0.347
4.963	1.5000	1.4927	0.735
5.094	1.5200	1.5188	0.122
5.224	1.5500	1.5449	0.510
5.355	1.5800	1.5710	0.898
5.486	1.6000	1.5971	0.286
5.616	1.6300	1.6233	0.673
5.747	1.6500	1.6494	0.061
5.878	1.6800	1.6755	0.449
6.008	1.7100	1.7016	0.837
6.139	1.7300	1.7278	0.224
6.269	1.7600	1.7539	0.612
6.400	1.7800	1.7800	0.000

Table 15: E-L1-1 gap magnitudes by discount factor over the full (λ, r, B) grid at that β. Every block is at machine precision; the threshold rule is never beaten. Canonical data (deterministic).

β	cells	max |gap|	mean |gap|
0.5	1200	$6.9e-18$	$2.3e-20$
0.8	1200	$3.6e-15$	$6.9e-17$
0.9	1200	$3.6e-15$	$7.7e-17$
0.95	1200	$1.4e-14$	$3.3e-16$
0.99	1200	$1.4e-14$	$6.0e-16$

Table 16: E-L1-3: the ten largest side-test deviations across the canonical grid, with their cells. The largest is $|z| = 3.37$, below the amended any-cell threshold of 5 and consistent with the folded-normal extreme-value scale for this many tests; wrong-side cells behave like the rest. Canonical data.

seed	γ	d	w_{rel}/ang	$\lvert w\rvert^2/\gamma$	wrong	$\lvert z\rvert$
11	1.0	5	1.0/0.523	1.000	no	3.37
13	1.0	10	0.5/1.047	0.250	no	3.28
13	1.0	5	1.0/1.570	1.000	no	3.11
13	9.0	2	1.5/0.785	2.250	yes	3.09
17	25.0	2	1.5/1.047	2.250	yes	3.00
13	9.0	5	0.5/0.0	0.250	no	2.97
13	9.0	5	0.25/0.0	0.062	no	2.92
13	4.0	10	1.5/1.047	2.250	yes	2.91
13	1.0	2	0.25/1.047	0.062	no	2.88
13	9.0	2	1.0/1.047	1.000	no	2.87

Table 17: E-L1-3: deviation statistics by (γ, d) block over all side-tests and seeds. No block's mean drifts from zero; the exactness is uniform over the grid rather than carried by favorable cells. Canonical data.

γ	d	side-tests	mean $\lvert z\rvert$	median	max
1.0	2	168	0.83	0.68	2.88
1.0	5	168	0.75	0.60	3.37
1.0	10	168	0.74	0.58	3.28
4.0	2	168	0.80	0.68	2.87
4.0	5	168	0.82	0.72	2.55
4.0	10	168	0.83	0.71	2.91
9.0	2	168	0.80	0.65	3.09
9.0	5	168	0.80	0.67	2.97
9.0	10	168	0.83	0.73	2.79
25.0	2	168	0.74	0.57	3.00
25.0	5	168	0.67	0.58	2.50
25.0	10	168	0.70	0.54	2.78

Table 18: E-L1-4: the twelve asymptotic cells (seed 11) with the largest descriptive z against the constant. The residual gap between MC and a tracks the benchmark exactly (the $O(1/m^2)$ term), which is why the amended falsifier tests the machinery against the benchmark and the constant by extrapolation. Canonical data.

γ	d	m	m/d	MC $m\varepsilon$	a	bench $m\varepsilon$	z_a (descr.)
1.0	32	1024	32.00	2.71868	2.75051	2.71803	20.76
1.0	16	512	32.00	1.32954	1.34225	1.32865	11.85
16.0	32	1024	32.00	0.11898	0.11811	0.11898	11.65
1.0	32	2048	64.00	2.73304	2.75051	2.73421	11.31
9.0	32	1024	32.00	0.36106	0.35887	0.36125	10.52
4.0	32	1024	32.00	0.97239	0.96788	0.97162	8.27
16.0	16	512	32.00	0.06459	0.06411	0.06454	7.88
9.0	16	512	32.00	0.18736	0.18618	0.18733	7.57
1.0	8	256	32.00	0.63320	0.63812	0.63377	6.66
1.0	16	1024	64.00	1.33518	1.34225	1.33544	6.58
4.0	16	512	32.00	0.48604	0.48394	0.48580	5.45
9.0	32	2048	64.00	0.35997	0.35887	0.36006	5.32

Table 19: E-L1-5: all fits, exact-excess and per-seed Monte Carlo, both regimes. Every confidence interval sits inside the amended band $[0.48, 0.52]$ with the small negative offset predicted by the finite-demand correction. On the exact variant the interval is a fit-residual band, not a sampling interval. Canonical data.

variant	γ	d	seed	exponent	95% CI	n
exact	16.0	32	na	0.49876	(0.49854, 0.49898)	13
exact	9.0	8	na	0.49971	(0.49961, 0.49981)	13
mc	9.0	8	11	0.49968	(0.49959, 0.49977)	13
mc	9.0	8	13	0.49975	(0.49968, 0.49981)	13
mc	9.0	8	17	0.49974	(0.49967, 0.49982)	13

Table 20: E-L1-4: representative cells in the continual-incremental regime $m < d$ (seed 11). These cells cannot falsify (registration, A-1): they map where the asymptotic departs, and the benchmark column shows the departure is real geometry (the $O((d/m)^2)$ terms), not measurement error, with the MC arm tracking the benchmark throughout. Canonical data.

γ	d	m	m/d	MC $m\varepsilon$	bench	z_a	z_b
1.0	4	2	0.50	0.1892	0.1888	364.11	1.40
1.0	8	2	0.25	0.2596	0.2593	2004.93	1.49
1.0	8	4	0.50	0.3781	0.3781	715.73	0.04
1.0	16	4	0.25	0.5225	0.5227	3112.82	0.79
1.0	16	8	0.50	0.7579	0.7573	1159.52	1.13
1.0	32	8	0.25	1.0495	1.0497	4582.11	0.35
1.0	32	16	0.50	1.5151	1.5162	1749.35	1.50
4.0	4	2	0.50	0.1367	0.1367	66.54	0.05
4.0	8	2	0.25	0.2589	0.2591	63.00	0.79
4.0	8	4	0.50	0.2708	0.2707	90.39	0.33
4.0	16	4	0.25	0.5197	0.5200	94.00	1.02
4.0	16	8	0.50	0.5395	0.5396	125.17	0.15
4.0	32	8	0.25	1.0433	1.0425	140.36	1.51
4.0	32	16	0.50	1.0759	1.0778	173.70	2.97

Table 21: E-L1-4: every asymptotic cell at seed 11 (the other seeds are indistinguishable at this precision). The MC arm tracks the exact benchmark (z_b within noise) across all twenty-four (γ, d) configurations at $m/d = 32$ and 64, and the benchmark extrapolates to the stated constant (Table 12); the residual MC-versus-a gap is the second-order term of Appendix J.3, shrinking in m. Canonical data.

γ	d	m	m/d	MC $m\varepsilon$	a	bench $m\varepsilon$	z_b
1.0	1	32	32	0.02196	0.02200	0.02189	0.99
1.0	1	64	64	0.02182	0.02200	0.02195	1.75
1.0	2	64	32	0.11117	0.11002	0.11117	0.01
1.0	2	128	64	0.11082	0.11002	0.11058	0.81
1.0	4	128	32	0.28502	0.28605	0.28593	1.85
1.0	4	256	64	0.28577	0.28605	0.28600	0.46
1.0	8	256	32	0.63320	0.63812	0.63377	0.76
1.0	8	512	64	0.63540	0.63812	0.63596	0.75
1.0	16	512	32	1.32954	1.34225	1.32865	0.83
1.0	16	1024	64	1.33518	1.34225	1.33544	0.25
1.0	32	1024	32	2.71868	2.75051	2.71803	0.43
1.0	32	2048	64	2.73304	2.75051	2.73421	0.76
4.0	1	32	32	0.03018	0.03025	0.03013	0.54
4.0	1	64	64	0.03023	0.03025	0.03019	0.40
4.0	2	64	32	0.06074	0.06049	0.06061	0.96
4.0	2	128	64	0.06060	0.06049	0.06055	0.34
4.0	4	128	32	0.12147	0.12099	0.12140	0.36
4.0	4	256	64	0.12088	0.12099	0.12119	1.61
4.0	8	256	32	0.24281	0.24197	0.24288	0.24
4.0	8	512	64	0.24222	0.24197	0.24243	0.76
4.0	16	512	32	0.48604	0.48394	0.48580	0.63
4.0	16	1024	64	0.48517	0.48394	0.48487	0.76
4.0	32	1024	32	0.97239	0.96788	0.97162	1.43
4.0	32	2048	64	0.96979	0.96788	0.96976	0.06
9.0	1	32	32	0.02426	0.02428	0.02425	0.11
9.0	1	64	64	0.02423	0.02428	0.02427	0.51
9.0	2	64	32	0.03508	0.03508	0.03513	0.61
9.0	2	128	64	0.03513	0.03508	0.03510	0.34
9.0	4	128	32	0.05694	0.05666	0.05688	0.59
9.0	4	256	64	0.05663	0.05666	0.05677	1.43
9.0	8	256	32	0.10048	0.09984	0.10036	0.94
9.0	8	512	64	0.10011	0.09984	0.10010	0.09
9.0	16	512	32	0.18736	0.18618	0.18733	0.19
9.0	16	1024	64	0.18649	0.18618	0.18675	1.70
9.0	32	1024	32	0.36106	0.35887	0.36125	0.90
9.0	32	2048	64	0.35997	0.35887	0.36006	0.43
16.0	1	32	32	0.01348	0.01350	0.01352	1.03
16.0	1	64	64	0.01345	0.01350	0.01351	1.54
16.0	2	64	32	0.01695	0.01687	0.01692	0.77
16.0	2	128	64	0.01698	0.01687	0.01689	1.99
16.0	4	128	32	0.02366	0.02362	0.02372	1.28
16.0	4	256	64	0.02363	0.02362	0.02367	0.95
16.0	8	256	32	0.03737	0.03712	0.03733	0.87
16.0	8	512	64	0.03722	0.03712	0.03722	0.04
16.0	16	512	32	0.06459	0.06411	0.06454	0.77
16.0	16	1024	64	0.06426	0.06411	0.06433	1.15
16.0	32	1024	32	0.11898	0.11811	0.11898	0.11
16.0	32	2048	64	0.11848	0.11811	0.11854	0.81

Table 22: E-L1-5: exact-benchmark stopping counts at every registered Λ in both regimes, with the asymptotic-validity flag. The counts march up the square root through two decades of demand; the Monte Carlo seeds (in the canonical CSV, not shown) sit on the same curve to within a fraction of a percent. Canonical data.

variant	γ	d	Λ	m^*	m^*/d	asymptotic
exact	9.0	8	$1.000e+05$	1001	125	True
exact	9.0	8	$1.468e+05$	1212	152	True
exact	9.0	8	$2.154e+05$	1468	184	True
exact	9.0	8	$3.162e+05$	1778	222	True
exact	9.0	8	$4.642e+05$	2154	269	True
exact	9.0	8	$6.813e+05$	2609	326	True
exact	9.0	8	$1.000e+06$	3161	395	True
exact	9.0	8	$1.468e+06$	3829	479	True
exact	9.0	8	$2.154e+06$	4639	580	True
exact	9.0	8	$3.162e+06$	5620	702	True
exact	9.0	8	$4.642e+06$	6809	851	True
exact	9.0	8	$6.813e+06$	8249	1031	True
exact	9.0	8	$1.000e+07$	9993	1249	True
exact	16.0	32	$1.000e+05$	1094	34	True
exact	16.0	32	$1.468e+05$	1324	41	True
exact	16.0	32	$2.154e+05$	1603	50	True
exact	16.0	32	$3.162e+05$	1940	61	True
exact	16.0	32	$4.642e+05$	2349	73	True
exact	16.0	32	$6.813e+05$	2844	89	True
exact	16.0	32	$1.000e+06$	3444	108	True
exact	16.0	32	$1.468e+06$	4171	130	True
exact	16.0	32	$2.154e+06$	5052	158	True
exact	16.0	32	$3.162e+06$	6119	191	True
exact	16.0	32	$4.642e+06$	7412	232	True
exact	16.0	32	$6.813e+06$	8978	281	True
exact	16.0	32	$1.000e+07$	10875	340	True

Table 23: E-L1-6: realized-ratio statistics per policy over all adversarial family streams (canonical seeds pooled; the constructive unboundedness instances are excluded here and reported separately). Break-even's maximum is exactly the classical bound approached from below; the fixed-count heuristics and the recency-frequency rule reach higher maxima on the same streams, though their medians sit below break-even's: the guarantee is paid for in typical cost. Canonical data.

policy	streams	min	median	mean	max
break-even	156	1.000	1.672	1.543	2.000
promote-after-2	156	1.000	1.483	1.662	4.754
promote-after-5	156	1.000	1.510	1.547	2.663
promote-after-10	156	1.000	1.441	1.432	2.055
recency-freq	156	1.057	1.373	1.541	3.111

Table 24: E-L1-6: break-even realized ratios by stream family (stop-just-before-break-even, its minus-one variant, interleaved, and the randomized smoothings). The stop-just-before family is the binding one, driving the ratio to exactly the classical bound; the smoothings reach it at their worst and sit below it at the median. Canonical data.

stream family	streams	min	median	max (break-even)
interleaved	3	1.114	1.114	1.114
sJB	18	2.000	2.000	2.000
sJBm1	15	1.000	1.000	1.000
smooth	120	1.021	1.672	2.000

Table 25: E-L1-7: the first matched instances (fixed intersection set), greedy allocation and objectives against the DP. The difference column is exactly zero throughout, as it is on all 72 instances. Canonical data.

instance	seed	M	greedy allocation	greedy obj	DP obj	diff
fixed	11	8	[1,1,2,4]	2.550000	2.550000	$0.0e+00$
fixed	11	12	[1,1,3,7]	2.971429	2.971429	$0.0e+00$
fixed	11	16	[1,2,4,9]	3.166667	3.166667	$0.0e+00$
fixed	11	20	[1,2,5,12]	3.280000	3.280000	$0.0e+00$
random_0	11	20	[3,10,5,2]	2.206715	2.206715	$0.0e+00$
random_1	11	20	[5,3,6,6]	3.990678	3.990678	$0.0e+00$
random_2	11	20	[3,6,5,6]	4.493464	4.493464	$0.0e+00$
random_3	11	16	[4,5,4,3]	2.218678	2.218678	$0.0e+00$

N THE REGISTRATIONS AND THEIR AMENDMENT LOGS

Three registration documents govern the experimental record and ship in the supplement: `experiments_design_L1_FROZEN.md` (Tier 1, frozen 2026-08-16 before the first canonical run), `experiments_design_L1_T1B.md` (Tier 1b, registered 2026-08-20 before its first smoke run), and `experiments_design_L1_T2ER1.md` (Tier 2, registered 2026-08-16, with its execution gate). Each fixes the design, the grids, the seeds, the tolerances, and the falsifiers before execution; each carries a numbered amendment log in which every post-registration change is recorded with its motivation and the smoke evidence that produced it, always before the canonical run it affects. The standing rules shared by all three: smoke seeds 0, 1, 2 are never canonical; canonical seeds are 11, 13, 17; numbers enter the manuscript only through the fill scripts from manifest-backed CSVs; every falsifier outcome, fired or not, is reported, a fired falsifier as a stated negative result; a falsifier that would fire on true formulas is a broken falsifier and is amended, on the record, before execution. The Tier 1 file is shipped as a marked redaction (internal session labels removed) with the SHA-256 of the unredacted file recorded for de-anonymized release. Four later registrations ship the same way and are read in their own appendices: Tier 3 (registered 2026-08-20, amendments T3-A1 to T3-A8, Appendix Q), Tier 4 (Appendix R), the parallel line's E-R2 (registered 2026-09-09, Appendix S) and E-R2C (registered 2026-09-11, amendment ER2C-A1, Appendix T); the SHA-256 of every unredacted registration is listed in `REGISTRATIONS_sha256.txt` of the supplement.

Tier 1 amendments (A1 to A8 on 2026-08-16, each before the run it affects; A9 and A10 written down on 2026-09-23 before their reruns, and entered in this log afterwards; the frozen registration file is unchanged).

A1 (tolerances). K-L1-1 numerical tolerance 10^{-8}; K-L1-3 three Monte Carlo standard errors; K-L1-4 asymptotic-cell criterion $m/d \geq 32$; K-L1-6 break-even ceiling 2.02; K-L1-7 tie tolerance 10^{-9}. Fixed at freeze and recorded in every manifest.

A2 (K-L1-3 multiplicity). The drafted any-cell three-sigma rule has an expected false-fire count of about 5.4 cells on the registered grid (336 cells per seed, 1008 rows over the three canonical seeds, two side-tests each, 2016 tests in all), so it would fire on exact formulas. Amended: fires iff any cell shows $|z| > 5$ or the pooled chi-square over all side-tests rejects at 10^{-3}; per-cell three-sigma exceedances are descriptive.

A3 (K-L1-4 multiplicity). The same amendment for the asymptotic cells of E-L1-4; small-m cells remain descriptive.

A4 (K-L1-4 second order). At $m/d \geq 32$ the $O(1/m^2)$ correction still exceeds Monte Carlo error for large d, so a raw MC-versus-constant comparison conflates the second-order term with a wrong leading constant. Amended: the machinery fires on MC-versus-exact-benchmark disagreement (any cell $|z| > 5$ or pooled rejection at 10^{-3}); the constant fires iff the OLS intercept of $m\epsilon$ on $1/m$ over asymptotic cells deviates from the stated $a(\gamma, d)$ by more than 0.5 percent (the retracted d-version deviates by 3 percent and more).

A5 (K-L1-5 band). The conditional law carries a multiplicative $1 + O(1/m^*)$ correction, so the finite-Λ log-log slope sits at $1/2$ minus a small offset and a near-deterministic fit's interval cannot contain exactly $1/2$. Amended: fires iff a fit's 95 percent interval lies entirely outside $[0.48, 0.52]$.

A6 (E-L1-8 added). The flow-dichotomy experiment at $N = 2$: base instance $r = 1$, $B = 10$, $\varphi = 0.5$, the 3×3 grid over $(B/r, \varphi/r) \in \{5, 10, 20\} \times \{0.25, 0.5, 1\}$, two-stage floods $Q_1 = 20$, $Q_v \in \{10^2, 10^3, 10^4\}$, alternating cycles; policies one-way break-even, FBE, WFA with a pinned tie-break convention (argmin of $w_t(s) + \hat{d}(\mathrm{cur}, s)$, ties broken toward moving, true costs charged); comparator the exact two-way DP optimum. Falsifiers: one-way ratio at $Q_v = 10^4$ fails to exceed 50; FBE exceeds 5 on a designed stream; WFA $\hat{d}$-cost ratio exceeds $14 + 80$ on any registered stream.

A7 (attack rows). Two parametric adversary families against FBE, overshoot ($B = 10$, $\varphi = 100$, $r_v = 1$) and ping-pong ($B = 10$, $\varphi = 25$, $r_v = 0.2$), plus the FBE-preemptive policy; the FBE falsifier is re-scoped to the non-attack streams; a new falsifier fires iff FBE's realized ratio on any attack row fails to exceed 20 (the counterexample must reproduce).

A8 (relabel). The runner's internal falsifier tags F1–F4 are reported as K-L1-8a–d; the runner and its manifests are untouched.

A9 (E-L1-8 single charge). Found in review on 2026-09-23, before any rerun: the attack-row runner `fbe_general` charged the rent twice on every rented query (the rent step, then the serving cost, which also bills the rent while the item is outside W), while `fbe`, the DP optimum and the WFA charge each query once. Amended: one charge per query; the attack rows are recomputed deterministically; the falsifier and its threshold are unchanged.

A10 (E-L1-5 interval). Found in review on 2026-09-23, before any rerun: the slope standard error divided the residual variance by $n-2$ twice, so every interval was $\sqrt{n-2}=\sqrt{11}$ times too narrow. Amended: one division; on the exact variant the interval is a fit-residual band, not a sampling interval. The band test of A5 is unchanged.

Tier 1b amendments (all before the first smoke run).

T1B-A1. The E-L1b-1 mean difference is placed along e_1, the heterogeneity-aligned direction, because along the all-ones direction the alternating covariance pattern makes both projected variances exactly γ and the population deviation vanishes identically.

T1B-A2. The E-L1b-3 slope interval uses the parametric bootstrap over cell means (a configuration has at most seven validity cells).

T1B-A3. A descriptive tail exponent (between the two largest cells) is added to the rate fits; falsifiers unchanged.

T1B-A4. After the smoke run: at $(d,\gamma)=(64,4)$ the validity region holds a single registered cell and the harness had required two to evaluate a slope, so the exact cell there was unevaluable and counted as a fire; the slope is now evaluated from one cell when only one exists, and the machinery falsifier is evaluated over evaluable exact cells with the unevaluable count recorded.

Tier 2 amendments (before any canonical real-data run).

T2-A1, T2-A2. A finite-baseline correction to the Arm-1 prediction was proposed from a synthetic slope of 0.81 and then withdrawn when the synthetic harness showed the differenced estimator conditions on the draw pool (the full-train plug-in baseline is the conditional truth), the 0.81 coming from unfiltered small-m cells in the continual-incremental regime; the prediction is $a(\hat\gamma, d_{\text{eff}})(1/m_c + 1/m_j)$ as registered and the validity filter stands. Both entries are kept because the record of the misdiagnosis is the validation's value.

T2-A3. Logged after E-L1b-3 and before any canonical real-data run. The registered code path whitened by the total covariance of the training features and fed the constant the participation ratio of that raw covariance; on synthetic data generated by the exact model the registered falsifier F-A1 fires (slope 3.4, interval $(3.2, 3.5)$), so it was a broken falsifier. Amended: whitening by the pooled within-class covariance; effective dimension the participation ratio of the whitened within-class spectrum (the raw ratio still recorded); the debiased margin $\hat\gamma - d_{\text{eff}}(1/n_c + 1/n_j)$ recorded and used; the F-A1 regression on the model's finite-m prediction (antithetic exact-conditional-error Monte Carlo at the measured $(\hat\gamma, d_{\text{eff}}, m)$) over all cells, the asymptotic form reported descriptively in the validity region, which is empty at $m \le 128$ once d_{eff} is the whitened-noise dimension of ViT features; pairs with non-positive debiased margin excluded and counted; a new end-to-end synthetic gate running the real code path on anisotropic Gaussians (amended path: slope 0.99, interval $(0.97, 1.03)$, d_{eff} exact). Band, backbones, datasets, grids, seeds, bootstrap, and the predeclared branches unchanged; the pre-amendment sources ship for the record.

O TIER 1B IN FULL: THE MISSPECIFICATION STRESS PROGRAM

Purpose and registered reading. Tier 1 verifies the exact ledger on its own generating process. Tier 1b (registered 2026-08-20, `experiments_design_L1_T1B.md`) breaks the three assumptions a real feature space violates and measures what survives: (i) shared covariance, (ii) Gaussian tails, (iii) exact whitening with a known margin in an isotropic space. The claim under stress is the one the stocking law rests on, that the plug-in excess decays as a/m, the rate $1/m$ being what the square-root law needs and the constant $a(\gamma, d)$ being what the exact ledger supplies. The registered reading, stated before any run: the rate survives (it follows from a smooth error map and

finite-variance mean noise by a second-order expansion that uses neither Gaussianity nor shared covariance) and the constant does not (it is the exact ledger's curvature and depends on the tails and the covariance). Tier 1b therefore tests the rate with a falsifier and reports the constant's degradation as a curve against the misspecification parameter; for (iii) the registered object is the Tier 2 adequacy regression itself, replicated on synthetic data where the truth is known, so that Tier 2's effective-dimension proxy and margin estimator are calibrated before any real-data number is read.

Common machinery. Pair (c, j) in d dimensions; the practitioner's rule is the paper's rule (pooled whitening, plug-in means from $m_c = m_j = m$ exemplars, the perpendicular bisector). Conditional errors given the estimated rule are computed exactly (Gaussian: Φ of the signed distance to the bisector in the class's own metric; multivariate t: the univariate t_ν CDF, since a unit-norm projection of the standard multivariate t_ν is univariate t_ν). Monte Carlo is over the mean noise only, with antithetic pairs $(w, -w)$ for both classes, which cancels the odd part of the error map exactly and leaves the even (excess) part with the variance of the quadratic term alone. The excess of a cell is the mean over pairs of the balanced error minus the balanced error of the population plug-in rule. Rate fits: OLS of $\log$ excess on $\log m$ over the asymptotic cells $m\gamma/d \geq 30$ (Appendix J.4), exponent $p = -\text{slope}$; constant ratio ρ: the OLS intercept of $m\epsilon$ on $1/m$ over the same cells (the Amendment A4 extrapolation) over $2a(\gamma, d)$; both with 95 percent parametric-bootstrap intervals (1000 replicates over the cell means). The m grid is $\{8, 16, 32, 64, 128, 256, 512\}$ throughout.

O.1 E-L1B-1: COVARIANCE HETEROGENEITY

Design. In pooled-whitened coordinates $\Sigma_c = \text{diag}(e^{\eta u_i})/\cosh\eta$, $\Sigma_j = \text{diag}(e^{-\eta u_i})/\cosh\eta$ with the alternating pattern $u_i = (-1)^i$, so the pooled covariance is the identity; the mean difference is $\sqrt{\gamma}\, e_1$, the coordinate where class c is widest and class j narrowest (Amendment T1B-A1, the heterogeneity-aligned worst case in which the bisector is not Bayes). Grid: $d \in \{4, 16\}$, $\gamma \in \{4, 9\}$, $\eta \in \{0, 0.25, 0.5, 1\}$; $R = 20000$ antithetic pairs per cell, mean noise sampled from its exact law $\mathcal{N}(0, \Sigma_c/m)$, $\mathcal{N}(0, \Sigma_j/m)$. Falsifier K-L1b-1: for some (d, γ, η) the rate interval lies entirely outside $[0.90, 1.10]$. Falsifier K-L1b-1m (machinery): at $\eta = 0$ the constant-ratio interval misses $[0.98, 1.02]$.

Result. K-L1b-1 fired; K-L1b-1m not fired. The rate falsifier fired, on every canonical seed, at 1 of 16 configurations, $(d, \gamma, \eta) = (4, 9, 1)$: the fitted exponent in the registered window was 1.154 to 1.160 (intervals above 1.10), while the descriptive tail exponent between the two largest cells was 0.988 to 1.054 and the constant ratio had collapsed to 0.25 to 0.26 of the shared-covariance value. The reading is the one the registration anticipated for a fire of this shape and is reported as a stated negative result: where heterogeneity collapses the leading constant, the $O(1/m^2)$ term dominates the registered window $m\gamma/d \geq 30$, the local exponent overshoots one, and the asymptotic regime of the square-root law recedes; the tail exponent near one says the rate is recovered at larger m, but the Tier 1 validity coordinate does not transfer to heterogeneous covariances without a constant-dependent correction. At every other configuration the exponent stayed inside the band (0.985–1.082). The constant shifts in both directions: at $\eta = 0.5$ the ratio is 0.83–1.20 and at $\eta = 1$ it is 0.25–1.17 across cells, rising at $\gamma = 4$ and collapsing at $\gamma = 9$, while the population-level deviations at $\eta = 1$ run +0.062 to +0.067 on the own side and -0.138 to -0.066 on the neighbor side, so the neighbor-side relief of the 3:1 split is amplified by heterogeneity aligned with the margin.

O.2 E-L1B-2: HEAVY TAILS

Design. Class conditionals $\mu + \sqrt{(\nu-2)/\nu}\, t_\nu(0, I)$ (unit covariance), $\nu \in \{3, 5, 10, \infty\}$; $d \in \{4, 16\}$, $\gamma \in \{4, 9\}$; $R = 4000$ antithetic pairs, the mean noise sampled as the mean of m multivariate-t vectors (no Gaussian shortcut). Registered caveat: at $\nu = 3$ the fourth moment of the noise is infinite, so the second-order remainder is not $O(1/m^2)$ and the local exponent in the window may sit below one. Falsifiers K-L1b-2 and K-L1b-2m as above.

Result. K-L1b-2 not fired; K-L1b-2m not fired. At every $\nu \geq 3$ the rate exponent stayed inside the band (0.987–1.090 over all finite ν; 0.994–1.090 at $\nu = 3$, where the intervals are widest, up to 0.107), so the $1/m$ rate survives tails as heavy as t_3 in the registered window, the registered caveat notwithstanding. The constant does not: the ratio falls monotonically with the tail index, 0.86–1.00

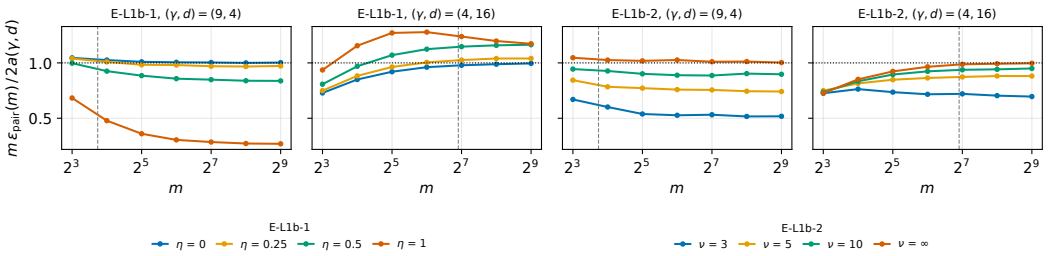

Figure 23: Tier 1b, canonical seeds 11, 13, 17: $m\,\epsilon/(2a(\gamma, d))$ against m (log m) for E-L1b-1 (left, by η) and E-L1b-2 (right, by ν), at $(\gamma, d) = (9, 4)$ and $(4, 16)$; a flat curve at one is the exact ledger, a flat curve below or above one is a changed constant with the rate intact, and a curve still moving in the registered window is the asymptotic regime receding. Vertical line: the validity cut $m\gamma/d = 30$. Canonical data.

at $\nu = 10$, 0.71–0.97 at $\nu = 5$, 0.46–0.85 at $\nu = 3$, because at unit covariance the t law puts less mass at the boundary than the Gaussian, which also lowers the population-level errors (own-side deviation -0.068 to -0.027 at $\nu = 3$). Heavy tails make the plug-in head cheaper to stock than the Gaussian ledger says, by up to a factor of two at $\nu = 3$; the square-root law is unchanged and its constant is conservative.

Table 26: Tier 1b rate and constant fits, canonical seed 11 (the other seeds agree within the intervals; all seeds in the supplement), at $(\gamma, d) = (9, 4)$: the rate exponent p with its interval, the descriptive tail exponent between the two largest cells, the constant ratio ρ against the shared-covariance Gaussian value, and the exact population deviations of the two sides from $\Phi(-\sqrt{\gamma}/2)$. Generated from the canonical CSVs.

Experiment	level	rate p (95% CI)	tail p	ρ (95% CI)	$D_c,\ D_j$
E-L1b-1	$\eta = 0$	1.005 (1.001, 1.008)	0.995	1.006 (1.000, 1.012)	+0.000, +0.000
E-L1b-1	$\eta = 0.25$	1.010 (1.007, 1.014)	1.007	0.966 (0.960, 0.972)	+0.023, -0.025
E-L1b-1	$\eta = 0.5$	1.029 (1.025, 1.032)	0.993	0.833 (0.827, 0.838)	+0.041, -0.046
E-L1b-1	$\eta = 1$	1.154 (1.145, 1.163)	0.988	0.258 (0.253, 0.262)	+0.062, -0.066
E-L1b-2	$\nu = \infty$	1.000 (0.993, 1.008)	1.020	1.019 (1.005, 1.033)	+0.000, +0.000
E-L1b-2	$\nu = 10$	1.010 (1.001, 1.018)	0.999	0.889 (0.878, 0.901)	-0.005, -0.005
E-L1b-2	$\nu = 5$	1.015 (1.006, 1.024)	1.013	0.742 (0.732, 0.753)	-0.012, -0.012
E-L1b-2	$\nu = 3$	1.051 (1.021, 1.081)	0.994	0.503 (0.480, 0.527)	-0.027, -0.027

O.3 E-L1B-3: ANISOTROPY, ESTIMATED WHITENING, AND THE EFFECTIVE-DIMENSION PROXY

Design. Shared-covariance Gaussian classes with $\Sigma = Q\,\mathrm{diag}(\lambda)Q^\top$, $\lambda_i \propto i^{-\alpha}$, $\alpha \in \{0, 0.5, 1, 1.5\}$, normalized to trace d, Q a seeded random orthogonal matrix; $d \in \{16, 64\}$; true Mahalanobis margin $\gamma \in \{4, 9, 16\}$ along a seeded random direction. A pooled training sample of $n \in \{256, 1024, 4096, \infty\}$ vectors ($n/2$ per class; $n = \infty$ the population objects) determines, exactly as in the Tier 2 runner: the Ledoit–Wolf shrunk covariance S, the whitening $T = S^{-1/2}$, the estimated margin $\hat{\gamma} = |T(\hat{\mu}_c - \hat{\mu}_j)|^2$ from the $n/2$-sample means, and the effective-dimension proxies. Two whitening variants are registered because the Tier 2 code path whitens with the total (uncentered) covariance of the training features while the theory's Σ is the within-class covariance: W (within-class, class-centered) and Tt (total, uncentered). Three proxies for the d entering $a(\hat{\gamma}, d)$: P0 the ambient d; P1 the participation ratio of S (the Tier 2 code path); P2 the participation ratio of the whitened within-class covariance $T\hat{\Sigma}_W T$ (the spectrum the theory's d refers to). Mean noise is sampled from its exact law $\mathcal{N}(0, T\Sigma T/m)$ in the whitened space, $R = 20000$ antithetic pairs, conditional errors exact with the true whitened covariance; the excess baseline is the plug-in rule at the $n/2$-sample means (Tier 2's full-train baseline). Metric: the origin slope of measured excess on $a(\hat{\gamma}, d_{\mathrm{proxy}})(2/m)$ over the validity cells $m\hat{\gamma}/d_{\mathrm{proxy}} \geq 30$, with a parametric-bootstrap interval (T1B-A2); descriptively the same slope with the debiased margin $\hat{\gamma} - d_{\mathrm{proxy}}(1/n_c + 1/n_j)$,

the finite-sample inflation of a squared distance between two independent sample means. Falsifier K-L1b-3 (machinery): at $(\alpha = 0, n = \infty, \mathrm{W}, \mathrm{P0})$ the slope lies outside $[0.8, 1.25]$. Pre-declared readings, both results: (R1) if P1's slope leaves the Tier 2 band $[0.5, 2]$ at $\alpha \geq 1$ while P0 or P2 stays inside, the Tier 2 effective dimension as coded is mis-specified for anisotropic features and a Tier 2 amendment must be logged before the canonical real-data runs; (R2) if the Tt slope departs from the W slope by more than the interval width at the same cell, total-covariance whitening distorts the margin and the same amendment switches Tier 2 to within-class whitening.

Result. K-L1b-3 not fired. At the exact cell ($\alpha = 0$, $n = \infty$, within-class whitening, ambient d) the adequacy slope is 0.972–1.034 across seeds, d, and γ, so the machinery is sound. The pre-declared readings both fired. (R1) fired: in 138 of 288 configurations with $\alpha \geq 1$ the Tier 2 proxy P1 left the band $[0.5, 2]$ while P0 or P2 stayed inside it: the participation ratio of the raw feature covariance is 0.07–0.51 of the ambient dimension at $\alpha \geq 1$ and predicts an excess smaller by the same factor, giving P1 slopes of 1.63–16.21 where P0 gives 0.82–2.40 and P2 gives 0.83–2.44; the theory's d is the dimension of the whitened noise, which anisotropy of the raw spectrum does not reduce (P2 stays at 0.80–1.00 of d under within-class whitening). (R2) fired in 588 of 864 cells: the total-covariance slope departed from the within-class slope by more than the interval width: with two classes the total covariance contains the between-class direction and whitening by it shrinks the estimated margin to 0.20–0.50 of its within-class value at $n = \infty$, the two-class worst case ($\gamma/(1 + \gamma/4)$ exactly), and the P0 slopes under total whitening are 0.99–2.96 against 0.82–2.40 under within-class whitening. Finite n matters as well: at $n = 256$ ($n/2$ per class) the ambient-d slope under within-class whitening ranges over 0.82–2.40 with the estimated margin and 0.70–2.07 with the debiased margin, against 0.97–1.04 at $n = \infty$; debiasing narrows the gap without closing it, because the asymptotic form $a(\hat{\gamma}, d)(2/m)$ is itself off at the registered counts, which is why the amended Tier 2 regresses on the model's finite-m prediction rather than on the asymptote.

Table 27: E-L1b-3 at canonical seed 11, whitening variant W unless marked, margin $\hat{\gamma}$ as estimated: the estimated margins under the two whitening variants, the Tier 2 proxy P1 against the ambient d, and the adequacy slopes by proxy (a slope of one is the exact ledger; the Tier 2 falsifier band is $[0.5, 2]$). Selected rows; every cell and the debiased margins are in the supplement CSVs. Generated from the canonical CSVs.

(d, γ, α, n)	$\hat{\gamma}_{\mathrm{W}}$	$\hat{\gamma}_{\mathrm{Tt}}$	P1	slope P0	slope P1	slope P2
$(16, 4, 0, 256)$	5.13	2.71	16.0	1.03	1.03	1.10
$(16, 4, 0, \infty)$	4.00	2.00	16.0	0.98	0.98	0.98
$(16, 4, 1.5, 256)$	3.99	2.03	3.9	0.96	3.79	0.96
$(16, 4, 1.5, \infty)$	4.00	2.00	3.7	0.98	3.98	0.98
$(16, 16, 0, 256)$	17.85	3.64	16.0	1.26	1.26	1.30
$(16, 16, 0, \infty)$	16.00	3.20	16.0	1.03	1.03	1.03
$(16, 16, 1.5, 256)$	17.73	3.25	4.6	1.22	3.17	1.22
$(16, 16, 1.5, \infty)$	16.00	3.20	3.7	1.03	3.15	1.03
$(64, 4, 0, 256)$	5.39	4.54	64.0	1.00	1.00	1.25
$(64, 4, 0, \infty)$	4.00	2.00	64.0	0.97	0.97	0.97
$(64, 4, 1.5, 256)$	4.89	2.03	5.1	1.13	9.94	1.17
$(64, 4, 1.5, \infty)$	4.00	2.00	4.6	0.97	11.52	0.97
$(64, 16, 0, 256)$	15.62	5.99	63.9	1.09	1.09	1.33
$(64, 16, 0, \infty)$	16.00	3.20	64.0	1.02	1.02	1.02
$(64, 16, 1.5, 256)$	17.51	3.35	5.0	1.71	13.59	1.75
$(64, 16, 1.5, \infty)$	16.00	3.20	4.6	1.02	9.64	1.02

What Tier 1b changes. Tier 2 cannot be read as registered. Before any canonical real-data run, amendment T2-A3 (Appendix N) replaces the code path's effective dimension by the participation ratio of the whitened within-class spectrum (P2), whitens by the within-class pooled covariance, and reports the debiased margin alongside the estimated one; the registered F-A1 band, grids, seeds, and the falsifier's logic are unchanged. Without the amendment F-A1 would fire on real features for a reason that has nothing to do with the theory, which is exactly the class of broken falsifier the standing rules require to be repaired on the record before execution. For the theory, Tier 1b's three results are: the $1/m$ rate and hence the square-root stocking law survive heavy tails outright and survive covariance heterogeneity except where the heterogeneity collapses the leading constant,

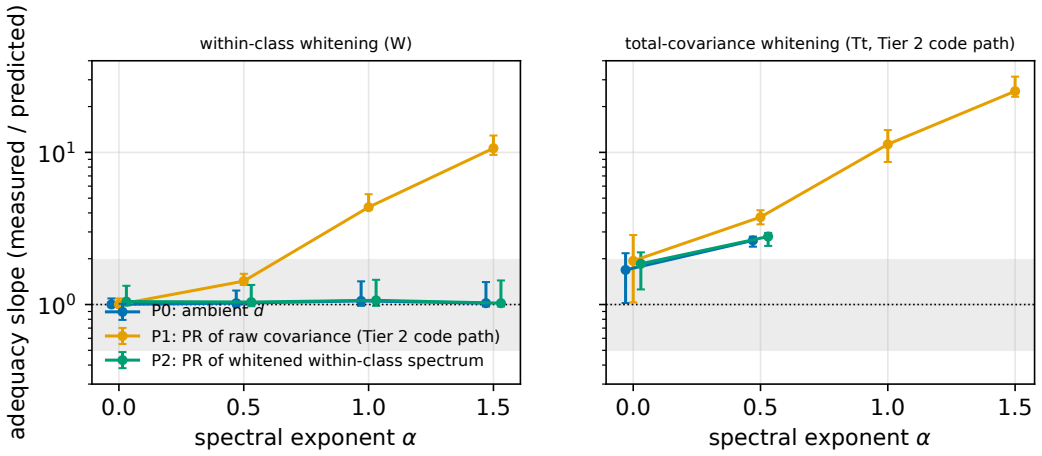

Figure 24: E-L1b-3, canonical seeds 11, 13, 17: adequacy slope by proxy (P0 ambient d, P1 the Tier 2 code path, P2 the whitened within-class spectrum) against the anisotropy exponent α at $d = 64$, for the within-class (left) and total-covariance (right) whitening variants: the median over γ, n, and seeds with the 10th to 90th percentile range; the shaded band is the Tier 2 falsifier band $[0.5, 2]$. Canonical data.

in which case the asymptotic regime recedes beyond the Tier 1 validity coordinate (the one fired falsifier); the constant $a(\gamma, d)$ is conservative under heavy tails and moves in either direction under heterogeneity, so the exact ledger should be read as an anchor with a misspecification band rather than as a point prediction outside the Gaussian substrate; and the quantity the real-data regression must be fed is the whitened noise dimension, not the raw spectrum's participation ratio.

P TIER 2: THE REAL-DATA ANCHOR (E-R1+)

This appendix records the registration of the real-data program and its status. By the registration's own rule, real-data numbers enter this paper only after the canonical runs execute, land as manifest-backed CSVs in the build, and pass through the fill script; the Tier 1 and Tier 1b results stand independently of Tier 2's outcome.

Arms. *Arm 1, model adequacy.* Frozen-backbone features (ViT-B/16 the registered anchor; ViT-L/16 and DINOv2-B/14 as robustness arms; MiniLM-L6 for a non-vision twin) over CIFAR-100 (registered) and Tiny-ImageNet-200; shrinkage covariance, whitening, per-pair $\hat{\gamma}$ and effective dimension; realized pairwise plug-in excess against the full-train plug-in baseline at buffer counts $m \in \{4, \ldots, 128\}$, $R = 50$ resamples per cell; origin regression of measured on $a(\hat{\gamma}, d_{\text{eff}})(1/m_c + 1/m_j)$ in the validity region $m\hat{\gamma}/d_{\text{eff}} \geq 30$, cluster bootstrap over pairs. Falsifier F-A1: the slope interval lies entirely outside $[0.5, 2]$ on the registered pair. The claim under test is the constant as an anchor, not synthetic-level exactness. A registered spot-check compares the per-side excesses to the 3:1 split on the 200 pairs with the largest validity coordinate. *Arm 2, selection.* Theorem 15's computable ϕ-weighted rule against herding and random at equal budgets $k \in \{1, \ldots, 32\}$, paired over draws, both outcomes pre-declared as results: favorable turns the "weights the heuristic lacked" clause into a measured claim; unfavorable isolates the oracle half of the computable/oracle split as the selection signal. Falsifier F-A2 fires only on machinery failure (random beats both policies at every k). *Arm 3, stocking.* The marginal rule under Zipf demand with the *measured* excess function; exponent interval against $[0.45, 0.55]$ in the validity region; plus the registered 3Λ sign test on the most confusable head/tail pairs.

Validation before execution. The pipeline is backbone-agnostic, so it was validated on synthetic shared-covariance Gaussians where the theory is exact, before any real run: in the validity region the origin slope is 1.12 with 95 percent interval $(0.83, 1.41)$, containing 1, the residual being the positive $O(1/m^2)$ correction; the all-cells slope (0.82) exhibits the small-m deficit of the continual-

incremental regime. The validation caught and withdrew a wrong finite-baseline amendment before it could contaminate a real run (T2-A1, T2-A2, Appendix N). Tier 1b's E-L1b-3 (Appendix O.3) extends this validation to anisotropic features and estimated whitening, where the isotropic synthetic gate is blind by construction.

Governance. Smoke seeds 0, 1, 2 validate the GPU path and are never canonical; canonical seeds 11, 13, 17; manifests carry code hashes, library and CUDA versions, and the GPU model; the runner package and the verbatim registration document ship in the supplement's `colab/` directory. The full registered matrix (six Arm-1 cells including the twin, two Arm-2 cells, two Arm-3 cells) is an estimated 6–12 GPU-hour program.

Status and results. Tier 2 as registered was subsumed by the Tier 3 registration (experiments_design_L1_T3.md, registered 2026-08-20, amendments T3-A1 to T3-A8 logged before the runs they affect), which carries Arms 1 to 3 forward with T2-A1 to T2-A3 in force and adds five arms. Tier 3 executed canonically on seeds 11, 13, 17; Appendix Q records the outcomes. This appendix remains the registration record of the smaller program Tier 2 was.

Q TIER 3: THE SCALED REAL-DATA PROGRAM, EXECUTED

Tier 3 subsumes the Tier 2 registration (Appendix P) with amendments T2-A1 to T2-A3 in force, and adds five arms; the registration (`experiments_design_L1_T3.md`, registered 2026-08-20, in the supplement) fixed grids, falsifier bands, pre-declared branches, a priority order, and an execution gate before any real run, and its amendment log (T3-A1 to T3-A8) records every design change with the evidence that motivated it, each logged before the run it affects. Canonical seeds 11, 13, 17; the synthetic gate passed on the execution machine (NVIDIA A100-SXM4-80GB) immediately before the canonical matrix, with the pre-amendment analysis path recovering $d_{\text{eff}} = 10.1$ against a truth of 64 where the amended path recovers 64.0, which is the failure T2-A3 exists to prevent. Every CSV carries a manifest with code hashes, versions, the GPU model, and the seed class; every number below enters through the fill script.

The registered anchor cell is ViT-B/16 on CIFAR-100 ($d_{\text{eff}} = 766.5$ of 768 ambient); every falsifier is evaluated there and nowhere else, a scoping fixed at registration, and all other cells are descriptive. Of the seven falsifiers, one fired.

Datasets and encoders. The real-feature programs (Tiers 2 to 4, E-R2 and E-R2C) use CIFAR-10 and CIFAR-100 (Krizhevsky, 2009), Tiny-ImageNet (Le & Yang, 2015), Food-101 (Bossard et al., 2014), Flowers-102 (Nilsback & Zisserman, 2008), Oxford-IIIT Pets (Parkhi et al., 2012), STL-10 (Coates et al., 2011) and EuroSAT (Helber et al., 2019) for vision, and AG News (Zhang et al., 2015), 20 Newsgroups (Lang, 1995), Banking77 (Casanueva et al., 2020), CLINC150 (Larson et al., 2019) and MASSIVE (en-US) (FitzGerald et al., 2023) for text. The frozen encoders are the ImageNet-pretrained (Deng et al., 2009) ViT-B/16 (AugReg weights (Steiner et al., 2022)) and ViT-L/16 (Dosovitskiy et al., 2021), DINOv2-B/14 and DINOv2-L (Oquab et al., 2024), CLIP ViT-B/16 and ViT-L/14 (Radford et al., 2021), ConvNeXt-T (Liu et al., 2022), ResNet-50 (He et al., 2016), Swin-B (Liu et al., 2021), MAE ViT-B (He et al., 2022), and the sentence encoders MiniLM-L6 (Wang et al., 2020; Reimers & Gurevych, 2019), MPNet (Song et al., 2020; Reimers & Gurevych, 2019), BGE (Xiao et al., 2024) and E5 (Wang et al., 2022); within-class covariances use Ledoit–Wolf shrinkage (Ledoit & Wolf, 2004).

Q.1 ADEQUACY OF THE CONSTANT ON LEARNED FEATURES (ARM 1, C1)

Per cell, all class pairs, nested per-class buffers shared across pairs (T3-A1), $R = 200$ replicates per seed; the regression target is the finite-m prediction of the shared-covariance model at $(\hat{\gamma}_{\text{db}}, d_{\text{eff}}, m_c, m_j)$ per T2-A3(d), and the origin slope of measured on predicted excess is bootstrapped over replicates.

Three disclosures scope the claim, stated here rather than left for a reader to compute. First, the weighting: the origin fit is concentrated, with half of the fit weight in the top 87 of 59,400 cells and 99 percent in the top 850, because the median predicted excess at the anchor is 2.2×10^{-15}

Table 28: Real-data adequacy on every Tier 3 cell: the origin slope of measured on predicted excess, with the effective dimension d_{eff} at which the prediction is evaluated; the anchor row (bold) is the registered cell. Every pooled vision slope is below one, so the ledger over-charges on these cells on average; per count cell it can under-charge by up to a factor of about two (Table 40), so budget with the finite-m prediction at the counts in use. The two Banking77 text cells (slopes above one) underflow the measurement protocol (Appendix Q); Flowers-102 sits outside the validity fence of Remark 3.

Backbone	Dataset	Slope	95% interval	d_{eff}
ViT-B/16	CIFAR-100	0.907	[0.904, 0.911]	766.5
ViT-L/16	CIFAR-100	0.855	[0.852, 0.859]	1017.1
DINOv2-B/14	CIFAR-100	0.742	[0.739, 0.744]	766.0
ResNet-50	CIFAR-100	0.846	[0.846, 0.847]	2038.9
ConvNeXt-T	CIFAR-100	0.867	[0.865, 0.869]	765.3
CLIP ViT-B/16	CIFAR-100	0.701	[0.698, 0.704]	509.9
ViT-B/16	Tiny-ImageNet	0.621	[0.616, 0.626]	766.8
ViT-B/16	Food-101	0.805	[0.802, 0.807]	766.8
ViT-B/16	Pets	0.753	[0.747, 0.760]	741.9
ViT-B/16	Flowers-102	outside the regime fence (no measurable excess)		
MiniLM	20 Newsgroups	0.786	[0.784, 0.787]	376.3
MiniLM	Banking77	4.386	[4.346, 4.427]	377.6
mpnet	Banking77	5.432	[5.366, 5.497]	747.6
mpnet	AG News	0.923	[0.922, 0.925]	763.1

against a measurement floor of 10^{-2}; the honest statement is that on the pairs where the closed form predicts a measurable excess it predicts it to within about ten percent (the most-confusable decile alone gives 0.875), and on the rest it predicts, correctly, nothing measurable. Second, the two-sided law: the asymmetric cells $(m_c, m_j) \in \{(4, 64), (8, 128), (16, 256), (4, 256)\}$, which test the additivity $a(1/m_c + 1/m_j)$ directly, give slope 0.893 ([0.890, 0.897]), indistinguishable from the diagonal. Third, a stated limitation: the per-side 3:1 split regression is not estimable on real features, returning coefficients 10.59 and -441.76 where both are predicted 1, because over the few cells carrying fit weight the two predictors are nearly collinear; on synthetic data, where every cell is measurable, the same regression recovers both coefficients at 1. The per-side content that survives on real features is Arm 4's, below. 0 pairs were excluded for non-positive debiased margin at the anchor.

Q.2 Selection: the computable half of the promotion value (Arm 2, C2)

The pre-declared favourable branch obtained: the registered composite (the Theorem 15 direction plus herding, equally weighted after standardization) beats herding on the paired one-sided interval at $k = 4$, $k = 8$ on the anchor, by -0.0223 ([-0.0237, -0.0208]) at $k = 4$, and the favourable branch replicates on every vision cell run. The shape is the theory's: the edge is largest where the buffer mean is noisiest and hands back to +0.0031 at $k = 64$ once the mean is well estimated. Two honest companions: the boundary-aligned direction alone degrades badly at large budgets (0.227 against herding's 0.141 at $k = 64$), so it is the composite and not the raw direction that is deployable; and the oracle bound targeting the evaluation-set mean (T3-A6) still beats herding by +0.0111 at $k = 64$, so oracle-only signal remains that no computable rule here reaches.

Q.3 The stocking exponent is not resolved (Arm 3, C3)

F-3 did not fire, and we decline to claim support: the anchor interval [0.308, 1.415] contains 0.5 but is 1.11 wide and contains exponents from well below one half to well above it, and the descriptive cells are no tighter. The design, the marginal rule simulated on the measured excess of each class against its nearest neighbor, does not have the power on real features to confirm or refute the square-root law, and is reported as exactly that. Tier 4 re-ran the arm at ten times the replication under a registered power gate and the question closed as unresolvable at this scale (Appendix R.2).

Q.4 THE EXTERNALITY IS REAL; ITS CONSTANT IS NOT (ARM 4, C3)

This is the fired falsifier, and the reading is registered (T3-A4): the measured own-to-neighbor marginal ratio $\hat{\rho}$ is compared to the finite-m prediction of the model at the same $(\hat{\gamma}_{\mathrm{db}}, d_{\mathrm{eff}}, m, m)$, with $3 + 4/(\varrho - 1)$ its population limit.

The qualitative core of the 3Λ rule lands emphatically. On the 200 most confusable anchor pairs at full resolution (measurable fraction 0.94 at $m = 32$), 94 percent show a measurably negative neighbor-side change from one more exemplar of the other class, 98 percent by the fitted slope on the registered fixed-m_j grid: tightening one class's mean measurably hurts its confusable neighbor on learned features. The marginals themselves track the model where they are largest: the selection-free regression of measured on predicted marginal at $m = 8$ gives slope 1.04 on the own side and 0.70 on the neighbor side, each predicted 1.

The quantitative flip point does not reproduce. The measured $\hat{\rho}$ across the balanced grid is 1.34, 1.58, 1.53, 1.29 at $m = 8, 16, 32, 64$ against finite-m predictions of 0.55, 1.40, 2.22, 2.73 and a median asymptote of 3.27; F-4 fired on its sign-agreement clause (sign agreement 0.682 against the registered 0.70, interval [1.302, 2.135] against $[1.5, 8]$). Reported as a stated negative result with this reading: on real features the neighbor-side harm is larger relative to the own-side gain than the constant says, so the demand ratio at which promotion turns net-negative sits nearer 1.5 than 3. The direction of the externality is confirmed and its magnitude is underestimated by the closed form, which therefore errs on the permissive side; a system using the 3Λ threshold as a promotion guard is, on this evidence, not conservative enough rather than too conservative. Tier 4 extended the grid to $m = 256$ at $R = 400$ and the registered reading is flat: the measured ratio does not track the finite-m prediction's climb where measurement is intact, a reading that is uninformative beyond $m = 128$, where measurement reaches its floor (Appendix R.3).

Q.5 THE DICHOTOMY WITH MEASURED PRICES (ARM 5, C4)

Every price measured: crowding from realized pairwise error rates, rents from the measured multi-class head error, on the internally confusable set of T3-A8. The one-way realized ratio at the longest registered flood is 107,496 against the falsifier threshold of 20, and between 10^4 and 10^5 on every descriptive cell; the work-function comparator stayed inside its metrical-task-system bound. The pre-declared realistic-stream reading came out on the favourable branch: the irreversibility gap survives on 45 of 45 Zipf-stream cells at the anchor, with a worst-case one-way ratio of 70x under demand no adversary chose. The impossibility is not an artifact of adversarial scheduling. Nor of stationarity: under drifting and bursting demand the gap survives on 9 of 9 registered cells, with the exact two-way optimum at $N = 10$ and the work-function policy on $4{,}096$ states at $N = 12$ (Appendix R.4).

Q.6 THE PRICED ALLOCATION LOSES, AND WHY (ARM 6, C4)

The one experiment in the program in which the theory could lose to a heuristic on the heuristic's own ground. It lost it decisively, on the pre-declared unfavourable branch, and the mechanism is diagnosable from the canonical CSVs.

Table 29: Class-incremental CIFAR-100, final demand-weighted test error, three canonical seeds, herding selection throughout so that allocation is isolated. The classical square-root law is competitive; it is specifically the externality correction, this paper's own addition, that fails as a global allocator.

Budget	Uniform	$\sqrt{\text{demand}}$	Priced	Priced $-$ $\sqrt{\text{demand}}$
1000	0.170	0.133	0.531	+0.398 [0.344, 0.475]
2000	0.156	0.135	0.309	+0.174 [0.147, 0.208]
5000	0.151	0.139	0.218	+0.079 [0.068, 0.089]

The mechanism, computed from the canonical geometry through the released code: at $d_{\mathrm{eff}} = 766.5$ the repulsion term dominates the crowding term by 7.6 to 1 at the median pair, so the priced score S_c is carried by each class's few confusable neighbors and is near zero, or negative, elsewhere; 34

of 100 classes receive $S_c \leq 0$ and are clamped to the floor, the surviving weights span a ratio of 10^7, and the allocation hands one class its cap of 500 exemplars while others sit at 3. Starving a class costs far more error than the pairwise ledger charges, because the ledger prices the marginal value of adding to a class the system already holds, not the value of holding it at all. That is the honest content of the registration's caveat that the priced rule is a worst-case, population-level object: the completed stocking law (Corollary 4) is a statement about marginal promotion, and using it as a global allocator under a hard budget is an extrapolation the theory does not license and the data now refuse. Tier 4 ran the diagnosis inverted, a starvation floor and a positive own-side price, and the registered branch fired not repaired on both cells: the floor removes the catastrophe, the prices never beat demand-aware baselines, and the question is closed as a stated negative result (Appendix R.1).

Q.7 The jointly refit head (Arm 7, C1)

F-7 did not fire: the ridge rate is [0.913, 0.957] and the cross-entropy rate [0.684, 0.716] against the registered band $[0.7, 1.3]$, with fitted constants 0.64 and 0.61 of the plug-in constant. The scoping sentence the paper states without measurement, that the refit head shares the $1/m$ form with a different constant, is now measured, with the cross-entropy head at the band's edge under the registered fixed-budget protocol and a convergence diagnostic at triple the step budget recorded in the CSVs.

Q.8 The non-vision twin (Arm 8)

Descriptive by registration; no text result can fire or clear a falsifier. On MiniLM over 20 Newsgroups the adequacy slope is 0.786 ([0.784, 0.787]), inside the vision range, and the selection composite's favourable branch replicates. Banking77, registered as the closest public analogue of an agent's memory over confusable topics, reads slope 4.386 [4.346, 4.427], five times the vision range and outside the descriptive band. That number is not evidence against the constant and we do not read it as such: Tier 4's diagnostic (Appendix R.5, Table 34) shows that sentence encoders separate fine-grained intents so widely that the predicted pairwise excess on Banking77 underflows the resolution of this protocol, leaving the adequacy ratio a measurement floor over a vanishing denominator. Banking77 is therefore reported, and excluded from the adequacy evidence, on a criterion fixed by the prediction's magnitude rather than by its answer.

R Tier 4: the hard-scenario program, executed

Tier 3 left four open questions, each named in its own subsection above: do the diagnosed repairs fix the priced allocator (Appendix Q.6), does the stocking law resolve at ten times the replication (Appendix Q.3), does the measured externality ratio track its finite-m prediction upward (Appendix Q.4), and does the flow-accounting gap survive nonstationary demand (Appendix Q.5). Tier 4 answers all four. Its registration (`experiments_design_L1_T4.md`, in the supplement) was written after Tier 3's results were known and before any Tier 4 run, an ordering it discloses; grids, falsifiers, readings, and branch language were fixed there first, and its amendment log (T4-A1 to T4-A4) records every subsequent change with its motivating evidence, each logged before any successful run of the cells it affects. The package is additive by construction: the Tier 4 self-test records the SHA-256 of every Tier 3 source file at registration and fails if any of them changes, so the Tier 3 record above is frozen under everything below. Anchor discipline is unchanged: every Tier 4 falsifier and pre-declared branch is evaluated on ViT-B/16 on CIFAR-100 and nowhere else, and every other cell is descriptive. The Tier 4 execution gate (G5 to G7) passed on the execution machine immediately before the canonical runs: the six allocators separate on synthetic features with zero floor violations, the synthetic externality ratio at m=32 reads 1.829 against a finite-m prediction of 2.990 (ratio 0.61, inside the registered factor-two band, read at m=32 per amendment T4-A1), and the N=12 state system builds all $4{,}096$ states with the work function inside its bound.

R.1 The repaired allocator does not repair the loss (Arm 9)

Arm 6's diagnosis (negative S_c clamping, a weight span of seven orders, class starvation) implies specific repairs, and Arm 9 runs exactly the diagnosis inverted, registered before execution: a per-class floor at half the uniform share under the Tier 3 priced weights, the own-side price $w_c =$

$\sqrt{\Lambda_c \sum_j (3A_{cj} + B_{cj})}$, which is positive by construction, and their combination, against the Tier 3 allocators under the identical class-incremental protocol.

Table 30: Arm 9, final demand-weighted error by allocator and budget on both cells, canonical seeds. The registered branch fired *not repaired* on the anchor and replicates on Tiny-ImageNet; F-9 (machinery) clean, floors respected everywhere.

	CIFAR-100 (anchor)			Tiny-ImageNet		
Allocator	M=1000	2000	5000	M=1000	2000	5000
Uniform	0.152	0.145	0.140	0.115	0.111	0.108
$\sqrt{\text{demand}}$	0.119	0.124	0.133	0.089	0.093	0.100
Priced	0.584	0.420	0.236	0.406	0.296	0.156
Priced + floor	0.203	0.145	0.132	0.204	0.119	0.110
Own-price	0.675	0.423	0.307	0.416	0.318	0.189
Own-price + floor	0.212	0.153	0.138	0.207	0.120	0.110

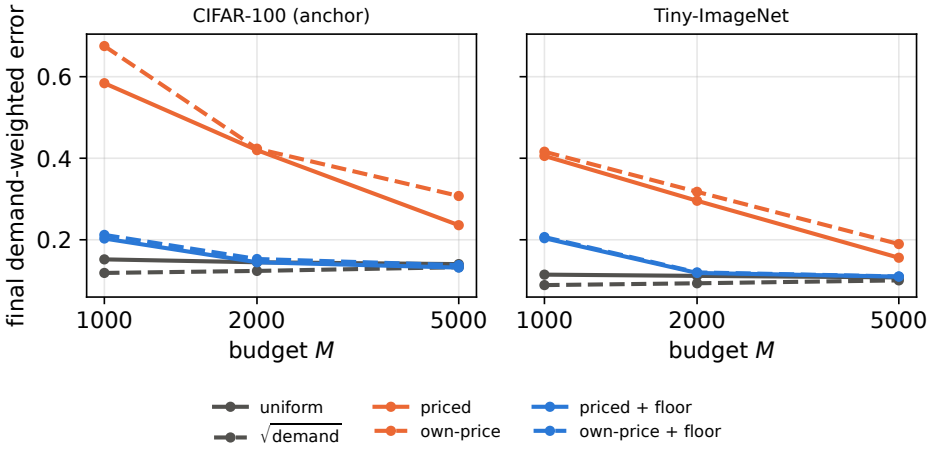

Figure 25: Arm 9: unfloored priced allocation is catastrophic at small budgets, the floored variants collapse to baseline level, and no repaired variant beats both baselines at a majority of budgets. Solid against dashed distinguishes members within each family.

The registered verdict is **not repaired** on the anchor and **not repaired** on Tiny-ImageNet (Figure 25): no repaired variant beats both uniform and $\sqrt{\text{demand}}$ at a majority of budgets. The table still earns its place, because the failure changes character. Unfloored priced allocation is catastrophic at small budgets (0.584 against uniform's 0.152 at M=1000; the pure own-side price is worse still at 0.675). The floor removes the catastrophe, cutting the priced allocator's error by a factor of 2.9 at that budget, and what remains is a gap to the better baseline that closes as the budget grows: 0.0844 at M=1000, 0.0212 at M=2000, -0.0009 at M=5000, where the priced floor is ahead on the point estimate but well inside the noise (paired difference -0.0086 against uniform, 95 percent interval [-0.0144, 0.0011], straddling zero). The repairs buy back the pathology and nothing more: at no budget does a repaired variant clear both baselines with an interval that excludes zero. The reading is the one the registration wrote for this branch: the floor removes the starvation pathology, but the pairwise ledger does not carry enough information for global allocation even after the diagnosed corrections, and the prices add nothing over demand-aware baselines. The question Arm 6 opened is closed as a stated negative result.

R.2 The stocking law is unresolvable at this scale (Arm 10)

Arm 3 at ten times the replication ($R = 200$, all 100 classes, counts to $m = 256$), with a power gate registered so the verdict cannot be fished: F-3 is read only if the 95 percent interval at the registered configuration is at most 0.30 wide. The realized interval is [0.252, 1.077] (width 0.825), so the F-3

verdict is NOT read, and the pre-declared outcome stands: *underpowered at ten times the replication*, reported as a design limit of measured-excess stocking experiments on learned features, and the question closes as unresolvable at this tier rather than being retried. The landscape is still informative as description: censoring at the registered configuration is 98 percent even at this replication, 5 of nine grid cells are estimable at all, and where estimable the point slopes span 0.50 to 0.66, bracketing the square-root law's 0.5 without the power to test it.

R.3 The externality ratio does not track its prediction (Arm 11)

F-4's verdict is already recorded above and this arm cannot alter it; Arm 11 determines which reading of the fire the paper carries, on the balanced grid extended to $m \in \{32, 64, 128, 256\}$ with $R = 400$ and the 300 most confusable eligible pairs.

Table 31: Arm 11, the externality ratio against its finite-m prediction as m grows. A pair is measurable at a cell when its neighbor-side marginal has the theory's sign; medians are over measurable pairs.

m	pairs	measurable	of which zero	median measured $\hat{\rho}$	median predicted
32	300	258	9	1.531	2.249
64	300	213	17	1.000	2.762
128	300	137	23	0.400	3.026
256	300	107	37	0.000	3.157

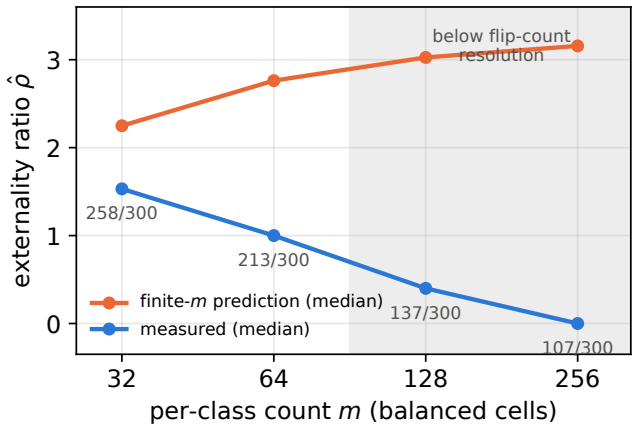

Figure 26: Arm 11: the finite-m prediction climbs while the measured median falls, and beyond $m = 128$ the one-extra-sample marginals sit below the evaluation set's flip-count resolution (shaded; fractions are measurable pairs of 300). The reading rests on the unshaded points.

The registered reading is **flat** (Figure 26). The honest gloss has two halves. At the well-measured end, $m = 32$, where 258 of the 300 pairs measure cleanly, the measured median sits at 1.531 against a predicted 2.249, a fraction of 0.68, and it does not track the prediction's climb thereafter. At the top of the grid the one-extra-sample marginals fall below the evaluation set's flip-count resolution: at $m = 256$ only 107 of 300 pairs are measurable at all, 37 of those measure an own-side marginal of exactly zero, and across those measurable pairs the median own-side marginal is 8.3×10^{-6} against a neighbor-side 2.5×10^{-5}, so the zero median at $m = 256$ is a statement about the measurement floor, not about the constant. The origin slope κ reads 8.3×10^{4} (bootstrap interval 5.2×10^{3} to 2.1×10^{5}, over 715 measurable cells) and is reported because it was registered, with the warning that a ratio statistic with near-zero denominators is dominated by its tail; the medians above carry the reading, as the registration anticipated. Both halves point the same way: the finite-m prediction overshoots the measured externality on learned features, the error is on the permissive side, and the paper's standing guidance, read 3Λ as permissive, holds where measurement is intact, up to $m = 128$.

R.4 THE DICHOTOMY SURVIVES NONSTATIONARY DEMAND (ARM 12)

Arm 5's system under demand no stationary model produces: a Zipf stream whose rank permutation is redrawn every τ queries, a Markov modulated burst stream, and the registered two-stage flood as the adversarial reference, at $N = 10$ with the exact two-way optimum, and at $N = 12$ the work-function policy on $4{,}096$ states, exact, no approximation.

Table 32: Arm 12 at $N = 10$ on the anchor, canonical seeds: realized cost ratios against the exact two-way optimum. The registered reading counts the drift and burst cells; the flood rows are the adversarial reference.

Stream	C_{con} mult.	one-way	FBE	WFA (two-way)
flood (reference)	1	107495.7	17.2	1.00
drift, $\tau = 500$	1	40.0	9.8	1.31
drift, $\tau = 2000$	1	35.0	9.5	1.29
burst	1	211.7	10.1	1.25
flood (reference)	10	10749.1	3.4	3.87
drift, $\tau = 500$	10	7.5	2.5	1.14
drift, $\tau = 2000$	10	6.8	2.5	1.09
burst	10	37.3	2.4	1.34
flood (reference)	100	1074.8	2.6	2.08
drift, $\tau = 500$	100	6.3	2.2	1.22
drift, $\tau = 2000$	100	6.2	2.4	1.06
burst	100	21.5	2.0	2.04

The registered reading: the flow-accounting gap survives nonstationarity on 9 of 9 drift and burst cells, with a maximum nonstationary one-way ratio of 211.7 (burst). The two policies that may demote separate here and the table reports both: the work function tracks the exact optimum closely, at most 2.04 times it on the drift and burst cells and 3.87 including the flood reference, while flow break-even, which may demote but decides greedily, reaches 10.09 and is therefore not a substitute for the work function under nonstationarity. Both stay inside the Theorem's bound at $N = 12$, where the one-way flood reference reaches 35004. Drifting and bursting demand does not rescue one-way accounting; it feeds it.

R.5 SCALE AND BREADTH CELLS (ARM 13, DESCRIPTIVE)

Table 33: Tier 4 vision scale cells (descriptive): the language-supervised scale point and two large backbones off the anchor dataset.

Backbone	Dataset	d_{eff}	Arm 1 slope	Arm 4 $\hat{\rho}$ (m=32)	Arm 6 branch
CLIP ViT-L/14	CIFAR-100	695.2	0.736 [0.732, 0.739]	1.29	unfavourable
ViT-L/16	Food-101	1016.7	0.861 [0.859, 0.863]	1.51	unfavourable
DINOv2-B/14	Food-101	766.6	0.914 [0.911, 0.917]	1.46	unfavourable

CLIP ViT-L/14 is the first strongly sub-ambient backbone in the program ($d_{\text{eff}} = 695.2$ of 768), and the adequacy anchor, the unfavourable Arm 6 branch, and the descriptive externality behaviour all replicate on it and on the Food-101 large backbone cells.

The diagnostic in that table changes what the text cells are worth, and the correction is worth stating plainly. Sentence-embedding *intent* datasets do not test the constant at all. Their classes are separated by debiased margins in the hundreds (240 to 790 at the median, against 42 on the topic datasets), and at margins that size the closed form predicts a pairwise excess of 10^{-5} down to 10^{-32}: an interference the theory says is, for practical purposes, absent. The measured excess in those cells is not a measurement of that quantity but the floor of the protocol itself, and the adequacy ratio is then a floor divided by a vanishing prediction. That is why Banking77 reads 4.386 and 5.432, far outside the descriptive band, while CLINC150 reads 0.899 and 0.771, apparently excellent: both are the same arithmetic accident, and neither is evidence about the constant. Of the 8 text cells only 2, the topic datasets 20 Newsgroups and AG News, sit where the prediction is resolvable, and those two

Table 34: Every text cell of the program, Tier 3 and Tier 4 together (descriptive; no text cell can fire or clear anything), with the diagnostic that decides whether its adequacy ratio is interpretable: the median debiased margin and the median finite-m excess that margin predicts, against the median measured excess. A cell is marked *not interpretable* when the predicted excess is below the floor this protocol can resolve.

Encoder	Dataset	median $\hat{\gamma}$	median pred.	median meas.	slope	reading
MiniLM	20 Newsgroups	42	4.4×10^{-2}	5.8×10^{-2}	0.786	in band
mpnet	AG News	36	1.6×10^{-1}	1.5×10^{-1}	0.923	in band
MiniLM	Banking77	345	1.7×10^{-14}	2.3×10^{-3}	4.386	not interpretable
mpnet	Banking77	538	1.4×10^{-20}	1.4×10^{-3}	5.432	not interpretable
MiniLM	CLINC150	449	6.9×10^{-20}	0	0.899	not interpretable
mpnet	CLINC150	790	1.1×10^{-32}	0	0.771	not interpretable
MiniLM	MASSIVE	240	1.2×10^{-5}	2.3×10^{-2}	2.007	not interpretable
mpnet	MASSIVE	349	1.8×10^{-6}	3.8×10^{-2}	2.611	not interpretable

read 0.786 and 0.923, inside the band and inside the vision range. The descriptive adequacy check registers a fire on 3 of the 8 cells, all of them in the degenerate group, reported as uninterpretable rather than as evidence against the theory; the honest content of the text twin is therefore narrower than the cell count suggests, and the program's real text evidence is two cells, not eight. We have not diagnosed why sentence encoders separate fine-grained intents this widely, and we do not speculate here; the practical consequence is that an adequacy test of this ledger needs a domain where classes actually interfere.

Provenance for every text cell is signature-checked: under `datasets` version 4 the script-era hub repositories retired unevenly, so amendments T4-A2 to T4-A4 register an ordered candidate chain of sources per dataset, and a candidate is accepted only if the loaded data matches the registered signature, the realized class count and the official split sizes. The check did real work: the maintained mirror of Banking77 silently deduplicates 14 rows (9,993 train rows against the official 10,003) and was rejected, and the cells run instead on the primary-source CSVs of the official PolyAI repository, the exact files the retired hub loader used to read. CLINC150 loads from the official repository under its official `plus` configuration, and MASSIVE from a maintained mirror whose realized signature equals the official en-US counts exactly. The winning source, verified signature, and label decoding of every cell are recorded with the data.

S E-R2: A WIDE MATRIX RUN IN A PARALLEL LINE, ITS FIRED FALSIFIERS, AND A POST HOC RE-ANALYSIS

This appendix reports a registered program that was run by a parallel development line of this project and whose prediction column was built on the pipeline this paper replaced. It is here because every registered falsifier outcome is reported, fired ones as stated negatives, and because the measurements themselves, 85,802 pairwise-excess cells across 11 backbones and 10 datasets, are the widest real-feature matrix the project owns. Nothing in it changes a number elsewhere in the paper: no Tier 3 or Tier 4 falsifier depends on it, and the post hoc part below is labeled as such throughout.

Provenance. E-R2 (`experiments_design_L1_T2ER2.md`, registered 2026-09-09, in the supplement as a hash-anchored marked redaction) was written and run by a development line that branched from this project at the Tier 1 state of 2026-08-16 and never received amendment T2-A3 of the canonical line (Appendix Q): its runner whitens by the total covariance, takes the raw participation ratio of the total spectrum as the effective dimension, uses the undebiased plug-in margin, and predicts with the asymptotic constant $a(\gamma, d)(1/m_c + 1/m_j)$, the combination whose failure the Tier 3 gate documents ($d_{\text{eff}} = 10.1$ recovered against a truth of 64). The matrix is 65 backbone-dataset combos at native width (the eight vision backbones on seven vision datasets, the three sentence encoders on three text datasets; seed `11`, R=12, pair cap 240), a deep tier of 12 combos at R=50 with pair cap 600, and a PCA ladder on 2 combos; arms 2, 3 and 5 of the Tier 2 design run on the deep-tier combos over the canonical seeds. One reading of the manifests needs care: the deep-tier arm-1 files list the seed set $\{11, 13, 17\}$, but the runner draws every arm-1 resample stream from the first listed

seed, so arm 1 is a seed-11 measurement at every replication level, and only arms 2, 3 and 5 iterate over the three seeds. Execution was on Colab on 2026-09-09 and 2026-09-10 under two runner versions, both recorded by hash in every manifest: the first failed on five dataset or backbone loaders, the parallel line patched them mid-program and re-ran exactly the failed combos and the deep tier under the second. That patch and a staging change were logged in the line's working notes but not appended to the registration's amendment log, which stops at T2E-A3; the registration, both runner versions' hashes, the second version's source, the patch, the failure notes, a labeled rerun of the synthetic gates, and all 85,802 cells with their manifests ship in the supplement (`data_er2/`), and every number here enters through `fill_er2.py`; the first version's source itself was not archived and cannot be reconstructed from the patch, which the README records. The manifests record the numpy and scipy versions only; the runtime's torch, timm, sentence-transformers and datasets versions, its Python and CUDA versions and its GPU were not recorded and are reported here as unrecorded rather than reconstructed. Two label collisions are worth a sentence: the parallel line's copy of the Tier 2 registration logs a different amendment under the name T2-A3 (a bootstrap-key fix), and its own amendments carry the prefix T2E; in this paper T2-A3 always means the canonical repair of 2026-08-20.

The registered outcomes, as they fired. Table 35 evaluates the five registered falsifiers with the program's own definitions on the program's own prediction column; every point estimate that the parallel line's own analysis record contains is reproduced to floating-point precision, and the cluster-bootstrap intervals are recomputed here under a fixed seed. Four of the five fired.

Table 35: E-R2 registered falsifiers, evaluated as registered (validity region $m\hat{\gamma}/d_{\text{eff}} \geq 30$ on the program's own d_{eff}, cluster bootstrap by combo and pair, band $[0.5, 2.0]$). Every verdict is the registered reading; the fired ones are stated negatives for the program as run. K-R2-1 is registered as the pool of the wide and deep tiers; per tier it reads 65.2 (63.4, 67.3) on 15,540 cells (wide) and 17.9 (17.0, 18.9) on 14,965 cells (deep), and the parallel line's own record, whose tier filter matched every combo, reports the all-tier pool, 44.4 (42.3, 46.6) on 32,964 cells; it fires under every reading.

Falsifier	Registered test	Outcome	Verdict
K-R2-1	origin slope, wide and deep tiers pooled; the 95% CI must meet $[0.5, 2.0]$	63.9 (62.2, 65.6) on 30,505 cells	fired
K-R2-2	the same per rung of the PCA ladder	Table 37, registered columns	fired
K-R2-3	herding beats random at k=8 in at least 70% of combo-seed pairs, one-sided $p < 0.05$	22 of 36 (12 combos, three seeds), $p = 0.12$	fired
K-R2-4	pooled log-log exponent of m^* on Λ inside $[0.3, 0.7]$	0.18 over 10,080 decisions	fired
K-R2-5	at least 80% of stream decisions within a factor of four of the 3Λ count	94.7% of 660 decisions	not fired

What the prediction column measured. The fired slopes are not small disagreements. On the anchor combo, ViT-B/16 on CIFAR-100, the median plug-in margin is $\hat{\gamma} \approx 96$ and the program's effective dimension is $d_{\text{eff}} = 70$, the participation ratio of a total spectrum dominated by a few between-class directions, where the whitened within-class noise of the same features occupies essentially the full width (Tier 3 measures $d_{\text{eff}} = 766.5$ of 768 on the anchor, Appendix Q.1). Across the wide tier the native width exceeds the program's d_{eff} by a factor of 2.1 to 547 (median 13). Two consequences follow. First, the registered validity filter admitted cells that are nowhere near the asymptotic regime: of the 30,505 validity cells, 249 (0.8%) still clear $m\gamma/d \geq 30$ when d is the native width, and on the anchor the largest value of $m\hat{\gamma}/d$ at native width is 11. Second, the asymptotic constant is exponentially small in the margin (its leading factor is the Gaussian density at $\sqrt{\gamma}/2$), so at m=4 on the anchor it predicts an excess 6×10^4 times smaller than the measured 0.178, whereas the finite-m prediction of T2-A3(d) at the native width gives 0.023 for the same cells. The plug-in bias of the margin, the third ingredient of T2-A3, is small here: $d(1/n_c + 1/n_j)$ is 3.07 on the anchor ($n = 500$ per class). The ladder makes the same point in one row: the registered slope climbs from 0.65 at sixteen dimensions to 2×10^3 at full width because the asymptotic formula leaves its regime as d grows, not because the features change character. K-R2-4 and K-R2-5 are computed from stocking counts that use the same constants, and their outputs sit at the floor of the search rather than on any law: the measured-ε stopping rule returned a single exemplar in 6,840 of 10,080 arm-3 decisions (68%) and never more than 10 against a clamp of 128 on the large-class combos, so the pooled exponent of 0.18 is the slope of a floor, and in arm 5 the 3Λ count is one in 90% of the 660 stream decisions (never above 6) with both counts equal to one in 390 of them,

so the 94.7% inside the factor-four band is agreement between two floors. Neither verdict carries information about the stocking law, in either direction. K-R2-3 compares two selection policies on measured error and does not involve the prediction column, so its fire stands on its own: herding beats random selection at k=8 in 22 of 36 combo-seed pairs, short of the registered 70 percent, a comparison Tier 3 does not make (its Arm 2 compares the registered composite against herding, Appendix Q.2).

Post hoc re-analysis, labeled as such. The measured excess in every cell is a legitimate measurement of the pairwise nearest-mean excess in the program's whitened space, so the same cells can be re-predicted with the canonical finite-m Gaussian prediction (T2-A3(d), `finite_m_prediction`) evaluated at the native width of the backbone (an upper bound on the whitened noise dimension; the program did not record the whitened within-class spectrum, and Appendix T measures it on the preserved caches: at least 0.88 of the width on every vision combo) or at the PCA target on the ladder, and at the plug-in margin minus $d(1/n_c + 1/n_j)$ with n the smallest class of the combo. Each of these choices errs toward predicting more excess, so the slopes below are, under the Gaussian model, generous to the prediction rather than to the measurement. One protocol fact runs the other way: the archived runner draws exemplar buffers with replacement (inherited from the Tier 2 machinery; Tier 3 replaced it with nested buffers under T3-A1), so a buffer of nominal size m holds fewer distinct points, about 6% fewer at m=64 on the anchor's classes of 500 and 16% fewer at m=8 on Flowers-102's classes of twenty, which raises the measured excess by a comparable relative amount. Cells whose repaired prediction falls below the Tier 4 floor of 10^{-3} are set aside as not interpretable, as in Appendix R, and a combo is summarised only when at least 20 of its cells remain. This was done after the registered outcomes were known and is not a registered test; it has no falsifier and no verdict.

Table 36: Post hoc, wide tier, by dataset: combos, smallest class per combo, the ratio of the bias correction to the median plug-in margin, the registered slope range (with the number of combos that had any validity cell where fewer than all did), the repaired slope (median and range over combos with at least 20 interpretable cells), and the count inside $[0.5, 2.0]$. The daggered datasets are the registration's own small-class cases, where the correction is a large fraction of the margin and the repaired numbers are correspondingly fragile.

Dataset	Combos	$n_{\min}$	Bias/margin	Registered slope	Repaired slope	In band
CIFAR-10	8	5000	0.01 to 0.05	29.44 to 53.57 (3 of 8)	1.29 (1.13 to 1.46)	8 of 8
STL-10	8	500	0.11 to 0.44	29.27 to 61.98 (4 of 8)	1.24 (1.04 to 1.44)	8 of 8
EuroSAT	8	1600	0.04 to 0.15	45.86 to 1×10^2	1.32 (1.28 to 1.51)	8 of 8
CIFAR-100	8	500	0.02 to 0.09	1×10^2 to 3×10^5	5.56 (1.31 to 8.95)	2 of 8
Food-101	8	750	0.02 to 0.10	54.80 to 3×10^3	2.62 (1.01 to 7.09)	3 of 8
Flowers-102†	8	20	0.38 to 1.26	5×10^4 to 5×10^7 (2 of 8)	3.10 (0.40 to 26.35)	2 of 8
Tiny-ImageNet	8	500	0.02 to 0.06	2×10^2 to 2×10^9	3.20 (1.41 to 21.53)	1 of 6
AG News†	3	100	0.96 to 1.65	no validity cell	37.97 (1.88 to 41.78)	1 of 3
20 Newsgroups	3	377	0.11 to 0.20	no validity cell	0.92 (0.86 to 0.92)	3 of 3
banking77†	3	35	0.26 to 0.46	5×10^3 to 5×10^5	1.12 (0.91 to 1.30)	3 of 3

Table 37: The PCA ladders, registered and repaired. Registered: validity cells and the origin slope on the program's prediction column, the K-R2-2 test. Repaired: interpretable cells and the origin slope on the finite-m prediction at the rung's dimension. The full-width rung is the anchor protocol of the wide tier at R=25.

Combo	d	Validity cells	Registered slope	Interp. cells	Repaired slope
ViT-B/16, CIFAR-100	16	451	0.65 (0.51, 0.80)	1184	0.58 (0.51, 0.64)
	64	556	3.79 (2.73, 5.43)	666	0.90 (0.78, 1.16)
	256	502	18.72 (8.35, 141.57)	347	1.39 (0.82, 3.91)
	768	476	2×10^3 (1×10^3, 4×10^3)	535	4.61 (4.20, 5.08)
BGE, 20 Newsgroups	16	310	0.73 (0.58, 0.89)	950	0.58 (0.54, 0.62)
	64	164	2.63 (2.27, 2.99)	950	0.86 (0.82, 0.89)
	256	0	no validity cell	950	0.95 (0.93, 0.98)
	768	0	no validity cell	950	0.92 (0.90, 0.94)

Read with those caveats, the re-analysis says three things (Tables 36 to 38 and Figure 27; the full per-combo table ships with the supplement as `generated/er2_combos.tex`). On the 51 combos

Table 38: The deep tier against its wide-tier twins, repaired slopes with interpretable cell counts. Twelve combos at $R{=}50$ (seed 11 stream) with pair cap 600 and m up to 128 reproduce the wide-tier values of the same combos to within a factor of 0.82 to 1.25; the deep tier is the more replicated measurement and changes no reading.

Backbone	Dataset	Registered	Wide, repaired	Cells	Deep, repaired	Cells	Deep/wide
BGE	20 Newsgroups	16.22	0.92 (0.90, 0.93)	950	0.91 (0.90, 0.93)	1140	1.00
CLIP ViT-B	CIFAR-100	65.49	2.77 (1.96, 4.29)	519	2.67 (2.09, 3.71)	1277	0.96
DINOv2-B	CIFAR-100	$2{\times}10^3$	6.85 (5.42, 9.29)	305	7.19 (6.14, 8.64)	766	1.05
DINOv2-B	Flowers-102	none	4.71 (4.28, 5.17)	276	4.91 (4.68, 5.17)	680	1.04
DINOv2-L	Food-101	$7{\times}10^2$	7.09 (5.70, 9.56)	361	5.83 (4.83, 7.25)	922	0.82
E5	banking77	$4{\times}10^5$	0.91 (0.85, 0.99)	673	0.89 (0.85, 0.93)	1713	0.97
MiniLM	20 Newsgroups	none	0.92 (0.90, 0.95)	950	0.92 (0.90, 0.95)	1140	1.00
ResNet-50	CIFAR-100	$4{\times}10^2$	1.99 (1.85, 2.15)	993	1.96 (1.87, 2.05)	2632	0.99
Swin-B	CIFAR-100	$1{\times}10^4$	8.48 (6.20, 12.12)	367	10.60 (9.32, 12.33)	917	1.25
ViT-B/16	CIFAR-100	$1{\times}10^3$	4.88 (4.36, 5.50)	536	4.13 (3.85, 4.45)	1398	0.85
ViT-B/16	Tiny-ImageNet	$2{\times}10^4$	11.63 (8.25, 20.57)	100	10.94 (7.87, 17.00)	249	0.94
ViT-L/16	CIFAR-100	$1{\times}10^3$	6.23 (5.23, 7.55)	372	5.53 (4.59, 6.93)	928	0.89

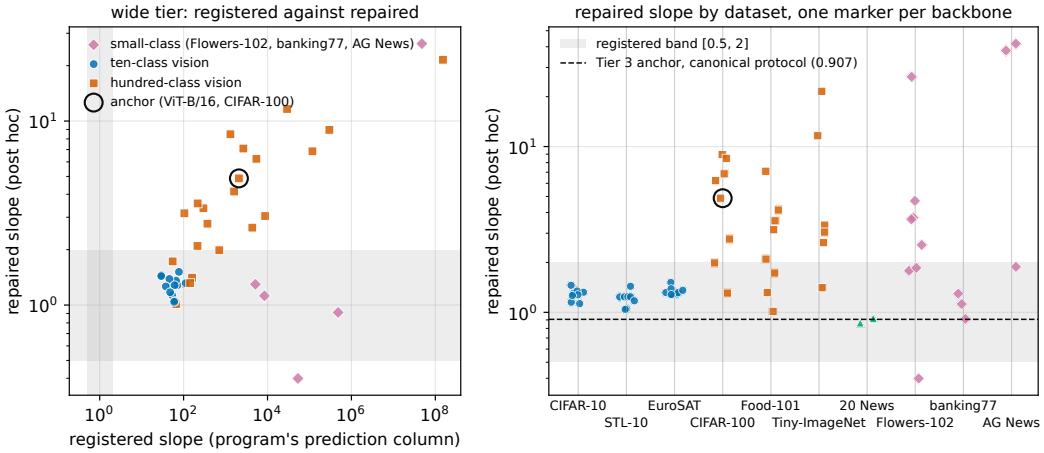

Figure 27: The wide tier, post hoc. Left: registered against repaired slope per combo (log-log; combos with a registered value and at least 20 interpretable cells), the registered band shaded on both axes. Right: the repaired slope of every wide-tier combo by dataset, one marker per backbone, the band shaded and the canonical Tier 3 anchor reading as the dashed line. The ten-class vision datasets and 20 Newsgroups sit in the band on every backbone; the hundred-class datasets do not, and the anchor (circled) reads 4.88 here against 0.907 under the canonical protocol.

outside the small-class datasets, 49 are interpretable and 33 of them sit inside $[0.5, 2.0]$ (36 within a factor of three; median 1.34, range 0.86 to 21.5), against registered slopes on the same combos from 29.27 to $2{\times}10^9$. The agreement is not uniform, and its structure is the useful part: all 24 combos on the ten-class vision datasets (CIFAR-10, STL-10, EuroSAT) are in band, between 1.04 and 1.51, and all three encoders on 20 Newsgroups are in band, whereas on the hundred-class vision datasets (CIFAR-100, Food-101, Tiny-ImageNet) only 6 of 22 combos are, with a median of 3.25 and a range of 1.01 to 21.5. The anchor reads 4.88 (4.36, 5.50) here against $2{\times}10^3$ registered and 0.907 under the canonical Tier 3 protocol on the same backbone and dataset (Table 28). The remaining factor on the many-class datasets is therefore a property of the program's protocol, not of the features, and the candidate cause we can name is the whitening: whitening by the total covariance compresses exactly the between-class directions in which class means differ, one per class up to a constant, so the Gaussian model with unit isotropic noise in that space misdescribes the noise where it matters most, and the more classes the more such directions. We recorded this as a hypothesis when the re-analysis was done; it cannot be tested on these outputs, and Appendix T tests it on the same features under a registration frozen for that purpose, with the canonical protocol (within-class whitening, the measured whitened noise dimension, the debiased margin) and a named falsifier, which did not fire. On the 14 small-class combos the correction is between 0.26 and 1.65 of the margin, 6 are in band,

and the range (0.40 to 41.8) says only that the plug-in margin is not resolvable at twenty to a hundred training points per class in several hundred dimensions, the regime fence of Remark 3 again.

What this appendix claims. Two things, and no more. The registered program fired four of its five falsifiers as stated, and those verdicts are reported here as the negative results they are for the program as run; a reader who wants the prediction column that was registered has it, with the code that produced it. And the same measurements, re-predicted post hoc with the repaired prediction of this paper, are consistent with the canonical Tier 3 reading on the datasets where the protocol difference is smallest and diverge from it where the candidate cause says they should; that consistency is evidence about the protocol, not a test of the theory, and the paper's registered real-feature evidence remains Tiers 3 and 4. The parallel line's feature caches, one file per combo with the full training and test splits at native width, are preserved, so the canonical protocol can be run on the same features under a new registration; that is Appendix T.

T E-R2C: THE CANONICAL PROTOCOL ON THE E-R2 FEATURE CACHES, A REGISTERED TEST OF THE WHITENING HYPOTHESIS

Appendix S ends with a hypothesis it could not test on the outputs it had (that the remaining factor on the many-class datasets was the parallel line's total-covariance whitening, not the features) and with the remark that the parallel line's feature caches were preserved. This appendix is that test. A third registration, E-R2C, was frozen on 2026-09-11 before any computation on the caches; it applies the Tier 3 arm-1 protocol, unchanged and byte for byte, to the 65 cached combos of E-R2, with three falsifiers written against the hypothesis and against the Tier 3 anchor; of the three, none fired. 64 of the 65 combos completed; the one that did not is a stated failure (Table 42). Every number below enters through `fill_er2c.py` from manifest-backed CSVs, and nothing here changes a number in Tiers 3 or 4.

Provenance. The registration is `experiments_design_L1_ER2C.md` (in the supplement as a hash-anchored marked redaction, `colab_er2c/`). Its data are the 65 `.npz` caches the parallel line wrote during E-R2, one per backbone-dataset combo, the full official splits at native width, produced by the parallel line's archived (second-version) extraction code of `data_er2/code_parallel_line/` for the combos that ran after its mid-program patch, and for the first pass (six square-image vision backbones on CIFAR-10, STL-10, EuroSAT and CIFAR-100, and the three encoders on AG News and 20 Newsgroups) by the first version, whose source is not archived and which, per the patch record, differs from the archived code only in loader paths those datasets do not exercise; the E-R2C manifests pin the features themselves by hash. The runner records the SHA-256, shapes, dtypes, class count and smallest class of every cache in the combo's manifest, never writes into the cache folder, and treats a cache that does not match the registered description as a stated omission. The runner, `er2c_run.py`, imports the Tier 3 library, arms and feature module (`t3_lib`, `t3_arms`, `t3_features`, `t3_policies`) as byte-identical copies whose hashes every manifest records beside the driver's, and adds only a cache reader in place of feature extraction; the protocol is therefore Tier 3 arm 1 as registered on 2026-08-20 with amendments T3-A1 and T3-A3 in force: the pool rule (a smallest training class below 256 pools train and test and re-splits 70/30 at the fixed split seed), the test cap of 100 per class, pooled within-class Ledoit-Wolf whitening (shrinkage 4×10^{-5} to 0.30 across the matrix), the participation ratio of the whitened within-class spectrum as d_{eff} with the raw participation ratio of the total spectrum recorded beside it, the plug-in margin debiased by $d_{\text{eff}}(1/n_c+1/n_j)$, nested per-class buffers at counts $\{2,\dots,256\}$ intersected with the smallest pool plus the four asymmetric cells, $R = 200$ replicates per seed on the canonical seeds 11, 13, 17, the finite-m Gaussian prediction of T2-A3(d), and the origin slope with the Tier 3 replicate bootstrap. The anchor, ViT-B/16 on CIFAR-100, runs on all 4,950 pairs as in Tier 3; every other combo on a seeded subsample of at most 1,000 pairs, logged in the manifest. The execution gate is the arm-1 component of the Tier 3 gate on synthetic features of known d_{eff}, and the driver refuses a canonical run unless a passing gate for the current code hashes is on disk; every canonical manifest records that precondition. One amendment exists, ER2C-A1, logged with its evidence before any smoke or canonical cell touched a cache: the gate as registered required the synthetic slope interval to contain 1, and on the run-time GPU it read $[0.971, 0.999]$, a miss of 0.001, where the same code read $[0.973, 1.002]$ on Tier 3's A100 and $[0.983, 1.019]$ on a CPU; the criterion was amended to an interval inside $[0.5, 2.0]$ within 0.02 of 1, the recovered d_{eff} tolerance unchanged,

and the registration's freeze hash before the entry is recorded in the entry. The gate then passed, the smoke run on the canon-excluded seeds 0, 1, 2 completed on the anchor (stamped non-canonical, in its own folder), and the canonical run took 552 seconds of wall clock on one NVIDIA RTX PRO 6000 Blackwell Server Edition (torch 2.11.0+cu128, a single code-hash set across all manifests; the anchor took 13 seconds). DINOv2-L on Flowers-102 failed inside the arm and is reported as it failed: the archived traceback locates the exception after the fit, in the descriptive per-side split ratio $(3A + B)/(B - A)$ at the median geometry, whose two asymptotic constants are both proportional to the Gaussian density at half the root margin and underflow to exactly zero in double precision once the debiased margin exceeds about $6{,}000$ (the other Flowers-102 combos reach median margins of 3,653); the regime fence reached numerically, on the one Flowers-102 combo the arm could not finish. The arm writes no output when it raises, the combo was not re-run under a patched arm, and the fill script lists it and writes no number for it.

Table 39: E-R2C registered falsifiers, evaluated as registered. The Tier 3 anchor slope enters K-C2 through its macro, 0.907 at registration and now. K-C3's interpretability rule is the registered one, at least 20 cells with a predicted excess above the Tier 4 floor of 10^{-3}; a stricter descriptive fence is reported in the text.

Falsifier	Registered test	Outcome	Verdict
K-C1	anchor adequacy, the Tier 3 F-1 form: the origin-slope interval must meet $[0.5, 2.0]$	0.874 (0.871, 0.877) on 59,400 cells	not fired
K-C2	replication: the anchor slope within a factor of 1.25 of the Tier 3 anchor, [0.726, 1.134]	0.874 against 0.907, ratio 0.963	not fired
K-C3	the whitening hypothesis: hundred-class median over ten-class median in $[0.5, 2.0]$, interpretable vision combos	0.882 (24 combos) over 0.905 (19 combos), ratio 0.974	not fired

The registered outcomes. Table 39. K-C1 and K-C2 concern the anchor and are read below. K-C3 is the test the previous appendix asked for: under within-class whitening the hundred-class vision combos (CIFAR-100, Food-101, Tiny-ImageNet) have a median slope of 0.882 and the ten-class ones (CIFAR-10, STL-10, EuroSAT) 0.905, a ratio of 0.974, where the same ratio under the parallel line's protocol, post hoc, was 2.54 (3.25 over 1.28, Appendix S); 23 of the 24 interpretable hundred-class combos and 19 of 19 ten-class combos sit inside $[0.5, 2.0]$. The many-class departure of Appendix S was the protocol's, as hypothesised, and the hypothesis stands as a registered result rather than a conjecture.

Table 40: The anchor by count cell, Tier 3 (its own extraction of ViT-B/16 on CIFAR-100, official splits, canonical seeds) against E-R2C (the parallel line's cache of the same backbone and dataset, canonical seeds), origin slope with the replicate-bootstrap interval; diagonal cells first, then the asymmetric cells.

(m_c, m_j)	Tier 3 slope (CI)	E-R2C slope (CI)
(2, 2)	1.809 (1.794, 1.823)	3.064 (3.047, 3.082)
(4, 4)	1.328 (1.312, 1.345)	1.596 (1.585, 1.609)
(8, 8)	1.260 (1.239, 1.280)	1.193 (1.181, 1.206)
(16, 16)	1.481 (1.449, 1.514)	1.214 (1.196, 1.233)
(32, 32)	1.944 (1.875, 2.011)	1.430 (1.395, 1.466)
(64, 64)	2.019 (1.865, 2.176)	1.596 (1.526, 1.667)
(128, 128)	0.284 (−0.052, 0.619)	1.525 (1.385, 1.670)
(256, 256)	−4.540 (−5.194, −3.936)	1.000 (0.773, 1.242)
(4, 64)	0.914 (0.910, 0.917)	0.878 (0.875, 0.881)
(4, 256)	0.931 (0.927, 0.934)	0.858 (0.855, 0.860)
(8, 128)	0.531 (0.528, 0.534)	0.737 (0.735, 0.740)
(16, 256)	0.411 (0.407, 0.415)	0.442 (0.439, 0.445)

The anchor, two extractions. On the parallel line's cache the anchor reads 0.874 (0.871, 0.877) over 59,400 cells, against 0.907 (0.904, 0.911) in Tier 3, a ratio of 0.963; the whitened-noise dimension is $d_{\text{eff}} = 766.8$ of 768 against 766.5 in Tier 3, and the raw participation ratio of the total spectrum on the same cache is 69.9, which is the d_{eff} of 70 that the parallel line's program used, reproduced on its own features (Tier 3's extraction gives 91.4 for the same pre-amendment quantity). The two programs do not share features. Both name the timm model `vit_base_patch16_224`

(Tier 3 resolves the pretrained tag `vit_base_patch16_224.augreg2_in21k_ft_in1k`; the parallel line's archived code takes timm's default for the name), but Tier 3 feeds it through the model's own timm data configuration and the parallel line through a fixed 224-by-224 bicubic resize with ImageNet normalization, and the whitened geometry differs accordingly: the median plug-in margin is 402 in Tier 3 and 248 here, 4.9% of Tier 3's cells have a predicted excess above the floor against 19.3% here, and the measured excess at $m{=}2$ has median 0.046 there and 0.130 here. Both caches hold 100 test images per class, the official CIFAR-100 test set below the cap, so the registration's reading clause for a fired K-C2 (which names test slices of different sizes) does not arise. The finite-m prediction tracks the measured excess at a slope within four percent of the same value on two feature sets whose margins differ by a factor of 1.6; that is what a replication of the anchor reading on independently extracted features can mean here, and it is more than the registration asked. Table 40 and Figure 28(c) put the two programs side by side per count cell: from $(4, 4)$ to $(64, 64)$ and on the four asymmetric cells the two readings agree to within a few tenths, at $(2, 2)$ they differ (3.06 here against 1.81), and at the two largest counts Tier 3's intervals are the widest of its grid and its $(256, 256)$ reading is negative where this cache reads 1.00 (0.77, 1.24); since the origin fit weights each cell by its predicted excess, the pooled slope is decided by the small-count cells in both programs. The per-decile slopes climb with the margin in both programs from below one in the most-confusable decile (0.875 in Tier 3, 0.711 here) to floor-dominated values in the top deciles, the concentration Appendix Q.1 discloses.

Table 41: E-R2C by dataset: completed combos, the smallest per-class training pool after the pool rule, the range of the whitened-noise d_{eff} and of the raw participation ratio across the backbones, the range of the percentage of cells whose predicted excess clears the floor, the origin slope over combos interpretable under the registered rule (median, and range), and the count of those inside $[0.5, 2.0]$.

Dataset	Combos	Pool	d_{eff}	Raw PR	% above floor	Slope, median (range)	In band
CIFAR-10	8	5000	763 to 2046	2 to 108	9 to 88	0.989 (0.897 to 1.105)	8 of 8
STL-10	8	500	732 to 1812	2 to 111	0 to 35	0.847 (0.717 to 1.256)	3 of 3
EuroSAT	8	1600	747 to 2042	1 to 23	48 to 60	0.895 (0.849 to 0.932)	8 of 8
CIFAR-100	8	500	763 to 2046	2 to 257	1 to 62	0.882 (0.694 to 1.057)	8 of 8
Food-101	8	750	766 to 2047	2 to 92	4 to 90	0.873 (0.781 to 1.008)	8 of 8
Tiny-ImageNet	8	500	766 to 2047	2 to 416	0 to 60	0.862 (0.678 to 4.007)	7 of 8
Flowers-102	7	28	724 to 1917	2 to 96	0 to 13	5.960 (2.735 to 14.006)	0 of 3
AG News	3	140	262 to 386	114 to 153	69 to 93	0.776 (0.731 to 0.792)	3 of 3
20 Newsgroups	3	377	376 to 741	76 to 131	75 to 81	0.829 (0.785 to 0.832)	3 of 3
banking77	3	52	378 to 744	38 to 57	1 to 2	4.379 (4.224 to 7.034)	0 of 3

The matrix. Table 41 and Figure 28(a, b); the full per-combo table ships with the supplement as `generated/er2c_combos.tex`. Under the registered interpretability rule 55 of the 64 completed combos are interpretable, 48 of them sit inside $[0.5, 2.0]$, the median slope is 0.898, and 40 of the 55 are below one, the direction Tier 3 reports (the ledger over-charges rather than under-charges). By family: every interpretable ten-class vision combo is in band (19 of 19), 23 of 24 hundred-class combos are (the exception is Swin-B on Tiny-ImageNet, with 0.6% of its cells above the floor), and of the 9 text combos 6 are: all three encoders on AG News and on 20 Newsgroups, between 0.731 and 0.832, and none on banking77, where 0.9% to 2.5% of the cells clear the floor and the slopes read 4.224 to 7.034, the arithmetic accident Tier 4 documents for the intent datasets (Appendix R, Table 34). The registered rule admits combos in which only a few dozen cells out of thousands carry a resolvable prediction, so we also report a stricter descriptive fence, at least 10 percent of a combo's cells above the floor, chosen after the outcomes were known and carrying no verdict: 38 combos clear it, 37 of them are in band, their median is 0.879 and their range 0.678 to 2.735 (the one outside the band is MAE ViT-B on Flowers-102, the only small-class combo above the strict fence); the 17 combos interpretable only under the registered rule have 0.4% to 9.1% of their cells above the floor and 11 of them are in band; and the K-C3 ratio under the strict fence is 0.929 (14 hundred-class against 17 ten-class combos), the same reading.

The fence, diagnosed. Which combos fall under the strict fence is not arbitrary, and the diagnostic is the one Tier 4 uses: the median debiased margin against the noise dimension it has to be resolved in. The 26 combos under the fence have median debiased margins between 0.31 and 4.91 times d_{eff} (348 to 3653 in whitened units); the 38 above it have 0.04 to 0.41 (median 0.14), and the two

sets overlap only between 0.31 and 0.41. At margins that large the Gaussian finite-m prediction is below the floor in nearly every cell (the median predicted excess of a fenced combo is at most 4×10^{-14}) while the median measured excess at $m=2$ runs from below 0.001 to 0.080 across them: the measurement is the protocol's floor and the slope is that floor over a vanishing prediction, which is why STL-10 on the strong backbones, Flowers-102 on every backbone but MAE, banking77 on all three encoders and the DINOv2 and Swin-B combos on the hundred-class datasets read slopes from 0.771 up to 3×10^{7} without any of them being evidence about the constant, in either direction. One consequence for the previous appendix is worth stating. Its post hoc analysis found the small-class datasets fragile because the bias correction was between 0.26 and 1.65 of the plug-in margin there; under within-class whitening the same datasets (28 to 140 training items per class after the pool rule) have median debiased margins of 28 to 3653 and a correction of 1% to 12% of the margin. The fragility was the protocol's as well: total-covariance whitening compressed the very directions the margin lives in, and the correction then looked large against a margin that had been shrunk.

The effective dimension. On every completed vision combo the whitened-noise d_{eff} is at least 0.88 of the ambient width, and the raw participation ratio of the total spectrum, the quantity the parallel line's program used, is 2.1 to 533 times smaller (median 14; Figure 28(d)), so the artifact diagnosed in Appendix S is measured on the features themselves rather than inferred. The three text combos on AG News are the one place the whitened-noise dimension sits well below the width (d_{eff} 262 to 386): that cache holds 140 training documents per class in four classes after the pool rule, fewer points than dimensions, so the within-class spectrum is finite-sample limited and the Ledoit-Wolf shrinkage is the largest of the matrix there (0.25 to 0.30); the slopes on those combos (in band, 0.731 to 0.792) are read with that in mind. At most 45.8% of any combo's cells lie in the asymptotic validity region, as in Tier 3, and the asymptotic prediction is not used for any verdict.

Table 42: E-R2C stated omissions: combos that did not complete, with the runner's recorded reason. No number for them enters the paper.

Combo	Recorded outcome
dinov2_l__flowers102	failed: ZeroDivisionError: float division by zero

What this appendix claims. Three things. The hypothesis Appendix S recorded survived a registered test with a named falsifier: under within-class whitening the hundred-class and ten-class medians agree to within three percent where the parallel line's protocol had them a factor of 2.54 apart, and 23 of 24 hundred-class combos sit in the band. The Tier 3 anchor reading replicates, within the registered factor and in fact within four percent, on features extracted by a different line through a different input pipeline, with the same d_{eff} to three figures. And the artifact that sank E-R2's prediction column, an effective dimension of 70 for features whose whitened noise fills 766.8 of 768 dimensions, is reproduced on the parallel line's own features. Not claimed: that the slope is one (it is 0.879 at the strict median and below one on 40 of 55 interpretable combos, as in Tier 3); anything about the fenced combos beyond the floor reading above; anything about arms 2 to 5, which E-R2C does not run; and that this program replaces Tier 3. It is one arm, re-run under registration on another line's caches in 552 seconds of GPU time, and the paper's registered real-feature evidence remains Tiers 3 and 4, now with this replication beside them.

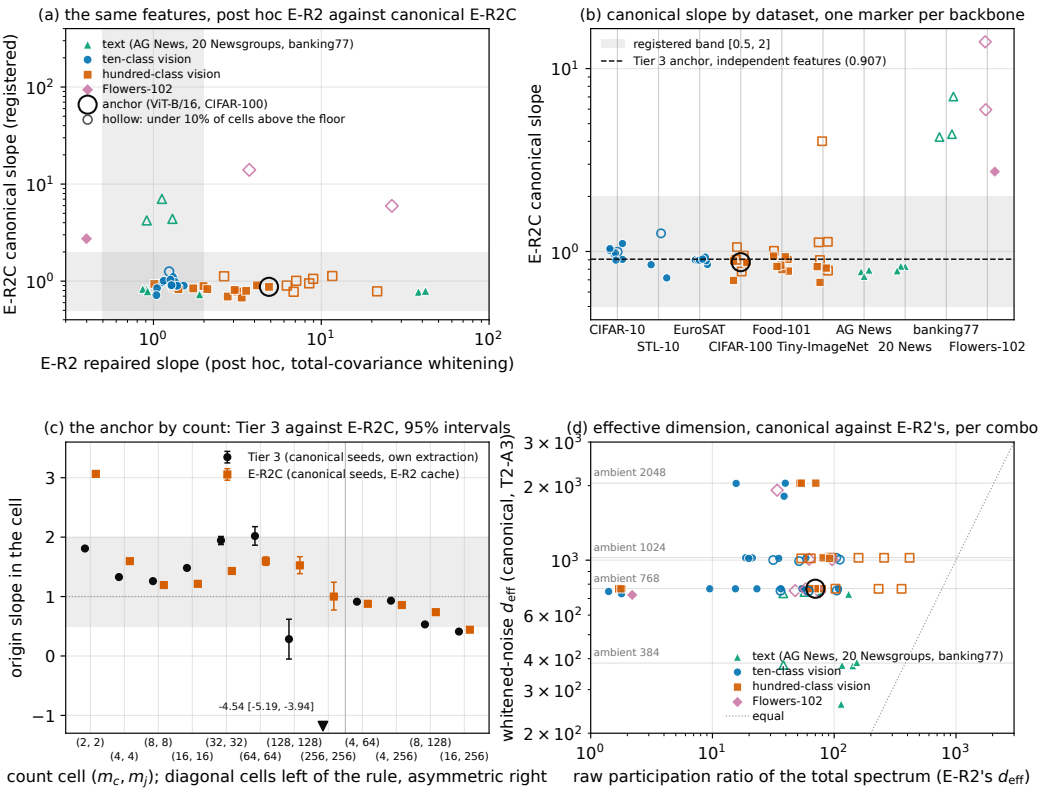

Figure 28: E-R2C. (a) The same features twice: the post hoc repaired slope of Appendix S (total-covariance whitening, native width) against the E-R2C canonical slope, one marker per combo, log-log, the band shaded on both axes; filled markers clear the descriptive strict fence, hollow ones are interpretable only under the registered rule, and combos with fewer than 20 cells above the floor are not drawn. (b) The canonical slope by dataset, one marker per backbone, the Tier 3 anchor reading dashed. (c) The anchor by count cell, Tier 3 against E-R2C, with 95 percent replicate-bootstrap intervals; the one off-scale value is printed. (d) The whitened-noise d_{eff} against the raw participation ratio of the total spectrum, per combo, the ambient widths as horizontal rules. The anchor is circled throughout.

U THREATS TO VALIDITY, STATED PLAINLY

Internal. The runners were validated on smoke seeds, three real bugs were caught and fixed before canonical execution (the MC threshold bug in E-L1-3's harness, the empty-stream division in E-L1-6, the two quadrature defects in the benchmark), and the canonical runs executed once, from the released code, on the recorded toolchain. The residual internal risk is shared harness error between runner and benchmark: both live in one library, so a common-mode formula error would pass the machinery test. Mitigations: the formulas are independently machine-verified at the symbol level (the sympy series and the released instruments), the $d = 1$ special case is exact in closed form (Appendix J.1), and the instrument bundle is released for independent re-execution.

Statistical. The falsifier amendments (A2–A5) exist because the drafted rules were statistically ill-posed at the registered grid sizes; each amendment is logged pre-execution with its smoke evidence. The reported descriptive exceedances (4 side-tests beyond three sigma in E-L1-3, against about 5.4 expected) are consistent with the amended analysis; a reader who prefers the raw rules can recompute them from the released CSVs and will observe exactly the fires the amendments predicted for perfect formulas.

External. Tier 1 and Tier 1b are synthetic by registration: Gaussian or near-Gaussian geometry, no learned representations, no LLM-scale systems. The registered frozen-feature anchor is a separate registration (Tier 2, Appendix P), so nothing in the synthetic tiers measures the theory's fit to deployed systems; the threat is the standard one for exact-in-a-model work, that the model's constants drift at scale. The scope paragraph of Section 2 states which claims survive verbatim outside the model (the online split's impossibility needs only the flow axioms) and which are instantiation-bound (every closed form); Tier 1b (Appendix O) measures the degradation of the closed forms under the three misspecifications a real feature space introduces.

Construction. All tables, figures, and quoted numbers in this manuscript are generated: canonical empirical content through the fill script and the results generators, exact content through the deterministic generators, code listings from the released sources. The generators are in the supplement; nothing in the results narrative was hand-computed.

V THE RELEASED SUPPLEMENT, FILE BY FILE

Reproduction protocol. With the released bundle: (1) verify the runner and library SHA-256 hashes against the manifests; (2) re-run each runner with its recorded seed list on the pinned toolchain (Python, NumPy, SciPy versions in the manifests); (3) re-run the fill scripts and the figure and table generators; (4) recompile. Steps (2) and (3) reproduce every number, table, and figure in this paper; step (1) is the trust anchor that the executed code is the released code, with two stated exceptions: E-R2's first runner and toolchain were not archived (Appendix S), and 91 of the 94 Tier 3 and 4 manifests record hashes of earlier versions of `t3_features.py`, `t3_run.py` or `t4_features.py` than the ones released (57 manifests for the first two scripts, 34 for the third). Every figure caption states whether its content is an exact closed form (deterministic) or canonical data (registered run); no figure mixes the two provenances in one panel.

Table 43: The released supplement's inventory. Hashes of the load-bearing sources appear in every manifest; the compilation sequence is runners, generators, fill scripts, `pdflatex`/`bibtex`.

Path	Contents
`experiments_design_L1_FROZEN_redacted.md`	the Tier 1 registration (marked redaction; SHA-256 of the unredacted file recorded), amendments A1–A8; A9–A10 in its addendum
`experiments_design_L1_T1B_redacted.md`	the Tier 1b registration, amendments T1B-A1 to T1B-A4
`colab/experiments_design_L1_T2ER1_redacted.md`	the Tier 2 registration with its execution gate and amendment log
`colab_t3/experiments_design_L1_T3_redacted.md`, `L1_T4_redacted.md`	the Tier 3 and Tier 4 registrations with their gates and amendment logs (marked redactions, SHA-256 of the originals in `REGISTRATIONS_sha256.txt`)
`data_er2/experiments_design_L1_T2ER2_redacted.md`, `colab_er2c/experiments_design_L1_ER2C_redacted.md`	the parallel line's E-R2 registration and the E-R2C registration (Appendices S and T)
`runners/l1_lib.py`	exact formulas, integral benchmarks (adaptive and batch), DP verifier, stream and policy harnesses, allocation solvers, frozen tolerances, manifest writer
`runners/l1_e1.py` … `l1_e8.py`	the eight registered Tier 1 runners
`runners/l1b_lib.py`, `l1b_e1.py` … `l1b_e3.py`	the Tier 1b library and its three runners
`colab/er1_lib.py`, `er1_run.py`, `er1_colab.py`	the Tier 2 package (feature extraction, the three arms, the synthetic validation gate, the Colab driver)
`colab_t3/t3_lib.py`, `t3_arms.py`, `t3_features.py`, `t3_policies.py`, `t3_run.py`, `t4_run.py`	the Tier 3 and Tier 4 package (whitening and geometry, the arms, the gates, the drivers)
`colab_er2c/er2c_run.py`	the E-R2C driver: the Tier 3 modules byte for byte, a cache reader in place of extraction
`runners/fill_l1_extended.py`, `fill_conference.py`, `fill_l1b.py`, `fill_t3.py`, `fill_t4.py`, `fill_er2.py`, `fill_er2c.py`	the only paths from canonical CSVs into this paper's macros
`runners/fig_gen.py`	deterministic closed-form figures
`runners/fig_results.py`, `fig_results_l1b.py`	canonical-data figures
`runners/gen_exact_tables.py`, `gen_results_tables.py`, `gen_manifests_tex.py`	table generators (exact, canonical, and manifest classes)
`runners/verify_reviews.py`, `verify_r3.py`, `verify_r4.py`	verification scripts (non-canonical; they reproduce every number quoted as such in the text)
`data/`	canonical CSVs and manifests, seeds 11, 13, 17 (Tier 1: `l1_e1` to `l1_e8`; Tier 1b: `l1b_e1` to `l1b_e3` with their fit and slope tables; Tier 2 when executed: `er1_*`)
`data_smoke/`	smoke outputs, seeds 0, 1, 2, manifests marked non-canonical
`data_t3/`	the Tier 3 and Tier 4 canonical CSVs, summaries and manifests, the gate records, the per-cell archive
`data_er2/`	the parallel line's E-R2 outputs as they landed (CSVs, manifests, failure notes, gate outputs), its own analysis record, the post hoc re-analysis summary, its archived runner code, a README of provenance facts
`data_er2c/`	the E-R2C outputs (64 canonical combos with manifests and summaries, the gate record, the one failure note, the smoke run), the fill summary, a README of provenance facts
`runners/audit/`	the independent recomputation battery (manifest hashes, citation cross-check, Tier 1 headline numbers refitted, theorem-by-theorem re-derivation, the E-R2 and E-R2C provenance checks, E7 brute force against the DP)
`figures/`	all PDF figures, deterministic and canonical, as cited
`iclr/`	this paper's sources, including `generated/` (all generated fragments)