The Minimal Coupling Principle for Budget and RoS Pacing: One Shadow Price is Enough (and Sometimes Necessary)

Liz Lemma Future Detective

January 16, 2026

Abstract

Modern autobidding stacks typically implement multiple pacing loops—most prominently spend pacing (budget) and efficiency pacing (target CPA / RoS). The survey literature shows that naively decoupling these loops can cause poor performance and constraint violations, while more coordinated designs can be near-optimal. We formalize this systems question as communication-constrained online optimization in repeated truthful auctions with a value-maximizing bidder facing both a budget constraint and a return-on-spend constraint. Our first contribution is a minimal-coupling sufficiency theorem: despite two constraints, the optimal bidding policy in truthful auctions remains uniform bid scaling, so the operational decision reduces to a single scalar multiplier. We construct a distributed pacing architecture where the budget and RoS controllers exchange only this scalar (a composite shadow price), and we prove it achieves $O(\sqrt{T})$ regret with tight feasibility guarantees, matching fully-coupled primal–dual baselines. Our second contribution is a necessity theorem: for a broad class of fully-decoupled architectures with separated feedback and no cross-message, there exist adversarial auction sequences where the controllers 'fight,' yielding $\Omega(T)$ welfare loss or persistent infeasibility. The results give actionable guidance for modular ad stacks in 2026: near-optimal performance requires only a single shared statistic per period, but removing even that coupling can be catastrophic. We complement theory with simulations on synthetic auction streams mimicking real pacing loops; no numerical methods are needed for the core characterizations beyond standard convex optimization steps.

Table of Contents

2. Related work: LP dual bidding rules in truthful auctions; online primal–dual pacing; empirical/engineering observations on coupling; contrasts with equilibrium-centric analyses.

3. Model: repeated truthful auctions, value maximizer, Budget and RoS constraints; uniform multiplier policies; information partitions across controllers; definition of architectures and communication budget.

4. Benchmark and objective: best fixed feasible multiplier κ^*; regret and constraint violation metrics; discussion of why one-dimensional control is economically meaningful.

5. Minimally-coupled algorithm: composite shadow price / single-message protocol; update rules; implementation details (hard budget via stopping, RoS via dual control).

6. Main theorem (sufficiency): regret and feasibility bounds; closed-form dependence on step size and horizon; discussion of extensions (partial feedback, delayed feedback).

7. Impossibility (necessity) for fully decoupled pacing: formal architecture class; adversarial construction; $\Omega(T)$ lower bound; interpretation as information bottleneck.

8. Simulations: synthetic auction streams; compare fully-coupled vs minimally-coupled vs decoupled; robustness checks (non-stationarity, feedback delays).

9. Extensions and discussion: multiple advertisers (external prices endogenous); multi-slot; non-truthful auctions (limits of the 1-scalar principle); engineering guidance and takeaways.

10. Conclusion.

1 Introduction and motivation

Modern advertising platforms increasingly implement bidding as a collection of modular "control loops" that operate as internal services. One loop paces spend against a budget, another enforces an efficiency or return-on-spend (RoS) requirement, and yet another may handle risk, experimentation, or delivery goals. This modular design is attractive in practice: each controller can be engineered, monitored, and tuned independently; each can consume a narrow slice of feedback (spend logs for budgeting, conversion value for RoS); and the overall system can be composed quickly across campaigns. Yet the constraints that these modules aim to satisfy are coupled in a fundamental economic sense. The budget constraint restricts the total quantity purchased, while the RoS constraint restricts the average quality of what is purchased relative to price. When the auction environment is volatile, these constraints compete for influence over the same one-dimensional action—how aggressively to bid—and this raises a basic architectural question: *how much coupling is actually required* between controllers to achieve near-optimal value while maintaining feasibility?

We study this question in a setting that intentionally strips away strategic complications in order to isolate the information and coordination problem inside the advertiser. The environment is a sequence of truthful single-slot auctions, so the advertiser faces a per-round threshold price and wins if and only if its bid clears that threshold. This reduces the auction-time decision to a cutoff rule: buy an impression if it looks "good enough" relative to price. In such settings, practitioners often implement bidding through a uniform multiplier (sometimes called a pacing multiplier), scaling the platform's estimated value by a single scalar. This design is not merely an engineering convenience; it is the natural degree of freedom in a truthful threshold model, and it is also the object around which budget and RoS controllers must coordinate.

The temptation, from a systems perspective, is to let each module compute its own preferred multiplier and then combine them mechanically, for instance by taking the minimum (to be conservative) or an average (to smooth noise). This "fully decoupled" design appears to respect modularity: the budget loop looks only at spend, the RoS loop looks only at value per dollar, and the final bid multiplier is an aggregator applied to the two outputs. The difficulty is that neither loop can infer, from its local feedback alone, whether the other constraint is currently binding, slack, or about to bind. In economic terms, each loop is attempting to learn a shadow price, but shadow prices are meaningful only jointly: raising the budget shadow price and raising the RoS shadow price both reduce bidding, but they do so for different reasons, and an agent that cannot distinguish those reasons may react in the wrong direction.

To see the intuition, consider two regimes that can alternate over time.

In one regime, impressions are plentiful and cheap but marginal in efficiency: spending is easy, but RoS is hard to maintain. In the other regime, high-efficiency opportunities arrive sporadically: they are exactly the impressions one should buy, yet doing so may require keeping the multiplier high enough not to miss them when they appear. A budget-only controller, observing that spend is below target, tends to increase the multiplier during the first regime, exacerbating RoS problems. A RoS-only controller, observing that efficiency is deteriorating, tends to decrease the multiplier, potentially pushing bids so low that the advertiser fails to capture the sporadic high-efficiency opportunities in the second regime. If the two modules are combined by a fixed monotone aggregator, an adversary (or simply a nonstationary market) can force persistent miscoordination: one module tightens when the other should loosen, and the system either violates constraints chronically or leaves substantial value on the table. The key point is informational: without some shared state summarizing the joint tradeoff, local residuals are not sufficient statistics for globally feasible bidding.

Our central message is that the required coupling is simultaneously *minimal* and *essential*. On the sufficiency side, we show that the relevant coordination object collapses to one scalar. Because the auction is truthful and allocation is a threshold rule, the offline optimum for a value maximizer subject to budget and RoS constraints can be implemented by a uniform multiplier—equivalently, by selecting a single cutoff on the ratio of price to value. This one-dimensional structure means that a distributed architecture does not need to exchange rich histories, gradient vectors, or separate multipliers. It needs only the effective multiplier itself (or, equivalently, the composite shadow price that induces it). We then construct a minimally coupled two-controller implementation of a primal–dual pacing algorithm: internally, the modules may maintain distinct dual variables (one for budget, one for RoS), but at decision time they communicate only one real number per round, the multiplier used to form the bid. The budget module can still enforce the hard budget mechanically by stopping when cumulative spend reaches the cap, while the RoS constraint is handled through online dual updates that ensure sublinear violation. In this sense, modularity is compatible with rigor: we can preserve the engineering separation of concerns while matching the performance of a fully centralized controller up to the standard online learning rates.

On the necessity side, we formalize the failure mode of fully decoupled designs. We consider a broad class in which each module updates using only its own feedback stream and the final multiplier is a fixed monotone combination of the two outputs. For this class, we show that there exist sequences of auction opportunities under which any such architecture must incur linear loss in at least one dimension: either large cumulative violations of budget and/or RoS, or linear regret relative to the best fixed feasible multiplier in hindsight. This is an information-theoretic impossibility: when the environment can switch between regimes in which different constraints bind, a decoupled system cannot reliably identify which shadow price should dominate without exchanging at least some coordinating message. The result clarifies why ad platforms that attempt strict service-level separation often reintroduce coupling through shared state, global pacing layers, or combined "effective bid" signals: these engineering patterns are not incidental, but rather reflect the minimal communication required by the economics of coupled constraints.

Finally, we address robustness considerations that matter for practice. Cross-service communication may be delayed, quantized, or noisy, and any useful architectural prescription should degrade gracefully under such frictions. Because the decision variable is effectively one-dimensional, we can leverage stability results from online optimization to show that modest delays or noise inflate regret and constraint violations by controlled additive terms, rather than causing qualitative breakdown. This reinforces the practical appeal of minimal coupling: a single scalar message is not only sufficient for coordination, but also easy to transmit reliably at scale.

Taken together, our contributions offer a simple organizing principle: *one scalar shadow price is the right unit of coordination* for budget-and-RoS-constrained bidding in truthful auctions. The remainder of the paper situates this principle relative to the duality-based bidding literature and to existing online pacing methods, and clarifies how our coupling results complement equilibrium-centered analyses.

2 Related work

Our results build on (and are meant to clarify) a line of work that uses linear programming duality to characterize optimal bidding in *truthful* auction environments when advertisers face long-run constraints. In threshold-price models—including second-price and posted-price variants—a bidder's decision in each round is essentially whether to accept a take-it-or-leave-it price. For value maximizers subject to a budget, the classical dual view says that an optimal policy can be expressed as a cutoff rule comparing value to a shadow price of budget, equivalently bidding a scaled value; see, e.g., the dual-bidding and pacing formulations in **???**. A related theme is that additional linear constraints (such as minimum ROI, CPA, or other average-efficiency requirements) can often be folded into the Lagrangian as extra dual variables, yielding bidding rules that remain simple but whose parameters reflect multiple shadow prices **??**. We view Proposition 1 as a sharpening of this perspective for the particular pair of constraints we study: despite the presence of two constraints, the induced allocation in truthful auctions is still implementable by a one-dimensional uniform multiplier (or a one-dimensional cutoff on a price-to-value ratio), which is precisely the statistic an implementation must coordinate on at decision time.

A second, closely related strand studies *online* pacing algorithms—typically primal–dual or mirror-descent updates—that adapt multipliers in response to realized spend and value. Online budget pacing in repeated auctions has been analyzed under i.i.d. inputs, adversarial inputs with bounded variation, and mixtures thereof; representative approaches include online primal–dual schemes and their connections to online convex optimization and stochastic approximation **???**. Many of these algorithms are presented in a centralized form: a single learner observes the necessary feedback and updates a multiplier (or a small set of dual variables) that governs bidding. Our contribution is not to improve the canonical $\sqrt{T}$-type rates per se, but rather to show that the same guarantees can be realized under an explicit *architectural* constraint motivated by how production systems are built: separate controllers with separate feedback streams and an explicit communication budget. In that sense, Proposition 2 can be read as an implementability result for online primal–dual pacing: the algorithmic object that matters for regret and feasibility is the *effective* multiplier κ_t, and the internal decomposition into "budget dual" and "RoS dual" can be kept local as long as the composed scalar is shared.

Our focus on RoS (or ROI) constraints connects to a broader literature on constrained bidding objectives in ad markets, including CPA/ROI optimization, value maximization with efficiency constraints, and variants with risk or delivery constraints **??**. Empirically, RoS-style requirements are often enforced at coarse time scales (daily or weekly) and are intertwined with budget pacing; much of the applied work emphasizes heuristics for stabilizing these coupled loops under delayed and noisy conversion signals. Our analysis abstracts from conversion delay by treating v_t as a realized per-impression value, but we do address delay and noise at the level of the cross-controller message (Proposition 4), which is the friction most directly tied to modular service architectures. This choice is deliberate: we aim to separate the economic question of *how much coordination is required* from the statistical question of *how accurately value can be predicted*, while acknowledging that the latter is first-order in many deployments.

The paper is also informed by engineering observations about how large-scale bidding stacks are actually organized. In practice, budget pacing commonly lives in one service, value/ROI optimization in another, and the serving system consumes a single "effective bid" (or multiplier) that is produced by composing multiple signals. It is therefore natural to ask whether one can maintain strict modularity (no shared state) without sacrificing performance. The negative result in Proposition 3 formalizes a failure mode that practitioners often report informally: when different constraints become binding in different market regimes, independent controllers that only see their own residuals can chase each other, producing either chronic constraint violations or excessive conservatism. While similar phenomena are studied in control theory under interacting feedback loops, our contribution is to cast the issue in an online economic decision problem and to provide an information-theoretic separation between "no coupling" and "one-scalar coupling."

Our viewpoint differs from equilibrium-centric analyses of sponsored search and auction markets, which study bidding incentives and equilibrium outcomes under generalized second price or VCG, often with multiple strategic advertisers and auction-specific features such as quality scores and reserve prices **??**. That literature explains how auction rules map bids into allocations and payments, and how strategic behavior shapes prices and welfare. We instead take the auction environment as exogenous and focus on a single advertiser's *internal* constrained optimization, motivated by the observation that large advertisers frequently delegate bidding to automated agents whose primary difficulty is not strategic deviation from truthfulness but rather real-time constraint management under uncertainty. The truthful threshold model is thus an intentional simplification that allows us to isolate the coordination role of shadow prices; extending the coupling/communication question to non-truthful settings is important, but it would require disentangling strategic considerations from architectural ones.

Finally, our results relate to work on decentralized and communication-constrained online learning. A growing literature studies how limited message passing affects regret in distributed optimization, and how delay/noise in gradients or dual variables degrades rates **???**. The novelty here is that the relevant decision space collapses to one dimension due to the truthful-auction structure, which makes the "right" message particularly simple: a single scalar multiplier that plays the role of a composite shadow price. This observation yields both a positive and a negative implication. Positively, a minimal protocol can match centralized learning guarantees up to standard stability terms. Negatively, if one insists on zero communication and fixed aggregation rules, then even very weak adversarial nonstationarity can force linear loss. We regard this as a useful bridge between algorithmic regret guarantees and the practical design question of what state must be shared across services in order for those guarantees to be attainable.

3 Model: repeated truthful auctions, value maximizer, Budget and RoS constraints; uniform multiplier policies; information partitions across controllers; definition of architectures and communication budget.

We evaluate any online pacing architecture by comparing it to the best *fixed* uniform-multiplier policy that is feasible in hindsight. For a scalar multiplier

$\kappa \in K$, let the induced win decision in round t be

$$x_t(\kappa) \;=\; \mathbf{1}\{\kappa v_t \geq p_t\},$$

so that total value and spend under κ are $\sum_t v_t x_t(\kappa)$ and $\sum_t p_t x_t(\kappa)$, respectively. The benchmark κ^* is then defined as an optimizer of the offline constrained problem over the same policy class,

$$\kappa^* \in \arg\max_{\kappa \in K} \sum_{t=1}^{T} v_t x_t(\kappa) \qquad \text{s.t.} \qquad \sum_{t=1}^{T} p_t x_t(\kappa) \leq B, \quad \sum_{t=1}^{T} (v_t - \tau p_t) x_t(\kappa) \geq 0.$$

This choice is deliberately austere: we do not compare against a dynamic oracle that can pick an arbitrary subset of auctions, nor against a policy that can condition on future realizations. Rather, we benchmark against the best *single* scalar that could have been used throughout the horizon. In our setting this is not merely a convenient restriction: Proposition 1 says that, in truthful threshold-price auctions, the offline optimum can be implemented by such a one-dimensional cutoff. Thus, by competing with κ^* we are effectively competing with the economically relevant offline optimum, while keeping the benchmark aligned with what a production bidding stack can plausibly represent (a single pacing/ROI multiplier applied to predicted value).

Given an online sequence $\{\kappa_t\}_{t=1}^{T}$, we measure performance along two axes: foregone value relative to κ^*, and the extent to which long-run constraints are violated. Regret is defined by

$$\mathrm{Reg}_T \;=\; \sum_{t=1}^{T} v_t x_t(\kappa^*) \;-\; \sum_{t=1}^{T} v_t x_t(\kappa_t),$$

while realized totals under the online policy are $V_T = \sum_t v_t x_t$ and $S_T = \sum_t p_t x_t$, yielding the violation magnitudes

$$\mathrm{Viol}_B = (S_T - B)_+, \qquad \mathrm{Viol}_\tau = (\tau S_T - V_T)_+.$$

These metrics separate "how much value we missed" from "how badly we missed the constraints." The separation is important in practice: a controller can trivially satisfy both constraints by bidding $\kappa_t \equiv 0$, but such a policy would incur linear regret whenever there exist profitable impressions; conversely, an aggressive policy can drive value but at the cost of chronic RoS shortfalls. Our target guarantees, $\mathrm{Reg}_T = o(T)$ and $\mathrm{Viol}_B, \mathrm{Viol}_\tau = o(T)$ (typically $O(\sqrt{T})$), formalize the intuitive goal that any average loss or average infeasibility should vanish with horizon.

The hard budget constraint admits a particularly sharp implementation: we may enforce $\sum_{u \leq t} s_u \leq B$ *pathwise* by stopping bidding once cumulative spend reaches B. In that case $\mathrm{Viol}_B = 0$ by construction, and the nontrivial constraint becomes RoS (or, more generally, any efficiency constraint that cannot be enforced myopically impression-by-impression). We nonetheless keep Viol_B explicit for two reasons. First, it allows a unified comparison across architectures that may not implement a hard stop (e.g., when spend is only observed with delay). Second, it clarifies what the online analysis is really doing: the algorithm is not "optimizing value subject to budget" in a static sense, but rather learning a scalar that balances an intertemporal scarcity (budget) against an intertemporal quality requirement (RoS), both of which are only revealed through realized allocations and payments.

The one-dimensionality of the decision variable κ is also economically meaningful. Because $x_t(\kappa) = \mathbf{1}\{\kappa \geq p_t/v_t\}$, choosing κ is equivalent to choosing a cutoff on the *price-to-value ratio* p_t/v_t, i.e., the cost per unit of predicted value. Put differently, κ^{-1} is the maximum "shadow cost" per unit value that the bidder is willing to accept. This interpretation aligns with how practitioners discuss pacing: lowering κ tightens the acceptance criterion, dropping marginally efficient (high p_t/v_t) auctions; raising κ relaxes it, buying more volume at potentially worse efficiency. The RoS constraint $\sum_t (v_t - \tau p_t)x_t \geq 0$ says that, in aggregate, the average value-per-dollar must be at least τ; in a ratio language, we must avoid systematically selecting impressions with low v_t/p_t. A single cutoff cannot guarantee $v_t/p_t \geq \tau$ *in every round*, but it can ensure that, over time, the induced mix of won auctions is sufficiently efficient while still spending budget in periods where good opportunities arrive.

Finally, the scalar benchmark and the scalar communication primitive coincide for a reason: in dual terms, the multi-constraint problem collapses at decision time to a single composite shadow price. If we write the per-round Lagrangian coefficient on allocating in round t as

$$(1 + \alpha_\tau)v_t \; - \; (\alpha_B + \tau\alpha_\tau)p_t,$$

then the decision "win if this coefficient is nonnegative" is equivalent to a threshold on p_t/v_t, hence implementable by bidding $b_t = \kappa_t v_t$ with κ_t proportional to $(1+\alpha_{\tau,t})/(\alpha_{B,t}+\tau\alpha_{\tau,t})$ (when the denominator is positive, with the usual truncations to keep $\kappa_t \in K$). The details of how the two dual variables are updated can remain internal to separate controllers; what must be shared with the bidding surface is exactly the composed scalar κ_t that encodes the *current marginal tradeoff* between value, budget scarcity, and RoS tightness. This is why the "one-scalar" design is not merely an engineering convenience: it is the minimal sufficient statistic for implementing the economically correct cutoff rule in truthful auctions, and it is the natural unit in which regret and feasibility can be stated and compared.

At the same time, we view the κ^* benchmark as an intentional abstraction with clear limitations. It leverages the threshold structure of truthful auctions and the proportional-bidding reduction; richer mechanisms (multi-slot, quality adjustments, non-truthful equilibria) may require higher-dimensional bidding policies, in which case one scalar may no longer be sufficient. Nonetheless, within the common production pattern where many services ultimately emit a single "effective bid" or multiplier, competing with κ^* and measuring $(\mathrm{Reg}_T, \mathrm{Viol}_B, \mathrm{Viol}_\tau)$ provides a concrete way to ask what modularity costs: if the economically relevant control variable is one-dimensional, then communication of that one dimension should be enough to recover centralized guarantees, and absence of it should be detectably harmful.

A minimally-coupled primal–dual control law. To operationalize the "one-scalar" idea, we implement the online policy as a distributed primal–dual method whose *only* cross-module artifact at auction time is the effective multiplier κ_t. The conceptual starting point is the Lagrangian relaxation of the two long-run constraints, with nonnegative dual variables α_B (budget scarcity) and α_τ (RoS tightness). Given duals (α_B, α_τ), the per-round contribution of winning auction t is weighted by the affine coefficient

$$(1 + \alpha_\tau)v_t - (\alpha_B + \tau\alpha_\tau)p_t.$$

In a truthful threshold-price auction, the allocation decision induced by bidding $b_t = \kappa v_t$ is exactly the indicator $\mathbf{1}\{\kappa v_t \geq p_t\}$. Hence, for fixed duals, the centralized primal step can be read as: choose a cutoff on p_t/v_t so as to win whenever the above coefficient is nonnegative. Rearranging yields the threshold condition

$$\frac{p_t}{v_t} \;\leq\; \frac{1 + \alpha_\tau}{\alpha_B + \tau\alpha_\tau},$$

which is implementable by a uniform multiplier

$$\kappa(\alpha_B, \alpha_\tau) \;=\; \Pi_K\left(\frac{1 + \alpha_\tau}{\alpha_B + \tau\alpha_\tau}\right), \tag{1}$$

with the understanding that if the denominator is nonpositive (an edge case early in learning), we set κ to the maximal aggressive bid $\kappa_{\max}$ before projecting onto K. Equation (1) is the key compositionality property: although the problem has *two* constraints, the auction-facing decision depends only on a *single* scalar ratio of dual variables.

Two controllers, one message. We implement (1) via two modules. The RoS controller $\mathcal{C}_\tau$ maintains an internal state intended to track $(\alpha_{B,t}, \alpha_{\tau,t})$ and computes the composite shadow price $\kappa_t = \kappa(\alpha_{B,t}, \alpha_{\tau,t})$. It then broadcasts the single real number

$$m_t \equiv \kappa_t$$

to the Budget controller $\mathcal{C}_B$. The Budget controller is deliberately simple: it converts the message into an auction bid $b_t = \kappa_t v_t$ so long as budget remains, and otherwise bids 0 (equivalently, sets $\kappa_t = 0$ locally) once cumulative spend reaches B. This hard stop enforces budget feasibility pathwise, independently of how the dual variables evolve:

$$\text{if } \sum_{u<t} s_u \geq B, \text{ then set } b_t = 0 \text{ and skip bidding thereafter.}$$

Crucially, no additional coordination is needed at bidding time: the only information that must traverse the module boundary is the scalar multiplier κ_t, which the budget layer can treat as an externally supplied "pacing/ROI knob."

Online dual updates from realized residuals. The remaining question is how $\mathcal{C}_\tau$ updates its internal dual variables using only realized feedback. A convenient choice is projected subgradient ascent on the dual, driven by per-round constraint residuals. Let $s_t = p_t x_t$ and $y_t = v_t x_t$ denote realized spend and realized value (both equal 0 when we lose). Then the instantaneous RoS slack is $\tau s_t - y_t$, while a natural budget pacing slack is $s_t - \rho_t$, where ρ_t is a per-round spend target. The simplest stationary choice is $\rho_t = B/T$, yielding the updates

$$\alpha_{B,t+1} = \Big[\alpha_{B,t} + \eta\,(s_t - B/T)\Big]_+, \tag{2}$$

$$\alpha_{\tau,t+1} = \Big[\alpha_{\tau,t} + \eta\,(\tau s_t - y_t)\Big]_+. \tag{3}$$

Here $[\cdot]_+$ denotes projection onto $\mathbb{R}_+$, and $\eta > 0$ is the step size (typically $\eta \propto 1/\sqrt{T}$). These updates are implementable with the prescribed information partition: they require only (s_t, y_t), which are observed ex post, and do not require observing the latent threshold price p_t when we lose. Moreover, they are interpretable: when realized efficiency falls short of the target (i.e., $\tau s_t > y_t$), the RoS dual $\alpha_{\tau,t}$ increases, which (via (1)) tightens the cutoff; when spending runs ahead of the pacing target, $\alpha_{B,t}$ increases, raising the effective "shadow cost of money" and likewise decreasing κ_t.

A practical refinement replaces the constant $\rho_t = B/T$ with an adaptive target based on remaining time and remaining budget, e.g.,

$$\rho_t \;=\; \frac{B - \sum_{u<t} s_u}{T - t + 1},$$

which improves transient behavior without changing the one-message nature of the protocol. With this choice, the same update (2) reacts to how quickly the advertiser is burning through budget relative to what is feasible to spread across the horizon.

Why this distributed protocol matches the centralized one. The salient architectural point is that, if a centralized algorithm maintains $(\alpha_{B,t}, \alpha_{\tau,t})$ and chooses bids through the composite map (1), then any decomposition that preserves the sequence $\{\kappa_t\}$ is behaviorally equivalent from the auction's perspective: allocations satisfy $x_t = \mathbf{1}\{\kappa_t v_t \geq p_t\}$ and the realized feedback (s_t, y_t) is identical. Thus, allowing $\mathcal{C}_\tau$ to own the dual variables while $\mathcal{C}_B$ owns the hard stop is not a conceptual compromise but an implementation of the same primal–dual recursion with a single shared sufficient statistic.

Guarantees and a constraint-tight variant. Under standard boundedness assumptions (e.g., v_t, p_t bounded and K compact), the primal–dual method with step size $\eta \propto 1/\sqrt{T}$ yields sublinear regret relative to the best fixed feasible multiplier and sublinear RoS violation; informally, the dual iterates cannot grow too quickly, and the average constraint residuals converge to feasibility. The hard budget stop ensures $\mathrm{Viol}_B = 0$ whenever spend is observed without delay and bidding can be halted instantaneously. If we seek tighter RoS control, we can strengthen the RoS update by adding an augmented penalty, for instance replacing (3) with an update on a smoothed potential that penalizes positive cumulative RoS deficit more aggressively. In effect, we trade off slightly more conservative bidding in early rounds for an $O(1)$ bound on Viol_τ (at the cost of constants in regret), while leaving the communication pattern unchanged: the bid still depends only on the single scalar κ_t.

Implementation takeaway. The minimally-coupled protocol can therefore be summarized as follows: the RoS side performs the economically meaningful accounting—tracking the shadow value of budget and the shadow tightness of RoS—and emits exactly one number, the composite multiplier κ_t; the budget side enforces the hard spend cap mechanically and otherwise acts as a thin execution layer. This division of labor mirrors production constraints (separate ownership of spend and ROI surfaces) while preserving the central theoretical insight: in truthful auctions with uniform-multiplier sufficiency, the only coordination needed at decision time is one real scalar that encodes the current marginal tradeoff across constraints.

Main sufficiency guarantee (regret–feasibility tradeoff). We now formalize the sense in which the minimally-coupled protocol—dual updates on $(\alpha_{B,t}, \alpha_{\tau,t})$ combined with the composite map $\kappa_t = \kappa(\alpha_{B,t}, \alpha_{\tau,t})$ and a mechanical budget stop—inherits the performance of a centralized primal–dual method. Throughout, we impose standard boundedness: $0 \leq v_t \leq \bar{v}$, $0 \leq p_t \leq \bar{p}$, and $K = [0, \kappa_{\max}]$. We also allow an explicit projection of the dual variables onto a compact set, $\alpha_{B,t} \in [0, A_B]$ and $\alpha_{\tau,t} \in [0, A_\tau]$, which is without loss for regret bounds (it only improves stability and simplifies constants).

Fix any horizon T. Run the distributed recursion with step size $\eta > 0$,

$$\alpha_{B,t+1} = \Pi_{[0,A_B]}\big(\alpha_{B,t} + \eta\,(s_t - \rho_t)\big), \qquad \alpha_{\tau,t+1} = \Pi_{[0,A_\tau]}\big(\alpha_{\tau,t} + \eta\,(\tau s_t - y_t)\big),$$

for some predictable pacing target ρ_t (e.g. $\rho_t = B/T$ or the remaining-budget rule), and bid using the single message $\kappa_t = \Pi_K\big((1+\alpha_{\tau,t})/(\alpha_{B,t} + \tau\alpha_{\tau,t})\big)$ as described above.

Theorem (sublinear regret and RoS violation). Assume the sequence $\{(v_t, p_t)\}_{t=1}^T$ is either i.i.d. or chosen by an oblivious adversary (independent of the learner's internal randomization, if any). Then there exist constants $C_{\text{reg}}, C_{\text{ros}}$ depending only on $(\bar{v}, \bar{p}, \tau, \kappa_{\max}, A_B, A_\tau)$ such that, for any feasible fixed multiplier $\kappa \in K$ satisfying the two constraints in expectation (or pathwise, in the oblivious-adversary model),

$$\mathbb{E}\left[\sum_{t=1}^T v_t x_t(\kappa) - \sum_{t=1}^T v_t x_t(\kappa_t)\right] \le \frac{C_{\text{reg}}}{\eta} + C_{\text{reg}}\eta T,$$

$$\mathbb{E}[(\tau S_T - V_T)_+] \le \frac{C_{\text{ros}}}{\eta} + C_{\text{ros}}\eta T.$$

In particular, choosing $\eta \asymp 1/\sqrt{T}$ yields

$$\mathbb{E}[\text{Reg}_T] = O(\sqrt{T}), \qquad \mathbb{E}[\text{Viol}_\tau] = O(\sqrt{T}).$$

Moreover, the budget layer enforces a hard spend cap in the usual platform sense: if the platform (or execution layer) stops serving once cumulative spend reaches B, then $\text{Viol}_B = 0$ pathwise; if instead we can only stop at the *next* round after observing spend, then $\text{Viol}_B \le \bar{p}$ (one-auction overshoot).

Discussion of constants and the role of η. The bounds above have the familiar "$\frac{1}{\eta}$-vs-ηT" shape from online mirror descent: the $\frac{1}{\eta}$ term is an initialization/diameter cost (how quickly the duals can move to the right scale), while the ηT term is an accumulated noise/variability cost (how much the algorithm jitters in response to per-round residuals). Economically, larger η makes the shadow values $\alpha_{B,t}$ and $\alpha_{\tau,t}$ react aggressively to recent overspending or RoS shortfalls, improving short-run feasibility but potentially inducing oscillations in κ_t; smaller η smooths bidding but slows the adjustment of the composite shadow price and raises regret. The recommended $\eta \propto 1/\sqrt{T}$ balances these forces and is the canonical choice when the horizon is known; standard doubling tricks recover the same rates when T is unknown.

Why "one message" is enough for the theorem. The theorem's proof is a bookkeeping exercise once we recognize that the auction outcome depends on $(\alpha_{B,t}, \alpha_{\tau,t})$ only through the scalar cutoff κ_t. A centralized primal–dual algorithm and the distributed implementation generate the same sequence $\{\kappa_t\}$ (hence the same x_t, s_t, y_t) as long as the RoS module can compute κ_t from its internal state and the budget module can enforce the stop. Thus, we analyze the *central* saddle-point recursion, but implement it with a single scalar broadcast; there is no additional approximation error introduced by modularization.

Constraint-tight variants. If we require tighter RoS feasibility than $O(\sqrt{T})$, we can replace the linear dual update for $\alpha_{\tau,t}$ with an augmented penalty that grows faster when the cumulative RoS deficit becomes positive (e.g. adding a quadratic penalty on the cumulative slack, or using an adaptive stepsize that increases when $(\tau S_t - V_t)_+$ grows). Under the same boundedness assumptions, such "augmented Lagrangian" style variants typically yield $\mathrm{Viol}_\tau = O(1)$ while preserving sublinear regret (often still $O(\sqrt{T})$ up to constants). Importantly, the architecture remains minimally coupled: the bid still depends only on the single number κ_t.

Extensions: partial feedback and missing prices. A practical advantage of the update rules is that they do *not* require observing the latent threshold price p_t when the advertiser loses. The RoS update uses only realized outcomes (s_t, y_t), which are typically observable from billing and conversion logging. Hence, partial observability of counterfactual prices does not affect implementability or the stated rates. If realized value is observed with noise—for instance, conversions arrive as an unbiased proxy $\hat{y}_t$ for y_t—the same analysis carries through in expectation, with an additional variance-dependent term of order $O(\sigma\sqrt{T})$ in regret/violation bounds (the standard stochastic-gradient phenomenon).

Extensions: delayed feedback and delayed messages. If spend/value feedback arrives with delay d, the dual updates become "stale" and effectively use residuals from $t - d$. The resulting performance degradation is graceful: the bounds above acquire an additive term scaling linearly in d (or, more precisely, proportional to d times a bound on the per-round residual magnitude), matching standard delayed online optimization guarantees. Similarly, if the one-scalar message κ_t is delivered with delay or corrupted by mean-zero noise, the regret and RoS-violation bounds worsen by at most $O(d)$ and $O(\sigma\sqrt{T})$, respectively, reflecting that the control problem remains one-dimensional at the decision boundary.

Takeaway. From a mechanism-design perspective, the content of the theorem is that truthful auctions collapse the bidder's action space to a single cutoff, so the "coordination rent" between budget pacing and ROI enforcement can be captured by one real number per round. With η tuned at the usual $\Theta(1/\sqrt{T})$ scale, this suffices to achieve centralized-grade guarantees while respecting the organizational constraint that budget and RoS live in separate modules.

A more explicit view of the regret–feasibility bounds. The stated inequalities are most transparently read as a bound on the *dualized* performance gap between our online sequence $\{\kappa_t\}$ and any fixed comparator $\kappa \in K$ that is feasible for the two long-run constraints. Concretely, define the per-round (random) residuals

$$r_t^B := s_t - \rho_t, \qquad r_t^\tau := \tau s_t - y_t,$$

and note that, under boundedness and the stop rule, they are uniformly bounded by constants depending only on $(\bar{v}, \bar{p}, \tau)$ and the choice of ρ_t (e.g. $|r_t^B| \le \bar{p} + \max_t \rho_t$ and $|r_t^\tau| \le \tau\bar{p} + \bar{v}$). The dual recursion is then simply projected online gradient ascent on the dual variables, driven by these residuals, while the primal action κ_t is the maximizer of the current Lagrangian surrogate (equivalently, a one-dimensional cutoff rule). Standard mirror-descent algebra yields a decomposition of the form

$$\sum_{t=1}^{T} \big(\ell_t(\kappa_t) - \ell_t(\kappa)\big) \;\le\; \underbrace{\frac{\|\alpha_1 - \alpha\|^2}{2\eta}}_{\text{diameter/initialization}} \;+\; \underbrace{\frac{\eta}{2}\sum_{t=1}^{T} \|g_t\|^2}_{\text{variance/instability}} ,$$

where ℓ_t is the (negative) per-round Lagrangian reward as a function of κ, g_t is the (bounded) gradient in the dual space induced by (r_t^B, r_t^τ), and α is any fixed dual vector. Translating this generic inequality back into our economic primitives gives exactly the advertised "$\frac{1}{\eta}$ versus ηT" bounds for regret and RoS violation, with constants proportional to $(A_B^2 + A_\tau^2)$ (dual diameter) and to $(\sup_t |r_t^B|^2 + \sup_t |r_t^\tau|^2)$ (gradient magnitude). This makes clear that the price of modular control is not statistical but purely parametric: once the bid is summarized by the single scalar κ_t, we are effectively solving a one-dimensional primal problem with a two-dimensional dual.

Choosing η and what the constants mean economically. Because the upper bounds scale like $\frac{1}{\eta} + \eta T$, the canonical choice $\eta \simeq c/\sqrt{T}$ yields $O(\sqrt{T})$ regret and $O(\sqrt{T})$ RoS shortfall, with the constant c shrinking when the residuals are small (stable environment) and growing when residuals are large (high variability in realized spend/value). Economically, this is exactly the expected sensitivity tradeoff: if the market is volatile, the shadow prices must react (larger η) to prevent persistent RoS drift, but reacting too quickly also makes the cutoff κ_t oscillatory and can waste learning opportunity on "wrong-side" adjustments. Conversely, in a stable market, smaller η is beneficial: it reduces churn in κ_t and makes the effective cutoff closer to a fixed κ^*, which is precisely the benchmark in our regret definition. When the horizon is unknown, the same rates follow from standard horizon-free schedules (e.g. $\eta_t \propto 1/\sqrt{t}$) or doubling; the key point is that none of these require cross-module communication beyond broadcasting the current κ_t.

Budget feasibility as a hard constraint (and the overshoot issue). The budget layer deserves a separate comment because it is not merely a Lagrangian penalty in practice: platforms typically implement budgets by suppressing delivery when remaining budget is insufficient. Under such a "mechanical stop," the algorithm's *pathwise* spend is capped at B up to execution granularity. In the idealized model in which we can stop immediately upon reaching B, we obtain $\mathrm{Viol}_B = 0$ deterministically. If instead we only learn s_t at the end of the round (the common logging convention), the worst-case overshoot is at most one payment, so $\mathrm{Viol}_B \leq \bar{p}$. This is not a failure of the primal–dual method; it is simply the discrete-time nature of auctions. Importantly, the stop rule is entirely local to the budget module and requires no additional information from the RoS module besides the common bid multiplier already being broadcast.

Partial feedback: why missing counterfactual prices do not break implementability. A frequent operational concern is that the bidder does not observe the threshold price p_t when it loses (and sometimes even when it wins, one only sees an invoiced amount with delays and adjustments). The minimally-coupled recursion is robust to this informational friction because the updates are written in terms of *realized* spend and realized value, (s_t, y_t), which are exactly the quantities recorded by billing and conversion measurement systems. When the advertiser loses, both s_t and y_t are zero, so the residuals (r_t^B, r_t^τ) are still well-defined and require no imputation of the counterfactual p_t. Conceptually, we are not estimating demand curves or price distributions; we are controlling aggregate feasibility via dual variables, and the stochastic approximation uses only realized constraint slack. If realized value is noisy—for instance, we observe $\hat{y}_t = y_t + \xi_t$ with $\mathbb{E}[\xi_t \mid \mathcal{F}_{t-1}] = 0$ and $\mathbb{E}[\xi_t^2] \leq \sigma^2$—then the same mirror-descent proof goes through with an added martingale-noise term, producing an additive $O(\sigma\sqrt{T})$ degradation in both regret and RoS-violation bounds. This is the standard stochastic-gradient phenomenon and does not alter the one-message nature of the protocol.

Delayed feedback and delayed messages. Delays matter in two distinct places: (i) the dual updates may use residuals from earlier rounds, and (ii) the scalar message itself may be delivered late, so the budget module executes κ_{t-d} rather than κ_t. In either case, the control is "stale" but still one-dimensional. Under bounded residuals, standard delayed online-optimization results imply an additive performance loss that scales at most linearly with the delay. A representative bound (suppressing constants) takes the form

$$\mathbb{E}[\mathrm{Reg}_T] \;\le\; O\!\left(\frac{1}{\eta} + \eta T\right) + O(d), \qquad \mathbb{E}[\mathrm{Viol}_\tau] \;\le\; O\!\left(\frac{1}{\eta} + \eta T\right) + O(d),$$

and if the delivered κ is further corrupted by mean-zero noise with variance σ^2, the usual $\sigma\sqrt{T}$ term appears additively. The economic interpretation is straightforward: delays temporarily decouple the shadow price from current feasibility, so we may briefly overshoot or undershoot RoS, but the system recenters once updates catch up. Crucially, there is no qualitative breakdown because the relevant decision boundary is still a scalar cutoff; we are not attempting to coordinate a high-dimensional action across modules.

What the sufficiency result does *not* claim. Finally, we emphasize a limitation that clarifies the scope of the theorem. The regret benchmark is the best *fixed* multiplier κ^* in hindsight (subject to feasibility), so the $O(\sqrt{T})$ rate should be read as: we compete well with the best static cutoff rule, not with an omniscient dynamic policy that can condition on future scarcity of high-v_t/p_t opportunities. In many advertising environments, this is exactly the right benchmark because the actionable degree of freedom is a stationary pacing factor; nevertheless, in highly non-stationary markets one would want tracking guarantees against a drifting comparator, which would require additional variation terms in the bounds. The point that survives these refinements is the architectural one: regardless of which comparator class we choose, truthful auctions reduce the real-time decision to a one-dimensional cutoff, and hence a single scalar message is sufficient to implement the centralized update rule without loss.

Why "no message" is genuinely restrictive. To make the necessity claim precise, it is useful to pin down what we mean by a *fully decoupled* pacing architecture. Fix a feasible multiplier set $K \subset \mathbb{R}_+$ and a *fixed* monotone aggregator $h : K \times K \to K$ (nondecreasing in each coordinate; canonical examples are min, product, and convex combinations). In each round t, the Budget controller outputs

$$\kappa_t^B \;=\; f_t\big(v_t;\, s_1, \ldots, s_{t-1}\big),$$

while the RoS controller outputs

$$\kappa_t^\tau \;=\; g_t\big(v_t;\, (y_1, s_1), \ldots, (y_{t-1}, s_{t-1})\big),$$

possibly using internal randomization, and the executed multiplier is

$$\kappa_t \;=\; h\big(\kappa_t^B, \kappa_t^\tau\big), \qquad b_t = \kappa_t v_t.$$

The defining constraint is that neither controller may condition on the other controller's internal signal (in particular on κ_t^B or κ_t^τ), and there is no cross-module message beyond what is mechanically revealed through the realized outcome (x_t, s_t, y_t).

The economic content of this restriction is subtle: *each* module is reacting to a long-run constraint, but the constraints are coupled through the same binary allocation decision $x_t = \mathbf{1}\{\kappa_t v_t \geq p_t\}$. In a centralized primal–dual method, this coupling is resolved by forming a single composite "shadow price" that trades off the marginal value of relaxing budget against the marginal value of relaxing RoS, and *then* applying a one-dimensional cutoff on p_t/v_t. In a fully decoupled architecture, by contrast, the cutoff is a deterministic function of two independently-evolved signals whose interaction is frozen ex ante by h.

A two-type adversary that forces linear loss. We now describe an adversarial construction (oblivious or adaptively chosen against a fixed deterministic algorithm; randomized algorithms can be handled by Yao's principle) showing that, for any such decoupled design, there exists a bounded sequence $\{(v_t, p_t)\}_{t=1}^T$ for which one *must* pay $\Omega(T)$ either in RoS shortfall or in regret relative to the best fixed feasible multiplier κ^*. The simplest intuition uses two impression types that alternately make *one* constraint locally informative and the *other* constraint locally misleading.

Fix $\tau > 0$ and choose constants (v^+, p^+) and (v^-, p^-) satisfying

$$v^+ - \tau p^+ \;=\; +\Delta, \qquad v^- - \tau p^- \;=\; -\Delta,$$

for some $\Delta > 0$, with both pairs bounded (so $\bar{v}, \bar{p}$ are finite). Type "+" impressions generate *positive* RoS slack when won; type "−" impressions generate *negative* RoS slack of equal magnitude. We additionally set the price/value ratios so that winning "−" impressions requires a strictly larger multiplier:

$$\frac{p^+}{v^+} \;<\; \frac{p^-}{v^-}.$$

Thus there exists an intermediate cutoff κ^{mid} that wins all "+" impressions and loses all "−" impressions, while a larger cutoff κ^{high} wins both.

Now arrange the horizon into alternating blocks. In *RoS-stress* blocks, we present only type "−" impressions. In *budget-stress* blocks, we present only type "+" impressions but with prices scaled so that spend is large relative to the remaining budget (so budget, not RoS, should be treated as binding). The offline benchmark can be made clean by selecting block lengths so that the globally optimal fixed multiplier is $\kappa^* = \kappa^{\text{high}}$: it wins both types, accumulates approximately zero net RoS slack over the full horizon (because the number of "+" and "−" wins balance), and respects the budget by construction (we choose B to match the spend induced by κ^{high}). Under κ^*, the advertiser collects $\Theta(T)$ value.

Why decoupling breaks: the "which constraint is binding?" ambiguity. Consider what a fully decoupled protocol must do on such a sequence. In a RoS-stress block, any win produces $r_t^\tau = \tau s_t - y_t = +\Delta$, pushing the RoS controller toward a more conservative κ^τ. But whether we *should* actually be conservative depends on the global plan: if ample "+" slack will arrive later and sufficient budget remains to buy it, then taking some "−" impressions now is efficient in the aggregate; if not, then those wins are truly infeasible. The centralized algorithm resolves this by a single scalar shadow price that already internalizes both the accumulated RoS deficit and the opportunity cost of budget.

In a decoupled protocol, however, the Budget controller cannot condition on the RoS deficit at all: its signal κ_t^B is a function of past spend, and in RoS-stress blocks the adversary can make spend uninformative (e.g. by choosing p^- small so that the budget residual is near zero even while RoS deteriorates). Symmetrically, in budget-stress blocks, the adversary can make RoS residuals uninformative (choose "+" impressions with very large $v^+ - \tau p^+$ so RoS appears extremely safe) while driving spend rapidly toward B, so that only the budget module "sees the fire." The result is that in alternating blocks the two modules are systematically *out of phase*: one is tightening while the other is loosening.

Because the executed action is $\kappa_t = h(\kappa_t^B, \kappa_t^\tau)$ with fixed monotone h, this out-of-phase behavior cannot be neutralized without direct coupling. If h is conservative in the sense of being bottlenecked by the smaller coordinate (as with min or product on $[0, 1]$), then the RoS module's tightening during RoS-stress blocks drives κ_t below κ^{high} for a linear number of rounds, so the bidder systematically fails to win the "−" impressions that κ^* would have won, implying $\text{Reg}_T = \Omega(T)$. If, instead, h is aggressive (e.g. closer to max or a large weight on κ_t^B), then in RoS-stress blocks the budget module's loosening keeps κ_t high and the bidder continues to win "−" impressions, accumulating $\Omega(T)$ RoS deficit before sufficient "+" slack can be realized, implying $\text{Viol}_\tau = \Omega(T)$ (and, depending on parameters, potentially exhausting budget in the wrong blocks as well).

Lower bound as an information bottleneck. The deeper interpretation is not that any particular update rule is flawed, but that *the needed statistic is joint.* The correct cutoff at time t depends on a composite shadow price—a single number encoding how costly additional spend is *given* the current RoS slack (and vice versa). In a fully decoupled design, the budget module and the RoS module each maintain only a partial view and then combine those views through a fixed aggregator that cannot represent the joint Lagrange multiplier except in special cases. An adversary can exploit this by alternating segments where one constraint should dominate the composite price and segments where the other should dominate. Hence, without transmitting at least one real scalar that aligns the two modules on a common shadow price, no architecture in this class can guarantee simultaneously sublinear regret and sublinear constraint violations: for some bounded auction stream, it must incur $\Omega(T)$ loss in feasibility or in value.

Simulation design and what we want to learn. The theoretical results above isolate a one-dimensional decision statistic (an effective multiplier κ_t) and argue that (i) a single scalar of coupling is sufficient to implement the same sequence of bids as a centralized primal–dual method, while (ii) removing that scalar can create linear losses under adversarial inputs. We complement those statements with simulations for two reasons. First, in practice auction streams are neither worst-case nor perfectly stationary, so we want to understand whether the minimally-coupled protocol is merely a proof device or a quantitatively reasonable engineering choice. Second, the impossibility result for fully decoupled designs is existential; simulations help illustrate *how* the failure mode manifests (persistent RoS drift versus chronic underbidding), and how sensitive it is to details like non-stationarity and feedback delays.

Synthetic truthful-auction environment. We generate sequences $\{(v_t, p_t)\}_{t=1}^T$ with T between 10^4 and 10^6. Values v_t are drawn from a bounded distribution (e.g. $v_t \sim \mathrm{Unif}[0, 1]$ or truncated lognormal), and threshold prices p_t are generated to mimic competition and reserve effects: we use $p_t = \rho v_t + \epsilon_t$ with $\rho \in (0, 2)$ and ϵ_t mean-zero noise truncated to keep $p_t \geq 0$, as well as a misspecified case in which p_t is independent of v_t. The bidder observes v_t before bidding, submits $b_t = \kappa_t v_t$, wins if $b_t \geq p_t$, and pays p_t upon winning. We set a budget $B = \beta\,\mathbb{E}[p_t]T$ with $\beta \in (0, 1)$ so that the budget binds but does not trivially stop the bidder immediately, and we choose an RoS target τ so that the feasibility region is nonempty but meaningful (e.g. τ set to a quantile of v_t/p_t). Performance is reported via realized value V_T, spend S_T, regret proxy against the best fixed multiplier in hindsight, and RoS violation $\mathrm{Viol}_\tau = (\tau S_T - V_T)_+$.

Architectures compared. We compare three bidding stacks that share the same base action space $K = [0, \kappa_{\max}]$ and the same step size η. (i) *Fully coupled*: a centralized controller runs a primal–dual update with dual variables for budget and RoS and outputs κ_t directly. Budget is enforced by halting bids once cumulative spend reaches B. (ii) *Minimally coupled*: the RoS-side computation maintains the dual state and broadcasts a single scalar message $m_t = \kappa_t$ each round; the budget module simply applies that κ_t while enforcing the hard stop at B. By construction, when the message is delivered instantly and without noise, this reproduces the centralized κ_t trajectory. (iii) *Fully decoupled*: a budget-only module outputs κ_t^B from spend history and an RoS-only module outputs κ_t^τ from RoS residuals; the executed multiplier is $\kappa_t = h(\kappa_t^B, \kappa_t^\tau)$ for fixed monotone h (we report min, product on $[0, \kappa_{\max}]$, and a convex combination). Each module uses the same style of update (mirror/gradient) but without access to the other module's state.

Baseline findings: one scalar is enough in benign regimes. Across i.i.d. environments, the fully coupled and minimally coupled systems are empirically indistinguishable up to Monte Carlo error: they spend at essentially the same rate until the hard budget stop, achieve the same V_T, and exhibit similar $O(\sqrt{T})$ scaling of RoS shortfall when we use a standard $\eta \propto 1/\sqrt{T}$. This is a useful sanity check: the point of the distributed design is not to improve on the centralized benchmark, but to match it under a stringent communication constraint. The decoupled designs, by contrast, are highly sensitive to the aggregator. With $h = \min$, the system tends to be overly conservative whenever either module becomes cautious, producing systematically lower win rates and value; with more aggressive aggregators, it tends to chase value in a way that appears locally RoS-safe (because high-v_t wins look good) but can accumulate a persistent global RoS deficit when the price environment shifts.

Non-stationarity and tracking. To stress adaptation, we consider piecewise-stationary streams in which (v_t, p_t) distribution parameters drift every L rounds (e.g. ρ increases, making auctions more expensive relative to value), as well as smoothly drifting processes. In these settings, the coupled and minimally coupled controllers degrade gracefully: the multiplier κ_t adjusts in the correct direction (down when prices rise relative to value; up when the environment becomes more favorable), and constraint violations scale with the variation budget of the sequence, consistent with standard online tracking intuition. The decoupled stacks often exhibit a lag mismatch: the budget module reacts quickly to increased spend (tightening), while the RoS module remains optimistic if realized wins still have high v_t, so the aggregator either (a) over-tightens (value loss) or (b) remains too loose (RoS drift), depending on h. Importantly, this happens even when the stream is not adversarial in any strong sense; mild regime changes suffice to expose the absence of a shared composite shadow price.

Feedback delays and message corruption. We next impose a delay d on the minimally coupled message, implementing bids using $\kappa_t = \kappa(m_{t-d})$ with the natural convention for $t \le d$. We also consider additive mean-zero noise $\tilde{m}_t = m_t + \xi_t$ with $\mathbb{E}[\xi_t^2] = \sigma^2$ and projection back to K. In both cases we observe a smooth deterioration: delays introduce oscillations in spend pacing (the controller tightens after the environment has already changed), while noise introduces jitter that slightly increases both regret and RoS violation. Quantitatively, the incremental loss appears roughly linear in d and proportional to $\sigma\sqrt{T}$ over the ranges we tested, aligning with the stability bounds one would expect from delayed/noisy one-dimensional online updates. Practically, this suggests that the single-scalar interface is not brittle: moderate latency or quantization does not destroy performance, though it does reduce the effective aggressiveness with which the system can track fast-moving conditions.

Takeaway from simulations. The simulations reinforce a simple engineering interpretation of the theory. When the auction is truthful and the bidder's action is well approximated by a uniform multiplier, the *only* statistic that needs to cross module boundaries at bid time is the effective cutoff κ_t (or an equivalent scalar shadow price). Implementing this scalar makes the distributed system behave like the centralized controller, including under moderate non-stationarity and delayed feedback. Conversely, eliminating even this minimal coupling forces the system to combine partial views through a fixed aggregator, and the resulting behavior can be systematically miscalibrated in realistic streams, not only in worst-case constructions.

Extensions: when prices are endogenous (many advertisers). Our model treats $\{p_t\}$ as exogenous thresholds, which is appropriate for understanding an individual bidder's pacing logic when its market share is small or when competition is sufficiently diffuse. With multiple strategic advertisers, however, each bidder's κ_t affects allocation and therefore the realized price process faced by others. In that regime the interpretation of the dual variables shifts: $\alpha_{B,t}$ and $\alpha_{\tau,t}$ remain meaningful as *private* shadow prices of constraints, but the mapping from those shadows to realized spend and value is mediated by equilibrium responses. Two implications follow. First, regret-style guarantees become equilibrium-sensitive: an online algorithm may be no-regret relative to a best fixed κ under a counterfactual *fixed* environment, yet still induce cycling if all agents simultaneously adapt. Second, the one-scalar interface can remain a useful *control* abstraction even when it is no longer an exact sufficient statistic: if each advertiser's feasible action is well approximated by a scalar multiplier, then market interaction reduces to a low-dimensional dynamical system in the multipliers, which can often be stabilized with damping (smaller η, averaging, or explicit inertia). Practically, this is the regime where platform-level safeguards (throttling, bid caps, pacing floors) and bidder-side smoothing jointly matter; our analysis informs the bidder-side compression, not the existence of a unique global optimum.

Multi-slot auctions and richer allocation rules. Single-slot truthfulness yields the clean cutoff $\mathbf{1}\{\kappa v_t \geq p_t\}$, which is the technical source of the one-dimensional reduction. With multiple slots and position effects, the allocation becomes rank-based and the payment depends on the next highest bid. Even in truthful multi-unit mechanisms (e.g. VCG in certain forms), the bidder's decision is not always characterized by a single threshold on p_t/v_t; one must account for the marginal value of moving from position j to $j-1$. That said, in many ad systems the bidder chooses a single bid that is then interpreted through an auction-specific scoring rule, and a uniform multiplier remains the dominant degree of freedom for budget pacing. A useful way to reconcile these facts is to view κ_t as selecting an *expected* tradeoff curve between spend and value induced by the auction: the multiplier picks a point on that curve, while the auction maps that point into a distribution over positions and prices. When that induced curve is sufficiently smooth and stable, the minimally-coupled architecture still captures the essential coordination problem (budget versus RoS) with one scalar. When the curve has kinks (due to discrete position jumps) or heavy-tailed price responses, one should expect larger variance and slower convergence, and may need either conservative step sizes or additional state (e.g. per-position multipliers) beyond the one-scalar principle.

Non-truthful auctions: limits of the one-scalar principle. First-price auctions, non-truthful generalized second price variants, and mechanisms with complex discounts fundamentally change the role of κ. In a first-price auction, the bid is not a truthful report, so a value-multiplier $b_t = \kappa_t v_t$ simultaneously controls *allocation* and *shading.* The optimal shading depends on beliefs about competition, which can vary across impressions and over time; therefore, even offline optimal behavior need not be implementable by a single κ applied uniformly to v_t. Conceptually, the sufficient statistic becomes a function $b_t = \beta_t(v_t, \text{features})$, and compressing that function into one scalar inevitably creates approximation error. Our view is that the one-scalar architecture remains defensible only to the extent that shading can itself be parameterized by a low-dimensional family (e.g. $b_t = \gamma_t \cdot v_t$ with a scalar γ_t that absorbs both pacing and shading, or $b_t = \gamma_t \cdot \hat{v}_t$ where $\hat{v}_t$ is already a calibrated "bid value" incorporating auction competitiveness). When competitiveness varies materially across contexts, a single global multiplier can be too blunt: the right remedy is not more communication between budget and RoS modules per se, but a richer action space K (e.g. segment-wise multipliers) so that the scalar message applies within a segment rather than across the entire stream.

Beyond one constraint pair: what scales and what does not. An important limitation of the minimal-coupling claim is that it leverages a collapse to one decision dimension. If the advertiser faces multiple simultaneous constraints—say separate budgets across campaigns, frequency caps, or multiple RoS targets across objectives—then the natural centralized primal–dual method maintains a vector of dual variables, and there is no general reason for those duals to compose into a single cutoff that is optimal for all impressions. In such settings, we should expect a corresponding increase in the necessary communication: the minimally sufficient message is typically the dimension of the effective decision statistic (often the number of binding constraints after accounting for the structure of the allocation rule). From an engineering standpoint, this suggests a design heuristic: invest in finding the *smallest* decision representation that preserves the intended policy class (e.g. one multiplier per campaign, or one multiplier per major segment), and then communicate only that representation across modules. The benefit is not merely bandwidth reduction; it is organizational clarity about what information must be consistent at bid time.

Engineering guidance: implementing the scalar interface robustly. The simulations suggest that the scalar channel is tolerant to delays and noise, but implementation details matter. We have found three practical guardrails to be especially important. First, treat κ_t as a *control signal* and enforce stability explicitly: project onto K, limit per-round changes $|\kappa_t - \kappa_{t-1}|$, and optionally smooth via an exponential moving average. Second, separate *hard* enforcement (budget stop) from *soft* enforcement (RoS), mirroring the theory: when budget is binding, hard stopping is simpler and safer than relying on a dual to avoid overspend. Third, calibrate the step size η and any penalty augmentation to the time scale of non-stationarity: if the environment drifts on a horizon L, then η tuned for $\sqrt{T}$ regret may be too sluggish, and one should instead tune for tracking (larger η plus damping) while accepting slightly higher variance. None of these modifications alter the one-scalar communication pattern; they merely make the scalar behave like a reliable pacing knob under realistic latency and measurement error.

Takeaways. The broader message is that the "one-scalar" result is best understood as a statement about *where* coupling is essential, not that all bidding problems are one-dimensional. Truthful single-slot auctions plus a uniform-multiplier policy class place the bidder in a particularly clean regime where budget and RoS interact only through a composite cutoff, making one scalar sufficient and, in a precise sense, necessary. As we move to endogenous prices, multi-slot effects, or non-truthful mechanisms, the same architectural idea survives as a design pattern: identify the smallest statistic that determines auction-time actions, and ensure that statistic is shared across constraint controllers. When that statistic ceases to be one-dimensional, the right response is to expand the statistic, not to revert to fully decoupled control. This perspective also clarifies the role of our negative result: it is not a claim that modularity fails, but that modularity without a shared decision statistic can be systematically miscalibrated.

Conclusion. We can now summarize the economic message of the paper in the most operational terms: in the repeated truthful single-slot environment, the advertiser's bid-time decision is effectively one-dimensional, and therefore the *right* form of modularity is not to isolate budget and RoS (ROI) logic completely, but to make them consistent through a shared scalar pacing signal. The substantive content is not that budgets and RoS are easy—they are coupled long-run constraints, and in general coupled constraints are difficult to manage online—but that, under the allocation rule $x_t = \mathbf{1}\{\kappa_t v_t \geq p_t\}$, the coupling can be *compressed* into a single statistic without sacrificing the classical learning guarantees that one would obtain from a fully centralized primal–dual controller.

On the positive side, the sufficiency claim says that if we restrict attention to the uniform-multiplier class $b_t = \kappa v_t$, then there exists a scalar κ (or cutoff on p_t/v_t) that implements an optimal offline policy, and there exists an online update that tracks the best fixed feasible κ^* with sublinear regret while keeping long-run RoS violations sublinear (and budget enforceable exactly via stopping). This is conceptually important: it reconciles two facts that practitioners often hold in tension. First, pacing is commonly implemented as a multiplicative adjustment to a predicted value. Second, feasibility involves *two* constraints that appear to require *two* prices (one for budget, one for RoS). Our reduction clarifies why these are not contradictory in the truthful single-slot case: the two shadow prices can be composed into a single effective multiplier that determines the allocation decision, so the auction interface need not expose more than one real number per round.

On the negative side, the necessity claim disciplines a popular engineering instinct: "let the budget module pace down when spend is high; let the ROI module pace down when ROI is low; combine the two multiplicatively or by taking the minimum." Without an explicit coupling channel, such a design can fail in a strong sense. The underlying economic reason is informational rather than computational. Each module sees only its own residual and therefore cannot infer whether tightening is socially efficient given the other constraint's status; an adversary (or simply a non-stationary environment) can present sequences where the correct action alternates between being budget-driven and RoS-driven. In such cases, fixed monotone aggregators $h(\kappa_t^B, \kappa_t^\tau)$ induce systematic underreaction or overreaction, producing either linear regret relative to κ^* or linear cumulative constraint violation. The upshot is that modularity per se is not the problem; modularity without a shared decision statistic is.

The broader design principle we take from these results is a communication–control equivalence: the minimal cross-module message should have the same dimension as the effective action statistic that determines allocation in the auction. In our baseline model that statistic is one scalar κ_t, hence one scalar of coupling is both sufficient and (up to constants) necessary. This framing is useful because it separates two levers that are often conflated. If a single global multiplier is too blunt for performance reasons (heterogeneous competitiveness, context-dependent shading, multiple objectives), the remedy is primarily to enrich the action space K (e.g. segment-wise multipliers) rather than to keep K one-dimensional but attempt to "coordinate harder" through more elaborate internal heuristics. Conversely, if the action space truly is one-dimensional, then insisting on fully decoupled control is not a virtue; it is an avoidable information constraint.

From a policy and platform perspective, the one-scalar interface suggests a concrete way to align bidder-side autonomy with system-side stability. Platforms often prefer bidders to implement pacing and ROI controls in a manner that is predictable, smooth, and easy to audit for compliance with spend limits and efficiency constraints. A scalar pacing knob naturally supports such governance: it can be rate-limited, bounded to K, and logged; and it interacts with platform safeguards (caps, floors, throttles) in a transparent way. At the same time, our analysis highlights a limitation of purely bidder-side solutions in thick markets: when many advertisers adapt simultaneously, the mapping from κ_t to outcomes becomes equilibrium-dependent, so no-regret properties relative to a fixed environment do not automatically translate into stable joint dynamics. This does not negate the value of the scalar interface; rather, it clarifies what it can and cannot guarantee without additional damping, coordination mechanisms, or platform-level stabilization.

We also want to be explicit about where the theory is tight and where it is an approximation. The truthful single-slot threshold structure is what makes the cutoff rule exact, and thus what makes one-dimensional control exact. Multi-slot rank-based allocations, first-price shading, and context-dependent competition all weaken the reduction: they can make the optimal policy genuinely multidimensional, and then one scalar cannot be information-theoretically sufficient. Likewise, when the advertiser faces multiple budgets, multiple RoS targets, or other non-separable constraints, the composite statistic typically becomes a vector. The main lesson persists, but with a different dimension: communicate the minimal vector that actually determines bid-time actions, and do not expect zero-communication modularity to succeed when constraints are coupled through allocation.

Finally, the paper is also a statement about organizational design. In many real systems, budget control and ROI control are owned by different teams, rely on different data pipelines, and operate under different latency constraints. Our results provide a disciplined contract between them: the shared object must be the scalar (or low-dimensional) decision statistic that pins down auction-time behavior, not a collection of loosely aligned heuristics. When that contract is respected, decentralization can preserve the performance of a centralized primal–dual controller; when it is not, decentralization can be predictably fragile. The model thus illuminates a tradeoff that is as practical as it is theoretical: we can have modularity with strong guarantees, but only if we pay the small price of communicating the right statistic.