How You Build Superintelligence: From Scaling Laws to Construction Laws

Anonymous authors
Paper under double-blind review

Abstract

A capability target does not specify the system that should be built to attain it. Scaling laws predict performance along a chosen development path, yet even exact knowledge of a scaling trajectory can leave the optimal resource allocation unidentified. We formulate *construction laws*: target-conditioned maps from identified mechanisms and physical resource constraints to a complete executable recipe, a forecast, and an auditable qualification test. We formalize superintelligence as an expert-relative capability specification giving a task population, within-task reliability, across-task breadth, and common information conditions. Our main inverse theorem converts uncertainty in forward predictions, resource costs, and task breadth into a guarantee on a qualifying construction; the controlling quantity is the *price of target tightening*, not prediction error alone. We then derive explicit construction laws for resource allocation, numerical precision and training data, reusable information, and conditional specialization, including an exact rule for jointly choosing sharing strength and specialist data. Prospectively frozen neural constructions reduce packed-weight storage by 1.33–$1.59\times$ while satisfying a predeclared 0.010-nat noninferiority criterion in all six target–seed comparisons. In a controlled hidden-support neural task, reusable structure reaches the same 80% reliability target with 32 rather than 256 task queries, while task-private structure removes the advantage. In external multi-hop retrieval, inherited transition structure consistently improves low-budget recovery, but a preregistered $2\times$ budget-crossing forecast misses validation by 0.66 percentage points, illustrating the instability of discrete construction decisions near a target boundary. Construction laws turn capability scaling from a forward forecast into an inverse question: *which system reaches the target, at what resource cost, and with what evidence?*

1 The target should determine the recipe

The same memory can hold one precise predictor, several specialists, or a reusable representation that makes future problems cheaper. The same computation can train a larger model, acquire more evidence, or improve a verifier. These choices produce different systems even when their nominal resource stocks agree. The construction question is therefore not only how performance changes with scale. More interesting is *which complete system should be built for the desired capability*.

Empirical scaling has already made important construction decisions. Model–data allocation, numerical precision, and conditional computation have been optimized using multivariate response models (Kaplan et al., 2020; Hoffmann et al., 2022; Kumar et al., 2024; Clark et al., 2022; Ludziejewski et al., 2025). We develop the additional structure needed when the destination itself is the input: a capability target, separate physical constraints, imperfect mechanism knowledge, and the cost of obtaining that knowledge. A forward model evaluates a recipe. A construction law returns one:

$$(\text{target},\ \text{resource limits},\ \text{identified mechanism}) \longmapsto (\text{recipe},\ \text{forecast},\ \text{qualification test}). \tag{1}$$

The distinction matters most when several development paths remain possible. Accurate prediction along one path may leave the useful conversion unidentified. Even accurate predictions over a family may lead to an expensive inverse when a small error changes whether a recipe qualifies.

We treat superintelligence as a construction destination: reliable performance beyond a declared expert reference over a declared breadth of tasks. Performance and breadth are distinct requirements,

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.

consistent with operational treatments of general intelligence (Morris et al., 2024). The recipe may share information across tasks and specialize its execution, but it must be one implementable system under the same information conditions. Taskwise access to an oracle-selected model would define a different problem.

Three results organize the paper. First, we characterize the evidence needed to distinguish constructions and give a stable target-to-cost certificate (section 3). Second, we derive explicit resource inverses for allocation, inheritance, and specialization (sections 4 and 5). In particular, sharing strength and specialist data can be chosen by a strictly convex scalar problem rather than an architecture sweep. Third, we test forecasts at different levels: new numerical-precision training recipes, hidden-support neural worlds, and conditional second-hop retrieval (section 6). Each test retains its own target and units. Their common focus is the prospective recipe, with gains measured against the same capability specification.

2 A CAPABILITY SPECIFICATION AND ITS RESOURCE INVERSE

A target fixes the task population ν, allowed information and tools, evaluation horizon, and taskwise reference. Let $x \in \mathcal{X}$ be a complete executable recipe, including training, retained artifacts, routing, and deployment. Its nonnegative task burden is $f_x(\tau)$. For expert-relative success, take $f_x(\tau) = 1 - p_\tau(x)$, where $p_\tau(x)$ is the probability of exceeding the prescribed reference. A loss or a resource requirement gives a narrower measurable criterion with its own units.

Define $F_x(t) = \nu\{\tau : f_x(\tau) \le t\}$. The system qualifies at tolerance ε and breadth β when

$$F_x(\varepsilon) \ge \beta. \tag{2}$$

For failure burden, $1 - \varepsilon$ is within-task reliability and β is across-task breadth. Equation (2) handles ties and atoms directly. It is equivalent to a lower-quantile requirement; the full capability-requirement formulation is given in Supplement A.

Let $c(x) \in \mathbb{R}^m_+$ price the complete remaining construction. Consumables such as training work sum across stages; retained memory is charged by peak simultaneous capacity. Numerical formats include scales and indexing overhead. With positive reference stock B, write

$$k_B(x) = \max_j \frac{c_j(x)}{B_j}, \qquad \mathsf{C}^\star_B(\varepsilon, \beta) = \inf_{x : F_x(\varepsilon) \ge \beta} k_B(x). \tag{3}$$

The inverse asks how much the reference stock must be dilated to admit a qualifying recipe. An explicit witness with $k_B(x) \le 1$ establishes affordability. Prices can instead scalarize genuinely interchangeable costs; a vector of architectural knobs is not itself physical resource accounting.

Why a path is insufficient. Consider two purchasable ingredients and the monotone, strictly convex burden $f_\theta(x_1, x_2) = (1-\theta)/x_1 + \theta/x_2$, $0 < \theta < 1$, at cost $x_1 + x_2$. Every environment has the same complete scaling path $f_\theta(t, t) = 1/t$. Yet its least-cost allocation satisfies

$$\frac{x_1^\star}{x_2^\star} = \sqrt{\frac{1-\theta}{\theta}}. \tag{4}$$

Knowing the path exactly leaves this ratio undetermined. The missing evidence concerns the direction in which resources should be converted, not the accuracy of the observed scaling exponent. Supplement B proves this example and the corresponding linear criterion: recipe contrasts are identifiable precisely when their feature differences lie in the measurement span. This connects construction to decision-directed experimental design (Soare et al., 2014; Fiez et al., 2019).

3 THE PRICE OF AN UNCERTAIN CONSTRUCTION

A useful inverse must return a qualifying recipe despite error in both performance and cost. Fix a finite recipe family and n evaluation tasks. Let $F_{x,n}$ be the empirical CDF of the true burdens. Suppose one simultaneous event gives

$$|\widehat{f}_x(\tau_i) - f_x(\tau_i)| \le \eta, \quad e^{-\xi}\widehat{c}_j(x) \le c_j(x) \le e^{\xi}\widehat{c}_j(x), \quad \sup_{x,t} |F_{x,n}(t) - F_x(t)| \le u. \tag{5}$$

These bounds express measurement and mechanism assumptions. Their justification must cover the recipes actually considered; finite held-out residuals provide diagnostics of that coverage. For fixed recipes and independent task draws, a finite-family CDF bound follows from the Dvoretzky–Kiefer–Wolfowitz inequality (Massart, 1990). Adaptive experiments require conditional or time-uniform coverage (Howard et al., 2021).

Use the pessimistic burden $\widehat{f}_x + \eta$, optimistic burden $\max(0, \widehat{f}_x - \eta)$, and upper/lower costs $e^{\pm\xi}\widehat{c}(x)$. The conservative set $\mathcal{X}^+$ requires empirical qualification at $(\varepsilon, \beta + u)$ using pessimistic burdens. The possible set $\mathcal{X}^-$ requires $(\varepsilon, \beta - u)$ using optimistic burdens. Let C^+ and C^- be their minimum upper and lower resource dilations.

Theorem 3.1 (Stable target-to-recipe inversion)**.** *Assume* (5)*,* $0 < \beta - u$*, and* $\beta + 2u \le 1$*. Then*

$$C^- \le \mathsf{C}^\star_B(\varepsilon, \beta) \le C^+ \le e^{2\xi}\mathsf{C}^\star_B(\varepsilon - 2\eta, \beta + 2u). \tag{6}$$

A conservative recipe $\widehat{x}$ *with upper cost at most* $e^\omega C^+$ *satisfies*

$$F_{\widehat{x}}(\varepsilon) \ge \beta, \qquad k_B(\widehat{x}) \le e^{\omega+2\xi}\mathsf{C}^\star_B(\varepsilon - 2\eta, \beta + 2u). \tag{7}$$

An empty tightened target has inverse $+\infty$*.*

Proof idea. A conservative empirical qualifier is a true population qualifier. Every true qualifier is a possible empirical qualifier. Finally, any recipe meeting the population target with 2η tighter burden and $2u$ greater breadth belongs to the conservative set; its estimated upper cost is at most $e^{2\xi}$ times its actual cost. Taking infima gives the bounds. Supplement C supplies each inclusion and the boundary cases. □

The controlling quantity is the *price of target tightening*,

$$\log \frac{\mathsf{C}^\star_B(\varepsilon - 2\eta, \beta + 2u)}{\mathsf{C}^\star_B(\varepsilon, \beta)}. \tag{8}$$

When this price is bounded, forward uncertainty becomes a bounded construction overhead. When it is unbounded, arbitrarily accurate forecasts can still force an arbitrarily costly safe recipe. A two-world example establishes this necessity in Supplement C. Thus correct ranking and stable qualification are different scientific achievements.

Reading the certificate as a decision. The bounds distinguish three outcomes. An upper cost $C^+ \le 1$ supplies an affordable qualifying recipe; a lower cost $C^- > 1$ excludes such a recipe within the declared family. Between them, the target remains unresolved. When $0 < C^- \le C^+ < \infty$, C^+/C^- bounds the cost of the exactly optimized conservative choice relative to the unknown optimum. This gap directs measurement toward recipes whose cost or qualification can change the decision. Prediction accuracy on already-qualified or clearly inadequate recipes may improve a global fit while leaving this gap unchanged.

For a concrete reading of (8), consider the single-burden, common-exponent allocation model in section 4. Its exact inverse is $C^\star(\varepsilon) = K\varepsilon^{-1/\alpha}$ for a positive mechanism constant K (Supplement E). With exact costs and known task population, tightening the target by 2η incurs the factor

$$\frac{C^\star(\varepsilon - 2\eta)}{C^\star(\varepsilon)} = \left(\frac{\varepsilon}{\varepsilon - 2\eta}\right)^{1/\alpha}, \qquad \varepsilon > 2\eta. \tag{9}$$

For an assigned $\alpha = 1$ mechanism, a burden-error bound $\eta = 0.05$ gives a factor $10/9 \approx 1.11$ at target $\varepsilon = 1$, but a factor 6 at target $\varepsilon = 0.12$. The same forward accuracy therefore supports very different construction decisions. A stringent target demands evidence precise enough relative to its remaining allowance.

Evidence is part of the construction. If a pilot consumes e, optimize against the remaining consumable stock and include the pilot in total cost. A finite-family learner with cost error ζ and direct contrast error ϵ_d achieves quality at least $\Phi(B - e - 2\zeta) - \epsilon_d$ when it selects safely from $e + \widehat{c}(x) + \zeta \preceq B$. The precise statement and proof appear in Supplement D. This separates the price of learning from the loss of choosing incorrectly, consistent with resource-constrained identification

Step	Target-to-recipe operation
Identify	Fix (ε, β), the information conditions, and the resource accounting. Measure interventions that distinguish feasible recipe contrasts; identify the response and its domain of validity.
Bound	On the simultaneous event in (5), form the conservative and possible qualifying sets $\mathcal{X}^+, \mathcal{X}^-$. Compute their upper and lower minimum costs C^+, C^-.
Select	Minimize upper cost on $\mathcal{X}^+$. If $C^+ \leq 1$, retain the qualifying witness. If $C^- > 1$, report the target as unaffordable in the family. Otherwise obtain evidence that resolves the boundary alternatives.
Freeze	Record the executable recipe, numerical forecast, complete resource accounting, and qualification rule before observing the completed construction's evaluation outcomes.
Qualify	Build the frozen system and evaluate its declared target on reserved tasks. Report target attainment, resource use, and forecast error in their respective units.

Table 1: **From identified mechanisms to a complete construction.** The confidence-qualified branch uses theorem 3.1; empirical forecasts follow the same identify–freeze–build–evaluate order. Pilot expenditure is charged before the remaining resource decision.

and costly-observation learning (Li & Cheung, 2024; Tucker et al., 2023). Reusable artifacts produced by a pilot remain available to subsequent constructions.

Table 1 makes the construction law operational. Identification determines which alternatives can be compared; the inverse selects one under a common specification; the qualification test checks the completed system. The execution order preserves the distinction between choosing a recipe from a fitted model and establishing its performance after construction.

4 Constructors for allocation and inheritance

Allocation: the target sets the ingredient schedule. An identified separable burden $E(x) = \sum_i a_i x_i^{-\alpha_i}$ with positive prices w_i has a complete inverse. The least-cost recipe meeting $E(x) \leq \varepsilon$ is

$$x_i^\star = \left(\frac{\lambda \alpha_i a_i}{w_i}\right)^{1/(\alpha_i+1)}, \qquad \sum_i a_i (x_i^\star)^{-\alpha_i} = \varepsilon. \tag{10}$$

One strictly decreasing scalar equation determines λ. With common exponent α, write $h_i = (a_i w_i^\alpha)^{1/(\alpha+1)}$ and $\pi_i^\star = h_i / \sum_j h_j$. The exact cost penalty of remaining on a fixed spending path π is

$$\frac{C_\pi(\varepsilon)}{C^\star(\varepsilon)} = \left[\sum_i (\pi_i^\star)^{\alpha+1} \pi_i^{-\alpha}\right]^{1/\alpha}. \tag{11}$$

The expression equals an exponentiated Rényi divergence. It quantifies the opportunity in a misallocated path, while showing why a strong already-balanced comparator may leave little allocation gain. Unequal exponents shift spending toward the slowest-improving ingredient as the target tightens (Supplement E).

For numerical model design, the identified response can instead be

$$E(N, D, b) = E_\infty + A(N/N_0)^{-\alpha} + C(D/D_0)^{-\gamma_D} + Qe^{-\gamma_b b}. \tag{12}$$

Here D counts processed training tokens. For a proposed architecture and precision, define the remaining loss allowance

$$R(N, b; \varepsilon) = \varepsilon - E_\infty - A(N/N_0)^{-\alpha} - Qe^{-\gamma_b b}, \qquad D_{\text{req}} = D_0 \left(\frac{C}{R(N, b; \varepsilon)}\right)^{1/\gamma_D}. \tag{13}$$

The expression applies when $R > 0$; otherwise no finite training allocation meets the target in this model. Round D_{req} upward to an executable batch count, reject choices exceeding the work cap, and select the least packed storage among the remaining architecture–precision pairs. Lower precision is worthwhile when its storage saving can pay for the extra training needed to recover the target. This is the constructor tested in section 6.1; its exponents and residuals belong to that measured family.

Inheritance: a reusable asset changes the cost of future targets. Let a construction cost P once and c_F per future task, against c_C for a qualifying task-local construction. The reusable program costs $P + Jc_F$ over J tasks and wins exactly when

$$J > \frac{P}{c_C - c_F}, \qquad c_C > c_F. \tag{14}$$

Reusable precomputation also motivates sleep-time computation (Lin et al., 2025). The substantive premise is transferable task structure: the reusable artifact must still qualify on future tasks. Transfer validation supplies this premise; the amortization identity prices it. In section 6.2, historical contexts identify where task-relevant observations tend to lie; a task-private perturbation tests the premise directly.

Information provenance also limits what reuse can buy. If a new computation Z is generated from the available transcript S and private randomness independent of an unresolved target W conditional on S, then $I(W; Z \mid S) = 0$. Computation can improve a decision from existing evidence; a new external observation can change the evidence itself. Supplement F gives the adaptive transcript form, a target-qualified information bound, and paid composition across successive constructions. This accounting distinguishes repeated computation from additional independent evidence (Cover & Thomas, 2006).

5 A TARGET-CONDITIONED LAW OF SPECIALIZATION

Allocation can choose not only how much of an ingredient to buy, but how much should be shared. We solve an explicit conditional-prediction model. Task class $I = i$ has probability $p_i > 0$ and desired vector θ_i. A router produces $R = r$ with probability w_r, and $q_{ir} = \mathbb{P}(I = i \mid R = r)$. Let

$$\mu = \sum_i p_i \theta_i, \quad \mu_r = \sum_i q_{ir} \theta_i, \quad H_0 = \sum_i p_i \|\theta_i - \mu\|^2, \quad H = \sum_r w_r \|\mu_r - \mu\|^2.$$

Thus H is the task variation explained by routing. Unbiased training estimates $\widehat{\theta}_i = \theta_i + e_i$, independent of the new query, have error Gram matrix $\Gamma_{ij} = \mathbb{E}\langle e_i, e_j \rangle$. Expert errors may be correlated. A shared–specialized recipe predicts

$$\widehat{m}_t(r) = \widehat{\mu} + t(\widehat{\mu}_r - \widehat{\mu}), \qquad 0 \le t \le 1. \tag{15}$$

Define $S = \sum_r w_r (q_r - p)(q_r - p)^\top$, $V_0 = p^\top \Gamma p$, and $V = \operatorname{tr}(S\Gamma)$.

Theorem 5.1 (Sharing strength and its intervention signature)**.** *The squared-error risk of* (15) *is*

$$\mathcal{R}(t) = H_0 + V_0 - 2tH + t^2(H + V). \tag{16}$$

For $H + V > 0$, its minimizer is $t^* = H/(H + V)$ and its gain over pooling is $H^2/(H + V)$. If task–route correspondence is replaced by independent routing with the same w_r, while keeping the trained templates fixed, the risk increases by exactly $2tH$.

The proof uses conditional expectation and the covariance of the compiled templates (Supplement G). Full specialization helps only when $H > V$; an intermediate construction can outperform both full sharing and full specialization. The intervention changes correspondence, not route loads or the stored templates. For general nonlinear experts, (15) is a functional mixture; compiling it into one active expert requires the stated linear/template structure.

Jointly purchasing the evidence. Suppose $\Gamma = \operatorname{diag}(v_i/n_i)$, with $v_i > 0$, sample prices $c_i > 0$, and $\sum_i c_i n_i \le T$. Set $s_i = S_{ii}$, $a_i(t) = v_i(p_i^2 + t^2 s_i)$, and $Z(t) = \sum_i \sqrt{c_i a_i(t)}$. Square-root allocation gives

$$n_i^*(t, T) = \frac{T\sqrt{a_i(t)/c_i}}{Z(t)}, \qquad R^*(T) = \min_{t \in [0,1]} \left[B(t) + \frac{Z(t)^2}{T} \right], \quad B(t) = H_0 - 2tH + t^2 H. \tag{17}$$

The inner allocation is the classical structure of stratified sampling (Neyman, 1934); its coupling to learned conditional structure makes the recipe target dependent.

When $H > 0$, the scalar objective is strictly convex, has a unique minimizer, and selects increasing specialization as T increases.

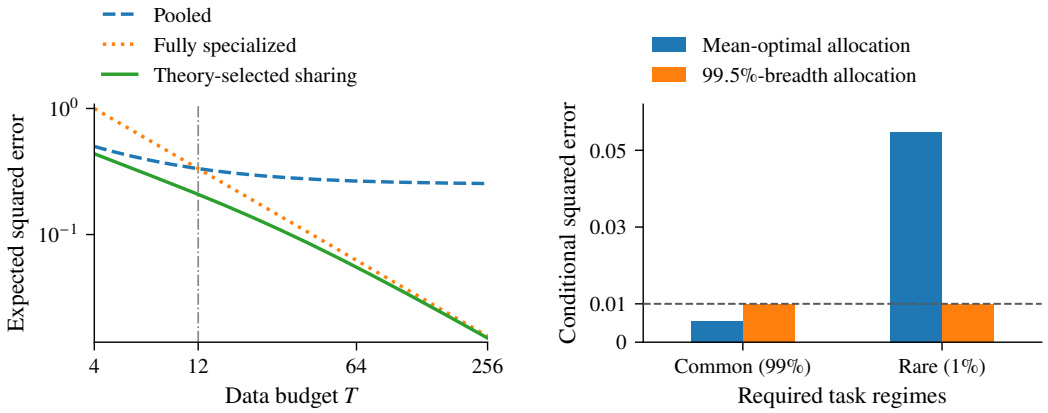

Figure 1: **Different targets produce different constructions.** Left: exact risks for four balanced regimes with $H = H_0 = 1/4$ and unit noise. Full specialization crosses pooling at $T = 12$; the constructed sharing rule improves on both. Right: with task masses $(0.99, 0.01)$ and $T = 200$, mean-optimal allocation leaves the rare regime above the 0.01 error target. A 99.5%-breadth requirement allocates 100 observations to each regime. Both panels are specified mathematical models, evaluated by the displayed laws.

Its target inverse is

$$\boxed{T^*(\varepsilon) = \inf_{t:B(t)<\varepsilon} \frac{Z(t)^2}{\varepsilon - B(t)}.} \tag{18}$$

These statements, integer allocation, and zero-signal cases are proved in Supplement H. They expose an irreducible routing floor $H_0 - H$ and the data cost of approaching it.

Precision and breadth change the inverse. For an identified task burden $v_i/n_i + \kappa_i e^{-kb_i}$ and target τ_i, the required data are $n_i = v_i/(\tau_i - \kappa_i e^{-kb_i})$. The minimum-data construction under $\sum_i d_i b_i \le M$ is strictly convex in continuous precisions on its feasible domain; a scalar dual root determines the bit allocation (Supplement I). The target must exceed each precision floor. Breadth can reverse mean-optimal spending by requiring a rare task to cross its own threshold, as figure 1 illustrates. Sharing, numerical format, and task coverage therefore enter the same inverse, rather than three unrelated tuning rules.

6 Prospective constructions and their boundaries

The experiments separate three evidentiary questions: does an untrained recipe meet a measured target, does inherited structure reduce task-local acquisition, and does a selected budget crossing replicate? Table 2 makes their qualification rules and charged resources explicit. Procedures, complete resource accounting, seeds, and source hashes are in Supplement J; all reported results are recomputed from the bundled records.

6.1 Numerical precision: build the recipe before observing its outcome

Byte-level Transformer models are trained on TinyStories (Eldan & Li, 2023). Eleven calibration configurations vary model size and processed tokens independently and identify (12); four separate mechanism checks have mean absolute error 0.02075 nats/byte. The fitted law is then fixed. For each of three loss targets and training-work caps, its inverse chooses the minimum declared packed storage in the architecture/precision family. The six resulting configurations are trained at two seeds, without modifying their recipes after evaluation. Table 3 gives the target and complete training allocation returned by each inverse.

Experiment	Qualification	Resource quantity	Evaluation scope
Numerical inverse	Paired loss degradation ≤ 0.010 nats/byte	Packed bits; processed parameter–tokens	Three pairs, two seeds, fixed evaluation stream
Blind Foundry	Private-sign and response-loss completion	Task-local queries; declared tariff	1,024 controlled worlds plus 128 private-shift worlds
Transition inheritance	Second support's shard reached	Conditional second-hop shard probes	379 calibration, 362 validation; final test closed

Table 2: Each experiment tests a target-conditioned recipe. Loss, completion probability, and shard coverage retain distinct meanings and resource measures.

Target	Loss target	Work cap	Model	Tokens, baseline/constructed	Bits
E1	0.882532	3.5×10^{14}	B	69,992,448/71,041,024	8/5
E2	0.818168	5.0×10^{14}	C	69,992,448/70,189,056	8/6
E3	0.786164	6.5×10^{14}	D	69,992,448/71,041,024	8/5

Table 3: Frozen inverse recipes. Work is dense parameters times processed tokens. Packed-size ratios are 1.591803, 1.329958, and 1.593442. Recipes meet the nominal fitted-loss target; the evaluation rule permits a 0.010-nat degradation relative to the trained comparator.

The comparator uses eight-bit weights; the constructor selects five, six, and five bits across the three targets. Within each pair the architecture is identical, and the constructor buys 0.28% or 1.50% more processed training tokens within the common cap. Packed storage includes row scales and retained high-precision parameters. Figure 2 reports the six paired loss contrasts. All meet the specified marginal noninferiority rule, upper 95% interval ≤ 0.010 nats/byte, at storage ratios 1.592, 1.330, and 1.593.

The measured conclusion is viable target-conditioned compression at this tolerance. The inverse enforces a nominal fitted-loss target, while qualification allows 0.010 nats. Empirical minimum-storage optimality would additionally require qualifying or excluding neighboring constructions, including the untested five-bit recipe at the middle target. All quantized values execute in floating-point kernels. Training work is recorded in parameter–tokens and storage in packed-weight bits, with hardware latency and energy outside these measurements. Supplement J.1 gives the exact targets, conditional intervals, corpus reuse, and qualification scopes.

6.2 Inheritance: pay once for information that future tasks can reuse

Blind Foundry is a controlled neural-world experiment. Each task has an eight-dimensional public context, public embeddings for $32{,}768$ modes, and eight hidden relevant modes. Their private signs modulate an unknown shared neural field. Historical tasks provide support identities from which a context-to-relevance map is learned. A width-64, depth-two network learns the shared field from grounded measurements. Querying a mode reveals its sign and a noisy relevance score. The task-local comparator chooses among predeclared $8/16/32/64$-probe ridge estimators and receives the same public metadata. Its probes are included in its query budget.

The forecast selects 32 queries with reuse and 256 for the task-local comparator. On $1{,}024$ new worlds, the systems complete 958 and 941 tasks, respectively; both exceed the common 80% reliability target using the experiment's simultaneous binomial lower bounds. The paired accuracy comparison has 76 wins and 59 losses for reuse (exact two-sided $p = 0.168$). Thus the result is an eightfold query reduction at the common target, rather than a demonstrated increase in accuracy.

The acquisition-and-primitive-training tariffs are $2{,}562.125$ and $6{,}146.125$ per task, a $2.40\times$ ratio. Charging historical support under two declared tariffs predicts break-even after 37 or 586 future tasks, as in (14). The private-structure intervention yields $3/128$ successes for reuse and $118/128$ for the cold controller. This establishes the conditional nature of inheritance: context-predictable relevance can be retained; task-private relevance must be reacquired. The selected reusable policy is a ridge representation, and same-reuse-class controls are reported in Supplement J.2.

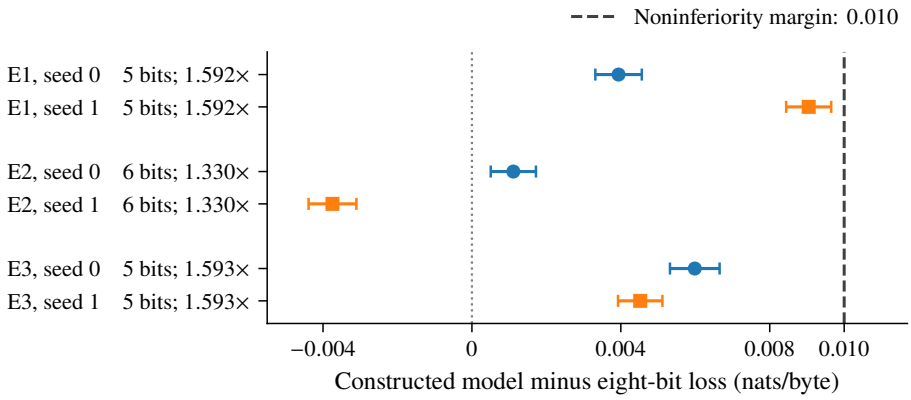

Figure 2: **A frozen inverse produces viable lower-storage recipes.** Points show test loss of the constructed model minus its eight-bit comparator; bars are the protocol's paired 95% evaluation-batch intervals. The vertical boundary is the declared 0.010-nat noninferiority margin. Forecast differences are approximately zero. The 1.33–1.59× ratios price packed representation; the empirical test is whether the preselected recipes preserve loss to the declared tolerance. Two training seeds are reported separately.

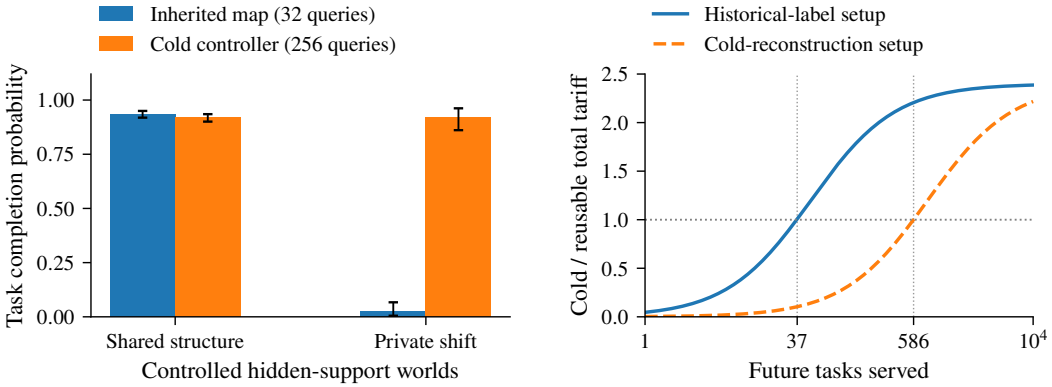

Figure 3: **Inheritance has both a reuse regime and a failure regime.** Left: completion in controlled hidden-support neural worlds. The inherited map uses 32 task queries and the cold controller 256. A private relevance component removes the inherited map's usefulness while preserving the cold controller's access to current-task observations. Error bars are pointwise exact binomial 95% intervals. Right: the declared acquisition-and-primitive-training tariff, including either historical-label or cold-reconstruction setup, crosses break-even at 37 or 586 future tasks. It prices the experimental mechanism, not hardware speed.

6.3 External retrieval: a useful shift can leave the target crossing unresolved

The external test uses HotpotQA bridge questions (Yang et al., 2018) and Qwen3-Embedding-0.6B (Zhang et al., 2025). It evaluates a *conditional second hop*: one gold support document is supplied as an anchor, and the objective is to probe the shard containing the other support. History estimates support-shard transitions; the cold control uses question and anchor semantics. The evaluation retains tasks whose support documents occupy distinct shards, yielding 379 calibration and 362 validation cases.

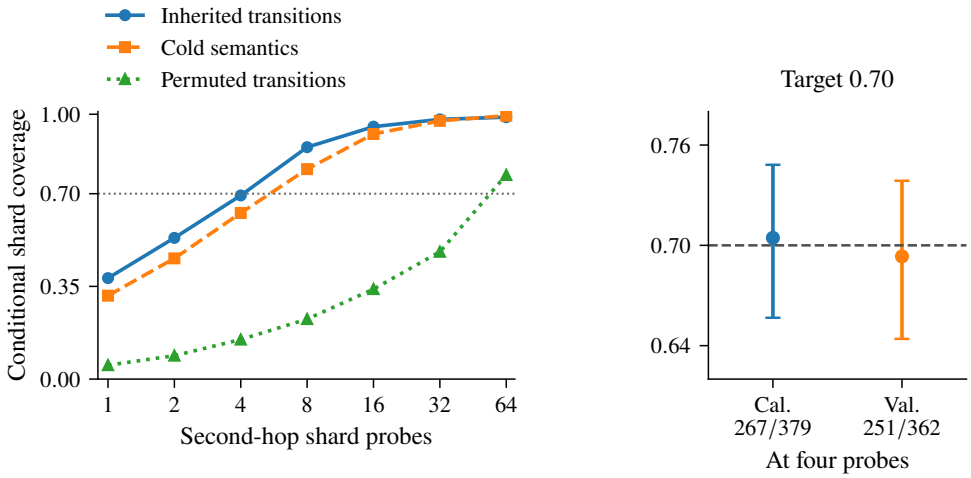

Figure 4: **A small forward error can change a discrete inverse.** The externally selected inheritance policy improves low-budget conditional shard coverage on validation, but its four-probe value falls below the 70% target. The right panel isolates the calibration/validation crossing with pointwise Wilson intervals. Historical-transition permutation reduces four-probe coverage to 14.92%; this damaging-history control is reported alongside the stronger no-history comparator. The declared exact-crossing validation remains a failure.

Calibration selects history weight 0.75, target 0.70, and crossing budgets four probes with inheritance versus eight without it. At four probes, coverage moves from $267/379 = 70.45\%$ in calibration to $251/362 = 69.34\%$ in validation. The cold eight-probe crossing reproduces; the inherited crossing moves to eight. The exact $2\times$ prediction therefore fails its declared validation gate, and final QA evaluation is not opened.

At matched budgets one, two, four, and eight, the selected policy's validation advantages are 6.63, 7.73, 6.63, and 8.29 percentage points. Permuting historical transition correspondences lowers four-probe coverage to $54/362 = 14.92\%$. These are validation effects on conditional shard retrieval, not an end-to-end QA or paired-significance claim. The result illustrates (8): a forecast 0.45 percentage points above a target is a fragile basis for an exact discrete cost ratio. Supplement J.3 supplies the complete curves and gate.

7 Conclusion

Scaling laws answer a forward question: how does a chosen system change as resources grow? Construction laws solve the complementary inverse: given a capability specification and available resources, what system should be built? The results here make that inverse operational. Decision-relevant mechanisms must be identifiable, forecast uncertainty must be translated into the price of target tightening, and the returned recipe must be frozen and tested under the same specification.

Across precision, inheritance, and specialization, the common theme is a resource conversion that must survive prospective qualification: representation can be exchanged for training, historical information for lower future acquisition cost, and task heterogeneity for structured sharing. The external retrieval result shows why a small forward error can overturn a discrete budget decision near a target boundary. For increasingly broad intelligence, this yields an iterative program: declare the target, identify the limiting mechanism, compile the next conversion, execute within the justified region, remeasure, and repeat. Scaling remains an essential instrument inside this loop; the construction law determines how it is used. The central result is therefore an auditable target-to-recipe map.

REPRODUCIBILITY STATEMENT

The appendix gives complete proofs, model assumptions, resource accounting, and experimental procedures. The accompanying source includes the recorded outcome tables, forecast hashes, figure-generation scripts, and deterministic reanalysis and theorem tests. Reported neural and retrieval results are reanalyses of the included experiment records; the new mathematical checks evaluate explicitly specified mechanisms. Supplement J distinguishes evaluation-batch uncertainty, training-seed replication, controlled-world inference, and the external validation stop.

ETHICS STATEMENT

The work studies resource allocation under explicit task and information specifications. Broad capability and safe deployment are separate requirements: an application can restrict the admissible recipe family by safety, privacy, and oversight constraints before optimization. The experiments use synthetic worlds and public language/question-answering corpora. The task-breadth analysis emphasizes that optimizing frequent tasks can leave required rare tasks unserved. Resource savings are reported in their measured or declared units, with no conversion into a general-intelligence score.

AI USE STATEMENT

AI tools were used for candidate-theorem exploration, literature retrieval, coding, drafting, figure preparation, and editing. Wolfram Language supported symbolic checks; Python supported numerical verification and reanalysis. The claims in the manuscript are tied to the stated assumptions, written proofs, and archived outcomes. The supplementary verification records report the symbolic and numerical checks supporting the written proofs.

REFERENCES

Aidan Clark, Diego de Las Casas, Aurelia Guy, Arthur Mensch, Michela Paganini, Jordan Hoffmann, Bogdan Damoc, Blake Hechtman, Trevor Cai, Sebastian Borgeaud, et al. Unified scaling laws for routed language models. In *International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 4057–4086, 2022. URL https://proceedings.mlr.press/v162/clark22a.html.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2 edition, 2006. doi: 10.1002/047174882X.

Ronen Eldan and Yuanzhi Li. TinyStories: How small can language models be and still speak coherent english? *arXiv preprint arXiv:2305.07759*, 2023. URL https://arxiv.org/abs/2305.07759.

Tanner Fiez, Lalit Jain, Kevin G. Jamieson, and Lillian Ratliff. Sequential experimental design for transductive linear bandits. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://arxiv.org/abs/1906.08399.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://arxiv.org/abs/2203.15556.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021. doi: 10.1214/20-AOS1991. URL https://arxiv.org/abs/1810.08240.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. URL https://arxiv.org/abs/2001.08361.

Tanishq Kumar, Zachary Ankner, Benjamin F. Spector, Blake Bordelon, Niklas Muennighoff, Mansheej Paul, Cengiz Pehlevan, Christopher Ré, and Aditi Raghunathan. Scaling laws for precision. *arXiv preprint arXiv:2411.04330*, 2024. URL https://arxiv.org/abs/2411.04330.

Zitian Li and Wang Chi Cheung. Best arm identification with resource constraints. In *International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 253–261, 2024. URL https://proceedings.mlr.press/v238/li24c.html.

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. URL https://arxiv.org/abs/2504.13171.

Jan Ludziejewski, Maciej Pióro, Jakub Krajewski, Maciej Stefaniak, Michał Krutul, Jan Małaśnicki, Marek Cygan, Piotr Sankowski, Kamil Adamczewski, Piotr Miłoś, and Sebastian Jaszczur. Joint MoE scaling laws: Mixture of experts can be memory efficient. *arXiv preprint arXiv:2502.05172*, 2025. URL https://arxiv.org/abs/2502.05172.

Pascal Massart. The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283, 1990. doi: 10.1214/aop/1176990746.

Meredith Ringel Morris, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg. Position: Levels of AGI for operationalizing progress on the path to AGI. In *International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 36308–36321, 2024. URL https://proceedings.mlr.press/v235/morris24b.html.

Jerzy Neyman. On the two different aspects of the representative method: The method of stratified sampling and the method of purposive selection. *Journal of the Royal Statistical Society*, 97(4): 558–606, 1934. doi: 10.1111/j.2397-2335.1934.tb04184.x.

Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. In *Advances in Neural Information Processing Systems*, volume 27, 2014. URL https://arxiv.org/abs/1409.6110.

Aaron D. Tucker, Caleb Biddulph, Claire Wang, and Thorsten Joachims. Bandits with costly reward observations. In *Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pp. 2147–2156, 2023. URL https://proceedings.mlr.press/v216/tucker23a.html.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Conference on Empirical Methods in Natural Language Processing*, pp. 2369–2380, 2018. URL https://aclanthology.org/D18-1259/.

Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*, 2025. URL https://arxiv.org/abs/2506.05176.

SUPPLEMENTARY MATERIAL: HOW YOU BUILD SUPERINTELLIGENCE: FROM SCALING LAWS TO CONSTRUCTION LAWS

Anonymous authors
Paper under double-blind review

SUPPLEMENTARY MATERIAL

This supplement is the technical companion to *How You Build Superintelligence: From Scaling Laws to Construction Laws*. It is organized to let a reader verify the paper from definitions through proofs to the recorded experiments. The notation and theorem numbering continue those of the main paper; every empirical quantity reported here is tied to the included recorded outcomes and reproduction code.

Reading guide. Sections A and B formalize the capability specification and show what evidence is needed to distinguish recipes rather than only forecast a path. Sections C and D give the full stable-inversion proof and price decision-relevant measurements. Sections E and F derive exact resource-allocation, information-provenance, and amortization laws. Sections G to I develop conditional specialization, sharing/data allocation, precision, routing, and breadth. Section J specifies the three experimental programs, frozen targets, recorded outcomes, and reproduction commands. Section K positions the construction-law results relative to complementary scaling, experimental-design, MoE, quantization, and information-theoretic work.

Reproducibility convention. The accompanying ZIP contains this PDF, anonymized source, vector figures, recorded numerical outcomes, deterministic reanalysis scripts, mathematical checks, and output-cleared resumable notebooks for the GPU experiments. The recorded-result analyses are fully local. Re-running the neural experiments additionally requires the public datasets/models and accelerator resources specified in the included notebooks.

A CAPABILITY SPECIFICATIONS AND COMPLETE RESOURCE ACCOUNTING

A.1 TASKWISE RELIABILITY AND THE CAPABILITY-REQUIREMENT DISTRIBUTION

Fix a task population ν, a taskwise reference, and an implementable policy π that receives only the information permitted by the specification. Let $p_\tau(\pi, h)$ be its task success probability at common resource allowance $h \geq 0$. Assume this function is nondecreasing in h. At within-task reliability r, define

$$a_\tau(\pi; r) = \inf\{h \geq 0 : p_\tau(\pi, h) \geq r\}. \tag{19}$$

An empty set gives $+\infty$. Suppose its finite thresholds are attained, so the qualifying set is $[a_\tau, \infty)$. Then

$$\nu\{\tau : p_\tau(\pi, h) \geq r\} \geq \beta \quad \Longleftrightarrow \quad \nu\{\tau : a_\tau(\pi; r) \leq h\} \geq \beta. \tag{20}$$

The least common allowance is the lower β-quantile of a_τ. More generally, for a real-valued burden with CDF F_x, $Q_\beta(f_x) = \inf\{t : F_x(t) \geq \beta\}$ and right continuity imply $Q_\beta(f_x) \leq \varepsilon$ if and only if $F_x(\varepsilon) \geq \beta$. This retains discrete masses and ties. If a taskwise resource infimum is not attained, equality at that threshold requires separate checking: costs $1 + 1/n$ have infimum one but no feasible witness at budget one. Strict slack or a finite catalog resolves the issue.

The shared-policy quantifier is essential. Let two equiprobable tasks have two admissible recipes with burdens $(0, 1)$ and $(1, 0)$. Each task separately admits zero burden; neither shared recipe attains zero burden on breadth greater than $1/2$. A selector can close the gap only when the information and cost needed to identify the task and invoke its recipe are admissible. The construction space must contain

that selector before its performance can enter the inverse. A superintelligence specification further specifies the human reference, success margin, autonomous information permissions, reliability, and breadth; the definition alone does not supply those measurements.

A.2 Boundary interventions and retraining

A smooth local version of the inverse identifies which tasks matter to the next intervention. Let $F(\theta, h) = \nu\{a_\tau(\theta) \le h\}$ and $q(\theta) = Q_\beta(a_\tau(\theta))$. Suppose F is continuously differentiable near (θ_0, q_0), $F(\theta_0, q_0) = \beta$, and $\partial_h F(\theta_0, q_0) > 0$. The implicit function theorem gives

$$\nabla_\theta q(\theta_0) = -\frac{\nabla_\theta F(\theta_0, q_0)}{\partial_h F(\theta_0, q_0)}. \tag{21}$$

If the joint distribution and differentiation-under-integration assumptions additionally justify the transport identity $\nabla_\theta F(\theta, h) = -f_{a,\theta}(h)\mathbb{E}[\nabla_\theta a_\tau(\theta) \mid a_\tau(\theta) = h]$, then

$$\nabla_\theta q = \mathbb{E}[\nabla_\theta a_\tau \mid a_\tau = q]. \tag{22}$$

This is the conditional-quantile derivative under transport regularity (Tasche, 2001), not merely existence of a positive density. It emphasizes tasks at the qualification boundary rather than average task improvement.

When $a_\tau = a_\tau(\theta, w^\star(\theta))$ includes retraining, suppose the local training optimum satisfies $\nabla_w L(w^\star(\theta), \theta) = 0$ and its Hessian $H = \nabla^2_{ww} L$ is nonsingular. Differentiating the stationarity equation yields $\mathrm{d}w^\star/\mathrm{d}\theta = -H^{-1}\nabla^2_{w\theta} L$, hence

$$\frac{\mathrm{d}a_\tau}{\mathrm{d}\theta} = \partial_\theta a_\tau - (\nabla_w a_\tau)^\top H^{-1}\nabla^2_{w\theta} L. \tag{23}$$

The second term is the implicit training response (Lorraine et al., 2020); omitting it evaluates a different, frozen-model intervention. These statements concern a differentiable local branch; they do not require all trained models to have a unique global optimum.

A.3 Paid sequences are complete recipes

For a construction program, additive expenditures sum and reusable capacities take the maximum simultaneous demand. A retained model, data store, or optimizer state occupies capacity for as long as it is retained. A training token, a byte of storage, and a shard probe remain separate quantities unless a declared price connects them.

One can also track a sequence of construction states K_t and remaining consumable stocks B_t. Let $\rho_t = \mathsf{C}^\star_{B_t}(\varepsilon, \beta; K_t)$ be positive and finite. Define gross technological gain and resource erosion by

$$g_t = \frac{\mathsf{C}^\star_{B_t}(\varepsilon, \beta; K_t)}{\mathsf{C}^\star_{B_t}(\varepsilon, \beta; K_{t+1})}, \qquad e_t = \frac{\mathsf{C}^\star_{B_{t+1}}(\varepsilon, \beta; K_{t+1})}{\mathsf{C}^\star_{B_t}(\varepsilon, \beta; K_{t+1})}. \tag{24}$$

Then $\rho_{t+1} = \rho_t e_t / g_t$. If $B_{t+1,j} \ge r_t B_{t,j} > 0$ for every consumable resource, every fixed recipe's normalized cost increases by at most $1/r_t$, so $e_t \le 1/r_t$. Thus

$$\log \rho_k \le \log \rho_0 - \sum_{t<k} \log(g_t r_t). \tag{25}$$

An identified cumulative net gain that exceeds the initial log gap, together with executable intermediate actions and an attained final inverse, supplies affordability. No independence of the successive gains is required. Each is evaluated after its predecessor and after paying its own cost. The identity accounts for discovered improvements; it does not discover them. Preparatory actions may be useful even when their immediate normalized cost increases.

B Identifying a construction rather than only a curve

B.1 A complete path can leave the allocation unknown

For $f_\theta = (1-\theta)/x_1 + \theta/x_2$ with positive x_i, its Hessian is diagonal with entries $2(1-\theta)/x_1^3$ and $2\theta/x_2^3$. It is strictly convex and decreasing in each coordinate. Along $x_1 = x_2 = t$, the burden is $1/t$, independently of θ.

Let $a_1 = 1 - \theta$, $a_2 = \theta$. By Cauchy–Schwarz,

$$(x_1 + x_2)\left(\frac{a_1}{x_1} + \frac{a_2}{x_2}\right) \geq (\sqrt{a_1} + \sqrt{a_2})^2. \tag{26}$$

At target $\varepsilon > 0$, equality is attained by $x_i^\star = \sqrt{a_i}(\sqrt{a_1} + \sqrt{a_2})/\varepsilon$. This proves (4) and the exact minimum cost. Varying θ changes the optimum while leaving the whole measured path fixed. It is a separation for this family; a different parametric path could span all decision-relevant features.

B.2 Exact contrast and decision identifiability

Let $\mathcal{X}$ be a finite feasible catalog with known costs. Assume $q_\vartheta(x) = \phi_x^\top \vartheta$, $\vartheta \in \mathbb{R}^d$, and noiseless interventions $y = Z\vartheta$.

Proposition B.1 (Decision-relevant span)**.** *The contrast $q_\vartheta(x) - q_\vartheta(z)$ is identifiable for every ϑ if and only if $\phi_x - \phi_z \in \mathrm{range}(Z^\top)$. Exact optimal selection for every unrestricted ϑ is possible if and only if every feasible contrast lies in this range.*

Proof. If $\phi_x - \phi_z = Z^\top a$, the contrast equals $a^\top y$. If it lies outside the row span, there is $h \in \ker Z$ with $(\phi_x - \phi_z)^\top h \neq 0$. Parameters ϑ and $\vartheta + th$ produce the same measurements but different contrasts. This proves the first assertion. When all contrasts are known, compare each recipe with a fixed reference and choose the maximizer. Conversely, a missing contrast produces a null vector h for which the values $\phi_x^\top h$ are not all equal. Worlds th and $-th$ have identical data, whereas their maximizing sets are disjoint. A single data-dependent choice cannot be optimal in both. For any randomized selection, the two expected regrets sum to $t(\max_x \phi_x^\top h - \min_x \phi_x^\top h)$, so the worst uniform regret is unbounded as $t \to \infty$. □

Known shape constraints or a bounded parameter set can reduce this ambiguity. Full parameter recovery is unnecessary: an unknown common offset does not change the best recipe, although it can change whether that recipe meets an absolute target.

An analogous finite-budget example makes the possible allocation loss explicit. For unknown $j \in \{1, \ldots, d\}$, set $q_j(x) = a\sum_i x_i + \kappa x_j$ on $x \succeq 0$, $\sum_i x_i \leq B$. All balanced paths agree, $q_j(t\mathbf{1}) = (ad + \kappa)t$. The informed optimum has quality $(a + \kappa)B$. For any choice, some coordinate receives at most B/d; uniform allocation attains the resulting minimax regret $\kappa B(1 - 1/d)$. The same proof applies to expected allocations under randomized choices. This example quantifies missing directional information without attributing a neural scaling exponent to it.

C Stable qualification and the cost of uncertainty

C.1 Proof of the stable inversion theorem

Take $\varepsilon \geq 0$, $\beta - u > 0$, $\beta + 2u \leq 1$. Define $f_x^+ = \widehat{f}_x + \eta$, $f_x^- = \max(0, \widehat{f}_x - \eta)$, and $c_j^\pm = e^{\pm\xi}\widehat{c}_j$. Write $k_B^\pm = \max_j c_j^\pm / B_j$ and

$$\mathcal{X}^+ = \{x : \nu_n(f_x^+ \leq \varepsilon) \geq \beta + u\}, \tag{27}$$

$$\mathcal{X}^- = \{x : \nu_n(f_x^- \leq \varepsilon) \geq \beta - u\}. \tag{28}$$

On the simultaneous event (5), the pointwise ordering at evaluation tasks is $f_x^- \leq f_x \leq f_x^+$. If $x \in \mathcal{X}^+$, then $F_{x,n}(\varepsilon) \geq \beta + u$, and $F_x(\varepsilon) \geq F_{x,n}(\varepsilon) - u \geq \beta$. Hence conservative recipes truly qualify. If $F_x(\varepsilon) \geq \beta$, then $F_{x,n}(\varepsilon) \geq \beta - u$ and $\nu_n(f_x^- \leq \varepsilon) \geq \beta - u$, so $x \in \mathcal{X}^-$. Since lower cost never exceeds true cost and upper cost never falls below it, infimizing proves $C^- \leq \mathsf{C}_B^\star(\varepsilon, \beta) \leq C^+$.

Now take any x with $F_x(\varepsilon - 2\eta) \geq \beta + 2u$. Its empirical true qualification is at least $\beta + u$. Wherever $f_x \leq \varepsilon - 2\eta$, we have $f_x^+ = \widehat{f}_x + \eta \leq f_x + 2\eta \leq \varepsilon$. It therefore belongs to $\mathcal{X}^+$. Moreover $c_j^+ \leq e^{2\xi} c_j$ for every coordinate, so $k_B^+(x) \leq e^{2\xi} k_B(x)$. Infimizing over tightened qualifiers proves the last inequality of (6). The selected recipe belongs to $\mathcal{X}^+$ and has actual cost bounded by its upper cost, giving (7).

For a nonempty finite catalog every finite minimum is attained. For an infinite class the same inequalities hold for infima, and an approximate witness must be supplied for a recipe statement. Empty qualifier sets have value $+\infty$. If $\varepsilon - 2\eta < 0$, nonnegative burdens make the tightened set empty. The ratio in (8) is used only for positive finite inverses; the feasibility inequalities do not need this ratio.

C.2 Obtaining the CDF event

For K fixed recipes independent of n i.i.d. task draws, the DKW–Massart bound and a union bound yield

$$\mathbb{P}\left\{\sup_{x,t} |F_{x,n}(t) - F_x(t)| > u\right\} \le 2Ke^{-2nu^2}. \tag{29}$$

Thus $u = \sqrt{\log(2K/\delta)/(2n)}$ is sufficient. The recipes may be random after conditioning on an independent training stage. Repeatedly redesigning the recipe on these same task outcomes requires a larger uniform class, conditional fresh blocks, or valid sequential inference. If task burdens themselves are estimated from repeated stochastic executions, their uncertainty also enters the response envelope. A CDF bound does not remove within-task measurement error.

C.3 Why accurate forecasts alone do not bound inverse cost

Fix arbitrarily small $\eta > 0$ and $\varepsilon > \eta$. There are two recipes: x_c costs one and x_e costs $M > 1$. Both observational worlds report $\widehat{f}(x_c) = \varepsilon$ and $\widehat{f}(x_e) = \varepsilon - \eta$. In world $-$, true burdens are $(\varepsilon - \eta, \varepsilon - \eta)$; in world $+$, they are $(\varepsilon + \eta, \varepsilon - \eta)$. Predictions have uniform error at most η. The expensive recipe qualifies in both worlds, whereas the cheap recipe qualifies only in world $-$. Any single surely qualifying construction based on these data must pay M, although world $-$ admits cost one. Since M is arbitrary, no approximation ratio follows from η alone. The price of target tightening is a necessary part of a stable inverse.

A multiplicative envelope gives a useful specialization. Replace the additive response bound by $e^{-\eta}\widehat{f} \le f \le e^{\eta}\widehat{f}$; the same inclusions give the stricter target $\varepsilon e^{-2\eta}$. Suppose an admissible improvement map sends any qualifier at ε to one with burden reduced by $e^{-\mu s}$ at cost increased by at most $e^{\nu s}$, without changing breadth. Then $s = 2\eta/\mu$ yields construction cost ratio at most $e^{\omega + 2\xi + 2\eta\nu/\mu}$ for $u = 0$. This is conditional on an implementable improvement map, not a license to extrapolate a response law.

D Paying for decision-relevant evidence

Let $q(x)$ be a quality to maximize, and $\Phi(B) = \max_{x:c(x)\preceq B} q(x)$ on a finite catalog. Let a pilot consume vector e from consumable resources. On one simultaneous event, suppose $|c_j(x) - \widehat{c}_j(x)| \le \zeta_j$ and

$$|[\widehat{q}(x) - \widehat{q}(y)] - [q(x) - q(y)]| \le \epsilon_d \quad \text{for all considered pairs.} \tag{30}$$

Select an approximate maximizer $\widehat{x}$ of $\widehat{q}$ over $e + \widehat{c}(x) + \zeta \preceq B$, with objective tolerance ϵ_o. Then

$$e + c(\widehat{x}) \preceq B, \qquad q(\widehat{x}) \ge \Phi(B - e - 2\zeta) - \epsilon_d - \epsilon_o. \tag{31}$$

Indeed the first statement follows from the upper cost bound. An optimizer x_0 at $B - e - 2\zeta$ is in the safe predicted set because $\widehat{c}(x_0) + \zeta \preceq c(x_0) + 2\zeta$. Therefore $\widehat{q}(\widehat{x}) \ge \widehat{q}(x_0) - \epsilon_o$, and the contrast envelope proves the second statement. The regret to the original-stock oracle separates as

$$\Phi(B) - q(\widehat{x}) \le \Phi(B) - \Phi(B - e - 2\zeta) + \epsilon_d + \epsilon_o. \tag{32}$$

A capacity such as peak memory is not consumed by each pilot observation. Its feasibility is represented by peak simultaneous use instead of subtracting e from capacity. Any reusable artifact produced by the pilot is included in the terminal recipe state. This result uses no convexity or differentiability; it requires valid envelopes and a feasible smaller-stock comparator. It does not assert that learning the full surface is always worth its bill.

E EXPLICIT TARGET-TO-RESOURCE ALLOCATION

Assume $a_i, w_i, \alpha_i > 0$, $x_i > 0$, and $\varepsilon > 0$. The objective is $\sum_i w_i x_i$ subject to $\sum_i a_i x_i^{-\alpha_i} \le \varepsilon$. Feasible points exist by increasing every x_i. On any finite objective sublevel, positivity of each response term bounds x_i away from zero, while positive prices bound it above. A minimum therefore exists. The constraint is active: otherwise a small proportional reduction of all coordinates reduces cost and remains feasible.

The constraint is strictly convex and admits a Slater point. Its Lagrange stationarity is $w_i = \lambda \alpha_i a_i x_i^{-\alpha_i - 1}$, with $\lambda > 0$. This gives (10). Substitution into the active constraint yields a continuous strictly decreasing function of λ, from $+\infty$ to zero; exactly one root exists. The construction is globally optimal and unique. These are standard convex-optimization arguments applied to the declared mechanism (Boyd & Vandenberghe, 2004).

If every $\alpha_i = \alpha$, define $S = \sum_i (a_i w_i^{\alpha})^{1/(\alpha+1)}$. Direct substitution gives

$$C^{\star}(\varepsilon) = S^{(\alpha+1)/\alpha} \varepsilon^{-1/\alpha}, \qquad x_i^{\star} = \frac{h_i}{w_i} \left(\frac{S}{\varepsilon} \right)^{1/\alpha}. \tag{33}$$

A fixed spending path has $x_i = C\pi_i / w_i$, with $\pi_i > 0$ and $\sum_i \pi_i = 1$. Its burden is $C^{-\alpha} \sum_i a_i w_i^{\alpha} \pi_i^{-\alpha}$. Inverting and dividing by $C^{\star}$ proves (11). For $\alpha + 1 > 1$, the definition of Rényi divergence identifies this ratio as $\exp D_{\alpha+1}(\pi^{\star} \| \pi)$. If a positive-effect ingredient receives zero spending, the required path cost is infinite. For unequal exponents, stationarity implies $x_i \asymp \lambda^{1/(\alpha_i+1)}$ as $\lambda \to \infty$. Thus the cost shares of ingredients with the smallest exponent dominate as $\varepsilon \downarrow 0$; ties retain coefficient-dependent shares.

For (12), let $R = \varepsilon - E_{\infty} - A(N/N_0)^{-\alpha} - Qe^{-\gamma_b b}$. For $R > 0$, isolating the data term gives (13), the minimum continuous processed-token allocation. If $R \le 0$, no finite D reaches the target in this response model. For an allowed finite architecture/bit family, compute (13), round upward to an executable training-batch quantum, and reject points exceeding the token or training-work cap. Price the remaining points using the full packed representation. This is exact for the fitted model and declared family; approximation error and empirical qualification are separate inputs to theorem 3.1.

F INHERITANCE, INFORMATION, AND AMORTIZATION

F.1 THE SCOPE OF A REUSABLE CONSTRUCTION

For fixed target-qualified per-task costs $c_F < c_C$, total reusable cost $P + Jc_F$ is smaller than Jc_C precisely when $J > P/(c_C - c_F)$. The first strictly beneficial integer is $\lfloor P/(c_C - c_F) \rfloor + 1$. Equality is break-even. If $c_F \ge c_C$ and $P \ge 0$, no positive number of future tasks repays this setup by that tariff alone. Different cost coordinates should be compared separately or under an explicit common price.

Suppose observations have revealed every variable except U independent, unbiased private signs that enter a task's required response. Conditional on the transcript, any single guessed assignment has probability 2^{-U} of being completely correct. Randomizing the assignment cannot improve that probability. This elementary conditional calculation supplies one component of the controlled acquisition model. It needs the stipulated independence of the unresolved signs; correlated or predictable signs define a different information requirement.

F.2 PROVENANCE OF ACQUIRED INFORMATION

Let W be an unresolved finite-valued target, S the available information, and Z a result computed from S and private randomness independent of W given S. The Markov condition $W - S - Z$ gives $I(W; Z \mid S) = 0$. For any estimator from (S, Z),

$$\mathbb{P}(\widehat{W} = W) \le \mathbb{E}\left[\max_w \mathbb{P}(W = w \mid S) \right]. \tag{34}$$

Conditional on $S = s$, the estimator randomizes among answers with fixed posterior probabilities. Its average correctness is no greater than the largest posterior probability. This proves the bound,

including randomized computation. It concerns new target information, not the computational effort needed to approximate the best decision supported by existing information (Cover & Thomas, 2006).

For adaptive acquisition, let $H_t = (H_{t-1}, A_t, Y_t)$. Nonanticipation means A_t is conditionally independent of W given the available history H_{t-1}. The chain rule gives

$$I(W; H_T \mid H_0) = \sum_{t=1}^{T} I(W; Y_t \mid H_{t-1}, A_t). \tag{35}$$

To prove it, expand each increment into $I(W; A_t \mid H_{t-1}) + I(W; Y_t \mid H_{t-1}, A_t)$; the first term is zero. One must not replace the right side by an expression conditioning on the entire future action sequence. For instance, if $Y_1 = W$ and $A_2 = Y_1$, conditioning on A_2 erases the information in Y_1, even though the agent learned W. If an M-valued target must be recovered with error at most e, the conditional Fano bound additionally requires acquired information at least $H(W \mid H_0) - h_2(e) - e\log(M-1)$ when $e \le 1 - 1/M$. Useful prior structure reduces the unresolved entropy; internally copied observations do not count as independent acquisitions.

G THE CONDITIONAL SPECIALIZATION MECHANISM

There are K regimes, $p_i > 0$, $\sum_i p_i = 1$, with desired vectors $\theta_i \in \mathbb{R}^d$. The router's fixed channel is $C_{ir} = \mathbb{P}(R = r \mid I = i)$, with $w_r = \sum_i p_i C_{ir} > 0$ after deleting empty routes. Let $q_{ir} = p_i C_{ir}/w_r$, $\mu = \sum_i p_i \theta_i$, and $\mu_r = \sum_i q_{ir}\theta_i$. Define

$$H_0 = \sum_i p_i \|\theta_i - \mu\|^2, \quad S = \sum_r w_r (q_r - p)(q_r - p)^\top, \quad H = \sum_r w_r \|\mu_r - \mu\|^2. \tag{36}$$

Conditional expectation gives $0 \le H \le H_0$. Suppose training produces unbiased estimates $\widehat{\theta}_i = \theta_i + e_i$. The new deployment pair (I, R) is independent of the training errors, whose second moments are $\Gamma_{ij} = \mathbb{E}\langle e_i, e_j\rangle$. Correlation across these errors is allowed. Both S and Γ are positive semidefinite. Put $V_0 = p^\top \Gamma p$ and $V = \operatorname{tr}(S\Gamma) \ge 0$.

Store one compiled vector per route,

$$\widehat{m}_t(r) = \widehat{\mu} + t(\widehat{\mu}_r - \widehat{\mu}), \qquad 0 \le t \le 1. \tag{37}$$

The population vector is $\mu + t(\mu_r - \mu)$. Expanding its squared error and using the conditional-mean projection identity gives bias $H_0 - 2tH + t^2 H$. The estimation error equals $\sum_i [p_i + t(q_{ir} - p_i)] e_i$. Its expected squared norm, averaged over routes, is $V_0 + t^2 V$, because $\sum_r w_r (q_r - p) = 0$. Unbiasedness and independence of the new query eliminate the population-error cross term. Therefore

$$\mathcal{R}(t) = H_0 + V_0 - 2tH + t^2(H + V). \tag{38}$$

For $H + V > 0$, completing the square gives

$$t^* = \frac{H}{H+V}, \quad \mathcal{R}(t^*) = H_0 + V_0 - \frac{H^2}{H+V}, \quad \mathcal{R}(t) - \mathcal{R}(t^*) = (H+V)(t - t^*)^2. \tag{39}$$

For $H = V = 0$, all choices give the same risk; pooling avoids unnecessary route-specific storage. Full specialization improves over pooling exactly when $H > V$. Partial sharing gives positive risk reduction whenever $H > 0$, before charging any additional setup, storage, or routing cost. This is a conditional-template statement: nonlinear neural experts generally cannot be averaged offline into one expert of the original size.

If the templates are biased, a general fixed predictor $f_0 + tg$ still has quadratic risk $\mathcal{R}(f_0 + tg) = \mathcal{R}(f_0) - 2tA + t^2 B$ with $A = \mathbb{E}\langle Y - f_0, g\rangle$ and $B = \mathbb{E}\|g\|^2$. For $B > 0$, its optimum on $[0, 1]$ is $\operatorname{clip}(A/B, 0, 1)$. These moments must be identified independently of the final evaluation. The more interpretable H, V decomposition requires the stated unbiased-template model.

A conservative choice. On a simultaneous event $H \in [H_-, H_+]$ and $V \in [V_-, V_+]$ with nonnegative endpoints, choose $t_s = H_-/(H_- + V_+)$, with value zero at zero denominator. For each fixed $t \in [0, 1]$ the gain is $(2t - t^2)H - t^2 V \ge (2t - t^2)H_- - t^2 V_+$. Maximizing this quadratic yields gain at least $H_-^2/(H_- + V_+)$. Confidence in a point estimate is not implied by unbiasedness: if $\widehat{\Theta} = \Theta + e$, then $\mathbb{E}\operatorname{tr}(\widehat{\Theta}^\top S \widehat{\Theta}) = H + \operatorname{tr}(S\Gamma)$, so even the naive signal estimate requires a variance correction.

A load-preserving intervention. Keep the templates and their training-error distribution fixed. Replace C_{ir} by $C'_{ir} = w_r$. Route frequencies stay fixed, while the route becomes independent of the task. Prediction norms and estimation variance are unchanged. The original cross moment between $\theta_I - \mu$ and $\mu_R - \mu$ is H; the new cross moment is zero. Consequently

$$\mathcal{R}_{\text{shuffled}}(t) - \mathcal{R}_{\text{original}}(t) = 2tH. \tag{40}$$

The identity averages over both training and deployment. Relabeling routes and their corresponding templates consistently is an invariant transformation and is not this intervention. Recompiling the conditional means after shuffling would construct a different model.

H JOINT SHARING AND PRICED DATA ACQUISITION

Assume additionally that training errors are uncorrelated across regimes, $\mathbb{E}\|e_i\|^2 = v_i/n_i$, and $v_i, c_i > 0$. Counts are initially continuous, with $n_i > 0$ and $\sum_i c_i n_i \le T$. Define $s_i = S_{ii}$ and

$$a_i(t) = v_i(p_i^2 + t^2 s_i), \quad Z(t) = \sum_i \sqrt{c_i a_i(t)}, \quad B(t) = H_0 - 2tH + t^2 H. \tag{41}$$

Cauchy–Schwarz gives $(\sum_i a_i/n_i)(\sum_i c_i n_i) \ge (\sum_i \sqrt{c_i a_i})^2$. Equality, together with saturation of the positive budget, uniquely yields

$$n_i^*(t, T) = \frac{T\sqrt{a_i(t)/c_i}}{Z(t)}. \tag{42}$$

The resulting scalar objective is $\Phi_T(t) = B(t) + Z(t)^2/T$. This square-root allocation is related to classical optimum stratified sampling (Neyman, 1934); the present objective couples it to the amount of shared structure.

Every $\sqrt{p_i^2 + s_i t^2}$ is nonnegative and convex, so $U(t) = Z(t)^2$ is convex. For $H > 0$, $B''(t) = 2H$ makes Φ_T strictly convex. Its derivative at zero is $-2H < 0$. Since $H > 0$ implies some $s_i > 0$, the derivative at one is $U'(1)/T > 0$. Thus a unique minimizer lies in $(0, 1)$. Implicit differentiation of $-2H(1-t) + U'(t)/T = 0$ gives

$$\frac{\mathrm{d}t^*}{\mathrm{d}T} = \frac{U'(t^*)}{T^2[2H + U''(t^*)/T]} > 0. \tag{43}$$

The same equation gives limits zero as $T \downarrow 0$ and one as $T \to \infty$. Gaussianity is unnecessary. Controlled allocation must permit the chosen n_i; passive traffic implements a different design.

At fixed t, meeting ε requires $T \ge Z(t)^2/(\varepsilon - B(t))$ with positive denominator. Taking the infimum proves the target inverse. Bias is bounded below by $W = H_0 - H$. For $\varepsilon \le W$, positive finite variances preclude attainment. For $\varepsilon > W$, $t = 1$ and sufficiently large T give a finite construction. Integer counts obtained by rounding up cost less than $\sum_i c_i$ additionally and cannot increase the variance terms; the extra cost must fit the resource budget.

Balanced closed form. For perfect routing, $p_i = 1/K$, $v_i = v$, and $c_i = 1$, equal sample allocation $n_i = T/K$ gives

$$\mathcal{R}_{\text{pool}} = h + v/T, \qquad \mathcal{R}_{\text{full}} = Kv/T, \tag{44}$$

$$t^*(T) = \frac{h}{h + (K-1)v/T}, \qquad \mathcal{R}^*_{\text{share}} = h + v/T - \frac{h^2}{h + (K-1)v/T}, \tag{45}$$

where $h = H_0 > 0$. Full specialization beats pooling exactly at $T > (K-1)v/h$. To invert the shared rule, set $x = v/T$. The target equation is $(K-1)x^2 + [Kh - (K-1)\varepsilon]x - \varepsilon h = 0$. Its positive root gives

$$T^*(\varepsilon) = \frac{v}{2\varepsilon h}\left[B_\varepsilon + \sqrt{B_\varepsilon^2 + 4(K-1)\varepsilon h}\right], \quad B_\varepsilon = Kh - (K-1)\varepsilon. \tag{46}$$

For $K = 4$, $h = 1/4$, $v = 1$, the crossover is $T = 12$. At $T = 8$, the risks are 0.375, 0.500, and 0.275; at $T = 64$, they are 0.265625, 0.0625, and 0.055099. Target 0.1 requires 31.7764 under optimal sharing versus 40 under full specialization. These numbers evaluate the specified model.

I PRECISION, ROUTING, AND BREADTH

For a fixed required task set, let $r_i(n_i, b_i) = v_i/n_i + \kappa_i e^{-kb_i}$ with positive v_i, κ_i, k, and per-task target $\tau_i > 0$. This may be an exact additive channel or a validated upper envelope. Minimizing sample cost for given precision saturates each qualifying constraint:

$$n_i^*(b_i) = \frac{v_i}{\tau_i - \kappa_i e^{-kb_i}}. \tag{47}$$

The denominator must be positive. With prices c_i and storage weights d_i, the remaining optimization minimizes $\sum_i c_i v_i/(\tau_i - q_i)$ over $\sum_i d_i b_i \le M$, where $q_i = \kappa_i e^{-kb_i}$. On its finite domain,

$$f_i'(b) = -\frac{c_i v_i k q_i}{(\tau_i - q_i)^2}, \qquad f_i''(b) = \frac{c_i v_i k^2 q_i(\tau_i + q_i)}{(\tau_i - q_i)^3} > 0. \tag{48}$$

The continuous problem is strictly convex. For an interior active-budget solution, stationarity $-f_i'(b_i) = \lambda d_i$ reduces to $y_i/(1-y_i)^2 = z_i$, where $y_i = q_i/\tau_i$ and $z_i = \lambda d_i \tau_i/(c_i v_i k)$. The root in $(0, 1)$ is

$$y_i = \frac{2z_i}{2z_i + 1 + \sqrt{4z_i + 1}}, \qquad b_i^* = \frac{1}{k}\log\frac{\kappa_i}{\tau_i y_i}, \qquad n_i^* = \frac{v_i}{\tau_i(1-y_i)}. \tag{49}$$

Bits decrease monotonically with λ. Clipping to admissible bit bounds and solving the remaining monotone storage equation yields the constrained continuous solution. If all upper bounds are reached, storage may be slack. Integer widths require a finite branch comparison; continuous rounding is not an uncharged feasibility argument.

Execution channels. For arbitrary quantization perturbation δ with $\mathbb{E}\|\delta\|^2 \le q$, the L^2 triangle inequality gives $\mathcal{R}(f+\delta) \le (\sqrt{\mathcal{R}(f)} + \sqrt{q})^2$. An additive identity requires the cross term to vanish, for example conditional unbiasedness given the original prediction and target. If observations lie in $[-L, L]^d$, their sample means and convexly compiled templates do too. Uniform stochastic rounding to 2^b levels is conditionally unbiased without overload and has total variance at most $dL^2/(2^b - 1)^2$. This supplies a finite-range channel. For a deterministic error bound, a sufficient target for the unquantized constructor is $(\sqrt{\varepsilon} - \sqrt{q})^2$ when $\varepsilon > q$.

Router precision has another mechanism. Let its top-two score margin be $m(x)$. Perturbing every score by at most u cannot change the selected route when $m(x) > 2u$. The winner can decrease by u while a competitor increases by u, proving the statement and the sharp factor two. If changing route changes loss by at most L, its expected contribution is at most $L\mathbb{P}(m \le 2u)$. Unbounded losses require tail or second-moment control instead. Expert rounding and discrete router perturbations should not be assigned one error law without evidence.

A breadth reversal. For regime probabilities $(0.99, 0.01)$, unit noise and prices, and $T = 200$, minimizing mean specialist risk $\sum_i p_i/n_i$ gives $n_i \propto \sqrt{p_i}$, namely $(181.735, 18.265)$. Errors are $(0.00550, 0.05475)$ and mean error 0.005995. Only 99% of task mass has conditional error at most 0.01. A 99.5%-breadth requirement includes both regimes; $n = (100, 100)$ meets it at the same cost, with mean error 0.01. For general breadth, selecting which tasks must qualify is itself a discrete covering problem. A rare required task does not become cheaper because its deployment frequency is small.

J EXPERIMENTAL RECORDS AND REPRODUCIBILITY

The numerical results in section 6 use three distinct evaluation criteria. Table 2 records the unit of success, charged resources, and evidentiary scope. The analysis scripts recompute the reported values from stored outcomes. Cryptographic hashes link each stored forecast to its corresponding result bundle. They document file consistency but do not constitute external preregistration.

J.1 PRECISION–DATA INVERSE: MODEL, IMPLEMENTATION, AND INFERENCE

The corpus is the byte-encoded TinyStories dataset (Eldan & Li, 2023). The dataset revision is `f54c09fd23315a6f9c86f9dc80f725de7d8f9c64`; cached data contain 120 million training bytes and separate eight-million-byte development and evaluation regions of the official validation

file. Windows have length 256. Training draws random windows, so D denotes processed tokens, including repeated corpus observations. The byte vocabulary has size 256.

The models are causal Transformers with learned token and position embeddings, tied token/output weights, pre-normalized residual attention and MLP blocks, GELU activations, and MLP expansion four. B has six layers, width 256, and four heads; C has seven layers, width 288, and six heads; D has seven layers, width 320, and eight heads. Their parameter counts are 4,856,320, 7,123,392, and 8,775,040. The full fitting architecture family also contains the smaller A and larger E–G models specified in the machine-readable records.

Weight quantization is symmetric, row-scaled, and straight-through in training. The signed maximum is $2^{b-1}-1$; each scale is the maximum absolute row value divided by that maximum. Quantization begins after 5% of optimizer steps. Activations and matrix kernels use ordinary floating-point execution. Packed-size accounting charges row-aligned integer payload, a 16-bit scale per row, and 16-bit LayerNorm parameters; tied weights are stored once. This is a declared representation size rather than a measured integer-kernel memory footprint.

AdamW uses $(\beta_1, \beta_2) = (0.9, 0.95)$, $\epsilon = 10^{-8}$, weight decay 0.1 on matrix parameters and zero on normalization parameters. The initial learning-rate scale is $3 \times 10^{-4}\sqrt{256/d}$, with a capped linear warm-up and cosine decay to one tenth. Gradients are clipped at one. Global batches contain 32,768 tokens. Seeds are 20260922 and 20260923. The precise token allocations are multiples of this global batch; the metadata retain all counts and training-work products.

The response parameters, in the order $(E_\infty, A, \alpha, C, \gamma_D, Q, \gamma_b)$, are

$$(0.0998340,\ 0.488342,\ 0.304411,\ 0.359703,\ 1.067277,\ 19.99999997,\ 1.753955), \tag{50}$$

with $N_0 = 8{,}775{,}040$ and $D_0 = 40$ million. Calibration uses the Cartesian grid B/D/F at $20/40/80$ million tokens and eight bits, plus D at four/six bits and 40 million tokens. The four mechanism checks are B/6-bit/60M, F/6-bit/30M, D/4-bit/60M, and F/4-bit/20M. Their maximum absolute residual is 0.040094 and mean absolute residual 0.020753. The fitted coefficient Q reaches its imposed upper bound. The joint precision penalty is the usable fitted quantity; these measurements do not separately establish a precise physical interpretation of Q or provide a uniform response band.

The target inverse searches B–F, bits 4–8, and continuous token counts within 20–80M before upward batch rounding. The fixed targets and complete selected recipes are in table 3. Selection is based on the fixed fitted law, not the new model outcomes. The resulting twelve trained models are evaluated at the same 320 batches of 16 windows per model. For each seed and target, the interval is the paired mean batch-loss difference ± 1.96 times its batch standard error. It is an evaluation-sampling interval conditional on those trained models, not a confidence statement over a population of training seeds. Potential corpus-window dependence and reuse of the fixed validation corpus also limit a fresh-task interpretation. The recipe forecasts precede these trained configurations; the corpus is not newly sampled for each scientific question.

The largest evaluation upper bound is 0.009652 for E1, seed 1. All six comparisons pass their declared marginal rule. Absolute nominal-target attainment is a different statistic: E3 seed 0 has construction loss 0.792312 against target 0.786164, while its baseline has loss 0.786326. The implemented development gate tests relative noninferiority, not exact nominal-target attainment. At E2, the fitted five-bit prediction at the available work cap is approximately 0.820668, just 0.002499 above the nominal target; the more tolerant evaluation rule would not exclude it on that fact alone. Consequently the recorded experiment establishes viable preselected recipes, not minimal bit widths or a confirmed optimality switch. The study does not compare against a full human-tuned or Bayesian architecture search with a matched experimental bill.

J.2 Controlled inheritance: world generator, comparators, and costs

A world has public normalized mode embeddings in $\mathbb{R}^{64}$ and an eight-dimensional context. A fixed hidden teacher maps the context through a feature family and a latent matrix to a relevance direction. The r highest relevance modes form the private target support; their signs are independent private bits. A query reveals that mode's sign, a noisy relevance score, and a measurement of a shared scalar field. The primary case uses the tanh-field family, 32,768 public modes, and $r = 8$. The field model is a depth-two, width-64 MLP; its ground-truth field and noise channels are specified in the supplied engine.

Historical training contains 512 contexts with 16 labeled support identities each. Their support-embedding centroids train ridge maps using raw, tanh, Fourier, radial-basis, and concatenated features, with penalties 0.1, 1, or 10. A disjoint selector subset chooses the map; a separate forecast subset chooses query budgets. The selected primary map is `static:concat`. Public embedding geometry is available to all policies. The cold controller estimates the current relevance direction from 8, 16, 32, or 64 task-local relevance probes and then ranks remaining modes by that ridge estimate. The primary comparator uses 64 such probes within its total budget of 256. It receives current-task evidence but no historical relevance labels.

The shared primitive starts from 128 grounded observations and a fixed prescribed training procedure. The query controller predicts the target signs, obtains their signed field values where queried, and guesses unresolved independent signs. A completed task must satisfy the stored sign, primitive-error, and response-error criteria: every required sign is correct, primitive MSE is at most $1/(8r)$, and normalized task-response MSE is at most $1/(18r)$. Hidden-support scoring is vectorized; it is not a large foundation-model generation run. The implementation and per-world outcomes are included for exact metric reconstruction.

At the selected budgets, reuse has $958/1024$ completions and cold $941/1024$. The registered simultaneous binomial lower bounds are approximately 0.9027 and 0.8831, both above the 0.8 target. Figure intervals are separately identified as pointwise Clopper–Pearson intervals (Clopper & Pearson, 1934). The paired discordant counts are 76 and 59; an exact binomial test on the 135 discordances gives two-sided $p = 0.168259$. The same-reuse-class `concat_shared` control has the same outcomes as the selected map, and `raw_shared` completes $959/1024$. The effect belongs to useful reusable information, rather than a uniquely superior online search heuristic.

The declared tariff assigns price 16 to a field/relevance query. A common 128-query component costs 2048 and the normalized primitive-training contribution is 2.125. Total per-task tariffs are therefore $2048+16(32)+2.125 = 2562.125$ and $2048+16(256)+2.125 = 6146.125$. Setup priced by historical labels is $512(16)(16) = 131072$; charging each historical task as cold reconstruction instead gives $512(256)(16) = 2097152$. Dividing by the per-task saving 3584 gives first strictly beneficial task counts 37 and 586. This acquisition-and-primitive-training tariff does not include measured selector-fit wall time, full index construction, or validation execution; it is not a complete hardware speedup.

In the private-structure canary, a task-private relevance component breaks the predictive relation between reusable context and support while leaving current-task acquisition available. Reuse completes $3/128$ worlds and cold $118/128$. This is a distinct world-distribution intervention, not the exact load-preserving channel in (40). The signed forecast payload, protected-stream link, per-world outcomes, and full nine-case audit family are retained in the source package.

J.3 Conditional external retrieval and its stopped forecast

The corpus uses HotpotQA's distractor configuration with bridge questions having two labeled support documents. The pinned dataset revision is `1908d6a`. Documents are embedded using Qwen3-Embedding-0.6B at revision `97b0c614be4d77ee51c0cef4e5f07c00f9eb65b3`. Unsupervised clustering creates 128 document shards. Historical support pairs supply transition destinations. A fixed hash chooses the support anchor, rather than recovering a ground-truth temporal order of reasoning. The routing query combines the question and the supplied anchor at weight 0.5; the reusable component combines source-shard transitions and 96 semantically neighboring historical questions. The candidate historical weights are $0.25, 0.5, 0.75, 1$.

The nominal splits contain 8000 historical tasks and 512 each for calibration, validation, and protected evaluation, after applying a fixed task-exclusion list defined before calibration. The conditional evaluation keeps tasks whose two supports are in distinct shards; 379 calibration and 362 validation tasks remain. Its unit is the shard containing the second support, conditional on a supplied gold first support. Counting shard probes does not by itself measure document reads, embedding work, passage reranking, answer accuracy, or total latency.

The budget grid is $1, 2, 4, 8, 16, 32, 64$ and the target grid is $0.5, 0.6, 0.7, 0.8$. Calibration selects $\lambda = 0.75$, target 0.7, and budgets $(4, 8)$. Validation must reproduce both exact empirical crossings, preserve a canary drop of at least 0.15, and have calibration-to-validation coverage shift at most

0.07. Its maximum reported coverage shift is 0.03304 and canary drop 0.54420, but the inherited four-probe crossing misses. The final protected retrieval and downstream QA stages are consequently not executed. The complete selected validation counts are given in table 4.

Probes	Cold successes	Inherited successes	Shuffled successes	Inherited minus cold (pp)
1	114	138	19	6.63
2	165	193	32	7.73
4	227	251	54	6.63
8	287	317	82	8.29
16	335	345	123	2.76
32	353	355	174	0.55
64	360	358	279	-0.55

Table 4: Selected external validation policy, $n = 362$, history weight 0.75. These are conditional shard-recovery counts. Aggregates permit descriptive matched-budget contrasts, but the stored stop report does not contain the joint per-task outcomes needed for a paired significance test.

At four probes, the inherited calibration proportion is $267/379 = 0.704485$ and validation is $251/362 = 0.693370$. The latter's pointwise Wilson interval is $[0.644046, 0.738633]$ (Wilson, 1927). A discrete 70% empirical crossing needs 254 validation successes. Under an illustrative binomial model with true success probability fixed at the calibration point estimate, the chance of repeating that crossing is $\mathbb{P}\{\mathrm{Bin}(362, 267/379) \geq 254\} = 0.572638$. This calculation exposes the sensitivity of exact grid crossings; it is not a replacement success criterion or an estimate treating the unknown success probability as known.

Historical-transition permutation is a damaging-knowledge control. Its difference from the unpermuted policy can reflect harm from misleading transitions as well as value from correct ones. The cold no-history comparator is therefore essential. The exact $2\times$ prediction remains a failure even though the low-budget descriptive contrasts are positive. Changing the history weight or reliability target after seeing this validation set would be a different experiment.

J.4 Reproducing the tables and figures

The supplementary archive contains the anonymous manuscript and supplement, all plotted data, vector figures, and the programs used for reanalysis. From its root, run

```
python -m pip install -r requirements.txt
bash build.sh --all
```

The build reconstructs the experimental summaries, runs 43 mathematical checks, renders the six plot assets and compiles both anonymous PDFs. A TeX installation with pdfLaTeX, BibTeX, PGF, and the listed font packages is required. The included `README.md` maps each result to its proof, input record, and implementation.

The three notebooks in `reproduction/notebooks/` run the precision–data, controlled inheritance, and conditional retrieval experiments. Each includes its source payload and persistent resume logic. Model-level replay requires the specified public model/corpus downloads and GPU resources; the compact package instead bundles the recorded outcomes needed to reconstruct every reported table and numerical figure. The external records contain aggregate validation curves, while the numerical-precision and controlled-world records retain their respective batchwise and worldwise outcomes. File checksums, template checksums, and executable-test results are provided with the archive.

K Positioning and scope of the construction laws

Multivariate scaling laws already optimize model size, data, precision, and expert count (Hoffmann et al., 2022; Kumar et al., 2024; Clark et al., 2022; Ludziejewski et al., 2025). Construction here adds a fixed qualification criterion, decision-identifying evidence, complete resource pricing, and

sensitivity of the target inverse. A one-dimensional path is not intrinsically uninformative; section B isolates when its observation span lacks the required contrasts. Convex duality, Rényi divergence, information provenance, and square-root allocation are established mathematical tools. The results combine them into explicit target-to-recipe statements rather than claim them as new primitives.

Task sharing has a long history in multitask learning (Caruana, 1997). The exact template model makes the signal explained by the router and the covariance of estimation error explicit, then jointly optimizes sharing and evidence allocation. Its constant active template dimension follows from offline linear compilation. In a nonlinear neural mixture, both the risk law and execution savings require their own identification. The new neural measurements in this paper concern the separate precision–data mechanism, while figure 1 evaluates an exact assigned specialization model.

For a claimed resource gain, the comparator, target tolerance, task distribution, and accounting horizon must agree. A ratio of two lower confidence bounds is not a lower bound on a ratio of true resource optima. A small forward error can create a large target-cost error, and a correct law can recommend the existing baseline. A controlled construction establishes its declared target. Extending that target to broad expert-relative superiority requires the corresponding expert measurements.

REFERENCES

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. URL https://web.stanford.edu/~boyd/cvxbook/.

Rich Caruana. Multitask learning. *Machine Learning*, 28:41–75, 1997. doi: 10.1023/A:1007379606734.

Aidan Clark, Diego de Las Casas, Aurelia Guy, Arthur Mensch, Michela Paganini, Jordan Hoffmann, Bogdan Damoc, Blake Hechtman, Trevor Cai, Sebastian Borgeaud, et al. Unified scaling laws for routed language models. In *International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 4057–4086, 2022. URL https://proceedings.mlr.press/v162/clark22a.html.

C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. doi: 10.1093/biomet/26.4.404.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2 edition, 2006. doi: 10.1002/047174882X.

Ronen Eldan and Yuanzhi Li. TinyStories: How small can language models be and still speak coherent english? *arXiv preprint arXiv:2305.07759*, 2023. URL https://arxiv.org/abs/2305.07759.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://arxiv.org/abs/2203.15556.

Tanishq Kumar, Zachary Ankner, Benjamin F. Spector, Blake Bordelon, Niklas Muennighoff, Mansheej Paul, Cengiz Pehlevan, Christopher Ré, and Aditi Raghunathan. Scaling laws for precision. *arXiv preprint arXiv:2411.04330*, 2024. URL https://arxiv.org/abs/2411.04330.

Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1540–1552, 2020. URL https://proceedings.mlr.press/v108/lorraine20a.html.

Jan Ludziejewski, Maciej Pióro, Jakub Krajewski, Maciej Stefaniak, Michał Krutul, Jan Małaśnicki, Marek Cygan, Piotr Sankowski, Kamil Adamczewski, Piotr Miłoś, and Sebastian Jaszczur. Joint MoE scaling laws: Mixture of experts can be memory efficient. *arXiv preprint arXiv:2502.05172*, 2025. URL https://arxiv.org/abs/2502.05172.

Jerzy Neyman. On the two different aspects of the representative method: The method of stratified sampling and the method of purposive selection. *Journal of the Royal Statistical Society*, 97(4): 558–606, 1934. doi: 10.1111/j.2397-2335.1934.tb04184.x.

Dirk Tasche. Conditional expectation as quantile derivative. *arXiv preprint arXiv:math/0104190*, 2001. URL https://arxiv.org/abs/math/0104190.

Edwin B. Wilson. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212, 1927. doi: 10.1080/01621459.1927.10502953.