
THE POWER DIAGRAM AUCTION: A FORMALLY VERIFIED VCG MECHANISM FOR LLM ADVERTISING

June Kim
Independent Researcher
june@june.kim
ORCID 0009-0005-3153-9396

July 6, 2026

ABSTRACT

Large language model (LLM) chat assistants are becoming an advertising surface. The previous surface, keyword search, shipped with broken incentives, and an industry grew to monetize the breakage; the new surface is a chance to get the incentives provably right on day one. The conversation embeds as a point x in a continuous space of hundreds of dimensions, and each advertiser declares a center c (their customer). The conceptual distance between them measures how closely they match. A reach σ (how widely they serve) and a bid b (what a conversion is worth) compare that distance against a willingness to pay. The slot beside the reply goes to the advertiser with the highest $\log(b) - \|x - c\|^2 / \sigma^2$. The mechanism involves no model training or response modification. We prove in Lean 4, with zero sorry, that the rule’s argmax allocation with Clarke payments forms a VCG mechanism. Truthful reporting is weakly dominant, the allocation maximizes welfare at every query point, and its territories form a power diagram with keyword auctions as the degenerate case. This mechanism allows an auction to sell regions of embedding space.

1 One-Shot Bidding

Keyword auctions are strategically broken, and the industry monetizes the breakage. [Edelman, Ostrovsky and Schwarz \(2007\)](#) proved that Google’s Generalized Second-Price auction has no dominant-strategy equilibrium, so your optimal bid depends on bids you cannot see. First-price display auctions require shading by an amount that also depends on bids you cannot see. An autobidding industry answered, agents running machine learning against each other to approximate what [Vickrey \(1961\)](#) made exact sixty-five years ago in one shot: report your value, pay the externality, done.

Ad auctions for large language models (LLMs) restart the design problem. A user’s conversation embeds to a point x in a real inner product space, in practice a few hundred dimensions. Each advertiser declares a center c (who their customer is), a reach σ (how wide a neighborhood they serve), and a bid b (what a conversion is worth). The platform scores each advertiser at x and the highest score wins. The scoring rule predates this paper.¹ An [open-source exchange](#) implements it, [multi-agent simulations](#) probe its market dynamics, and an [interactive explorer](#) renders the allocation live. The mechanism lacked a proof that it deserves the trust the proposal asks for.

We supply the proof in Lean, so that trust in the mechanism reduces to [running a build command](#). The contribution is one bridge lemma, `score_eq_log_reportedVal`: the scoring rule is the logarithm of the value a report implies, unconditionally. Everything downstream is the classical VCG argument of [Vickrey \(1961\)](#), [Clarke \(1971\)](#), and [Groves \(1973\)](#), executed formally. We also prove that the allocation is a power diagram for arbitrary heterogeneous reaches, and that the mechanism collapses to Vickrey’s sealed-bid second-price auction at any keyword point.

2 Model

An advertiser’s report is a triple (c, σ, b) :

¹Proposal series: [june.kim/power-diagrams-ad-auctions](#), [june.kim/the-price-of-relevance](#), [june.kim/keywords-are-tiny-circles](#). Implementation: [github.com/kimjune01/vectorspace-adserver](#). Simulations: [june.kim/relocation-fees](#), [june.kim/relocation-fee-dividend](#). Demo: [june.kim/vectorspace-ads](#), archived as DOI 10.5281/zenodo.21365798.

Field	Name	The advertiser declares
$c \in E$	center	who their customer is, a point in embedding space
$\sigma > 0$	reach	how wide a neighborhood around c they serve
$b > 0$	bid	what a conversion is worth to them

Their private valuation is a triple (c^*, σ^*, v) of the same shape, and each report field may differ from its starred counterpart. The distinction is operational. Reported distance controls which impression the advertiser wins; true distance controls whether the person who sees the ad converts. A supplement advertiser can declare its center on running-shoe queries, but the declaration does not turn supplements into shoes. Its true value at query x remains Gaussian around its private center:

$$\text{trueVal}(x) = v \cdot \exp(-\|x - c^*\|^2 / \sigma^{*2})$$

True value reads as margin times conversion rate: v is the margin a conversion clears, and the Gaussian factor is the probability that the lead at x converts. σ is a length scale: the Gaussian is unnormalized, and the exponent carries no $1/2$. The Gaussian is strictly positive: even a badly matched impression can convert, only with a probability that decays rapidly with semantic distance. That is the object observed in impression advertising, not a binary relevance label.

The Gaussian family is an approximation, and the best available one. Embedding coordinates are not human-interpretable, so a peak and a width are all an advertiser can meaningfully declare, and the Gaussian is the [maximum-entropy](#) model of a conversion curve known only by those two. Any other decay imports structure for which the advertiser has no evidence. The theorems hold inside the family, and the Limits section says what stays outside it. The platform scores reports by

$$\text{score}(x) = \log(b) - \|x - c\|^2 / \sigma^2$$

and allocates each query to an argmax of score. Thus allocation uses the reported Gaussian while utility uses the fixed true Gaussian. Changing a report changes the impressions won, never the underlying conversion surface. The winner pays the Clarke pivot. With a single winner per query, the pivot has a closed form: losers pay nothing, and the winner pays the runner-up's implied value (zero when unopposed).

$$\text{payment}_w(x) = \max_{j \neq w} b_j \cdot \exp(-\|x - c_j\|^2 / \sigma_j^2)$$

Utility is quasilinear. Impressions are opt-in opportunities. The model does not force an advertiser into an auction it declined to enter. A publisher may separately restrict eligibility, but that policy defines the player set upstream and is not a second allocation mechanism in the model. All definitions follow Nisan's mechanism-design chapter in [Nisan, Roughgarden, Tardos and Vazirani, eds. \(2007\)](#).

The scoring rule is the embedding-space member of a known family. [Lahaie and Pennock \(2007\)](#) analyzed keyword scoring rules of the form $\text{bid} \times \text{quality}^s$. With the score's logarithm taken in base β , ranking by score is ranking by $b \cdot q^{\ln(\beta)}$ with quality $q = \exp(-\|x - c\|^2 / \sigma^2)$. The log base is therefore the squashing parameter, $s = \ln(\beta)$, and the squashing-parameter literature transfers intact. Sweeping the log base is the platform's revenue-relevance dial.

The formalization states everything over an arbitrary real inner product space. No finite-dimension hypothesis appears anywhere in the chain. The theorems hold for a 384-dimensional sentence embedding and for next year's wider one alike. Drawn in one and two dimensions, the figures show the structure that survives the trip up: cells, boundaries, argmax. Their proportions do not survive it. In three hundred dimensions distances concentrate and most of a cell's volume sits near its boundary. Pictures illustrate; theorems generalize.

3 The Bridge

The mechanism-design results hang on one identity. Define the value implied by a report, $\text{reportedVal}(x) = b \cdot \exp(-\|x - c\|^2 / \sigma^2)$. Then

$$\text{score}(x) = \log(\text{reportedVal}(x))$$




Figure 1: The allocation is the upper envelope of the reported-value curves (their logs are the score parabolas). Winning intervals along the axis are the power-diagram cells. A bid sets each curve’s height, a center its position, σ its width.

with no truthfulness hypothesis (`score_eq_log_reportedVal`). The disguise is a logarithm. Log is monotone, so the argmax of score is the argmax of reported value. The mechanism always maximizes reported welfare, which is the allocation rule VCG requires. When a report is truthful, reported value equals true value (`reportedVal_eq_trueVal_of_truthful`), and the same identity becomes `score = log(trueVal)` (`score_eq_log_trueVal`). The winner is the advertiser who values the impression most (`winner_maximizes_welfare`).

4 Weakly Dominant Strategy

4.1 The Incentive Theorem

`vcg_dsic` is the main incentive theorem: for any deviation report r' , truthful reporting yields at least the utility of reporting r' , regardless of what every other player reports. The Clarke payment depends only on others’ reported values, so a player’s report moves their allocation and never their price schedule (`welfareOthersWithout_invariant`); the four-case comparison then closes the argument. A false center can make the supplement advertiser win the running-shoe impression, but its utility still contains the small true supplement-conversion value, while its payment equals the value of the shoe advertiser it displaced. In the figure’s instance, the supplement advertiser S ’s books at the shoe query x :

S ’s ledger at x	truthful	false center
allocation	loses x	wins x
earns (true value at x)	0	0.08
pays (Clarke pivot)	0	1.25
utility	0	-1.17

With one rival the pivot is the rival’s reported value, so a winner’s utility is its maximum willingness to pay at x less what the competitor’s report says the impression is worth. The accounting generalizes: the same four rows settle any deviation, in any dimension, against any number of rivals, in one shot. `vcg_dsic` is this table quantified over all of them.



4.2 Observational Equivalence

Weak dominance does not by itself identify the literal report. In a thin market, an unopposed winner pays zero, and every center report that preserves its allocation gives the same utility. [The artifact](#) proves the limiting case explicitly, with pointwise and expected utility invariant under every report when only one player competes (`playerUtility_invariant_of_subsingleton`, `expectedPlayerUtility_invariant_of_subsingleton`).

Competition closes that gap. A center deviation that crosses a contested allocation boundary sacrifices true surplus and is strictly worse (`vcg_strict_at_contested_point`), and when the disagreement covers a positive-weight region the loss survives aggregation (`vcg_strict_expected_on_allocation_change_set`). Observational equivalence remains: reports that preserve allocation where competition matters. An expected-utility tie cannot place positive query weight on queries where the allocations disagree and the deviation is strictly worse (`expected_utility_tie_disagreement_not_positive`). Finite weighted query logs construct the required strictness witness whenever that set contains a logged query of positive weight.

4.3 Spatial Competition

The strictness theorems answer the most immediate Hotelling objection. In spatial competition from [Hotelling \(1929\)](#) onward, sellers drift toward the center of the market, and the analogue in embedding space is advertisers declaring centers near high-traffic regions. VCG makes truthful center reporting weakly dominant, so Hotelling drift cannot increase advertiser utility. Competition converts false spatial proximity from a costless declaration into a strictly costly allocation error. A reported center changes placement but not customer fit: the supplement advertiser from that ledger still earns its small true value while paying for the shoe advertiser it displaces. As rival coverage thickens, the allocation-equivalent class can contract toward the advertiser’s true center. Under the artifact’s explicit spatial-coverage condition, every selection from that surviving class converges to `c*` (`centers_converge_of_competitive_coverage`). The convergence result is conditional rather than a claim about adjustment dynamics, and profitable drift can return through what the model excludes: budgets, volume-dependent objectives, endogenous products or creatives, and non-Gaussian true values.

4.4 Comparison with Alternative Mechanisms

The qualifier earns its value in comparison. GSP gives keyword bidders no dominant strategy ([Edelman, Ostrovsky and Schwarz, 2007](#)), and even an exact second-price keyword auction makes only the bid truthful: keyword selection and match types remain strategic choices made against unobserved rivals, so targeting sits outside the incentive guarantee. In the present mechanism targeting is part of the report, and the guarantee covers it. The same comparison orders the LLM mechanisms:

Mechanism	Truthful report	Allocation	Welfare	Proof
GSP (keywords)	<i>none exists</i>	deterministic	<i>unguaranteed</i>	—
Vickrey (keywords)	bid only; <i>targeting</i> <i>strategic</i>	deterministic	<i>per conflated bin</i>	classical
MOSAIC	reward function	<i>randomized softmax</i>	<i>approximate</i>	paper
Retrieval-time	<i>bid only</i>	<i>probabilistic draw</i>	<i>unguaranteed</i>	paper
LLM-Auction (training)	<i>no report to state</i>	<i>weights, opaque</i>	<i>unstated</i>	<i>none</i>
Power diagram	bid, center, reach	deterministic argmax	exact, pointwise	Lean, zero sorry

Italics mark the concessions.

5 Geometry

The name “power diagram auction” is now a theorem: in E when advertisers share a reach, and in $E \times \mathbb{R}$ when they do not.

5.1 Equal Reach

When two advertisers share σ , comparing scores at x is comparing power distances $\|x - c\|^2 - w$ with sites at the centers and weights $w = \sigma^2 \log(b)$ (`score_le_iff_powerDist_le`). The score difference is an affine function of x (`score_sub_affine`), so cell boundaries are hyperplanes. That is the classical power diagram of [Aurenhammer \(1987\)](#), with bids setting the weights. When every advertiser shares σ , the winner rule is that cell assignment in E directly: the argmax of score minimizes power distance (`winner_minimizes_powerDist`). No lift required.

5.2 Heterogeneous Reach

With distinct σ s the in-space boundaries curve, and the folk description of the allocation as a power diagram fails in E . It succeeds in $E \times \mathbb{R}$. Lift each query to the paraboloid $(x, \|x\|^2)$. Each report (c, σ, b) determines a lifted site and weight under which maximizing score is minimizing lifted power distance. The auction’s winner rule is then the lifted diagram’s cell assignment (`liftedPowerDist_paraboloid`, `score_le_iff_liftedPowerDist_ge`, `winner_minimizes_liftedPowerDist`; the site and weight formulas are in `PowerDiagram.lean`). Aurenhammer proved diagrams of quadratic distance functions lift this way; the formalization instantiates his lift for this scoring rule and checks the algebra once. [The explorer](#) renders the cells live. Exact $O(\log N)$ point location is a fixed-low-dimension result and does not survive embedding dimension, but the lift still pays at serving time. Score comparisons are affine in the lifted query $(x, \|x\|^2)$. So finding the winner is a maximum-inner-product search against static lifted sites, which approximate-nearest-neighbor indexes and cache layers serve sublinearly in practice, with the $O(N)$ scan as the exact fallback.




Figure 3: Five advertisers with unequal reach. Tight specialists carve bubbles out of wider generalists, and adjacent cells meet in circular arcs. The boundaries are flat only upstairs: quadric bisectors in E are hyperplanes in $E \times \mathbb{R}$.

6 Keywords Recovered

The migration path follows from the geometry: a keyword is a tiny circle, a report whose reach has shrunk to a point. Keyword auctions are the degenerate case of the embedding auction, so the general mechanism is backward compatible: adopting it strands no existing buyer. The claim is now a theorem pair.

6.1 The $\sigma \rightarrow 0$ Limit

At any point other than its center, a report’s score diverges to $-\infty$ as $\sigma \rightarrow 0$, while its score at the center is $\log(b)$ independent of σ (`keyword_is_degenerate_limit`, `score_at_center`). The tiny circle collapses to its point.

6.2 Exact Vickrey

At a query point that is every advertiser’s shared center, the Gaussian factor is $\exp(0) = 1$ for every bidder, so reported value is the bid. The winner is a highest bidder (`winner_maximizes_bid_of_common_center`), the Clarke pivot equals the highest competing bid (`vcgPayment_common_center_second_price`), and every loser pays nothing (`vcgPayment_eq_zero_of_loser`). That is Vickrey’s sealed-bid second-price auction, allocation and payment both. A keyword auction is what the embedding auction does at a point.

7 Optimality

Pointwise welfare maximization integrates. For any measure over queries, expected welfare under the score-argmax allocation weakly dominates expected welfare under any rule the expectation operator evaluates, power-diagram or not (`integral_efficiency`, `gaussian_optimality`; the theorem statements carry the exact quantifiers).




Figure 4: Shrinking σ at a fixed center against a wide rival with equal bid. The value at the center is b for every σ , while territory off-center collapses; in the limit the advertiser competes at one point only, on bid alone.

`gaussian_vcg_weakly_dominates` conjoins the three properties: welfare-optimal, dominant-strategy incentive compatible, equilibrium-efficient. The last is the welfare guarantee evaluated at truthful play, which `vcg_strategyproof` certifies as a Nash equilibrium. Inside the game, a truthful report weakly dominates every deviation (`vcg_dsic`); across allocation rules, the mechanism’s expected welfare weakly dominates every rival the operator evaluates. The same order certifies the advertiser’s best strategy and the platform’s best mechanism.

The artifact separately proves the equilibrium-decomposition theorem of Ghani, Hedges, Winschel and Zahn (2018). In their open games, where large games compose from smaller ones, an equilibrium of the composite induces equilibria of the components (`composed_equilibria_decompose`). Whether DSIC composes the same way remains open, and the formalization records it as open instead of assuming it.

Welfare, as the `Model` defines it, means willingness to pay. A true value $b \cdot \exp(-\|x - c^*\|^2 / \sigma^{*2})$ is what a conversion at query x is worth to the advertiser, so pointwise welfare maximization routes every query to the bidder willing to pay the most there, and the incentive theorems keep the declared curves honest about where that willingness lives. The guarantee is also dimension-free, and the keyword language is not. A keyword is a point report whose territory vanishes in high dimension, so covering an embedding space with keywords takes exponentially many reports, while a Gaussian covers a region with three numbers.

The weak-order chain uses one modeling assumption, `QueryMeasure.integral_mono`: a monotone expectation operator respects pointwise inequality (no normalization or probability mass required). Strict expected comparisons additionally carry `PositiveAt` or `PositiveOn`, stating that the operator detects a contested point or region. These are structural hypotheses, not Lean axioms. `QueryMeasure.dirac` and `QueryMeasure.ofWeightedFinset` construct the operator. A finite weighted query log constructs the strictness witness whenever a contested region contains a logged query of positive weight. A measure-theoretic integral satisfies the corresponding interfaces on functions and positive-mass regions it can evaluate. Keeping the operator abstract lets the theorems quantify over query distributions instead of fixing one, and each theorem carries the exact strength it uses.

8 Related Work

8.1 LLM Ad Auctions

Feizi, Hajiaghayi, Rezaei and Shin (2025) survey the design space as four modules (modification, bidding, prediction, auction) and leave open how the pieces sustain an ad ecosystem. The idea of an auction inside the generation begins with Dütting, Mirrokni, Paes Leme, Xu and Zuo (2024), who aggregate advertisers’ token distributions into the reply and price the result through a monotonicity condition. Dubey, Feng, Kidambi, Mehta and Wang (2024) auction placements inside an LLM-written summary of the ads. MOSAIC and the retrieval-time mechanism are later points on that line; all of them keep the ad inside the text the model emits.




Figure 5: Simulated conversion values against distance from the true center, drawn along two directions of embedding space. Preferences need not be Gaussian for the model to be right about what a report can say: a peak and a width are all an advertiser can declare, and the isotropic Gaussian, equal weight in every dimension, is the maximum-entropy curve those two numbers justify.

[Soumalias, Curry and Seuken \(2025\)](#) give the prior LLM ad mechanism with dominant-strategy guarantees. Their MOSAIC auction allocates the reply itself: advertiser reward functions score candidate responses sampled from the model, and softmax selects one. Computing the exact welfare-maximizing reply is out of reach, so exact VCG is unavailable. The softmax allocation is instead cyclically monotone, and [Rochet \(1987\)](#) turns cyclic monotonicity into dominant-strategy payments through a convex potential. Auctioning an embedded intent point instead of the reply shrinks the outcome space from token sequences to a finite set of known advertisers. Exact VCG becomes one argmax plus one counterfactual argmax, and the guarantees become simple enough for Lean to check. Allocating a separate ad object also removes MOSAIC’s deployment costs: no multi-candidate generation per query, deterministic winners an advertiser can budget against, and a discrete placement that can be attributed and priced.

[Hajiaghayi, Lahaie, Rezaei and Shin \(2024\)](#) place the auction at retrieval time instead: their mechanism retrieves ads probabilistically per discourse segment, by bid and relevance, with incentive-compatible pricing. The outcome space is far smaller than MOSAIC’s, but the placement still lives inside the generated text, and the winner is a draw rather than an argmax. The geometric mechanism keeps the retrieval-time simplicity while making the allocation deterministic and the ad a separate object.

[Alaei, Makhdomi and Malekian \(2026\)](#) share the geometric premise: a user is an unknown preference vector, each advertiser a feature vector, match value their inner product. Their platform learns the vector as the conversation proceeds, trading a sharper estimate against the risk the user leaves. The static [Model](#) holds that learning dynamic fixed, and fixing the query point and the Gaussian family buys the closed-form diagram and the compiled proof.

The other family bakes advertising into the model by training. Alibaba’s [LLM-Auction](#) post-trains the model with preference alignment to balance ad revenue against user experience, encoding commercial incentives directly in the weights. Training is the wrong cadence for an ad market: campaigns change hourly, fine-tuning takes weeks, and retraining per campaign rotation costs more than the campaign.

It is also the wrong trust surface. Outsiders cannot audit bias learned in weights, and users cannot police it from inside either: [evaluators shown chatbot responses with embedded ads](#) failed to detect the ads and preferred the responses that contained them. The power-diagram mechanism is the opposite limiting case. The scoring rule is public, the allocation is its argmax, the incentive guarantee is a compiled theorem, and the audit reduces to a build command.

8.2 Semantic Matching

Advertisers competing in a vector space is not new. [Grbovic et al. \(2016\)](#) learned joint query and ad embeddings to match ads beyond literal keywords, and embedding-based retrieval is now standard ad-system infrastructure. In that

line the geometry supplies candidates or relevance features to a separate auction. In the power diagram the distance function is the valuation, and the auction is the geometry.

8.3 Formalized Auctions

Caminati, Kerber, Lange and Rowat (2013) proved soundness of combinatorial Vickrey auctions in Isabelle and generated verified executable code. Barthe, Gaboardi, Gallego Arias, Hsu, Roth and Strub (2015) verified incentive compatibility for VCG in a probabilistic relational calculus, and Jouvelot and Gallego Arias (2022) built `mech.v`, a Coq library that proves VCG’s incentive properties generically and refines an online-advertising VCG down to that specification. Formalized VCG is not new, and neither is a verified online-ad mechanism. What [this artifact](#) adds is the object: the verified allocation is a partition of a vector space, and the same artifact that checks truthfulness checks its identification as a power diagram. Kerber, Lange and Rowat (2016) formalized Vickrey’s auction in Isabelle and argued mechanized reasoning should be ordinary practice in economic theory; Ghani, Hedges, Winschel and Zahn (2018) rebuilt game theory compositionally so that mechanisms assemble from smaller games. The artifact extends the practice from a discrete allocation rule to a geometric one.

8.4 Filters and Reserves

Hartline, Hoy and Taggart (2023) show competitive efficiency survives reserve pricing. The same holds trivially for any pre-auction relevance filter, since welfare maximization restricted to a nonempty eligible set is welfare maximization on that set (`tau_preserves_efficiency_among_eligible`). Revenue also lives in reserves. The Clarke pivot collects only the runner-up’s implied value, zero when a winner is unopposed, and a reserve puts a price floor under exactly those impressions.

9 Limits

The theorems end where the model does, and the model is deliberately narrow.

Element	In the formalization
Embedding space E	any real inner product space; dimension unconstrained
Advertisers	arbitrary finite type; no bound on the count
Query distribution	universally quantified expectation operator; one assumption, monotonicity
Utility	quasilinear; no budget constraints
True valuations	isotropic Gaussian family
Deviations	any (center, σ , bid) triple
Rival allocation rules	any rule the expectation operator evaluates
Winners per query	exactly one; ties broken by fixed enumeration

Single winner. Slates, pacing, and cross-impression externalities are outside the model.

No budgets. Budget-constrained clearing has known structure (with fixed budgets it is [semi-discrete optimal transport](#), whose solution is again a power diagram), but its incentives and dynamics are open. [The simulations](#) suggest the dynamics are where the difficulty lives.

Gaussian truth. The Gaussian family is a bidding language, and every deployed auction clears one. The reigning language is the keyword, the $\sigma \rightarrow 0$ point mass of this same family. A bidding language buys expressiveness against clearing cost, the central tradeoff of the combinatorial-auction literature (Nisan, 2000). On that frontier the isotropic Gaussian is the most expressive language currently known to admit sublinear clearing, exact VCG, and a compiled proof. Anisotropic preferences (elliptical rather than spherical reach) stay computable at $O(N)$ per query but break the bridge lemma as stated; whether a lifted variant survives for quadratic-form preferences is the natural next theorem, since Aurenhammer’s lift already accommodates general quadrics. Mixtures turn the score into a log-sum-exp whose bisection admits no fixed-dimension lift, and the clearing speed and the proof leave with the geometry.

Frozen embedder. We hold the embedding model fixed. Retraining or replacing it moves every conversation point and every declared center at once, invalidating reports with no advertiser misreporting anything; re-declaration cadence is an operational question the model does not pose.

Statics only. Nothing in the artifact says bidding dynamics converge. [The simulations](#) study the dynamics empirically; a formal convergence or limit-cycle result would be a separate paper’s contribution.

10 Discussion

Four implications sit outside the theorems but follow from their shape; none is compiled, and each names its evidence. Together they answer why now: the agent literature prices coordination toward zero, that collapse runs through discovery, and this mechanism offers diversity in the firms that can be discovered.

Positioning becomes the bidding language. A report is a positioning statement made formal: c is who the customer is, σ how far the claim extends, b what the fit is worth. Keywords made both sides translate positioning into a vocabulary neither thinks in. In [Keywords Recovered](#) the keyword auction embeds as one point of the general mechanism, and `vcg_dsic` covers all three fields, so no misdeclaration profits and the allocation-changing ones cost (the series argues the sociological case). Richer bidding languages have paid before: [Sandholm \(2007\)](#) cleared \$35 billion of sourcing through expressive combinatorial bids and reported \$4.4 billion of hard-dollar savings. That evidence comes from procurement, and advertising is the untested extension.

The keyword tax shrinks. [Levin and Milgrom \(2010\)](#) call it conflation, pooling heterogeneous items into one auction. A keyword bin is conflation at the item-definition layer, where the climbing specialist pays to compete on pelvic-floor queries it will never convert. Heterogeneous reach is the un-conflation: a tight- σ specialist wins where its reported value leads without submitting a separate bid for every nearby keyword. The geometry section proves the carving; the simulations measure the surplus it recovers in a synthetic fifteen-advertiser market, and the bar those measurements have to clear is Levin and Milgrom’s own counterweight: that finer segmentation can thin markets and weaken price discovery. The ecological hypothesis: a finite keyword vocabulary routes demand only to businesses it has words for, while a continuous positioning space can host any firm a center and a reach can describe. So the richer language admits a more diverse population of discoverable firms. The pattern has retail precedent: when online catalogs lowered discovery costs, [Brynjolfsson, Hu and Smith \(2006\)](#) measured demand shifting into the niche tail.

The minimum viable firm shrinks. [Coase \(1937\)](#) set the boundary of the firm where the cost of transacting outside it exceeds the cost of organizing inside. A recent literature asks whether AI agents collapse those costs market-wide: [Shahidi, Rusak, Manning, Fradkin and Horton \(2025\)](#) call the limit the Coasean singularity. That collapse runs through discovery: small firms survive by specialization, so it needs an ad market where specialists are viable against large budgets. Keyword advertising is a scale game of vocabulary research, bidding infrastructure, and autobidding retainers that amortize only over large budgets, and training-based placement gates entry on the platform’s fine-tuning calendar. The power-diagram mechanism replaces the keyword portfolio with one report per campaign; a deployment would measure whether it lowers the minimum profitable campaign size and shifts entry toward smaller firms.

The collapse runs on elicited information. Cheap coordination does not supply the information coordination needs. Dispersed knowledge is elicited, never given ([Hayek \(1945\)](#)), and advertising is the market that produces the who-serves-whom half of it; [Nelson \(1974\)](#) founded the economics of advertising on its information content. The stakes run to the whole margin: the Model prices true value as margin times conversion rate, so the ceiling on what an advertiser will pay for discovery is everything the sale clears, and the mechanism sets which share changes hands. The incentives of the mechanism that clears it set the quality of that information: conflated auctions produce conflated knowledge, weights produce knowledge nobody can audit, and in the power diagram, truthful reporting is the compiled incentive. An agent-mediated market inherits the incentive properties of whatever ad mechanism its discovery clears through.

The formal agenda: whether DSIC composes across open games (recorded open in the artifact, with the equilibrium half proved, `composed_equilibria_decompose`), whether the bridge lemma reformulates for anisotropic preferences, and whether bidding dynamics converge once budgets enter the mechanism. The empirical agenda: what it costs an advertiser to estimate (c, σ, b) , whether the richer bidding language cuts campaign setup cost and misallocation, whether specialist surplus, market thickness, and small-firm entry hold up together, and whether allocations stay stable after embedder recalibration. The artifact verifies the mechanism as modeled; these questions ask whether the model describes a market worth running.

Artifacts

Formalization

The formalization is at github.com/kimjune01/auction-proof, archived as DOI 10.5281/zenodo.21214697, AGPL-3.0, Lean 4 (v4.29.0-rc6) with Mathlib. Verification:

```
lake exe cache get
lake build
```

Zero sorry, and no trusted axiom of our own. `#print axioms` on every theorem below returns only Lean’s three standard background axioms (`propext`, `Classical.choice`, `Quot.sound`), because the single modeling assumption is a `QueryMeasure` typeclass that concrete instances discharge by proof. The claims-to-theorems map:

Claim	Lean theorem	File
score = log(reported value)	score_eq_log_reportedVal	Efficiency.lean
winner maximizes true welfare	winner_maximizes_welfare	Efficiency.lean
truthful reporting is dominant	vcg_dsic	Strategyproof.lean
no allocation rule beats VCG	gaussian_optimality	GaussianOptimality.lean
capstone conjunction	gaussian_vcg_weakly_dominates	GaussianOptimality.lean
equal- σ power diagram, hyperplane bisectors, winner minimizes power distance	score_le_iff_powerDist_le, score_sub_affine, winner_minimizes_powerDist	PowerDiagram.lean
variable- σ power diagram via paraboloid lift	liftedPowerDist_paraboloid, winner_minimizes_liftedPowerDist	PowerDiagram.lean
keyword limit	keyword_is_degenerate_limit	VectorSpace.lean
exact Vickrey at a keyword point	vcgPayment_common_center_second_price	SecondPrice.lean
losers pay nothing	vcgPayment_eq_zero_of_loser	Strategyproof.lean

Exchange

[vectorspace-adserver](#) implements the mechanism as an open-source exchange, [live at vectorspace.exchange](#) and archived as [DOI 10.5281/zenodo.21365911](#): the scoring rule, the argmax allocation, and the Clarke pivot.

Simulations

[openauction/cmd/simulate](#), archived as [DOI 10.5281/zenodo.21366035](#), runs the market on real [BGE-small-en-v1.5](#) embeddings: fifteen advertisers, 62 queries, 50 randomized trials, keyword and embedding clearing compared head to head. [Keyword Tax](#) and [Relocation Fees](#) report the results.

Explorer

[vectorspace-ads](#) renders the allocation live, archived as [DOI 10.5281/zenodo.21365798](#): drag a center or raise a bid and the territories shift.

Acknowledgments

Conversations with Sébastien Lahaie and Mohammad Hajiaghayi sharpened the mechanism-design framing, and Sébastien pointed me to MOSAIC. Errors, and the claims, are mine.

Disclosures

LLM use. This paper and its artifact were produced with large language models via [Claude Code](#). The author directed the research and reviewed every claim; agents wrote the Lean proofs, generated the figures, and drafted the prose. A separate model ran an adversarial review of the prose against the formalization. No guarantee in this paper rests on any model’s judgment: the artifact type-checks with zero sorry, and verification reduces to lake build.

Funding. Self-funded independent research. No external funding, no employer direction, no advertiser or platform relationship.

References

- Alaei, S., Makhdoumi, A., and Malekian, A. (2026). Dynamic Learning and Optimal Advertising Mechanism for LLM Platforms. [SSRN working paper](#).
- Aurenhammer, F. (1987). Power Diagrams: Properties, Algorithms and Applications. [SIAM Journal on Computing](#) 16(1), 78–96.
- Barthe, G., Gaboardi, M., Gallego Arias, E. J., Hsu, J., Roth, A., and Strub, P.-Y. (2015). Computer-Aided Verification in Mechanism Design. [arXiv:1502.04052](#).
- Brynjolfsson, E., Hu, Y. J., and Smith, M. D. (2006). From Niches to Riches: Anatomy of the Long Tail. [Sloan Management Review](#) 47(4), 67–71.
- Caminati, M. B., Kerber, M., Lange, C., and Rowat, C. (2013). Proving Soundness of Combinatorial Vickrey Auctions and Generating Verified Executable Code. [arXiv:1308.1779](#).
- Caplan, P. C. (2021). Higher-Dimensional Power Diagrams for Semi-Discrete Optimal Transport. [arXiv:2106.14730](#).
- Clarke, E. H. (1971). Multipart Pricing of Public Goods. [Public Choice](#) 11, 17–33.

- Coase, R. H. (1937). The Nature of the Firm. *Economica* 4(16), 386–405.
- Dubey, K. A., Feng, Z., Kidambi, R., Mehta, A., and Wang, D. (2024). Auctions with LLM Summaries. [arXiv:2404.08126](#).
- Dütting, P., Mirrokni, V., Paes Leme, R., Xu, H., and Zuo, S. (2024). Mechanism Design for Large Language Models. [arXiv:2310.10826](#).
- Edelman, B., Ostrovsky, M., and Schwarz, M. (2007). Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review* 97(1), 242–259.
- Feizi, S., Hajiaghayi, M., Rezaei, K., and Shin, S. (2025). Online Advertisements with LLMs: Opportunities and Challenges. [arXiv:2311.07601](#).
- Ghani, N., Hedges, J., Winschel, V., and Zahn, P. (2018). Compositional Game Theory. *LICS ’18*; [arXiv:1603.04641](#).
- Grbovic, M., Djuric, N., Radosavljevic, V., Silvestri, F., Baeza-Yates, R., Feng, A., Ordentlich, E., Yang, L., and Owens, G. (2016). Scalable Semantic Matching of Queries to Ads in Sponsored Search Advertising. *SIGIR ’16*; [arXiv:1607.01869](#).
- Groves, T. (1973). Incentives in Teams. *Econometrica* 41(4), 617–631.
- Hajiaghayi, M., Lahaie, S., Rezaei, K., and Shin, S. (2024). Ad Auctions for LLMs via Retrieval Augmented Generation. [arXiv:2406.09459](#).
- Hartline, J., Hoy, D., and Taggart, S. (2023). Robust Analysis of Auction Equilibria. [arXiv:2310.03702](#).
- Hayek, F. A. (1945). The Use of Knowledge in Society. *American Economic Review* 35(4), 519–530.
- Hotelling, H. (1929). Stability in Competition. *The Economic Journal* 39(153), 41–57.
- Jouvelot, P., and Gallego Arias, E. J. (2022). mech.v: A Coq Formalization of Mechanism Design. [GitHub repository](#).
- Kerber, M., Lange, C., and Rowat, C. (2016). An Introduction to Mechanized Reasoning. *Journal of Mathematical Economics* 66, 26–39.
- Lahaie, S., and Pennock, D. M. (2007). Revenue Analysis of a Family of Ranking Rules for Keyword Auctions. *EC ’07*, 50–56.
- Levin, J., and Milgrom, P. (2010). Online Advertising: Heterogeneity and Conflation in Market Design. *American Economic Review* 100(2), 603–607.
- Nelson, P. (1974). Advertising as Information. *Journal of Political Economy* 82(4), 729–754.
- Nisan, N. (2000). Bidding and Allocation in Combinatorial Auctions. *EC ’00*.
- Nisan, N., Roughgarden, T., Tardos, É., and Vazirani, V. V., eds. (2007). Algorithmic Game Theory. [Cambridge University Press](#).
- Rochet, J.-C. (1987). A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context. *Journal of Mathematical Economics* 16(2), 191–200.
- Sandholm, T. (2007). Expressive Commerce and Its Application to Sourcing: How We Conducted \$35 Billion of Generalized Combinatorial Auctions. *AI Magazine* 28(3).
- Shahidi, P., Rusak, G., Manning, B. S., Fradkin, A., and Horton, J. J. (2025). The Coasean Singularity? Demand, Supply, and Market Design with AI Agents. In *The Economics of Transformative AI*. [NBER](#).
- Soumalias, E., Curry, M. J., and Seuken, S. (2025). Truthful Aggregation of LLMs with an Application to Online Advertising. [arXiv:2405.05905](#).
- Tang, B. J., Sun, K., Curran, N. T., Schaub, F., and Shin, K. G. (2024). Ads that Talk Back: Implications and Perceptions of Injecting Personalized Advertising into LLM Chatbots. [arXiv:2409.15436](#).
- Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16(1), 8–37.
- Zhao, C., Hu, Q., Song, S., Chen, D., Zhu, H., Xu, J., and Zheng, B. (2025). LLM-Auction: Generative Auction towards LLM-Native Advertising. [arXiv:2512.10551](#).