Platform Competition and Information Design in the Autobidding Era: Why Precision Can Reduce Revenue

Liz Lemma   Future Detective

January 16, 2026

Abstract

Autobidding has shifted online advertising from manual bids to target-based delegation, making advertiser behavior depend on platform-provided prediction signals (click/conversion/value estimates). Recent theory studies equilibria and efficiency of auction formats under RoS and budget constraints, and separately studies platform competition when advertisers allocate across channels. This paper connects these threads by introducing an explicit information-design lever: platforms choose not only auction format (first- vs second-price) and reserves, but also the precision of the value signals exposed through prediction APIs and reporting. We build a tractable two-platform model with RoS-constrained value-maximizing advertisers who allocate spend based on posterior net value. More informative signals are modeled via Blackwell ordering and shown to increase cross-platform elasticity—advertisers respond more strongly to differences in effective prices/auction formats. Our main characterization delivers (i) a threshold rule in sensitivity, competition intensity, and within-platform monetization under which second-price mechanisms dominate first-price mechanisms for revenue, and (ii) an information-garbling result: when precision substantially increases advertiser responsiveness, platforms optimally coarsen signals to soften competition, despite precision improving match quality. We discuss implications for auction format choice, prediction transparency, and antitrust/policy debates about strategic opacity in digital ad markets, and validate comparative statics in simulation on synthetic multi-channel auction environments.

Table of Contents

1. 1. Introduction: autobidding, platform-controlled ML signals, and why information policy is now a strategic lever in auction competition.
2. 2. Related literature and positioning: autobidding equilibrium/PoA; ML advice (reserves/boosts); platform competition in autobidding; information design/garbling in industrial organization.

3. 3. Model: two platforms, RoS-constrained advertisers, auction formats as price instruments, and Blackwell-ordered signal precision.

4. 4. Advertiser allocation rule and elasticity: posterior net value, RoS participation cutoff, and how higher precision increases cross-platform sensitivity.

5. 5. Format competition holding precision fixed: closed-form threshold where SPA dominates FPA (recovering and extending prior competition results).

6. 6. Joint choice of reserves and precision: equilibrium characterization; coarsening (garbling) incentives; welfare and revenue comparison to first-best precision.

7. 7. Extensions (clean but optional): heterogeneous advertisers/targets; partial commitment (platform can adjust reserves); more than two platforms K; adding budgets; regulator-imposed minimum precision or transparency cost.

8. 8. Numerical validation: synthetic multi-channel auction simulations; mapping reduced-form parameters to auction primitives; robustness checks.

9. 9. Policy and product implications: transparency mandates, interoperability, prediction API design, and competitive effects of signal coarsening.

1 Introduction

A growing share of display and search advertising is no longer mediated by human bidders choosing keywords, bids, and budgets in real time. Instead, advertisers increasingly delegate these decisions to platform-provided “autobidders” that optimize toward high-level objectives such as return on spend (RoS), cost per acquisition, or value-based conversion goals. In practice, the autobidder is only as good as the information it receives: the platform supplies machine-learning (ML) predictions, diagnostics, and targeting signals that shape the advertiser’s inferred value of an impression and, ultimately, the advertiser’s willingness to pay. As a result, what used to be an engineering choice about prediction accuracy has become an economic choice about how informative, granular, and comparable the platform makes its signals to the demand side.

This paper studies the strategic role of *information policy* in platform competition when demand is governed by autobidding under RoS-type constraints. Our starting point is a simple observation from marketplace practice. The dominant frictions are no longer limited to auction rules or reserve prices; rather, they include the platform’s control over the *measurement and prediction layer* that determines how advertisers perceive value and how readily they can compare alternatives across platforms. Platform signals include predicted click-through and conversion rates, value estimates, attribution and incrementality models, and sometimes opaque “quality” adjustments that can be interpreted as part of the information environment faced by advertisers. When signals are more informative, advertisers can target and allocate budgets more effectively within a platform. At the same time, more informative signals can make advertisers more confident about *relative* performance across platforms, thereby increasing their propensity to reallocate spend in response to small differences in fees, auction formats, or policy changes. This double-edged nature of information is the central tradeoff we formalize.

We build intuition before formalism. Consider two platforms offering impressions that are close substitutes for a typical performance advertiser. Each platform would like to convince advertisers that it delivers high value per dollar, but it also wants to extract revenue per impression. If advertisers can only imperfectly infer their realized value from a platform’s reporting and ML predictions, their cross-platform reallocation is sluggish: a platform can raise effective prices with limited immediate share loss. In contrast, when reporting and prediction are highly precise, advertisers can identify where their marginal dollar performs better and rapidly shift budgets. Precision therefore acts like a *competition amplifier*. This suggests an incentive for platforms to limit precision (or comparability) even when they possess the technical capability to improve it. The resulting logic resembles classic arguments about obfuscation and information disclosure, but the mechanism here is tailored to autobidding: information changes the *elasticity* of platform demand, not merely the level of willingness to pay.

The institutional details of modern ad markets make this mechanism particularly salient. First, platform auctions are repeated, fast, and algorithmic, so small changes in predictions propagate immediately into bidding behavior. Second, platform-side tools are typically bundled: the same system that provides targeting and value predictions also determines pacing, optimization, and sometimes eligibility. Third, advertisers often face explicit or implicit RoS constraints (e.g., a campaign must meet a target return over a rolling window), which induces threshold behavior: if the inferred net value of buying impressions on a platform falls below a profitability cutoff, spend is redirected elsewhere. Under these conditions, increased signal quality can make the demand response to price-like instruments steep. Thus, the platform's information policy is naturally strategic in a competitive environment.

Our model captures these ideas in reduced form while remaining close to the economic objects of interest. Each platform chooses an auction format (first-price versus second-price, represented as an expected price increment), a reserve/price instrument, and a signal precision parameter. Advertisers observe platform-provided signals and allocate spend shares across platforms via an objective consistent with RoS-constrained value maximization. The key modeling move is to link signal precision to cross-platform sensitivity: more informative signals reduce decision noise, so the allocation share becomes more responsive to differences in net value across platforms. We summarize this channel by a sensitivity parameter $\beta(\sigma)$ that increases with precision. Separately, precision may directly raise the value generated within a platform through better matching, improved optimization, or tighter within-platform bidding; we allow for such direct benefits as well. This separation clarifies what is at stake: platforms may want to improve ML accuracy for productive reasons, yet they may simultaneously fear the intensified competition that accuracy enables.

Two sets of results motivate the analysis. First, we show that the choice between first-price and second-price formats interacts sharply with advertiser sensitivity. A first-price format can mechanically raise expected payments per impression (captured by an increment $\Delta > 0$ holding other instruments fixed). But higher payments also lower advertiser net value, which reduces demand share when advertisers can easily substitute. When sensitivity is sufficiently high, the share loss dominates the per-unit gain, making the second-price format revenue-dominant against a rival second-price platform. The resulting condition yields a closed-form threshold β^*, highlighting a basic point: even if first-price auctions extract more per impression in a partial-equilibrium sense, they can be unprofitable in equilibrium when demand reallocates elastically across platforms. In this way, the equilibrium viability of "high-price" formats depends on the informational environment that determines how quickly advertisers react.

Second, and more centrally, we show that platforms may optimally choose *less* informative signals than would be socially efficient. The reason is not that information is intrinsically harmful; rather, information tightens competitive constraints by raising $\beta(\sigma)$. From a platform's perspective, increasing precision has a direct return (better within-platform monetization and/or higher match value) but also a strategic cost (tougher competition, lower markups). When the elasticity effect of precision is strong relative to the direct return, the platform's best response involves strategic coarsening, i.e., choosing a lower σ than the welfare benchmark. In symmetric equilibrium, this can produce systematically opaque information environments even when both platforms have access to high-quality prediction technology. This result speaks directly to contemporary debates about transparency, measurement quality, and the incentives of vertically integrated platforms that both run the auction and provide the tools used to bid in it.

These findings have practical and policy relevance. Industry participants often frame changes in reporting granularity, attribution windows, privacy-preserving aggregation, or "black-box" bidding tools as responses to privacy or complexity. While such factors matter, our analysis emphasizes an additional incentive: opacity can be a competitive strategy that relaxes price pressure. This perspective helps interpret why improvements in measurement and comparability (for example, standardized conversion APIs or independent attribution) may face resistance even when they appear efficiency-enhancing for advertisers. Conversely, it clarifies why platforms might selectively increase precision in dimensions that raise within-platform revenue while limiting precision in dimensions that facilitate cross-platform benchmarking. The model thus suggests that policies aimed at transparency or interoperability can have pro-competitive effects, but also that their incidence depends on how they alter advertiser sensitivity and the platforms' ability to extract surplus through format and reserve choices.

We also acknowledge limitations. Our reduced-form treatment of signal precision abstracts from many institutional details: heterogeneous advertiser objectives, multi-objective autobidders, learning dynamics, and the possibility that platforms commit to long-run measurement regimes while advertisers adapt over time. We focus on two platforms to keep the strategic logic transparent, though the forces we isolate plausibly strengthen with more competitors. Finally, we model format differences as an expected price increment; in reality, the welfare consequences of first- versus second-price auctions also involve bidding strategies, risk, and dynamic budget management. These simplifications are deliberate: they allow us to cleanly separate the direct value of information from its competitive externality through elasticity, and to derive transparent comparative statics.

The remainder of the paper situates our contribution in the related literature, develops the model, and characterizes equilibrium format and information choices. Throughout, our goal is not to claim that any particular platform deliberately degrades prediction quality, but to illuminate a general tradeoff: in autobidding markets, information is simultaneously a productivity-enhancing input and a strategic lever that shapes competitive intensity.

2 Related literature and positioning

Our analysis speaks to four adjacent literatures: (i) autobidding and constrained optimization in ad auctions, including efficiency and price-of-anarchy considerations; (ii) platform-side "auction design with ML advice," where platforms jointly choose pricing rules and prediction/quality adjustments; (iii) competitive platform design in advertising markets, especially format choice and multi-homing; and (iv) information design and strategic garbling in industrial organization. We briefly position our contribution within each strand and clarify what is deliberately abstracted away.

Autobidding, budget/RoS constraints, and equilibrium inefficiency. A large theoretical and empirical literature studies auctions and pacing when bidders are not standard quasilinear agents but face budget constraints or performance targets. In these environments, bidders (or their delegates) choose bids to satisfy constraints such as average cost-per-click, return on ad spend, or value-per-dollar objectives, often in repeated and high-frequency settings. Related models include analyses of budget-constrained bidding, pacing equilibria, and the efficiency consequences of constrained bidding, as in work on smooth mechanisms and welfare bounds, and in the ad-auctions tradition that links platform rules to equilibrium bids and allocations (e.g., **?**; **?**; **?**). A recurring theme is that constraints can distort bids away from true marginal values, generating welfare loss relative to a full-information, unconstrained benchmark.

We view our reduced-form RoS constraint as a tractable way to capture this empirically salient feature without committing to a specific pacing algorithm. Rather than characterizing within-auction bid shading or dynamic budget smoothing, we focus on the *cross-platform* implication: constrained bidders exhibit threshold-like participation and reallocation when their inferred net value falls relative to effective price. This emphasis connects to price-of-anarchy style reasoning in a different dimension. The inefficiency we highlight is not only allocative (misallocation across impressions within a platform), but also informational and strategic (platforms may choose to reduce the informativeness of the signals that drive the autobidder), which can lower total surplus even when each platform could technically improve prediction quality.

Platform "ML advice" as an instrument: reserves, quality scores, and bid multipliers. A second strand studies how platforms influence auction outcomes through the prediction layer: reserve prices, scoring rules, quality adjustments, and other transformations of bids or values that are justified as relevance, user experience, or ML calibration. In sponsored search, quality scores and ad-rank rules effectively reweight bids; in display, exchanges frequently apply bid multipliers, floors, and other policy levers. A related theoretical literature examines revenue-optimal mechanisms when the seller has access to information or can choose what to reveal, including the design of signaling schemes in auctions and the role of information in extracting surplus (e.g., **?**; **?**). In parallel, empirical work in ad markets documents how prediction improvements and scoring changes can affect both allocative outcomes and platform revenue.

Our contribution is to treat the prediction/reporting layer itself as a strategic object that interacts with competition. We abstract from the fine structure of the ad-rank formula and instead summarize the platform's "ML advice" by a one-dimensional precision parameter that is Blackwell ordered. This choice is not meant to deny the richness of platform ML, but to separate two conceptually distinct effects: (i) a *direct* return from better prediction (higher match value and tighter within-platform monetization), and (ii) an *elasticity* effect whereby better information makes advertisers more responsive in their cross-platform allocation. Much of the existing mechanism-design-with-ML perspective emphasizes the first channel; our main results rely on the second.

Format choice and competitive platform design in advertising. There is a substantial literature on auction format in ad markets, particularly the practical shift from second-price to first-price auctions and the resulting role of bid shading, equilibrium markups, and intermediary fees (e.g., **?**; **?**). At the same time, industrial organization work on platforms and multi-homing emphasizes that advertisers often allocate budgets across multiple channels, and that platforms compete through fees, targeting capabilities, and measurement. In such settings, a platform's attempt to raise per-unit revenue can be counteracted by share loss when advertisers substitute. This competition logic is familiar, but the ad-tech setting adds a distinctive feature: the objects that govern substitution are not only prices and product characteristics, but also *measurement and comparability.*

We build on this insight by making format choice and information policy jointly strategic. We model the difference between first-price and second-price as an expected payment increment holding other instruments fixed, which allows us to isolate the competitive tradeoff without taking a stand on the full bidding equilibrium under each format. The resulting format threshold result can be read as a reduced-form complement to the detailed auction-theoretic analyses: even if a first-price format mechanically raises payments in partial equilibrium, it may be dominated in a competitive environment once one accounts for induced reallocation. Importantly, because information affects the responsiveness of reallocation, format viability becomes endogenous to the information regime. This interaction is less emphasized in the format-choice literature, which typically takes bidder information and cross-platform outside options as given.

Information design, obfuscation, and strategic garbling in IO. Finally, our paper relates to the information design literature, which studies how an informed party optimally chooses what to reveal to influence downstream actions (e.g., **?**; **?**). In IO, a complementary literature on obfuscation and shrouded attributes shows that firms may strategically reduce transparency to soften competition or exploit behavioral frictions (e.g., **?**; **?**). In many of these models, disclosure affects demand either by changing perceived quality/price levels or by creating search and comparison frictions.

Our mechanism is closest in spirit but differs in economic object. We do not assume naive consumers or limited attention. Instead, advertisers are rational but constrained and delegate decisions to an optimization algorithm whose inputs are platform-provided signals. Garbling then operates through a rational elasticity channel: more informative signals reduce decision noise and steepen the demand response to relative net value. In this sense, we adopt an "as-if" discrete-choice representation of allocation in which information policy shifts the slope parameter. This reduced form can be microfounded using standard Blackwell comparisons and rational choice with noisy signals, but we do not require a particular behavioral foundation. The key point is that even with fully rational advertisers, the platform may prefer opacity because precision intensifies competition.

Positioning and scope. Relative to these literatures, we contribute a parsimonious framework in which platforms choose *three* instruments—format, reserve/price, and information precision—and where the distinctive role of information is to affect cross-platform sensitivity. This focus yields sharp comparative statics and transparent conditions under which (i) lower-price formats can dominate higher-price formats in revenue once demand reallocation is elastic, and (ii) equilibrium information can be inefficiently coarse because platforms do not internalize the surplus gains from improved allocation while they do internalize the competitive pressure it creates.

At the same time, our modeling choices impose limitations that guide interpretation. We abstract from dynamic learning by advertisers and from repeated-game incentives that might discipline or amplify information manipulation. We also compress a multidimensional reporting environment (attribution, incrementality, privacy-preserving aggregation, diagnostics) into a single precision index; in practice, platforms can selectively increase precision in some dimensions while decreasing comparability in others. Finally, we do not attempt to adjudicate the welfare ranking of first- versus second-price auctions in full generality; instead, we take the observed "price increment" effect as a sufficient statistic for how format changes translate into effective prices facing autobidders. These simplifications are intentional: they allow us to highlight a competition-relevant tradeoff that cuts across institutional details and motivates the model we develop next.

Roadmap. Guided by these connections, Section 3 formalizes a two-platform environment in which platforms choose auction format, a price instrument, and a Blackwell-ordered signal precision, while RoS-constrained advertisers allocate spend across platforms as a function of posterior net value and effective prices. The subsequent analysis uses this structure to characterize equilibrium format and information choices and to compare private incentives to a welfare benchmark.

3 Model

We model two competing advertising platforms, indexed by $k \in \{1, 2\}$, that sell a homogeneous unit of attention (an "impression") to a continuum of advertisers $i \in [0, 1]$. The goal is to isolate a simple but practically salient interaction: platforms can raise the effective price advertisers pay (through auction format and related pricing instruments), but they also control how *measurable and comparable* performance is through the information they provide to advertisers and their autobidders. More informative signals improve match quality and monetization within a platform, yet also make advertisers more willing to reallocate spend across platforms in response to small differences in net value.

Advertisers, values, and RoS constraints. Advertiser i has a platform-specific expected value v_{ik} from an impression on platform k. This value can be interpreted broadly: expected conversions times margin, incrementality-adjusted lift, or any internal value metric the advertiser would like the autobidder to maximize. We allow v_{ik} to differ across platforms even for the same advertiser, capturing differences in audience composition, creative fit, or measurement.

A distinctive feature of the environment is that advertisers are constrained by performance targets. Each advertiser i has a return-on-spend (RoS) requirement parameter $\tau_i \geq 0$. In reduced form, we treat the RoS requirement as a constraint that ties expected value to expected payments. Intuitively, if a platform becomes "too expensive" relative to the advertiser's perceived value, the advertiser (or its autobidder) will either scale back or stop allocating spend there. This captures a broad set of practical constraints used in autobidding systems (e.g., target ROAS, target CPA, or value-per-dollar pacing), while remaining agnostic about the exact algorithm.

Platform instruments: format, reserve/price, and information precision. Each platform k commits to a triple of instruments

$$(M_k, r_k, \sigma_k) \in \{\text{SPA}, \text{FPA}\} \times \mathbb{R}_+ \times [0, 1].$$

The first component M_k is the auction format, either second-price (SPA) or first-price (FPA). The second component $r_k \geq 0$ is a price instrument that the platform can directly control. We call it a "reserve" for concreteness, but it should be interpreted more broadly as any lever that shifts the expected payment per impression holding demand fixed: floors, take rates, scoring/multipliers that scale bids, or other implementation details that, from the advertiser's perspective, show up as a higher expected cost per unit of delivery.

The third component $\sigma_k \in [0, 1]$ indexes the platform's information policy. Higher σ_k corresponds to a more informative signal about advertiser-impression value and, crucially for our competitive logic, to a more *decision-relevant* signal for cross-platform allocation. While real reporting environments are high-dimensional (attribution windows, conversion modeling, incrementality diagnostics, privacy noise, aggregation, and so on), we summarize the informativeness of what the platform provides by a one-dimensional precision index that is Blackwell ordered.

Signal structures and the Blackwell order. For each platform k, precision σ_k selects a signal structure $S_k(\sigma_k)$. Advertiser i then observes a platform-specific signal $\hat{v}_{ik}$ drawn according to $S_k(\sigma_k)$. We assume that if $\sigma' > \sigma$, then $S_k(\sigma')$ Blackwell-dominates $S_k(\sigma)$: the lower-precision signal can be obtained by garbling the higher-precision one. This assumption is the standard way to represent a platform's ability to commit to being more or less informative without committing to a particular parametric noise distribution. It also matches practice: a platform can coarsen by adding noise, reporting at a higher level of aggregation, delaying feedback, censoring diagnostics, or otherwise reducing how sharply the advertiser can infer which placements, audiences, or platforms are performing best.

Advertisers use Bayes' rule to form posterior expected values $\mathbb{E}[v_{ik} \mid \hat{v}_{ik}]$. We emphasize that nothing in the model relies on advertisers being naive: the information design problem arises even when advertisers process signals optimally, because signal precision affects the *responsiveness* of subsequent allocation decisions.

Timing. The game has three stages.

1. **Stage 0 (platform commitment).** Platforms simultaneously choose (M_k, r_k, σ_k).

2. **Stage 1 (advertiser allocation).** Advertisers observe signals $(\hat{v}_{i1}, \hat{v}_{i2})$ and allocate spend shares across platforms. We denote advertiser i's share on platform k by $a_{ik} \in [0, 1]$, with (a_{i1}, a_{i2}) in the simplex (allowing an outside option if neither platform meets the RoS constraint).

3. **Stage 2 (auction execution and payments).** Auctions execute within each platform, and platform k collects expected payment per impression p_k times the allocated quantity it receives.

We interpret Stage 1 as the outcome of an autobidder that takes as inputs platform-provided value signals and effective prices and outputs campaign-level allocation across channels.

Reduced-form effective prices and the role of auction format. Rather than modeling the full within-platform bidding equilibrium under each format, we summarize the platform's expected per-impression payment by an effective price function $p_k(M_k, r_k, \sigma_k)$. This abstraction serves two purposes. First, it lets us treat "format" as a pricing rule that affects expected payment while also affecting advertiser net value, without committing to the detailed mechanics of bid shading or pacing. Second, it keeps the competitive logic transparent: what matters for cross-platform substitution is how platform choices shift the expected cost of buying a unit of delivery.

We adopt the following simple decomposition:

$$p_k(\text{SPA}, r, \sigma) = r + \phi(\sigma), \qquad p_k(\text{FPA}, r, \sigma) = r + \phi(\sigma) + \Delta,$$

with $\Delta > 0$ an exogenous increment capturing the idea that, holding other instruments fixed, first-price implementation tends to raise expected payments relative to second price. The term $\phi(\sigma)$ captures within-platform monetization effects of precision: better prediction and tighter matching can increase competition for valuable impressions, reduce waste, and (in expected terms) raise the price per impression even under a fixed reserve. We allow $\phi(\sigma)$ to be weakly increasing.

This specification does not claim that first-price *always* raises payments in all environments; rather, Δ is a sufficient statistic for the partial-equilibrium payment shift relevant for our comparative statics. In applications, Δ can be interpreted as incorporating bid shading frictions, intermediary fees, or other implementation differences that make the effective price higher under first price from the advertiser's perspective.

Advertiser payoffs and feasibility under RoS. Given effective prices, advertiser i chooses (a_{i1}, a_{i2}) to maximize expected value subject to meeting its RoS constraint:

$$\max_{(a_{i1},a_{i2})\in\Delta} \sum_{k=1}^{2} a_{ik}\,\mathbb{E}[v_{ik} \mid \hat{v}_{ik}] \quad \text{s.t.} \quad \sum_{k=1}^{2} a_{ik}\,\mathbb{E}[v_{ik} \mid \hat{v}_{ik}] \;\geq\; \tau_i \sum_{k=1}^{2} a_{ik}\, p_k(M_k, r_k, \sigma_k).$$

This formulation captures the idea that advertisers scale spend only where expected value is sufficiently high relative to expected cost. It also creates a natural participation cutoff: if a platform's expected net value falls below zero (relative to the target τ_i), the advertiser optimally assigns it zero weight. In the next section, we impose a tractable choice rule that aggregates these individual decisions into a smooth demand share system and makes the elasticity effects of σ_k explicit.

Platform payoffs and the cost of precision. Let D_k denote the aggregate demand/spend share that platform k receives in Stage 1. Platform k's payoff is

$$\Pi_k \;=\; p_k(M_k, r_k, \sigma_k)\, D_k \;-\; c(\sigma_k),$$

where $c(\sigma_k)$ is a convex cost of providing precision. This cost can represent engineering and data infrastructure costs, privacy and compliance burdens, increased fraud/measurement risk, or brand/trust concerns associated with highly granular reporting. Convexity captures diminishing returns and the idea that moving from moderately informative to highly informative signals may be disproportionately costly (technologically or institutionally).

Discussion and scope. Two modeling choices deserve emphasis. First, we treat information as a platform commitment about what the autobidder can observe, not as a passive byproduct of technology. This aligns with the practical reality that platforms choose what to measure and report, what to aggregate, and what to expose via APIs and dashboards. Second, we isolate a competitive channel through which higher precision can intensify platform competition: as advertisers can better infer relative performance, they can more finely arbitrage across platforms, making demand more sensitive to small price or format differences. Our subsequent analysis separates this elasticity channel from the direct effects of precision on match value and within-platform monetization, and it compares the private incentive to choose σ_k to a welfare benchmark in which increased comparability is socially valuable.

4 Advertiser allocation and demand elasticity

To study format and information competition we need a tractable mapping from platform choices (M_k, r_k, σ_k) into *aggregate* spend shares (D_1, D_2). The microfoundation in Section 3 delivers a natural object: advertiser i's posterior expected value $\mathbb{E}[v_{ik} \mid \hat{v}_{ik}]$, which is then compared to the platform's effective payment $p_k(M_k, r_k, \sigma_k)$ through the RoS requirement. The difficulty is that, taken literally, RoS feasibility generates corner solutions and discontinuous market shares: an arbitrarily small change in price can flip a campaign from "in" to "out." We therefore impose a smooth reduced form that (i) preserves the RoS participation cutoff logic, and (ii) makes explicit how signal precision increases cross-platform responsiveness.

Posterior net value and the RoS participation cutoff. Fix platform choices and consider advertiser i's posterior *RoS-adjusted* net value (or "surplus index") on platform k,

$$m_{ik} \;:=\; \mathbb{E}[v_{ik} \mid \hat{v}_{ik}] \;-\; \tau_i \, p_k(M_k, r_k, \sigma_k). \tag{1}$$

The RoS constraint implies a participation cutoff: if $m_{ik} < 0$, then allocating spend to platform k worsens expected performance relative to the target, so an optimizer sets $a_{ik} = 0$ unless it is needed to satisfy some other constraint not modeled here. Conversely, when $m_{ik} > 0$, platform k is "feasible" for advertiser i and can receive positive allocation. In the knife-edge deterministic benchmark in which advertisers perfectly rank platforms by (1) and face no adjustment frictions, each advertiser would assign all spend to $\arg\max\{m_{i1}, m_{i2}, 0\}$, where the outside option (no spend) has normalized payoff 0. This benchmark is conceptually useful—it highlights the role of RoS cutoffs—but it is too discontinuous for comparative statics in (M_k, r_k, σ_k).

A smooth allocation rule: random utility with an outside option. We smooth the deterministic cutoff by allowing for residual campaign-level dispersion in how advertisers translate posterior net values into realized allocations. One interpretation is that advertisers face additional unmodeled constraints (budgets, learning, attribution uncertainty, creative fatigue) that create "mistakes" or inertia relative to pure myopic maximization of (1). Formally, suppose advertiser i chooses among $\{0, 1, 2\}$ according to

$$\text{choose } k \in \{0,1,2\} \;\text{ to maximize }\; m_{ik} + \varepsilon_{ik},$$

where ε_{ik} are i.i.d. Type-I extreme value shocks and $m_{i0} \equiv 0$ is the outside option. The standard logit formula implies that the probability (and, with a continuum of infinitesimal campaigns, the spend share) assigned to platform k is increasing in its net value index and decreasing in rivals' indices. Importantly, the *scale* of the shocks governs how sharply shares respond to utility differences. Writing this scale as an "inverse noise" parameter $\beta \geq 0$ yields choice probabilities proportional to $\exp(\beta m_{ik})$.

To connect the allocation rule to platform instruments, we aggregate advertiser heterogeneity and signal realizations into a platform-level net value index. In our reduced form we write

$$u_k \;:=\; V(\sigma_k) \;-\; p_k(M_k, r_k, \sigma_k), \tag{2}$$

where $V(\sigma_k)$ captures the expected (RoS-normalized) value that advertisers can generate on platform k when precision is σ_k. This term is meant to absorb both (i) the direct match-quality value of precision (better targeting, fewer wasted impressions), and (ii) the effect of precision on posterior expected value formation. We then posit the market share system

$$D_k \;=\; \frac{\exp(\beta(\sigma_k)\, u_k)}{1 + \exp(\beta(\sigma_1)\, u_1) + \exp(\beta(\sigma_2)\, u_2)}, \qquad k \in \{1, 2\}, \tag{3}$$

with outside option share $D_0 = 1 - D_1 - D_2 = \big(1 + \sum_{j=1}^{2} \exp(\beta(\sigma_j) u_j)\big)^{-1}$. The outside option is the reduced-form representation of RoS-driven non-participation: when both platforms deliver negative net value indices, most mass goes to D_0; when a platform's u_k is strongly negative, its share vanishes. In the limit $\beta(\sigma) \to \infty$, (3) converges to the deterministic cutoff rule that places essentially all weight on the best nonnegative option, recovering the sharp RoS participation logic.

Elasticities and substitution patterns. Equation (3) implies transparent comparative statics. Let $\beta_k := \beta(\sigma_k)$ and define $z_k := \exp(\beta_k u_k)$ and $Z := 1 + z_1 + z_2$. Then $D_k = z_k/Z$ and the share ratio satisfies

$$\frac{D_1}{D_2} \;=\; \exp(\beta_1 u_1 - \beta_2 u_2), \tag{4}$$

so relative demand is exponentially sensitive to relative net values, with sensitivity governed by the β's. Differentiating delivers the familiar logit response:

$$\frac{\partial D_k}{\partial u_k} = \beta_k\, D_k(1 - D_k), \tag{5}$$

$$\frac{\partial D_k}{\partial u_j} = -\beta_j\, D_k D_j, \qquad j \neq k. \tag{6}$$

Thus higher β_j makes platform j's net value more "pivotal" for all shares: small changes in u_j trigger larger reallocations.

Because u_k subtracts the effective payment, the same expressions translate into price elasticities. Holding σ fixed, a marginal increase in p_k reduces u_k one-for-one, so

$$\frac{\partial \ln D_k}{\partial p_k} \;=\; -\beta_k(1 - D_k), \tag{7}$$

and raises rivals' (and the outside option's) shares. This is the key sense in which $\beta(\sigma)$ controls *cross-platform contestability*: when $\beta(\sigma)$ is large, the demand loss from a higher effective price (whether induced by r_k or by choosing FPA) is amplified.

Why precision increases cross-platform sensitivity. We now connect $\beta(\sigma)$ to the information policy. The assumption that $S_k(\sigma')$ Blackwell-dominates $S_k(\sigma)$ for $\sigma' > \sigma$ means that, with higher precision, advertisers can form posterior means that more accurately rank alternatives. In discrete choice terms, information reduces decision noise. A simple illustration is to suppose that the advertiser's perceived net value difference is a noisy signal

$$\tilde{d}_i(\sigma) \;=\; (u_1 - u_2) + \eta_i(\sigma),$$

where $\eta_i(\sigma)$ is mean-zero noise whose dispersion falls with σ. Then the probability of allocating to platform 1 is $\Pr(\tilde{d}_i(\sigma) \geq 0)$, whose slope with respect to $(u_1 - u_2)$ increases as σ increases. Under Gumbel noise this exactly yields a logit form with $\beta(\sigma)$ proportional to the inverse noise scale; under other common noise families (normal, logistic), the same monotonicity holds: more precision makes the choice probability function steeper. We therefore treat

$$\beta'(\sigma) \;>\; 0 \tag{8}$$

as the reduced-form implication of the Blackwell order for cross-platform allocation. For later comparative statics it is useful to parameterize this slope by $b := \beta'(\sigma)$ (e.g., $\beta(\sigma) = \beta_0 + b\sigma$), which measures how strongly improvements in measurement and reporting translate into more aggressive spend reallocation.

Two channels of σ in demand: levels versus elasticity. Finally, note that precision affects demand in two conceptually distinct ways. First, it shifts the *level* of net value through $V(\sigma_k)$ (and, through within-platform monetization, potentially through p_k via $\phi(\sigma_k)$). Second, it changes the *slope* of the demand system through $\beta(\sigma_k)$. To see the decomposition, observe that D_k depends on σ_k through the index $\beta(\sigma_k)u_k(\sigma_k)$; thus

$$\frac{\partial \ln z_k}{\partial \sigma_k} \;=\; \beta'(\sigma_k)\, u_k \;+\; \beta(\sigma_k)\, \frac{\partial u_k}{\partial \sigma_k}.$$

The first term is the "elasticity channel": holding the net value level fixed, making information more precise changes how sharply advertisers react to differences. The second term is the "direct value/monetization channel": precision changes u_k itself. Our main competitive results will turn on the first term: even when more information is intrinsically helpful (so $\partial u_k / \partial \sigma_k$ is positive), it can intensify substitution and thereby discipline prices and formats.

With the demand system (3) in hand, we can now examine format competition holding σ fixed. In particular, because format affects p_k through the increment Δ, the profitability of FPA versus SPA will hinge on the sensitivity parameter $\beta(\sigma)$: when advertisers can finely compare performance across platforms, they punish high effective prices more strongly, shifting equilibrium format choices.

5 Format competition holding precision fixed

We next isolate the platforms' *format* incentives by holding information and reserves fixed. This serves two purposes. First, it provides a clean closed-form condition under which a second-price auction (SPA) dominates a first-price auction (FPA) in a competitive environment with cross-platform substitution. Second, it makes transparent why the format margin will later interact sharply with the information margin: the same precision choice that raises match value can also raise the sensitivity parameter $\beta(\sigma)$ that disciplines high-price formats.

5.1 A reduced-form format wedge

Fix $\sigma \in [0,1]$ and $r \geq 0$ and suppose both platforms share these values. Format affects the effective expected payment per impression through the constant wedge $\Delta > 0$,

$$p(\mathrm{SPA}, r, \sigma) =: p_S, \qquad p(\mathrm{FPA}, r, \sigma) = p_S + \Delta.$$

We interpret Δ as the expected price increment from switching the auction rule while holding the platform's other instruments fixed (e.g., due to reduced bid shading under FPA relative to SPA, or, more generally, a format-induced change in the expected price impact faced by the autobidder). This is not meant to be a literal statement about auction theory under arbitrary information; rather, it is the minimal reduced form needed to compare formats in the presence of elastic cross-platform allocation.

Given σ, advertiser responsiveness is summarized by $\beta := \beta(\sigma)$. Because $V(\sigma)$ is held fixed in this section, a platform that switches from SPA to FPA reduces its net value index u one-for-one by Δ, i.e.,

$$u(\mathrm{FPA}) = u(\mathrm{SPA}) - \Delta,$$

while simultaneously raising its per-impression payment by Δ. The format problem is therefore a canonical tradeoff between *margin* (higher p) and *share* (lower D), with the share response governed by β.

5.2 Closed-form threshold: SPA versus FPA against an SPA rival

To obtain a closed form, we focus on the standard competitive benchmark in which the outside option is negligible at the margin of format choice (equivalently, in the limit in which both platforms deliver sufficiently positive net value that $D_0 \approx 0$). In that case, demand splits only across the two platforms and the logit rule reduces to the familiar binary formula. If platform

2 uses SPA and platform 1 deviates to FPA, the effective price difference is Δ, so platform 1's demand share is

$$D_1(\text{FPA} \mid \text{SPA}) \;=\; \frac{1}{1+\exp(\beta\Delta)}. \tag{9}$$

If instead platform 1 also uses SPA, symmetry implies $D_1(\text{SPA} \mid \text{SPA}) = 1/2$. Profits in the format subgame (holding (r, σ) fixed) are therefore

$$\Pi_1(\text{SPA} \mid \text{SPA}) = \frac{p_S}{2}, \qquad \Pi_1(\text{FPA} \mid \text{SPA}) = \frac{p_S + \Delta}{1+\exp(\beta\Delta)}.$$

Comparing these expressions yields a simple threshold in β.

Proposition (Format threshold with fixed (r, σ)). Fix (r, σ) and let $p_S = p(\text{SPA}, r, \sigma)$ and $p(\text{FPA}, r, \sigma) = p_S + \Delta$ with $\Delta > 0$. Suppose platform 2 uses SPA and the outside option is negligible. Then platform 1 (weakly) prefers SPA to FPA if and only if $\beta \geq \beta^*$, where

$$\beta^* \;=\; \frac{1}{\Delta} \ln\!\left(1 + \frac{2\Delta}{p_S}\right). \tag{10}$$

Derivation. The condition $\Pi_1(\text{SPA} \mid \text{SPA}) \geq \Pi_1(\text{FPA} \mid \text{SPA})$ is

$$\frac{p_S}{2} \;\geq\; \frac{p_S + \Delta}{1+\exp(\beta\Delta)}.$$

Rearranging gives $\exp(\beta\Delta) \geq 1 + 2\Delta/p_S$, which is equivalent to (10).

Interpretation. The right-hand side of (10) is decreasing in p_S and increasing in Δ in the natural way. A larger baseline payment p_S makes the incremental margin gain Δ from FPA relatively less important, so the share loss induced by a higher effective price becomes decisive at a lower level of sensitivity. Conversely, when Δ is large, FPA is more "extractive" in net value terms, so a smaller β suffices for advertisers to reallocate strongly enough to make SPA revenue-dominant.

The limiting cases are also intuitive. As $\beta \to 0$ (almost inelastic allocation), (9) yields $D_1 \approx 1/2$, so FPA strictly dominates by raising per-impression payment while barely affecting share. As $\beta \to \infty$ (almost deterministic ranking), $D_1(\text{FPA} \mid \text{SPA}) \to 0$, so the FPA deviator is essentially excluded, and SPA strictly dominates.

5.3 Equilibrium implications: when is SPA uniquely selected?

The threshold in (10) characterizes a best response to an SPA rival. To understand the full 2×2 format game, we also compare deviations when the rival uses FPA. Under the same covered-market approximation, if platform 2 uses FPA and platform 1 uses SPA, then platform 1's demand share is

$$D_1(\mathrm{SPA} \mid \mathrm{FPA}) \;=\; \frac{\exp(\beta\Delta)}{1+\exp(\beta\Delta)},$$

so

$$\Pi_1(\mathrm{SPA} \mid \mathrm{FPA}) = p_S \frac{\exp(\beta\Delta)}{1+\exp(\beta\Delta)}, \qquad \Pi_1(\mathrm{FPA} \mid \mathrm{FPA}) = \frac{p_S+\Delta}{2}.$$

A second threshold β^{**} solves $\Pi_1(\mathrm{SPA} \mid \mathrm{FPA}) \geq \Pi_1(\mathrm{FPA} \mid \mathrm{FPA})$, yielding (when $p_S > \Delta$)

$$\beta^{**} \;=\; \frac{1}{\Delta} \ln\left(\frac{p_S+\Delta}{p_S-\Delta}\right). \tag{11}$$

This implies three regimes. If $\beta < \beta^*$, each platform prefers FPA even against an SPA rival, so (FPA, FPA) is the unique equilibrium. If $\beta > \beta^{**}$, each platform prefers SPA even against an FPA rival, so (SPA, SPA) is the unique equilibrium. For intermediate $\beta \in [\beta^*, \beta^{**}]$, both (SPA, SPA) and (FPA, FPA) can be equilibria, reflecting a coordination problem: when both platforms commit to a high-price format, each faces a weaker incentive to unilaterally soften, but when the rival is already soft, matching softness is attractive because share is highly contestable.

This multiplicity is economically important for practice. It suggests that observed persistence of first-price formats need not indicate that FPA is intrinsically revenue-superior; it may instead reflect equilibrium selection in a setting where platforms anticipate that softening format would be met by aggressive responses elsewhere (e.g., via reserves, quality adjustments, or other fee instruments not explicitly modeled here).

5.4 Remarks: outside option and asymmetries

Two modeling choices deserve emphasis. First, incorporating the outside option D_0 generally strengthens the share-loss channel from raising effective prices, because demand can exit rather than merely reallocate. The closed-form expressions above are therefore best interpreted as conservative benchmarks: when RoS constraints bind tightly and nonparticipation is salient, the region in which SPA dominates expands.

Second, the clean thresholds rely on a symmetric comparison that fixes (r, σ) across platforms. With asymmetric reserves or asymmetric baseline net values, the same logic applies but the thresholds become state-dependent: a platform with higher baseline net value can sometimes sustain the high-price format for longer, because it starts from a position of higher share. This will matter in Section 6, where r_k and σ_k are themselves strategic choices and may be used to manipulate both the level u_k and the slope $\beta(\sigma_k)$ of demand.

Taken together, the results of this section formalize a simple competitive message: holding information fixed, high sensitivity $\beta(\sigma)$ makes demand highly contestable, so a format that raises effective payments by Δ can be revenue-inferior despite its higher per-impression margin. In the next section we allow platforms to choose σ_k and r_k jointly, and show that precisely because higher precision raises $\beta(\sigma)$, platforms may have an incentive to *coarsen* information to sustain higher markups and, in some parameter regions, to support high-price formats that would otherwise unravel.

6 Joint choice of reserves and precision

We now turn to the main strategic interaction in the model: platforms simultaneously choose a price instrument (the reserve r_k) and an information policy (precision σ_k). Unlike the format comparison in Section 5, the reserve margin cannot be analyzed under a fully covered market, because if the outside option were literally irrelevant then a symmetric increase in both platforms' reserves would leave shares unchanged while raising per-impression payments. To discipline pricing and to make precision consequential for participation, we therefore work with the baseline multinomial logit demand system that includes the outside option,

$$D_k \;=\; \frac{\exp(\beta(\sigma_k)u_k)}{1+\exp(\beta(\sigma_1)u_1)+\exp(\beta(\sigma_2)u_2)}, \qquad u_k \;=\; V(\sigma_k)-p_k,$$

and with the reduced-form payment rule

$$p_k \;=\; r_k+\phi(\sigma_k)+\mathbf{1}\{M_k=\text{FPA}\}\Delta.$$

In this section we treat M_k as given when characterizing (r_k,σ_k), and then interpret the resulting equilibrium $\beta(\sigma^*)$ through the format thresholds derived earlier.

6.1 Reserve choice conditional on precision

Fix (M_1,M_2) and (σ_1,σ_2). Platform k's profit is

$$\Pi_k \;=\; p_kD_k-c(\sigma_k).$$

Because p_k depends one-for-one on r_k and $u_k=V(\sigma_k)-p_k$, the reserve first-order condition (when interior) has the familiar logit "markup equals inverse elasticity" form. Using $\partial D_k/\partial u_k=\beta(\sigma_k)D_k(1-D_k)$, we obtain

$$\begin{aligned}\frac{\partial \Pi_k}{\partial r_k} &= D_k+p_k\frac{\partial D_k}{\partial r_k}=D_k+p_k\frac{\partial D_k}{\partial u_k}\frac{\partial u_k}{\partial r_k}\\ &= D_k-p_k\,\beta(\sigma_k)\,D_k(1-D_k).\end{aligned} \tag{12}$$

Thus any interior optimum satisfies

$$p_k \;=\; \frac{1}{\beta(\sigma_k)\,(1-D_k)}. \tag{13}$$

Equation (13) makes the central role of $\beta(\sigma)$ transparent: holding fixed demand share D_k, greater sensitivity (higher β) forces a lower equilibrium effective payment p_k, and hence a lower reserve r_k. In other words, the same informational improvement that helps advertisers allocate more sharply also disciplines platform pricing.

In a symmetric profile $(M_1, M_2) = (M, M)$ and $(r_1, \sigma_1) = (r_2, \sigma_2) = (r, \sigma)$, we can summarize reserve setting by the pair of fixed-point relationships

$$D(\sigma, r) \;=\; \frac{\exp(\beta(\sigma)u)}{1 + 2\exp(\beta(\sigma)u)}, \qquad p \;=\; r + \phi(\sigma) + \mathbf{1}\{M = \text{FPA}\}\Delta, \qquad u \;=\; V(\sigma) - p, \tag{14}$$

together with the pricing condition $p = 1/(\beta(\sigma)(1-D))$. While (14) typically does not deliver a closed-form solution for r, it provides a sharp comparative static: any change in σ that increases $\beta(\sigma)$ tends to reduce the feasible markup in (13), thereby pushing equilibrium reserves downward unless offset by a sufficiently strong direct shift in u through $V(\sigma)$ and $\phi(\sigma)$.

6.2 Precision choice and the incentive to garble

We next examine how a platform evaluates a marginal increase in σ_k. Precision affects profits through three conceptually distinct channels.

(i) Direct match-value channel. Higher σ_k increases advertisers' expected match value $V(\sigma_k)$, raising u_k and therefore D_k for any fixed effective payment p_k. This is the standard efficiency rationale for more information: better targeting and measurement make impressions more valuable to the advertiser side and expand demand.

(ii) Direct monetization channel. Higher σ_k may also increase within-platform monetization $\phi(\sigma_k)$ (e.g., tighter bid distributions or less wasteful allocations raise clearing prices). This increases p_k mechanically, holding r_k fixed, but it *reduces* advertiser net value u_k. Whether $\phi'(\sigma)$ is profit-increasing depends on how strongly the induced increase in per-impression payment outweighs the share loss from the lower u_k.

(iii) Elasticity (competition) channel. Crucially, by assumption $\beta'(\sigma) > 0$: more precise signals reduce decision noise and make cross-platform allocation more responsive to net-value differences. In the reserve condition (13), this channel operates like an endogenous increase in demand elasticity. Intuitively, as σ_k rises, advertisers become quicker to reallocate away from any platform that raises its effective price, so the platform internalizes that it cannot profitably sustain the same markup. This is the strategic force toward *coarsening* (garbling): lowering precision softens competitive pressure by keeping allocation relatively inert.

A convenient way to formalize this tradeoff is to consider the platform's *reduced* objective after optimizing the reserve, $\widehat{\Pi}_k(\sigma_k; \sigma_{-k}) := \max_{r_k \geq 0} \Pi_k(r_k, \sigma_k; r_{-k}, \sigma_{-k})$. The envelope theorem implies that, at an interior optimum $r_k^*(\sigma_k)$, the marginal value of precision is

$$\frac{d\widehat{\Pi}_k}{d\sigma_k} = \frac{\partial}{\partial \sigma_k}\big(p_k D_k\big)\Big|_{r_k = r_k^*(\sigma_k)} - c'(\sigma_k), \tag{15}$$

where $\partial(p_k D_k)/\partial \sigma_k$ captures the direct V' and ϕ' effects *and* the strategic effect working through $\beta'(\sigma_k)$. In symmetric equilibrium, this yields the qualitative implication emphasized in the global context: whenever the elasticity effect is strong (large $b := \beta'(\sigma)$), the equilibrium precision σ^* is depressed relative to a benchmark that ignores competitive spillovers.

Proposition (Strategic coarsening relative to first-best). Suppose $\beta(\sigma) = \beta_0 + b\sigma$ with $b > 0$, $V'(\sigma) > 0$, $\phi'(\sigma) \geq 0$, and $c(\sigma)$ convex. Consider an interior symmetric equilibrium (r^*, σ^*). Then σ^* solves (15) (symmetrically), and there exists a parameter region with sufficiently large b (relative to the direct gains $V'(\sigma) + \phi'(\sigma)$) in which $\sigma^* < \sigma^{FB}$, where σ^{FB} maximizes welfare.

The economic content is that platforms do not internalize the allocative benefits of making demand more price-sensitive; they only experience the resulting erosion of markups. Thus, even if precision improves match quality, platforms may rationally choose opacity to dampen substitution.

6.3 Welfare, revenue, and the transparency wedge

To compare equilibrium precision to a welfare benchmark, we need to specify what is socially valuable. In our environment, payments p_k are transfers between advertisers and platforms, while precision affects (i) the realized value of allocated impressions via $V(\sigma)$, (ii) participation via $D_1 + D_2$, and (iii) platform costs $c(\sigma)$. A natural reduced-form welfare objective therefore treats revenue as a transfer and focuses on surplus from matching net of information costs:

$$W(\sigma_1, \sigma_2, r_1, r_2) = \int_0^1 \sum_{k=1}^{2} a_{ik}\, \mathbb{E}[v_{ik} \mid \hat{v}_{ik}]\, di - \sum_{k=1}^{2} c(\sigma_k),$$

which, under the logit reduced form, is increasing in $V(\sigma_k)$ and in participation $D_1 + D_2$, and decreasing in $c(\sigma_k)$. The first-best σ^{FB} trades off the direct informational gains against $c'(\sigma)$ but does *not* assign independent value to keeping $\beta(\sigma)$ low. In contrast, the private platform optimum σ^* values low $\beta(\sigma)$ precisely because it relaxes the pricing discipline embodied in (13). This gap is the transparency wedge: information policies can be privately too opaque even when their direct allocative effects are positive.

Revenue comparisons further sharpen the distinction between private and social incentives. Because higher σ raises $\beta(\sigma)$, equilibrium reserves and effective payments can fall, so platform revenue pD may decline with transparency even as welfare rises (through higher $V(\sigma)$ and higher participation). This is the sense in which precision is a complement to competition: it benefits advertisers and efficiency, but it can make the market more contestable and reduce platform markups. In turn, through the format thresholds in Section 5, the equilibrium σ^* also indirectly shapes whether high-price formats (modeled as an FPA wedge Δ) are sustainable: greater precision pushes the environment into the region where SPA becomes revenue-dominant.

These results motivate the extensions in the next section, where we relax homogeneity and commitment assumptions and consider policy instruments such as minimum transparency (a lower bound on σ) or increased curvature in transparency costs.

7 Extensions

This section records a set of clean extensions that do not change the basic logic—precision raises allocative value but also raises cross-platform responsiveness—while clarifying when the strategic garbling force is stronger or weaker. Our goal is not to solve each variant exhaustively, but to indicate which objects in the baseline analysis are robust (e.g., a markup discipline channel through $\beta(\sigma)$) and which are knife-edge (e.g., full commitment to (r, σ)).

7.1 Heterogeneous advertisers and heterogeneous RoS targets

In practice, advertisers differ sharply in both value and constraints. A parsimonious way to incorporate this is to let RoS targets τ_i be distributed on $[0, \infty)$ and allow match-value gains from precision to differ by type, e.g. $V_i(\sigma)$ with $V_i'(\sigma) \geq 0$. In the reduced form, type i has utility index

$$u_{ik} = V_i(\sigma_k) - p_k,$$

and allocates according to the same noisy best-response rule (or a participation cutoff induced by the RoS constraint). Aggregate demand becomes an integral over types,

$$D_k = \int_0^1 \frac{\exp(\beta(\sigma_k)\, u_{ik})}{1 + \sum_{j=1}^{2} \exp(\beta(\sigma_j)\, u_{ij})}\, di,$$

possibly with the understanding that types with $\max_j u_{ij} < 0$ select the outside option with high probability.

Two forces are worth highlighting. First, heterogeneity in τ_i makes participation and elasticity state-dependent: advertisers with tight RoS targets are effectively "closer to the margin," so a given increase in p_k can move them to the outside option or to the rival platform. This amplifies the competitive discipline effect of $\beta(\sigma)$ precisely among types that generate large volume in equilibrium. Second, heterogeneity in $V_i'(\sigma)$ creates a targeting motive for transparency: if high-spend advertisers also benefit more from precision (larger V_i'), then the direct match-value channel is weighted toward inframarginal demand, potentially offsetting the strategic coarsening motive. Put differently, with heterogeneous types, the relevant comparison is no longer "b versus $V' + \phi'$" pointwise, but rather an equilibrium-weighted comparison between the incremental profit from attracting high-value/high-budget types and the incremental profit loss from making *all* types more willing to switch.

An empirically useful implication is that observed average precision may mask strong segmentation: a platform can provide high precision for categories where $V_i'(\sigma)$ is large (measurement-intensive verticals) while maintaining coarser reporting elsewhere to soften competitive pressure.

7.2 Partial commitment and reserve adjustment

The baseline timing assumes commitment to (r_k, σ_k) at Stage 0. A natural alternative is partial commitment: platforms choose σ_k up front (reflecting engineering and policy choices) but can adjust reserves r_k later as market conditions evolve. A minimal formalization is:

> Stage 0: platforms choose σ_k; Stage 1: advertisers observe signals and/or policies; Stage 2: platforms choose r_k; Stage 3: allocation and auctions occur.

In a subgame-perfect equilibrium, the reserve at Stage 2 is a best response given (σ_1, σ_2) and anticipated shares. The same markup condition (13) continues to characterize the reserve choice *conditional on* σ. What changes is the interpretation of the precision choice: when reserves are adjustable, increasing σ_k tightens the pricing constraint *in every future state* in which pricing occurs. This strengthens the strategic motive to keep $\beta(\sigma)$ low, and therefore can push equilibrium precision further below the first-best.

At the same time, partial commitment opens an offsetting possibility that is important in applications: if platforms can condition r_k on observables that correlate with advertiser posterior values (e.g. query categories, user cohorts, or measured conversion propensity), then higher σ can improve not only match quality $V(\sigma)$ but also the platform's ability to price-discriminate. In reduced form, this can be represented as a stronger $\phi'(\sigma)$ or as an expanded choice set $r_k(\cdot)$. The model then predicts an ambiguous net effect on transparency: precision raises $\beta(\sigma)$ (hurting markups) but also increases the feasible sophistication of pricing (helping extraction). Which force dominates is a quantitative question, and it provides one motivation for the simulation exercises in the next section.

7.3 More than two platforms K

The extension to $K \geq 2$ competing platforms is immediate under the multinomial logit demand with an outside option:

$$D_k = \frac{\exp(\beta(\sigma_k)u_k)}{1 + \sum_{j=1}^{K} \exp(\beta(\sigma_j)u_j)}.$$

In a symmetric profile with $u_k = u$ and $\beta(\sigma_k) = \beta(\sigma)$, we have $D = \exp(\beta u)/(1+K\exp(\beta u))$, so $1-D = (1+(K-1)\exp(\beta u))/(1+K\exp(\beta u))$. The reserve condition continues to take the markup form $p = 1/(\beta(\sigma)(1-D))$, but now the equilibrium share D is smaller and $1-D$ is larger for larger K. Holding $\beta(\sigma)$ fixed, this pushes the feasible markup down as competition becomes more crowded; conversely, holding the markup fixed, the implied β discipline is tighter. Both comparative statics strengthen the coarsening incentive: when there are more alternatives, any given increase in $\beta(\sigma)$ makes demand more contestable, so platforms have stronger private incentives to dampen precision.

The format logic also generalizes. If a unilateral move from SPA to FPA raises effective price by Δ, then the deviator's demand share falls more sharply when there are many close substitutes, and the revenue gain from the price increment is spread over a smaller base. Thus, the region in which SPA is revenue-dominant expands with K, consistent with the view that "competitive thickness" and transparency jointly discipline high-price mechanisms.

7.4 Budgets, pacing, and the endogeneity of elasticity

A salient feature of ad markets is that many advertisers face hard budgets B_i and pacing constraints. Budgets can be introduced by letting each advertiser choose total spend $s_i \leq B_i$ and allocate shares a_{ik} across platforms subject to the RoS constraint. In such an environment, the mapping from utility indices u_k to aggregate spend shares D_k is no longer purely substitution-driven: when many advertisers are budget-constrained, marginal changes in u_k reallocate *where* spend goes but do not increase total spend much, and the outside option becomes less relevant.

In reduced-form terms, budgets tend to compress the effective elasticity faced by platforms in the relevant region: when $s_i = B_i$ for many types, the outside-option margin is slack and cross-platform switching is limited by campaign-level constraints, learning lags, and pacing rules. This suggests a practical refinement of the baseline $\beta(\sigma)$ channel: we can interpret $\beta(\sigma)$ as an *effective* sensitivity that itself depends on how often advertisers are at interior margins. Precision may still increase responsiveness conditional on being at an interior margin, but the fraction of advertisers at that margin may shrink when budgets bind. The model then predicts that strategic garbling should be weaker in periods or segments where budgets are tight (e.g. peak season), and stronger when marginal dollars are actively re-optimized across platforms.

7.5 Minimum transparency regulation and transparency-cost instruments

Finally, the model provides a simple language for policy: regulation can either impose a floor $\sigma_k \geq \underline{\sigma}$ (minimum measurement/reporting precision) or increase the private cost of precision through a steeper $c(\sigma)$ (privacy compliance, auditing, or liability). A binding minimum precision directly eliminates the low-σ equilibrium that platforms may prefer for competitive reasons. Mechanically, it raises $\beta(\sigma)$, which (by (13)) reduces equilibrium effective payments and thus platform rents, while increasing match value and participation through $V(\sigma)$. It also pushes the environment toward the region where SPA is revenue-dominant, because the penalty from raising price (or adopting high-price formats) is larger when allocation is more responsive.

An increase in the curvature $c''(\sigma)$, by contrast, works in the opposite direction: it makes high precision privately expensive and can rationalize further coarsening even if transparency is socially valuable. This distinction matters for interpreting privacy regulation. Policies that raise the cost of data use (modeled as higher c) can unintentionally promote opacity; policies that mandate standardized reporting or interoperability (modeled as $\sigma \geq \underline{\sigma}$) can counteract the strategic garbling motive but may reduce platform revenue and induce other margin adjustments (e.g. higher reserves where feasible, or product redesigns that shift value away from measured channels).

These extensions underscore a common theme: the key empirical objects are (i) how precision shifts allocative value $V(\sigma)$ and within-platform monetization $\phi(\sigma)$, and (ii) how precision shifts effective sensitivity $\beta(\sigma)$ once budgets, frictions, and the number of competitors are accounted for. We now turn to numerical validation to map these reduced-form forces to auction primitives and to check robustness beyond the symmetric analytic benchmarks.

8 Numerical validation: mapping reduced-form forces to auction primitives

Our analytic results are intentionally reduced-form: we summarize transparency by a one-dimensional precision index σ, and we summarize the resulting cross-platform discipline by an effective sensitivity $\beta(\sigma)$. This abstraction is useful for proving clean comparative statics, but it naturally raises a practical question: when we embed the same economic ingredients into a more literal multi-channel auction environment with (i) noisy value prediction, (ii) within-platform auction competition, and (iii) RoS-constrained autobidding, do we recover the same objects $\{V(\sigma), \phi(\sigma), \beta(\sigma), \Delta\}$ and the same qualitative tradeoffs? We address this question with synthetic simulations designed to (a) make the mapping from primitives to reduced-form parameters explicit, and (b) stress-test the strategic coarsening and format-reversal logic outside the symmetric closed-form benchmark.

8.1 A synthetic multi-channel auction environment

We simulate two (or K) platforms that repeatedly run independent single-impression auctions. On each auction instance t, a set of advertisers $i \in \{1, \ldots, n\}$ draws platform-specific values $\{v_{ikt}\}$. We allow for rich correlation by writing

$$v_{ikt} \;=\; \mu_k + \eta_{it} + \xi_{ikt},$$

where η_{it} is an advertiser–impression component common across platforms (capturing, e.g., product-seasonality fit) and ξ_{ikt} is platform-specific fit (capturing, e.g., audience match). This structure makes platforms imperfect substitutes even under full information.

Platform k discloses to each advertiser a signal $\hat{v}_{ikt}$ generated by a precision-controlled signal structure. A convenient parametric implementation is additive noise,

$$\hat{v}_{ikt} \;=\; v_{ikt} + \varepsilon_{ikt}(\sigma_k), \qquad \varepsilon_{ikt}(\sigma_k) \sim \mathcal{N}\big(0,\; s^2(\sigma_k)\big)\,,$$

with $s'(\sigma) < 0$. Higher σ is Blackwell-more-informative in the Gaussian location family, and this monotone informativeness lets us interpret σ as "measurement quality" or "reporting granularity" in a way that is close to industry practice.

Each advertiser has an RoS target τ_i and delegates bidding to a simple value-based autobidder. Given $\hat{v}$, the bidder computes a posterior mean $m_{ikt} := \mathbb{E}[v_{ikt} \mid \hat{v}_{ikt}]$ and submits a bid proportional to the maximum cost-per-impression consistent with meeting the RoS target in expectation,

$$b_{ikt} \;=\; \frac{m_{ikt}}{\tau_i}.$$

This rule is not meant to be an equilibrium of the full dynamic bidding problem; rather, it captures the operational reality that autobidders translate predicted value into bids under a target constraint.

Within each platform, we run either a second-price auction (SPA) or a first-price auction (FPA) with reserve r_k. We record the realized payment and winner for each auction. Aggregating across auction instances yields an expected payment per impression $p_k(M_k, r_k, \sigma_k)$ and an expected gross value (to the advertiser) among allocated auctions. Importantly, σ affects p_k even holding r_k fixed, because better signals change bid dispersion and hence expected clearing prices.

Cross-platform allocation is implemented at the campaign layer. At the beginning of a simulation run, each advertiser observes platform policies (M_k, r_k, σ_k) and then allocates a share of its auctions to each platform based on estimated net value. To parallel our baseline, we use a smoothed choice rule with an outside option:

$$\Pr\{\text{choose } k\} \;=\; \frac{\exp(\beta_0\, \tilde{u}_{ik})}{1 + \sum_j \exp(\beta_0\, \tilde{u}_{ij})}, \qquad \tilde{u}_{ik} := \tilde{V}_{ik}(\sigma_k) - \tilde{p}_k(M_k, r_k, \sigma_k),$$

where $\tilde{V}_{ik}(\sigma_k)$ and $\tilde{p}_k(\cdot)$ are advertiser i's beliefs (computed from signal models and historical outcomes). This step makes explicit what is implicit in practice: advertisers treat channels as competing options, and more precise measurement makes those comparisons sharper. We then *measure* the effective $\beta(\sigma)$ induced by the environment rather than imposing it, as described next.

8.2 Recovering $V(\sigma)$, $\phi(\sigma)$, Δ, and $\beta(\sigma)$

The simulation produces primitives (signals, bids, payments, allocations) from which we back out reduced-form objects.

Within-platform monetization $\phi(\sigma)$ and format wedge Δ. Fix participation and allocation shares, and run a single-platform counterfactual in which only σ varies. Define

$$\phi(\sigma) \;:=\; \mathbb{E}\big[p(\text{SPA}, r, \sigma)\big] - r,$$

so that $p(\text{SPA}, r, \sigma) = r + \phi(\sigma)$ by construction. We then define the format increment as

$$\Delta(\sigma, r) \;:=\; \mathbb{E}\big[p(\text{FPA}, r, \sigma)\big] - \mathbb{E}\big[p(\text{SPA}, r, \sigma)\big].$$

In many simulated environments $\Delta(\sigma, r)$ is approximately constant in (σ, r) over relevant ranges, validating the convenient reduced-form approximation $\Delta > 0$. When it is not constant, we can still compute the relevant local wedge at the equilibrium (r^*, σ^*).

Allocative value $V(\sigma)$. We define $V(\sigma)$ as the (allocation-weighted) expected gross value generated when the platform uses precision σ, holding fixed the advertiser population and the platform's market share. Operationally, we run a counterfactual in which each advertiser is assigned to the platform with exogenous probability (so that competition effects are held constant), and we compute

$$V(\sigma) \;:=\; \mathbb{E}\big[v \mid \text{allocated to platform}, \sigma\big],$$

where the conditioning includes the effect of σ on bidding and thus on which impressions are won. This definition captures the idea that better measurement improves match quality and conversion efficiency, even absent any strategic share shifting.

Effective sensitivity $\beta(\sigma)$. To map σ into cross-platform responsiveness, we estimate an *effective* logit slope as follows. For each σ (symmetric across platforms for measurement), we run small perturbations to a platform's reserve r (or equivalently to its expected payment p) and record the resulting change in demand shares D_k. Under the logit structure,

$$\ln\left(\frac{D_k}{D_0}\right) \;=\; \beta(\sigma)\, u_k + \text{constant},$$

so we can recover $\beta(\sigma)$ by regressing $\ln(D_k/D_0)$ on u_k across perturbations, where u_k is computed using the simulated $V(\sigma)$ and $p(\cdot)$. Intuitively, higher σ reduces the variance of posterior errors in m_{ik}, which makes advertisers more confident about small net-value differences across platforms, and this steepens the empirical choice curve. In our experiments, $\beta(\sigma)$ is increasing in σ across a wide range of primitives, providing a microfoundation for the assumed $\beta'(\sigma) > 0$.

8.3 Equilibrium computation and what the simulations verify

Given the mapping above, we compute platform best responses on a grid over (M_k, r_k, σ_k) and search for pure-strategy Nash equilibria (and, where needed, mixed strategies over formats). The primary objects we track are: equilibrium precision σ^*, equilibrium format M^*, equilibrium payment p^*, advertiser surplus, and total surplus.

Three regularities emerge.

First, we replicate the *format threshold logic*: when the estimated $\beta(\sigma)$ is high, the demand-share penalty from increasing effective price dominates the mechanical revenue gain from FPA, so SPA is revenue-dominant for each platform against an SPA rival and typically also against an FPA rival. Empirically, the simulated $\beta(\sigma)$ threshold at which SPA becomes dominant is close to the analytic β^* computed using the local Δ and p_S, which supports the interpretation of Proposition 1 as a useful rule-of-thumb.

Second, we replicate *strategic coarsening*: as we increase the responsiveness of allocation to measured ROI—either by tightening the RoS distribution (more types near the margin) or by reducing campaign-level frictions—the estimated slope $b = \beta'(\sigma)$ rises and the equilibrium σ^* falls. Importantly, this occurs even though $\phi(\sigma)$ and $V(\sigma)$ are increasing: platforms reduce precision not because it is technologically unproductive, but because it intensifies cross-platform substitution and erodes markups.

Third, we observe an *endogenous complementarity between transparency and "softer" formats*: when σ is forced upward exogenously (e.g., by a simulated reporting mandate), FPA becomes less attractive and equilibrium selection shifts toward SPA. This mirrors the policy intuition that transparency can discipline high-price mechanisms by making advertiser switching more salient and credible.

8.4 Robustness checks and limitations

We stress-test these findings along several dimensions that matter in applications. We vary the number of bidders n, the degree of common-value components (which affects the informativeness of rivals' bids), the correlation of values across platforms (which affects substitutability), and the presence of budgets and pacing. Budgets, in particular, dampen the outside-option margin and mechanically lower estimated $\beta(\sigma)$, weakening (but not reversing) the coarsening force: opacity is most profitable precisely when marginal dollars are flexible and re-optimized across channels.

We also replace additive Gaussian noise with discrete coarsening (bucketed conversion reporting) and obtain similar qualitative results, which is reassuring because many platform measurement policies are naturally discrete. Finally, we allow partial commitment by letting r_k adjust after σ_k is chosen; this tends to strengthen the incentive to keep σ low, because the platform anticipates being "held up" by future elasticity in every pricing subgame.

These exercises are not a claim of literal calibration. They abstract from learning dynamics, multi-objective platform design, and general-equilibrium feedback from advertiser product-market competition. Their role is narrower: to demonstrate that the reduced-form objects we use—$V(\sigma)$, $\phi(\sigma)$, Δ, and especially $\beta(\sigma)$—can be recovered from plausible auction primitives, and that the key qualitative predictions persist under operationally meaningful deviations from the knife-edge symmetric benchmark. With this mapping in hand, we can interpret policy and product interventions as shifts in σ, in the curvature $c(\sigma)$, or in the feasible dependence of pricing on signals, which we take up next.

9 Policy and product implications: transparency mandates, interoperability, prediction API design, and competitive effects of signal coarsening

Our framework is stylized, but it isolates a practical mechanism: information policies that improve advertisers' measurement and targeting (σ higher) also raise cross-platform responsiveness ($\beta(\sigma)$ higher), which disciplines pricing and makes high-price formats less viable. This complementarity between transparency and competition creates a wedge between what is privately optimal for platforms (strategic coarsening) and what may be desirable from the standpoint of advertiser surplus and allocative efficiency. In this section we draw out implications for regulation and for product design choices that, in practice, are often justified in technical terms but have clear competitive content.

9.1 Transparency mandates and reporting requirements

A natural intervention is to impose minimum reporting quality—for example, standardized conversion reporting, disclosure of key auction statistics, or restrictions on discretionary "bucketing" of performance signals. In the model, such policies operate as a lower bound on precision, $\sigma_k \geq \underline{\sigma}$, or equivalently as a reduction in the platform's effective cost of transparency $c(\sigma)$ by clarifying legal safe harbors and compliance standards. Two effects follow.

First, mandating higher σ directly raises advertisers' match value $V(\sigma)$ (better targeting and measurement) and may also increase within-platform monetization $\phi(\sigma)$ through tighter bids. Second, and more distinctively, it raises $\beta(\sigma)$, steepening the demand-share response to differences in net value. This latter effect pushes equilibrium away from high-effective-price choices, including the format premium Δ associated with first-price implementations. Thus transparency mandates can have an indirect *mechanism-design* consequence: even if a mandate does not explicitly regulate auction format, it can shift equilibrium selection toward lower-price (or lower-price-impact) formats by making switching more credible and more sensitive to ROI differences.

A key limitation is that platforms will typically re-optimize on other instruments. If σ is forced up, a platform may attempt to restore margins through reserves r_k, fees, or product-level bundling. Our reduced form highlights when such recoupment is difficult: when $\beta(\sigma)$ is high, any increase in effective price p_k is met with substantial share loss. In that region, the policy-induced increase in σ tends to translate into lower equilibrium prices and higher advertiser surplus rather than being fully offset elsewhere. In contrast, when advertisers are relatively insensitive (low β), transparency can raise $V(\sigma)$ without much competitive discipline, and much of the surplus may be appropriated by the platform through higher r or other charges.

Practically, this suggests that regulators should evaluate transparency rules jointly with frictions that determine effective responsiveness: multi-homing costs, contractual restrictions, and measurement comparability all move β. A reporting mandate in a market with high switching frictions may improve measurement while leaving market power largely unchanged; the same mandate in a low-friction environment can materially discipline pricing and format.

9.2 Interoperability, portability, and common measurement standards

Interoperability policies—data portability, common attribution standards, and API-level access that lets advertisers compare performance across channels—are often analyzed as reducing lock-in. In our model they can be interpreted as increasing β holding σ fixed: even if each platform's internal signal about value is unchanged, the *decision noise* in cross-platform allocation falls because comparisons become cleaner and more trusted. Put differently, interoperability shifts the market toward the high-sensitivity region in which per-impression prices are more tightly disciplined.

This observation has two implications. The first is positive: interoperability can substitute for direct transparency regulation by increasing the effective elasticity that deters high-price formats and excessive markups. The second is strategic: when platforms anticipate that interoperability raises β, they may respond by reducing σ (garbling) to partially restore slack. This is exactly the strategic coarsening force: if higher σ increases $\beta(\sigma)$, and higher β intensifies competition, then platforms have an incentive to reduce the informativeness of the signals they control. Interoperability thus creates an endogenous "arms race" between policies that make platforms more comparable and platform choices that make their own performance harder to verify or interpret.

From a design perspective, interoperability requirements are therefore most effective when they pin down not only access, but also *semantics*: definitions of conversions, treatment of delayed attribution, and error bars or uncertainty quantification. Standardization reduces the scope for strategic coarsening that is formally "compliant" but economically obfuscatory. In the language of our reduced form, standardization effectively constrains the mapping from technical reporting to σ, preventing platforms from relabeling lower precision as a benign product change.

9.3 Prediction and measurement API design as an information policy

Many of the most consequential platform choices are not framed as "transparency" at all, but as engineering decisions: how granularly to report conversions, whether to provide user-level vs aggregated signals, whether prediction APIs expose calibrated probabilities or only coarse scores, how quickly signals are delivered, and whether privacy-preserving noise (e.g., differential privacy) is added. Each of these choices changes the Blackwell informativeness of the advertiser's signal structure $S(\sigma)$, and thus moves σ.

The model clarifies a tradeoff that product teams face but rarely articulate in competitive terms. Increasing σ can raise $V(\sigma)$ and $\phi(\sigma)$—advertisers bid more effectively and auctions may clear at higher expected payments—but it can also raise $\beta(\sigma)$, eroding margins through increased substitutability across platforms. Consequently, a platform with market power may rationally choose an API design that is *less informative than technologically feasible*, even absent any privacy cost. Conversely, when competitive pressure is strong, platforms may supply high-precision measurement as a competitive differentiator, but only up to the point where the induced elasticity makes further precision privately unprofitable.

This perspective suggests concrete diagnostics. If an API change reduces reporting granularity or increases noise while being justified as "simplification," we should expect (i) reduced cross-platform reallocation in response to price or format changes (a drop in estimated β), and (ii) increased ability to sustain higher effective prices p or higher-price formats. Conversely, exogenous increases in precision—for instance, mandated disclosure of calibrated conversion probabilities or standardized incrementality metrics—should predictably (a) steepen demand response, and (b) shift equilibria toward lower-price-impact formats, consistent with the format threshold logic.

9.4 Competitive effects of signal coarsening and the boundary with exclusion

Strategic coarsening has a competitive interpretation that is distinct from classic price-setting. By reducing σ, a platform makes advertisers less able to rank channels by true ROI, weakening discipline from multi-homing and thereby relaxing competition. This resembles a form of nonprice obfuscation: the platform does not prohibit switching, but it makes switching less informed. In settings where advertisers rely on platform-provided measurement to satisfy RoS constraints, lower σ can also shift spend toward the incumbent simply because the autobidder cannot confidently validate alternatives.

This raises a policy question: when should coarsening be treated as a legitimate privacy or user-experience choice (captured by a higher $c(\sigma)$), and when is it best viewed as an anticompetitive degradation of comparability? Our model does not adjudicate intent, but it suggests an evidence-based approach. One can estimate $\beta(\sigma)$ and its change following reporting modifications using quasi-experimental variation (reserve changes, fee changes, or format changes) and test whether reduced precision measurably lowers cross-platform responsiveness. If it does, then a reporting change has competitive impact even if it leaves average conversion counts unchanged.

A related implication concerns mergers or exclusivity. Consolidation that reduces the number of outside options effectively reduces contestability and can lower the competitive cost of raising σ, potentially increasing precision while also increasing markups. This ambiguity underscores that transparency is not a monotone proxy for competition: a dominant platform may be quite transparent about outcomes and still sustain high prices if advertisers have few credible alternatives. In our notation, high σ does not guarantee high β when alternatives are limited or highly differentiated.

9.5 Implementation guidance and limitations

Two practical recommendations follow from the reduced-form mapping. First, transparency and interoperability policies should be paired with monitoring of behavioral responsiveness, i.e., tracking β (or elasticity of spend shares with respect to net value) rather than focusing exclusively on disclosure checklists. Second, when privacy concerns necessitate noise addition, regulators and platforms can aim for *structured* transparency: preserving the information that is most important for cross-platform comparability (which drives efficient allocation) while allowing coarsening along dimensions that are privacy sensitive but less central for ROI ranking. In terms of $S(\sigma)$, this is a move from scalar σ to multi-dimensional precision, a natural extension in which only some components of information raise β sharply.

Finally, we emphasize what our model does not capture. Real systems feature learning dynamics, budget smoothing, and multi-objective platform constraints; these can dampen or delay the elasticity channel. Nonetheless, the core prediction is robust: whenever improved measurement makes advertiser substitution more sensitive to small net-value differences, platforms face a private incentive to limit informativeness, and policies that increase comparability can indirectly discipline pricing and format. This is precisely why "measurement" and "auction design" should be treated as jointly competitive choices rather than as separable engineering decisions.