Targeted Advertising Platforms: Pooling Audiences with Poaching Risks *

Klajdi Hoxha †

August 23, 2026

Abstract

E-commerce platforms are inviting merchants to share their platform-built audiences, promising to use the pooled audience to target customers most likely to click on their ads. Yet pooling raises customer-poaching concerns that current platform designs fail to address. We develop a model of targeted advertising platform design that incentivizes merchants to share audiences despite poaching risks. We characterize the mechanism that maximizes a weighted sum of platform revenue, merchant surplus, and customer engagement, measured by expected clicks. On sufficiently large platforms, the optimal indirect mechanism segments the market into *selling* and *buying* markets, in which the platform respectively buys and sells targeted ads at posted per-click prices, and *exchange* markets, in which merchants exchange audiences for potentially fewer ads with higher click-through rates. An endogenous vector of shadow prices—designer's *visible hand*—clears the markets against supply from the pooled audience, generalizing the classic third-degree price discrimination problem to a multiproduct setting with merchant-owned products. We show that merchant-efficient design is possible without outside subsidies and characterize when some degree of both merchant and customer efficiency can be implemented. We apply the model to Shopify Audiences, and show how Shopify could tailor its pricing by product-based tier and audience contribution. The model applies to other combinatorial exchanges such as community time banks and greenhouse gas credit markets.

*I am grateful to Michael Ostrovsky, Andrzej Skrzypacz and Weijie Zhong for their guidance and support. I am thankful to Heski Bar-Isaac, Ellen Muir, Spencer Pantoja, Bing Liu, Ilya Segal, Eric Tang, Frank Yang and seminar participants at Stanford University for helpful comments and suggestions. I owe special thanks to Varanya Chaubey for exceptional editing advice.

The author declares that he has no relevant or material financial interests that relate to the research described in this paper. AI tools assisted in generating examples and tikz figures, all of which were independently verified by the author. All em/en-dashes were human-generated.

†Graduate School of Business, Stanford University. Email: hklajdi@stanford.edu.

1 Introduction

E-commerce platforms are rolling out ambitious targeted advertising initiatives that rely on merchants to share their platform-built audiences. These initiatives are in response to recent privacy policies—such as California Consumer Privacy Act (CCPA), the European Union's General Data Protection Regulation (GDPR), and Apple's App Tracking Transparency (ATT)[1]—that have limited third-party tracking, the backbone of targeted online advertising. With Safari, Firefox, and Edge having phased out third-party cookies, Chrome is now the only major browser still relying on them as it seeks alternative solutions.[2] A promising direction involves pooling audiences owned by in-platform merchants. Doing so gives the platform access to a larger audience, allowing it identify and target customers with ads they are more likely to click than those from their initial merchants.

Yet pooling audiences exposes merchants to customer poaching. Participating merchants can compete for any customer in the pool typically through auctions based on bids and measures of ad quality, such as *click-through rates* (CTRs), defined as the number of clicks your ads receive divided by the number of impressions. Some merchants that contribute valuable audiences may therefore lose some of their existing customers to competitors that bid more aggressively or generate higher click-through rates. Unless the platform protects merchants from poaching or compensates them for the profits they forgo, they may prefer not to participate and instead advertise only to their own audiences.

Our running example is Shopify, an e-commerce platform that powers the online stores of millions of merchants. In 2022, it launched Shopify Audiences, a tool that pools participating merchants' audiences with the goal of improving ad targeting. Shopify's FAQ acknowledges concerns about customer poaching by addressing whether merchants can "exclude competitors from using their data".[3] Shopify states that "the data you contribute will be used to benefit other participating merchants," highlighting an important trade-off that merchants face when deciding whether to participate.

This paper develops a model of targeted advertising platform design and characterizes the optimal design. We first study platforms with a finite number of merchants ($i \in N$) and later extend the model to a continuum, grouping merchants into finitely many product-based tiers (henceforth, *tiers*) ($k \in K$) that capture observable differences in profit margin

[1]Following introduction of ATT in iOS 14.5, about 96% of iPhone users opted out of app tracking: https://www.imore.com/96-iphone-users-have-opted-out-app-tracking-ios-145-launched.

[2]Google's Privacy Sandbox initiative, a privacy-preserving alternative to third-party cookies, fell short of its intended goals; see https://privacysandbox.com/news/privacy-sandbox-update.

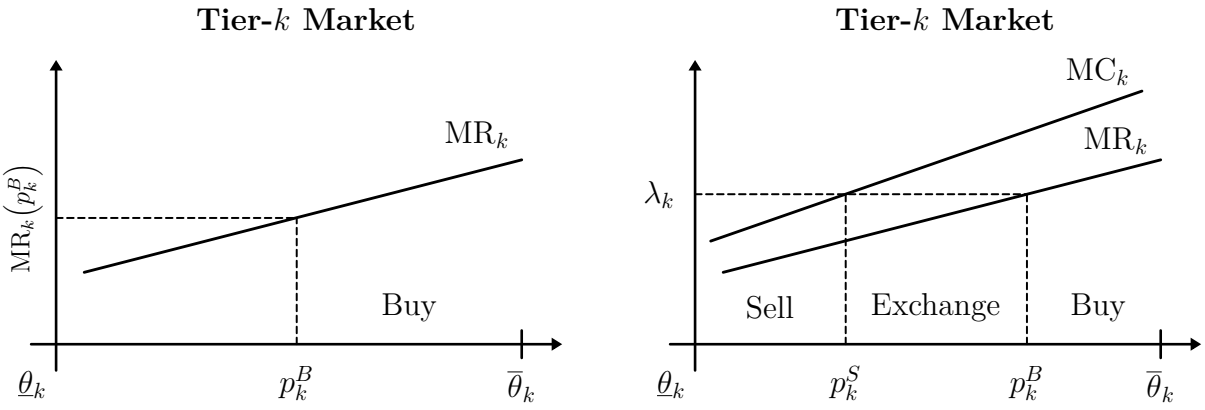
[3]See https://help.shopify.com/en/manual/promoting-marketing/shopify-audiences/faq.

(type) distributions and initial audience contributions. Merchants participate by contributing their customer audiences and then compete to advertise to customers in the resulting pool. Each customer is described by an observable vector of CTRs, $\omega = (\omega_i : i \in N)$, where each CTR ω_i is the probability that the customer clicks merchant i's ad. Each merchant is privately informed about their profit margin per click, θ_i. Absent product prices, customer welfare corresponds to customer engagement, measured by expected clicks. Because the model abstracts from product prices, we use customer engagement, measured by expected clicks, as a proxy for customer welfare. Our question therefore concerns the design of targeted ads and payments that maximize a weighted sum of platform revenue, merchant surplus, and customer engagement, with respective welfare weights η_r, η_v, and η_w summing to 1, subject to incentive and participation constraints.

Unlike position auction models, targeted advertising involves heterogeneous CTRs across customer-merchant pairs, the key novelty of our model. In position auction models, CTRs depend on merchant characteristics and advertising position, but not on customer identity. Varian (2007) assumes that the CTR of merchant i in position s is separable, $e_i p_s$, where e_i is the ad's "quality effect" and p_s is the "position effect", whereas Edelman et al. (2007) restrict their analysis to CTRs that depend only on position. Although a position can be interpreted as a customer seeking a match with a merchant, the relative ranking of merchants does not vary across positions. In our model, by contrast, each customer represents a distinct advertising opportunity and induces a potentially different ranking of merchants' ads. Lastly, customers initially belong to exclusive merchant audiences, which creates incentive participation constraints and customer-poaching concerns that are absent in position auction models.

Our main result shows that, in sufficiently large (limit) market, for each tier there exists unique posted per-click selling $\left(p_k^S\right)$ and buying $\left(p_k^B\right)$ prices that optimally segment the marketplace into selling, exchange, and buying markets. Within each tier, merchants with margins below the selling price $\left(\theta_k < p_k^S\right)$ give up their audiences, receive no targeted ads, and receive lump-sum transfers equal to the expected clicks they forgo, valued at the posted selling price. Merchants with margins above the buying price $\left(\theta_k > p_k^B\right)$ receive expected clicks up to capacity and pay the posted buying price per additional expected click. Those with margins between the two prices exchange access to their audiences for higher-CTR ads; optimally, they receive just enough expected clicks to reproduce their outside option from advertising to their own audiences, and zero lump-sum transfers.

Across tiers, a *visible hand* coordinates the marketplace: the designer determines endogenous shadow prices that clear the demand for targeted ads from each tier with the



(a) Canonical third-degree price discrimination

(b) Merchant-owned products third-degree price discrimination

Figure 1: Third-Degree Price Discrimination.

supply generated by the pooled audience. On the supply side, each tier-specific shadow price λ_k determines the value of an additional expected click in that tier. Consequently, each representative customer from the limiting pooled audience, $(\omega_k : k \in K)$, is assigned to the tier that maximizes its shadow value $\omega_k \lambda_k$. On the demand side, the shadow price determines both posted prices—and hence the total demand for targeted ads—by equating both the marginal cost of buying audiences from low types and the marginal revenue from selling targeted ads to high types to that shadow price. In equilibrium, the resulting shadow prices equate demand and supply in every tier.

Fundamentally, the designer's problem becomes equivalent to a multiproduct monopolist's problem engaging in third-degree price discrimination across tiers, except that here merchants initially own the products. This interpretation builds on Bulow and Roberts (1989), who show that the optimal auction problem is equivalent to a monopolist's third-degree price discrimination problem.[4] If customers were homogeneous, the optimal selling strategy is supplying quantities that equalize marginal revenues across tiers until the marginal revenue from assigning an additional customer reaches zero or the available supply is exhausted. With heterogeneous customers, initially part of merchants' audiences, this logic generalizes to a vector of shadow prices that jointly determines the optimal mechanism. In the standard setting, all shadow prices coincide and equal the marginal revenue generated by the last unit sold.

Fig. 1 illustrates these findings. Panel (a) illustrates canonical third-degree price discrimination, in which a single-product monopolist sells the same product across segmented markets, potentially charging a different price p_k^B in each. Because the monopolist does

[4]The study of third-degree price discrimination dates back to the classic work of Pigou (1920) and Robinson (1933).

not purchase the product from market participants, only marginal revenue matters in each tier-k market. The monopolist sells additional units in the market with the highest positive marginal revenue until either marginal revenue falls to zero or capacity is exhausted. If capacity is slack, marginal revenue is zero in every active market; if capacity binds, positive marginal revenues are equalized across active markets. Thus, for any two active markets k and ℓ with interior buying cutoffs,

$$\mathrm{MR}_k\left(p_k^B\right) = \mathrm{MR}_\ell\left(p_\ell^B\right) = \lambda,$$

where $\lambda = 0$ when capacity is slack and $\lambda > 0$ when it binds.

Panel (b) generalizes this logic to a multiproduct monopolist that sells heterogeneous products initially owned by merchants. Because the platform must first purchase the products from merchants, within each tier k it mediates trade until the marginal cost of buying from net sellers equals the marginal revenue from selling to net buyers. At interior buying and selling cutoffs,

$$\mathrm{MC}_k\left(p_k^S\right) = \mathrm{MR}_k\left(p_k^B\right) = \lambda_k,$$

where $\lambda_k \geq 0$ is the market-clearing shadow price in tier-k market. Because products are heterogeneous, market clearing generally requires a vector of tier-specific shadow prices rather than one common shadow price as in the canonical setting.

Efficiency. In the limit market, merchant-efficient design is always possible without outside subsidies. Within each tier, efficiency eliminates the exchange market and leaves a single cutoff price p_k^E: merchants below it sell, while those above it buy. Exchange markets are inefficient because they allocate ads to satisfy merchants' outside options, rather than according to the value of each customer-merchant pair, $\omega_k\theta$. At an interior solution, the tier-specific shadow price satisfies $\lambda_k^E = p_k^E$. Therefore, assigning each customer to the tier k that maximizes the expected profit of its cutoff type, $\omega_k p_k^E$, yields an efficient allocation. When the designer also values customer engagement—puts positive weight η_w on maximizing expected clicks—at an interior solution the shadow price becomes $\lambda_k^E = \eta_w + (1-\eta_w)p_k^E$. Thus, a higher weight η_w may direct ads toward lower marginal types. The resulting total payments from buyers may be insufficient to compensate sellers, in which case merchant- and customer-efficient design requires an outside subsidy.

The merchant-efficiency result complements the core convergence theorems for exchange economies (Debreu and Scarf (1963)) and the efficient dissolution of *any* sufficiently replicated partnership (Cramton et al. (1987)). As the market grows, individual merchants become negligible and the empirical distribution of private types converges to its popula-

tion distribution, creating aggregate certainty about total demand in each tier. In the limit, shadow prices play a role analogous to competitive prices in the Debreu-Scarf limit: they emerge as supporting prices that decentralize the efficient allocation. The result can also be interpreted as an extension of partnership dissolution to a setting with multiple heterogeneous assets. Cramton et al. (1987, Proposition 4) show that any partnership, regardless of its initial ownership shares, can be dissolved efficiently after sufficiently many replications—adding multiple copies of each partner and splitting that partner's share equally among them. Whereas their setting concerns the reallocation of ownership in a single asset, ours concerns the reallocation of advertising access to multiple *heterogeneous* customers.

Policy Implications. Applying our setting to Shopify, our results suggest that Shopify could expand access to all merchants on its platform by tailoring its pricing by product-based tiers and audience contributions. Currently, Shopify Audiences is available only to Shopify Plus merchants, effectively tying access to a single premium subscription tier rather than offering a separate pricing plan for pooling audiences. This uniform pricing scheme may exclude merchants that could contribute valuable audiences but have low demand for additional advertising.

We propose pricing Shopify Audiences separately from the plan subscription, offering it to all in-platform merchants. In the tier-pricing structure, net audience sellers will receive lump-sum transfers or advertising credits for future ad campaigns, whereas net buyers will pay only for the additional expected clicks they receive. The rest of the merchants contribute valuable audiences to the pool and are guaranteed to receive targeted ads that are just as valuable as advertising to their own audiences.

Technical Contributions. Our main technical contribution shows that, under partial exchangeability assumption on CTRs, the optimal finite market design converges to the limit market solution as the number of merchants grows large. An infinite sequence $(\omega_i : i \geq 1)$ is *partially exchangeable* if the joint law of any finite sample (N) is invariant with respect to permutations (π_k) of the indices within each observable tier k in K:

$$(\omega_{i,k} : i \in N, k \in K) \overset{\text{dist.}}{\sim} \left(\omega_{\pi_k(i),k} : i \in N, k \in K\right).$$

De Finetti's theorem states that a partially exchangeable sequence is i.i.d. conditional on latent probability measures $\mathbf{P} := (P_i : i \in K)$, that is,

$$(\omega_{k,i} : i \in N \mid \mathbf{P}) \overset{\text{i.i.d.}}{\sim} P_k \quad \text{for all } k \text{ in } K,$$

where $\mathbf{P}$ is drawn from the de Finetti *mixing measure* G.[5] In the limit market, the pooled audience is therefore characterized by representative customers $(\omega_k : k \in K)$ such that each tier-specific CTR ω_k is identified with the $\operatorname{ess\,sup}(\omega_{k,i} \mid \mathbf{P})$. The convergence results make these approximations precise: using perturbations of the optimal limit market allocation, we construct converging feasible finite market allocations to the optimum.

Finite Market Results. To build intuition for solving the limit market, we first solve the optimal design in the finite market. We show that a *scoring* mechanism that allocates ads in bundles is optimal: merchant i's report generates a single score $h_i(\theta_i)$, and each customer ω is assigned to a merchant with highest CTR-weighted score $\omega_i h_i(\theta_i)$; ties are resolved by customer-specific tie-breaking rule that ensures participation constraint of the worst-off merchant binds. Moreover, the scoring function coincides with the virtual cost function for low types and the virtual value function for high types. A scoring mechanism therefore pools merchants' countervailing incentives to overstate (understate) their types when selling customers (buying ads), which helps the designer reduce information rents.

The finite market delivers a general *direct* scoring implementation, while the limit model offers an *indirect* market implementation consisting of posted prices and market-clearing shadow prices. In the finite market, incentive compatibility and individual rationality are imposed at the interim stage, so allocations and transfers depend on beliefs about other merchants' types. As a result, a merchant with a given type may be a net seller or a net buyer ex-post, so strict cutoff types that split the market are not generally feasible ex-post.

Moreover, the scoring function is flat between the virtual cost and virtual value regions to ensure incentive compatibility as merchants switch from net sellers to net buyers. The scoring function is therefore *ironed* over an interval, with the optimal ironing parameter chosen so that types in the ironed region are worst-off: their interim expected clicks equal their outside option, and their participation constraints bind. The resulting ironing parameters have natural counterparts in the limit market: the tier-specific shadow prices. In the limit, these parameters play a dual role, determining the worst-off types while also clearing the markets.

Characterizing the optimal ironing parameters in the finite market is generally computationally intractable because determining customer-specific tie-breaking rules creates a multidimensional bottleneck. We first characterize the optimal scoring functions explicitly in a stylized setting with two merchants and binary CTRs, where customers are either inclusive or exclusive. In this simple yet rich setting, we show how the model connects

[5] De Finetti's theorem is a celebrated result in Bayesian statistics because the mixing measure provides a foundation for treating latent parameters as draws from some prior distribution.

to the classic settings of bilateral trade and partnership dissolution and how it generalizes these settings. We then consider continuously distributed CTRs and show that ties occur with probability zero. Avoiding the tie-breaking bottleneck provides greater tractability and allows us to derive an explicit system of nonlinear equations characterizing the optimal scoring functions.

Related Literature. Theoretically, the model can be recast as a generalization of settings like bilateral trade (Myerson and Satterthwaite (1983)) and partnership dissolution (Cramton et al. (1987), Loertscher and Wasser (2019)) to an environment in which trade occurs over multiple items. More generally, the model fits within a class of combinatorial exchange markets in which agents have privately known heterogeneous values for a bundle of items.[6] While studying optimal design in combinatorial exchanges quickly becomes intractable outside a few special settings,[7] we gain considerable tractability by having one-dimensional heterogeneity (type) and additive preferences over multiple heterogeneous and divisible items (customers) in a quasi-linear environment with side payments.

Our result that targeted ads are optimally allocated in bundles connects to the literature on optimal bundling. Yang (2025) studies nested bundling in a setting with one-dimensional heterogeneity and non-additive preferences. If preferences are additive, bundling strategies generally do not benefit the designer and the problem reduces to a simple one-dimensional screening problem. In our setting, preferences are additive across customers, but countervailing incentives due to initial endowments link merchants' incentives across customers, making bundling beneficial.

Our limit market model also connects to the literature on platform-mediated matching. Gomes and Pavan (2016, 2024) study many-to-many optimal matching between privately informed agents from two sides of a market, and ask how cross-subsidization and price discrimination shape matching patterns, targeting, and welfare. In our setting only one side of the market is privately informed: the platform and merchants observe customers' CTRs, while each merchant privately observes their type, and match value is the product of the two. There's an initial matching between merchants and customers, and the designer's objective is to improve upon this matching while leaving every merchant weakly better off.

[6] Following Babaioff et al. (2008), a combinatorial exchange generalizes combinatorial auctions by having agents act as both buyers and sellers: an agent can offer a bundle of items, provided she receives a different bundle and/or an appropriate side payment in return.

[7] The celebrated Shapley-Scarf housing market is a combinatorial exchange in which each agent initially owns an indivisible house, has preferences over all houses, and may exchange it for another without side payments. Shapley and Scarf (1974) propose the strategy-proof Top Trading Cycles mechanism, which selects an allocation in the core.

On the practical side, the paper contributes to the literature on digital platforms. Bergemann and Bonatti (2011) develop a model with many advertisers and many advertising markets, where each consumer is located in a specific advertising market and has preferences for one specific product. They study how targeting in advertising markets affects the allocation and price of ads and competition in the advertising market. Eliaz and Spiegler (2020) study personalized targeting when advertisers may strategically choose whom to target. Bergemann and Bonatti (2024) and Bergemann et al. (2024) study digital platforms where advertisers can reach customers on and off the platform with a focus on optimal ad pricing and efficient product matches. Bergemann and Bonatti (2015) and Elliott et al. (2022) study information design and pricing when the platform controls consumer data, whereas in our setting merchants initially control their own customer audiences. See also Bergemann et al. (2022), Galperti and Perego (2022) and Galperti et al. (2024), for settings in which customers control the data supplied to intermediaries, and Farboodi and Veldkamp (2023) for a comprehensive review of data and markets.

Roadmap. The remainder of the paper proceeds as follows. In Sec. 2 we illustrate finite and limit market examples that preview the main results. In Sec. 3 we introduce the benchmark model with finitely many merchants, characterize the optimal mechanism and provide proof sketches of the main results. In Sec. 4 we introduce and solve the limit market, study efficiency, and provide proof sketches of the main results. Lastly, in Sec. 5 we discuss model extensions and applications, and conclude in Sec. 6. Auxiliary proofs are in the appendix.

2 Illustrative Examples

This section presents illustrative examples for the finite and limit markets. The first draws parallels with bilateral trade and partnership dissolution models, while highlighting how the finite market model generalizes each of them. The second previews the three-market segmentation result and draws parallels with third-degree price discrimination, which our framework generalizes to a multiproduct market with merchant-owned products.

2.1 Illustrative Examples: Finite Market

Consider two merchants $i = 1, 2$ with *private* per-click profit margins (henceforth, *types*) θ_i drawn identically and independently from $F_i \sim \text{Uniform}[0, 1]$. Merchant i initially has exclusive access to an audience of mass $\alpha_i \in [0, 1]$, with $\alpha_1 + \alpha_2 = 1$. An audience consists

of customers with *observable* CTRs $\omega = (\omega_1, \omega_2)$, where ω_i denotes the probability that the customer clicks on merchant i's ad. If merchant i advertises to a customer of type ω, she earns an expected profit of $\omega_i \theta_i$.

The designer mediates trade among merchants by reallocating exclusive access to customers and setting payments to maximize a weighted sum of platform revenue, merchant surplus, and customer engagement, measured by the net increase in expected clicks.

2.1.1 Homogeneous Audiences

Suppose the audiences consist of a single customer type, $\omega = (1, \tau)$ for some $\tau \in [0, 1]$. Thus, merchant 2's *effective type*, $\tau\, \theta_2$, is uniformly distributed on $[0, \tau]$. Suppose further that the designer maximizes revenue.

Bilateral Trade with a Mediator. Suppose merchant 1 has access to no customers, that is, $\alpha_1 = 0$. Then the problem becomes virtually identical to bilateral trade with a mediator, with merchant 1 being the buyer (B) and merchant 2 being the seller (S) (Myerson and Satterthwaite, 1983, Section 5). Then let merchant 1's virtual value and merchant 2's virtual cost be defined as usual by

$$\phi_1^B(\theta_1) := \theta_1 - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} = 2\theta_1 - 1 \quad \text{and} \quad \phi_2^S(\tau\theta_2) := \tau\theta_2 + \frac{F_2(\tau\theta_2)}{f_2(\tau\theta_2)} = 2(\tau\theta_2),$$

respectively. The optimal mechanism therefore grants merchant 1 exclusive access to merchant 2's audience if and only if

$$\phi_1^B(\theta_1) \geq \phi_2^S(\tau\theta_2) \quad \Leftrightarrow \quad \theta_1 - \tau\theta_2 \geq 1/2.$$

If $\tau = 0$, the seller requires no compensation, and the designer acts as a monopolist selling merchant 1 exclusive access to a unit mass of customers of type $\omega = (1, 0)$. The optimal posted (per-click) price is therefore $p^* = 1/2$, so trade occurs if and only if $\theta_1 \geq 1/2$.

Partnership Dissolution with a Mediator. Now set $\tau = 1$, so that all customers have CTRs $\omega = (1, 1)$, and instead suppose $\alpha_i \in (0, 1)$ for each $i = 1, 2$. Then the problem becomes virtually identical to revenue-maximizing partnership dissolution with a mediator (Cramton et al., 1987; Loertscher and Wasser, 2019). Whenever trade occurs, one merchant obtains exclusive access to the pooled audience; otherwise, each merchant retains exclusive access to her initial audience. Loertscher and Wasser (2019) show that the optimal mechanism grants the merchant with the higher *ironed* virtual type function (henceforth, *scoring function*) $\overline{\phi}_i(\theta_i)$ exclusive access to the pooled audience, where $\overline{\phi}_i(\theta_i)$ is

obtained by ironing the virtual type function $\phi_i(\theta_i, \theta_i')$ at some *critical* type θ_i'.[8] Specifically, for each $i = 1, 2$ there exists an ironing parameter $z_i \geq 0$ such that

$$\overline{\phi}_i(\theta_i; z_i) = \begin{cases} 2\theta_i, & 2\theta_i \leq z_i, \\ z_i, & 2\theta_i - 1 < z_i < 2\theta_i, \\ 2\theta_i - 1, & 2\theta_i - 1 \geq z_i. \end{cases} \tag{1}$$

When the merchants' ironed virtual types are tied, the designer chooses a tie-breaking rule that makes each merchant's critical type a *worst-off type*—expected rents are zero. In the partnership dissolution setting, this is equivalent to ensuring that a merchant with the critical type retains her initial share of the pooled audience in expectation.

As a consistency check, suppose $\alpha_1 = 0$. One choice of optimal ironing parameters is $z_1 = 0$ and $z_2 = 1$,[9] which recover the optimal bilateral trade mechanism under which merchant 1 acquires exclusive access to the pooled audience if and only if $\theta_1 - \theta_2 \geq 1/2$. This event occurs with probability $1/8$, so expected trade volume is $1/8$. Conditional on trade, the average gain is $2/3$; hence, expected net gains from trade are $(1/8)(2/3) = 1/12$. Since the first-best expected gains are $(1/2)(1/3) = 1/6$, the optimal bilateral trade mechanism captures $1/2$ of the first-best gains.

When $\alpha_1 \in (0, 1)$, the winner of the pooled audience is determined only after the merchants report their types, making the buyer and seller roles *endogenous* and creating countervailing incentives: high types tend to understate, whereas low types tend to overstate, but both incentives are weaker than under the fixed roles in bilateral trade. These incentives are weakest for "intermediate" types, who retain their initial shares in expectation and earn zero rents.

For example, with symmetric initial audience shares, $\alpha_1 = \alpha_2 = 1/2$, the optimal ironing parameters are $z_1 = z_2 = 1/2$, so the merchants' scoring functions are flat for $\theta_i \in [1/4, 3/4]$. The mechanism always grants one merchant exclusive access to the pooled audience, raising expected trade volume from $1/8$ to $1/2$. The partnership dissolution mechanism differs from the first best only when both types are intermediate and uniform tie-breaking selects the lower type. This occurs with probability $1/8$, and the average type difference conditional on this event is $1/6$, so the expected efficiency loss is $(1/8)(1/6) = 1/48$. Subtracting this loss from the first-best gains gives expected net gains of $1/6 - 1/48 = 7/48$, compared with $1/12$ under bilateral trade. Thus, the optimal partnership dissolution mechanism captures

[8] The virtual type function at the *critical type* θ_i' is defined by $\phi_i(\theta_i, \theta_i') = \phi_i^S(\theta_i)$ for $\theta_i < \theta_i'$ and $\phi_i(\theta_i, \theta_i') = \phi_i^B(\theta_i)$ for $\theta_i \geq \theta_i'$.

[9] More generally, any $z_2 \in \left[1, \phi_2^S(1)\right] = [1, 2]$ implements the same allocation. Merchant 1's scoring function never exceeds 1, so changes in merchant 2's scoring function above 1 has no bite.

7/8 of the first-best gains, compared with 1/2 under bilateral trade.

2.1.2 Heterogeneous Audiences

We now generalize the preceding analysis by allowing CTRs to vary across customers. Unlike in bilateral trade and partnership dissolution, the optimal mechanism may target different customers with ads from different merchants.

Fix $\varepsilon \in [1/2, 1)$. Suppose that, for each customer in either merchant's initial audience,

$$\omega_1, \omega_2 \overset{\text{i.i.d.}}{\sim} \text{Uniform}[2\varepsilon - 1, 1].$$

The parameter ε governs both the average and the heterogeneity of CTRs: $\mathbb{E}[\omega_j] = \varepsilon$, and as ε increases, CTRs become higher and less dispersed, converging to the homogeneous audience case $\omega = (1, 1)$ as $\varepsilon \to 1$. Moreover, because both initial audiences have the same i.i.d. CTR distribution, the pooled audience has that same distribution regardless of the initial audience shares α. Lastly, let a_i^ε denote merchant i's total expected clicks from advertising to her initial audience:

$$a_i^\varepsilon := \alpha_i \mathbb{E}[\omega_i] = \alpha_i \varepsilon.$$

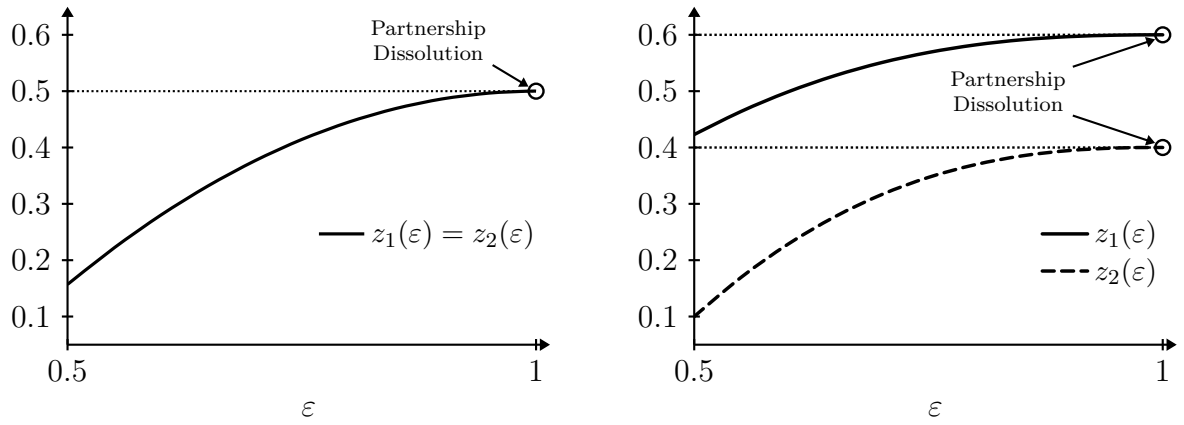
Her outside option expected profits are therefore $a_i^\varepsilon \theta_i = \alpha_i \varepsilon \, \theta_i$.

Even though different customers may receive ads from different merchants, the designer does not run a separate mechanism for each customer. Our first result (Theorem 1) is that a single merchant score governs her entire targeted advertising campaign, effectively bundling targeted ads across customers. Under revenue maximization ($\eta_r = 1$), this score is the scoring function in (1), with an ironing parameter z_i that varies with merchant i's initial audience contribution and the degree of CTR heterogeneity. Each customer with CTR vector ω is then targeted with an ad from the merchant with the highest CTR-weighted ironed score,

$$\omega_i \overline{\phi}_i(\theta_i; z_i).$$

A merchant's single report θ_i therefore determines both which customers in her initial audience continue to receive her ads and which customers from the other merchant's audience she begins to target. By linking these decisions, the designer pools countervailing incentives across customers as well, thus reducing further merchants' information rents.

In this example, continuously distributed CTRs imply that nonzero CTR-weighted scores tie only on a null set of customers. Hence, no nontrivial tie-breaking rule is needed, and the optimal mechanism is fully characterized by the ironing parameters (Proposition 1).



(a) $\alpha_1 = \alpha_2 = 1/2$.

(b) $\alpha_1 = 4/5$ and $\alpha_2 = 1/5$.

Figure 2: Optimal Ironing Parameters $z_1(\varepsilon)$ and $z_2(\varepsilon)$.

Case: $\alpha_1 = 0$**.** Here merchant 1 can only buy customer access and merchant 2 can only sell it. One choice of optimal ironing parameters is $z_1 = 0$ and $z_2 = 2$, giving merchant 1 the scoring function $(2\theta_1 - 1)_+$ and merchant 2 the scoring function $2\theta_2$.[10] Therefore, customers with CTRs ω are targeted with merchant 1's ad if and only if

$$\omega_1(2\theta_1 - 1) \geq \omega_2(2\theta_2) \quad \Leftrightarrow \quad \omega_1\theta_1 - \omega_2\theta_2 \geq 1/(2\omega_1).$$

This generalizes bilateral trade by making trade customer specific. As $\varepsilon \to 1$, the condition reduces to $\theta_1 - \theta_2 \geq 1/2$, recovering the optimal bilateral trade mechanism.

However, the optimal scoring functions are independent of ε: merchant 1, who has no initial audience, is always the buyer, while merchant 2 is always the seller. Once both audience shares are positive, however, the optimal ironing parameters generally depend on audience heterogeneity and hence, in this example, on the parameter ε.

Case: $\alpha_1 \in (0, 1)$**.** When both initial audience shares are positive, each merchant may buy and sell targeting access across customers. We show that there exists optimal ironing parameters $(z_1(\varepsilon), z_2(\varepsilon))$ such that each merchant i's critical worst-off type receives exactly her outside option expected clicks, $\alpha_i\varepsilon$. Using Eqn. (9) of Proposition 1, yields a system of two nonlinear equations (49) derived in Appendix C.1. Fig. 2 plots its numerical solutions for symmetric initial audience shares, $\alpha_1 = \alpha_2 = 1/2$, and asymmetric shares, $\alpha_1 = 4/5$ and $\alpha_2 = 1/5$.

With symmetric initial audience shares, symmetry implies $z_1(\varepsilon) = z_2(\varepsilon)$. For lower

[10] More generally, for $\varepsilon > 1/2$, any $z_2 \in \left[\min\left\{(2\varepsilon - 1)^{-1}, 2\right\}, 2\right]$ induces the same allocation almost surely; at $\varepsilon = 1/2$, only $z_2 = 2$ does. As $\varepsilon \to 1$, the lower bound converges to 1, recovering the homogeneous audience case.

values $\varepsilon < 1$, heterogeneity in CTRs creates additional "gains from trade" by targeting each customer with an ad from the merchant with whom she has the higher CTR. These gains allow the designer to provide each critical worst-off type her outside option expected clicks with a lower ironing parameter, so in the mechanism fewer types are sellers and more are buyers. As $\varepsilon \to 1$, these additional gains disappear and both ironing parameters converge to $1/2$, recovering the optimal partnership dissolution ironing parameters with symmetric shares and homogeneous audiences.

With asymmetric shares, $\alpha_1 = 4/5$ and $\alpha_2 = 1/5$, merchant 1 has the larger outside option:

$$a_1^\varepsilon = (4/5)\varepsilon \quad \text{and} \quad a_2^\varepsilon = (1/5)\varepsilon.$$

The optimal mechanism therefore sets $z_1(\varepsilon) > z_2(\varepsilon)$, making merchant 1 more likely to retain customers from her initial audience or target customers from merchant 2's audience. By assigning her a higher score, the designer reduces the risk that customers from her larger initial audience are poached, thereby preserving her incentive to participate. As $\varepsilon \to 1$, the ironing parameters converge to $z_1 = 0.6$ and $z_2 = 0.4$, which again correspond to the optimal partnership dissolution ironing parameters with asymmetric shares $\alpha = (4/5, 1/5)$.

2.2 Illustrative Example: Limit Market

We illustrate the limit-market mechanism in a stylized setting motivated by Shopify Audiences, with merchants grouped into product-based tiers.

Specifically, suppose Shopify segments the product market into three stylized product-based tiers $k \in K = \{\mathrm{L}, \mathrm{M}, \mathrm{H}\}$: low-tier products (L), medium-tier products (M), and high-tier products (H). There is a unit mass of merchants, with masses $\beta_\mathrm{L} = 1/6$, $\beta_\mathrm{M} = 1/6$, and $\beta_\mathrm{H} = 2/3$ in the respective tiers. Each merchant has a private type drawn independently from the tier-specific distribution F_k common to all merchants within tier k. Let θ_k denote the type of a tier-k merchant, and suppose

$$\theta_\mathrm{L} \sim \mathrm{Uniform}\left[0, \frac{1}{2}\right], \quad \theta_\mathrm{M} \sim \mathrm{Uniform}\left[\frac{1}{4}, \frac{3}{4}\right], \quad \theta_\mathrm{H} \sim \mathrm{Uniform}\left[\frac{1}{2}, 1\right].$$

There is a unit mass of customers in the market, initially uniformly distributed across merchants' platform-built audiences: any group of merchants with mass $\beta \in [0, 1]$ collectively contributes mass β to the pooled audience. Thus, tier k contributes audience mass β_k, and therefore $\sum_{k \in K} \beta_k = 1$ is the mass of the pooled audience when all merchants participate. By advertising to their own platform-built audiences, merchants in tiers L, M, and H generate total expected clicks $\mu_\mathrm{L} = 1/4$, $\mu_\mathrm{M} = 3/8$, and $\mu_\mathrm{H} = 7/10$, respectively.

In the limit market, we impose supply capacity constraints to rule out anomalous cases in which a single merchant sells products to an arbitrarily large number of customers relative to her market size. We normalize this capacity to one, meaning that each merchant can meet at most one unit of demand, measured by the expected clicks generated by her ads. By participating in the platform, a tier-k merchant can therefore increase her total expected clicks from μ_k to at most one, achieving maximum potential targeting gains $1 - \mu_k$.

Under partial exchangeability assumptions on the infinite sequence $(\omega_i : i \geq 1)$ with $K = \{\mathrm{L}, \mathrm{M}, \mathrm{H}\}$, our main technical result shows that there exists a distribution $G \in \Delta\big([0,1]^{|K|}\big)$ such that each customer in the pooled audience is represented by tier-specific CTRs (henceforth, *CSRs*) $\omega = (\omega_{\mathrm{L}}, \omega_{\mathrm{M}}, \omega_{\mathrm{H}}) \sim G$, which weakly outperform each tier's initial audience in terms of expected clicks:

$$\mathbb{E}_G[\omega_k] \geq \mu_k \quad \text{for all } k \in K.$$

Suppose for simplicity that the CSRs are mutually independent and satisfy

$$\omega_{\mathrm{L}} \sim \mathrm{Uniform}\left[\frac{1}{4}, \frac{1}{3}\right], \quad \omega_{\mathrm{M}} \sim \mathrm{Uniform}\left[\frac{1}{8}, \frac{2}{3}\right], \quad \omega_{\mathrm{H}} \sim \mathrm{Uniform}\left[\frac{1}{2}, 1\right],$$

where

$$(\mathbb{E}_G[\omega_{\mathrm{L}}], \mathbb{E}_G[\omega_{\mathrm{M}}], \mathbb{E}_G[\omega_{\mathrm{H}}]) = \left(\frac{7}{24}, \frac{19}{48}, \frac{3}{4}\right) > \left(\frac{1}{4}, \frac{3}{8}, \frac{7}{10}\right).$$

Optimal Mechanism. Suppose Shopify maximizes revenues ($\eta_r = 1$). Our main result (Theorem 2) shows that the optimal mechanism is characterized by shadow prices $\lambda = (\lambda_{\mathrm{L}}, \lambda_{\mathrm{M}}, \lambda_{\mathrm{H}})$ that clear, in each tier k in K, total supply $S_k(\lambda)$ generated from the pooled audience and total demand $D_k(\lambda)$ for expected clicks.

On the supply side, given shadow prices λ, each customer with CSR vector ω is targeted by a merchant from the tier k that maximizes the shadow value $\lambda_k \omega_k$. For $\lambda \neq 0$, the expected clicks supplied to tier k from the pooled audience are therefore

$$S_k(\lambda) := \mathbb{E}_G\left[\omega_k \mathbf{1}\left\{k = \operatorname*{argmax}_{j \in K} \lambda_j \omega_j\right\}\right], \tag{2}$$

where ties occur only on a null set of customers.

On the demand side, the shadow price λ_k determines posted per-click selling and buying prices p_k^S and p_k^B, segmenting tier-k merchants into selling, exchange, and buying markets. Since marginal revenue and marginal cost curves in each tier k are given by the virtual value and cost functions $\phi_k^B(\theta)$ and $\phi_k^S(\theta)$, respectively, whenever the prices are interior, they satisfy

$$\phi_k^S\big(p_k^S\big) = \lambda_k = \phi_k^B\big(p_k^B\big).$$

Now the optimal expected clicks and transfer for a type θ merchant in tier k, denoted by $q_k^*(\theta)$ and $t_k^*(\theta)$, respectively, are given by

$$q_k^*(\theta) = \begin{cases} 0, & \theta < p_k^S, \\ \mu_k, & \theta \in \left[p_k^S, p_k^B\right], \\ 1, & \theta > p_k^B, \end{cases} \qquad t_k^*(\theta) = \begin{cases} -\mu_k p_k^S, & \theta < p_k^S, \\ 0, & \theta \in \left[p_k^S, p_k^B\right], \\ (1-\mu_k) p_k^B, & \theta > p_k^B. \end{cases}$$

That is, sellers give up their initial audiences, receive no targeted ads, and are compensated at the per-click price p_k^S for the μ_k expected clicks they forgo. Exchange merchants receive targeted ads generating exactly their outside option expected clicks μ_k and make zero net payments. Buyers receive one unit of expected clicks which exhausts their capacity, thus achieving the maximum targeting gains $1 - \mu_k$ and paying the per click price p_k^B for the additional expected clicks. We interpret positive transfers as payments to Shopify, whereas negative transfers as (advertising) credits to merchants.

The resulting total expected clicks demand in each tier k is therefore

$$D_k(\lambda) := \beta_k \left[\mu_k\Big(F_k\Big(p_k^B\Big) - F_k\Big(p_k^S\Big)\Big) + 1 - F_k\Big(p_k^B\Big)\right]. \tag{3}$$

The first term in (3) assigns μ_k expected clicks to each exchange merchant, while the second assigns one unit to each buyer. Sellers demand no expected clicks.

The optimal shadow price vector λ^* is chosen so that the demand in (3) equals the supply in (2) in every tier:

$$D_k(\lambda^*) = S_k(\lambda^*), \quad \text{for all } k \in K. \tag{4}$$

Solving (4) gives

$$\lambda^* \approx (0.52, 0.35, 0.28).$$

The corresponding cutoff pairs are

$$\Big(p_{\mathrm{L}}^S, p_{\mathrm{L}}^B\Big) \approx (0.26, 0.50), \quad \Big(p_{\mathrm{M}}^S, p_{\mathrm{M}}^B\Big) \approx (0.30, 0.55), \quad \Big(p_{\mathrm{H}}^S, p_{\mathrm{H}}^B\Big) \approx (0.50, 0.64),$$

and the resulting markets are

tier	selling market	exchange market	buying market
low-tier (L)	$[0, 0.26)$	$[0.26, 0.50]$	$\varnothing$
medium-tier (M)	$[0.25, 0.30)$	$[0.30, 0.55]$	$(0.55, 0.75]$
high-tier (H)	$\varnothing$	$[0.50, 0.64]$	$(0.64, 1]$

Thus, the low tier has selling and exchange markets, the medium tier has all three markets, and the high tier has exchange and buying markets. Appendix C.2 derives the market-

clearing system and verifies this solution, while Theorem 2 establishes its optimality.

The example suggests offering Shopify Audiences separately from plan subscriptions and tailoring its pricing by product tier. Low-tier merchants, whose lower profit margins limit their willingness to pay for additional expected clicks, would either receive advertising credits for their audience contributions or obtain targeting that preserves their current targeting efforts. High-tier merchants could choose to pay for additional expected clicks, while medium-tier merchants would choose among getting advertising credits for the next round, or paying for expanding targeting reach.

3 Finite Market

We formally introduce the model with finitely many merchants in Sec. 3.1, state designer's problem in Sec. 3.1.2, and discuss underlying assumptions in Sec. 3.1.3. In Sec. 3.2.1 we characterize the optimal mechanism when CTRs are distributed over a finite support, and in Sec. 3.2.2 we derive an explicit system of equations that pins down the optimal mechanism when CTRs are continuously distributed. Sec. 3.3 provides proof sketches of both results.

3.1 Model

A designer operates an advertising platform serving a finite set of merchants, each of whom owns a *platform-built audience*. These audiences are pairwise disjoint and collectively constitute a unit mass of unit-demand customers.

Each merchant i in N has private per-click profit margin (henceforth, *type*) θ_i in $\Theta_i = \left[\underline{\theta}_i, \overline{\theta}_i\right] \subset \mathbb{R}_+$. Types are independently drawn from distributions F_i, each admitting a strictly positive, continuous density f_i. We assume that each F_i is *regular*, meaning that its virtual value and cost functions (also referred to as *marginal revenue* and *cost* curves[11])

$$\phi_i^B(\theta_i) := \theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \quad \text{and} \quad \phi_i^S(\theta_i) := \theta_i + \frac{F_i(\theta_i)}{f_i(\theta_i)},$$

respectively, are strictly increasing. Let $\Theta := \prod_{i \in N} \Theta_i$ denote the type space, and let $F := \bigotimes_{i \in N} F_i$ denote the joint type distribution.

Each customer is characterized by a vector of *click-through rates* (CTRs)

$$\omega = (\omega_i)_{i \in N} \in [0,1]^{|N|},$$

[11] Bulow and Roberts (1989) use these formulations to show that the optimal auction problem is equivalent to the monopolist's problem in third-degree price discrimination. Our main result in Sec. 4 admits a similar interpretation: the designer engages in third-degree price discrimination across tiers.

where the CTR ω_i is the probability that the customer clicks on merchant i's ad. CTRs are *observable*: the realization ω_i is commonly observed by the designer and merchant i. Let $\Omega \subseteq [0,1]^{|N|}$ be the (Borel) set of possible CTR realizations.

Each merchant i initially has exclusive access to a platform-built audience of mass $\alpha_i > 0$, where $\sum_{i \in N} \alpha_i = 1$. Conditional on belonging to merchant i's audience, customers are distributed according to Borel measures $\gamma_i \in \Delta(\Omega)$. Let $\alpha = (\alpha_i)_{i \in N}$ and $\boldsymbol{\gamma} = (\gamma_i)_{i \in N}$ denote the audience shares and conditional CTR distributions, respectively. Finally, let $\gamma_\alpha \in \Delta(\Omega)$ denote the *pooled audience*, defined as the mixture distribution

$$\gamma_\alpha(B) := \sum_{i \in N} \alpha_i \, \gamma_i(B) \quad \text{for every Borel set } B \subseteq \Omega.$$

Customers are *captive* to in-platform merchants and discover products only via targeted ads. Initially, merchant i advertises to their own platform-built audience, generating total expected clicks

$$a_i := \alpha_i \int_\Omega \omega_i \, d\gamma_i(\omega).$$

Merchant i's type-dependent *outside option* is therefore given by total expected profits $\theta_i a_i$.

3.1.1 Platform Design

We describe how the platform operates by specifying the policy on audience sharing and the design of targeted advertising campaigns.

Audience Sharing. Audience sharing requires merchants to waive their exclusive targeting rights over the audiences they contribute. Once the rights are waived, the platform may limit a merchant's ability to direct ads to its initial audiences, even if the merchant continues to hold the underlying customer records. The assumption therefore concerns control over platform-mediated advertising access, rather than ownership of the customer records.[12] Moreover, assuming the designer can costlessly replicate each merchant's initial advertising campaigns by targeting only that merchant's initial audience within the pooled audience, we can restrict attention, without loss of generality, to mechanisms in which every merchant participates.

[12] Shopify merchants can access and export customer records generated through their own stores. Shopify Audiences, however, controls the construction and delivery of advertising audiences: merchants cannot observe the identities or contact information of customers included in a particular audience or directly select those customers. Thus, a merchant may retain its customer records while having limited control over how those customers are reached through platform-mediated advertising.

Targeted Advertising Campaigns. Let $\mathcal{M}_+(\Omega)$ denote the space of finite nonnegative Borel measures on Ω, and let $(x,t) = \Big((x_i)_{i\in N}, (t_i)_{i\in N}\Big)$ denote a targeted advertising campaign (henceforth, *campaign*), consisting of

i) targeted ads allocation $x_i \in \mathcal{M}_+(\Omega)$ (henceforth, *targeted ads*) subject to feasibility constraints

$$\sum_{i\in N} x_i(B) \leq \gamma_\alpha(B) \quad \text{for every Borel set } B \subseteq \Omega,$$

ii) transfers $t_i \in \mathbb{R}$ (also referred to as *participation fees*).

A campaign specifies targeted ads, which determine the type of customers targeted by each merchant, and transfers, which specify the payments between the merchant and the platform. The allocation $x_i(B)$ specifies the mass of customers with CTR realizations in the (Borel) set $B \subseteq \Omega$ who are targeted by merchant i's ads. Feasibility constraints imply that, for each type of customers B in Ω, the total mass assigned across merchants cannot exceed the mass $\gamma_\alpha(B)$ available in the pooled audience.

In practice, e-commerce platforms commonly allow participating merchants to choose an advertising budget, which then determines the performance of their campaigns. In our setting, however, the designer must also account for merchants' participation incentives by guaranteeing each of them at least the value of their outside option. A campaign may therefore assign negative participation fees, implemented through instruments such as cash payments, fee rebates, or advertising credits.

Direct Mechanism. By the revelation principle, it is without loss of generality to restrict attention to direct mechanisms. A direct mechanism (x,t) consists of a targeted ads allocation rule (henceforth, *targeted ads*) $x = (x_i)_{i\in N}$ and transfer rule $t = (t_i)_{i\in N}$ where each $x_i : \Theta \to \mathcal{M}_+(\Omega)$ and $t_i : \Theta \to \mathbb{R}$ are measurable with respect to F, and satisfy the usual interim incentive-compatibility and individual-rationality constraints

$$\theta_i(Q_i(\theta_i) - a_i) - T_i(\theta_i) \geq \theta_i\Big(Q_i\Big(\widetilde{\theta}_i\Big) - a_i\Big) - T_i\Big(\widetilde{\theta}_i\Big) \quad \text{for all } \theta_i, \widetilde{\theta}_i \in \Theta_i \text{ and } i \in N, \tag{IC}$$

$$\theta_i(Q_i(\theta_i) - a_i) - T_i(\theta_i) \geq 0 \quad \text{for all } \theta_i \in \Theta_i \text{ and } i \in N, \tag{IR}$$

where the *expected clicks* $Q_i(\theta_i)$ and expected transfers $T_i(\theta_i)$ are given by

$$Q_i(\theta_i) := \mathbb{E}_{F_{-i}}\left[\int_\Omega \omega_i \, x_i(\theta_i, \theta_{-i}; d\omega)\right] \quad \text{and} \quad T_i(\theta_i) := \mathbb{E}_{F_{-i}}[t_i(\theta_i, \theta_{-i})],$$

and targeted ads must satisfy ex-post feasibility constraints

$$\sum_{i\in N} x_i(\theta; B) \leq \gamma_\alpha(B) \quad \text{for every Borel set } B \subseteq \Omega \text{ and } \theta \in \Theta. \tag{Fsb}$$

3.1.2 Designer's Problem

The designer designs a direct mechanism that maximizes a weighted sum of expected merchants' surplus, customer engagement, and platform revenue, subject to incentive compatibility, individual rationality and feasibility constraints.

Merchants' Surplus. Let $V(x)$ denote expected merchants' surplus, defined by

$$V(x) := \mathbb{E}_F\left[\sum_{i\in N} \theta_i\Big(\int_\Omega \omega_i\, x_i(\theta; d\omega) - a_i\Big)\right].$$

$V(x)$ measures total expected profits generated by targeted ads, net of the expected profits merchants forgo from not advertising to their own audiences. A platform may value merchants' welfare because satisfied merchants are more likely to continue participating and increase their future advertising expenditures.

Customer Engagement. Let $W(x)$ denote expected customer engagement, defined by

$$W(x) := \mathbb{E}_F\left[\sum_{i\in N}\Big(\int_\Omega \omega_i\, x_i(\theta; d\omega) - a_i\Big)\right].$$

$W(x)$ measures the expected net clicks generated by targeted ads x. It captures the potential improvement in targeting precision generated by reallocating customers toward merchants whose products they are more likely to click. Because the model abstracts from product prices, we use CTRs as a measure of customers' welfare from the advertised products. Poorly targeted or repetitive advertising may create ad fatigue and user attrition, whereas more precise targeting can improve the customer experience and sustain engagement with the platform.

Platform Revenues. Let $R(t)$ denote platform's expected revenues, defined as usual by

$$R(t) := \mathbb{E}_F\left[\sum_{i\in N} t_i(\theta)\right].$$

Now let $\eta = (\eta_v, \eta_w, \eta_r) \in \mathbb{R}^3_+$ be *welfare weights* satisfying $\eta_v + \eta_w + \eta_r = 1$ and $\eta_r > 0$, and let $\Psi_\eta(x,t)$ denote the designer's objective function for given direct mechanism (x,t), defined by

$$\Psi_\eta(x,t) := \eta_v V(x) + \eta_w W(x) + \eta_r R(t).$$

The designer therefore solves

$$\max_{(x,t)} \quad \Psi_\eta(x,t) \quad \text{subject to (IC), (IR) and (Fsb).}$$

Lastly, note that operating the platform may require costly infrastructure investments. We treat these additional costs as sunk and therefore do not model them explicitly. The weight on revenues (η_r) may therefore capture, in part, the necessity of recovering these costs.

3.1.3 Discussion of Assumptions

Observable CTRs. The assumption that CTRs are observable is motivated by the platform's transparency regarding merchants' campaign performances. Because platform-built audiences are managed by the platform, it can provide merchants with campaign-level performance metrics, including CTRs.[13] The platform and merchants can therefore share a common assessment of CTRs.

Per-Click Profit Margins. For a micro-foundation on types θ_i, assume customers fall into two categories: with probability ω_i, customers with CTR realizations ω are interested in merchant i's product, and otherwise are not interested. Conditional on being interested, each customer's willingness to pay v_i is drawn from some cumulative distribution V_i. Let θ_i therefore be defined as the *per-interested-customer* profit margin $\theta_i(c_i) := p(c_i) - c_i$, where $p(c_i)$ is an optimal price given private constant marginal cost $c_i \geq 0$ and market demand $D_i(p) := 1 - V_i(p)$, that is,

$$p(c_i) \in \operatorname*{argmax}_{p \geq 0} (p - c_i) D_i(p).$$

One Ad per Customer. (Fsb) requires that each customer be targeted at most once, ruling out cases in which a customer who does not click on one ad is subsequently e another. Such sequential targeting strategies creates a richer space of targeted ad campaigns, potentially increasing the value generated from each customer. This introduces additional incentive considerations that lie beyond the focus of this paper. Nevertheless, one could argue that ads create brand awareness "with or without the added benefit of click-through" Percy and Rossiter (1992), and repeated exposure may generate ad fatigue Craig et al. (1976), potentially reducing CTRs. Lastly, note that under the one ad per customer restriction, customers remain captive to some merchant. It is therefore without loss of generality to abstract from price competition when per-click profit margins are microfounded as above,

[13] Shopify Audiences, for example, makes data about benchmark metrics available, including CTR metric, so that merchants can compare their performance with relevant benchmark metrics; see *Advertising Campaign Benchmarks* on https://help.shopify.com/en/manual/promoting-marketing/shopify-audiences/benchmarks.

provided that winning an ad does not lead merchants to update their beliefs about market demand.

3.2 Optimal Mechanism

We show that the optimal mechanism assigns targeted ads using merchant-specific scoring functions, referred to as a *scoring mechanism*.

Definition 1 (Scoring Mechanism)**.** *A scoring mechanism* (h, π) *consists of merchant-specific, nondecreasing scoring functions* h_i *and a customer-specific tie-breaking rule* π*, defined by (Borel) measurable functions* $h_i : \Theta_i \to \mathbb{R}_+$ *for every* i *in* N *and* $\pi : \Theta \times \Omega \to \Delta(N \cup \{0\})$ *such that*

$$\begin{aligned} x_i(\theta; B) &= \int_B \pi_i(\theta, \omega)\, d\gamma_\alpha(\omega) \quad \text{for every Borel set } B \subseteq \Omega, \\ \text{where} \quad \pi_i(\theta, \omega) &= 0 \ \text{ if } \ i \notin \operatorname*{argmax}_{j \in N} \{\omega_j h_j(\theta_j)\} \quad \text{for every } i \in N, \theta \in \Theta, \end{aligned} \tag{5}$$

where $\pi_0(\theta, \omega)$ *is the share of unassigned customers with CTR realizations* ω *retained by the designer.*

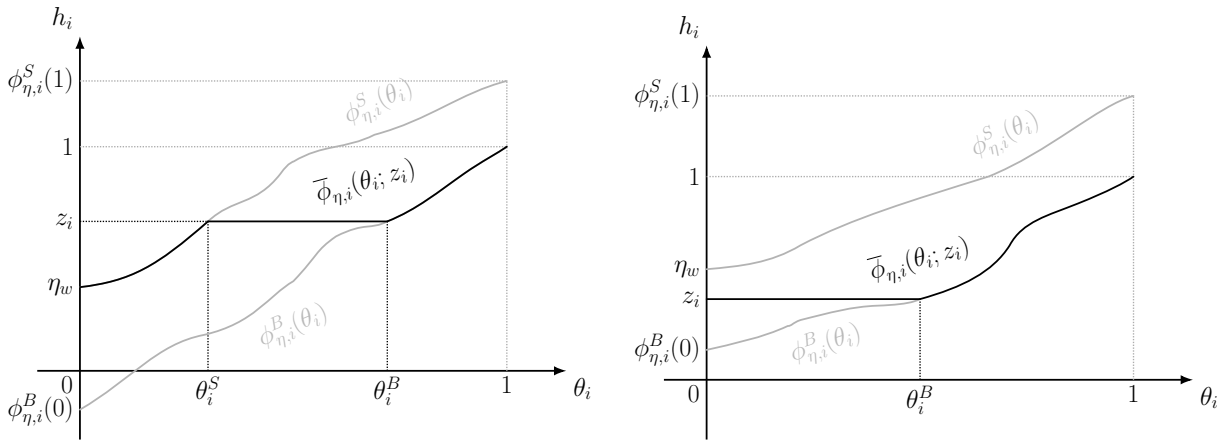
At each type profile θ, a scoring mechanism assigns customers with CTR realizations ω to merchants with the highest CTR-weighted score, $\omega_j h_j(\theta_j)$. The tie-breaking rule may split an atom of the audience across tied merchants. When ties occur only on γ_α-null sets, almost every customer with same CTR realizations are targeted by the same merchant.

Next define the weighted virtual value and cost functions by

$$\phi^B_{\eta,i}(\theta_i) := \eta_w + \eta_v \theta_i + \eta_r \phi^B_i(\theta_i) \quad \text{and} \quad \phi^S_{\eta,i}(\theta_i) := \eta_w + \eta_v \theta_i + \eta_r \phi^S_i(\theta_i),$$

respectively. Fig. 3 illustrates the weighted virtual cost and value functions along with an *ironed* function of the two, denoted by $\overline{\phi}_{\eta,i}(\theta_i; z_i)$ with z_i being the value of the function on the ironing interval. The ironed function coincides with the weighted virtual cost (value) function to the left (right) of the ironing interval.

We show that a scoring mechanism with scores constructed by ironing the weighted virtual cost and value functions is optimal. To state the result, for each j in $\{S, B\}$ and i in N, we use the inverse convention that $\left(\phi^j_{\eta,i}\right)^{-1}(z) = \underline{\theta}_i$ whenever $z < \phi^j_{\eta,i}(\underline{\theta}_i)$ and $\left(\phi^j_{\eta,i}\right)^{-1}(z) = \overline{\theta}_i$ whenever $z > \phi^j_{\eta,i}\left(\overline{\theta}_i\right)$.



(a) $\phi^B_{\eta,i}(0) < 0$. (b) $\phi^B_{\eta,i}(0) > 0$.

Figure 3: Examples of (optimal) scoring and weighted virtual functions.

3.2.1 Finite-Support CTRs

Let Ω be a finite set. For ease of notations, we write $\gamma_\alpha(\omega) := \gamma_\alpha(\{\omega\})$, $x^\omega_i(\theta) := x_i(\theta; \{\omega\})$, and $\pi^\omega_i(\theta) := \pi_i(\theta, \omega)$ for every $\omega \in \Omega$.[14]

Theorem 1 (Optimal Scoring Mechanism)**.** *There exists an optimal scoring mechanism such that, for each i in N, the scoring function h_i is given by $h_i \equiv \overline{\phi}_{\eta,i}(\cdot\,; z_i)$ for some ironing parameter $\left(\phi^B_{\eta,i}(\underline{\theta}_i)\right)_+ \le z_i \le \phi^S_{\eta,i}\left(\overline{\theta}_i\right)$, where*

$$\overline{\phi}_{\eta,i}(\theta_i; z_i) := \begin{cases} \phi^S_{\eta,i}(\theta_i) & \text{if } \theta_i < \theta^S_i, \\ z_i & \text{if } \theta^S_i \le \theta_i \le \theta^B_i, \\ \phi^B_{\eta,i}(\theta_i) & \text{if } \theta_i > \theta^B_i, \end{cases} \tag{6}$$

with cutoff types $\theta^j_i = \left(\phi^j_{\eta,i}\right)^{-1}(z_i)$ for each j in $\{S, B\}$. Moreover, the customer-specific tie-breaking rule π is chosen such that the critical types $\hat{\theta}_i$, defined uniquely for each i in N by

$$\hat{\theta}_i(z_i) := \mathbb{E}_{F_i}\left[\overline{\phi}_{\eta,i}(\theta_i; z_i) - \eta_w - \eta_v\theta_i\right]/\eta_r \in \left[\theta^S_i, \theta^B_i\right], \tag{7}$$

are worst-off types à la Cramton et al. (1987), that is, either $Q_i\left(\hat{\theta}_i\right) = a_i$, or $Q_i(y) < a_i$

[14] Note that (Fsb) becomes equivalent to

$$\sum_{i \in N} x^\omega_i(\theta) \le \gamma_\alpha(\omega) \quad \text{for every } \omega \in \Omega \text{ and } \theta \in \Theta.$$

for all $y < \hat{\theta}_i$ *and* $Q_i(y) > a_i$ *for all* $y > \hat{\theta}_i$*, where*

$$Q_i(\theta_i) = \mathbb{E}_{F_{-i}}\left[\sum_{\omega\in\Omega} \omega_i \, \pi_i^{\omega}(\theta_i, \theta_{-i}) \, \gamma_\alpha(\omega)\right].$$

Lastly, interim transfers are recovered by

$$T_i(\theta_i) = \theta_i(Q_i(\theta_i) - a_i) - \int_{\hat{\theta}_i}^{\theta_i} (Q_i(y) - a_i)dy.$$

Relative to the outside option a_i, expected clicks Q_i dictate whether a merchant is a net seller or a net buyer in the mechanism. As we show later, (IC) requires expected clicks to be nondecreasing in type, so in general low types are net sellers and high types are net buyers. The optimal scoring function in (6) reflects precisely this distinction: low-type scores are governed by the weighted virtual cost function $\left(\phi^S_{\eta,i}\right)$, whereas high-type scores are governed by the weighted virtual value function $\left(\phi^B_{\eta,i}\right)$.

Ironing pools intermediate types and pins down each critical worst-off type $\hat{\theta}_i(z_i)$ defined in (7). We say a critical type $\hat{\theta}_i$ is *interior* when the expected clicks she receives under the mechanism equal those generated by advertising on her own audience; that is, when $Q_i\left(\hat{\theta}_i\right) = a_i$. An interior critical type $\hat{\theta}_i$ is a worst-off type because, in expectation, she can neither be charged a positive participation fee nor require compensation to incentivize truthful reporting. This generalizes the worst-off-type condition in Cramton et al. (1987): whereas their worst-off types are pinned down by the expected quantity of a single traded good, ours are pinned down by the expected clicks generated by a heterogeneous pool of customers.

Importantly, the designer can replicate merchant i's initial expected clicks a_i without restricting targeting to her initial audience: potentially she designs targeted ads $x_i \not\equiv \alpha_i\gamma_i$ that target a smaller mass of customer in the pooled audience with higher CTRs. In fact, because all the types θ_i in the ironing interval $\left[\theta_i^S, \theta_i^B\right]$ have the same score $h_i(\theta_i) = z_i$, they are all worst-off types and satisfy $Q_i(\theta_i) = a_i$. This is precisely the finite market counterpart of the exchange market result in the limit market: all types in the ironing interval exchange their audiences while preserving their initial expected clicks a_i, though only in expectation—unlike in the limit market, where this also holds ex-post.

When the critical worst-off type is not interior, it must lie at an endpoint of the type space, $\underline{\theta}_i$ or $\overline{\theta}_i$, an outcome that does not arise in the partnership dissolution setting (Loertscher and Wasser, 2019). The critical worst-off type equals $\underline{\theta}_i$ when the pooled audience contains, relative to her initial contribution, many customers who generate high CTRs when matched with her. The same can also occur when the designer prefers targeting every customer rather than leaving some unassigned because she places positive weight on

customer engagement or merchant surplus. The lower endpoint $\underline{\theta}_i$ is therefore the worst-off type, since every higher type enjoys strictly larger rents. Conversely, when the designer's objective and the composition of the pooled audience favor assigning merchant i relatively few ads, she is a net seller regardless of her type, making $\overline{\theta}_i$ the critical worst-off type because she enjoys the lowest rents.

Now, since Ω is finite, tie-breaking is essential for determining the critical worst-off types. Because optimal scoring functions may be flat on intervals of positive measure, ties in CTR-weighted scores can occur with positive probability. The designer must therefore specify how ties are resolved at each type profile, and these choices affect, in particular, the expected clicks allocated to each critical type. Because every tied merchant contributes equally with the maximal CTR-weighted score to the objective, its value is unaffected by how the tie is resolved. We then prove that it suffices for the designer to choose tie-breaking rules under which each critical type's participation constraint binds, thereby making each one of them a worst-off type.

Determining the optimal tie-breaking rule in the finite-support case remains difficult. Because the critical worst-off type depends *endogenously* on the tie-breaking rule, the problem's complexity becomes inherently multidimensional, whose size grows with the cardinality of the set Ω. The exclusive and inclusive customer setting in Online Appendix illustrates how challenging characterizing the optimal tie-breaking rule can be even when ties occur only on inclusive customers. Nevertheless, as we show in the next section, under additional continuity assumptions on the mixture distribution γ_α, this tie-breaking bottleneck disappears: ties occur only on null sets, making the problem considerably more tractable.

3.2.2 Continuous CTRs

When ties occur only on null sets, we derive an explicit system of equations that determine the optimal ironing parameters, yielding a complete characterization of the optimal scoring mechanism.

Assumption 1 (Tie-Free CTRs)**.** *We say that the mixture distribution γ_α is tie-free if*

$$\operatorname*{argmax}_{i\in N}\{c_i\omega_i\}$$

is a singleton γ_α-almost everywhere for every $c=(c_i)_{i\in N}\in\mathbb{R}_+^N\setminus\{\mathbf{0}\}$.

Assumption 1 implies that, for every nonzero realization $c=(h_i(\theta_i))_{i\in N}$ of the merchants' score vector, the merchant with the highest CTR-weighted score is unique for almost every customer ω in Ω. In particular, the assumption holds when ω_i's are mutually

independent and each has an atomless distribution, and, more generally, whenever γ_α is absolutely continuous with respect to Lebesgue measure on $[0,1]^{|N|}$.

Proposition 1. *Suppose Assumption 1 holds. There exists an optimal scoring mechanism as in Theorem 1. If the vector of ironing parameters $z = (z_i)_{i\in N} \neq \mathbf{0}$, then, for F-almost every type profile, ties occur only on γ_α-null sets and the optimal customer-specific tie-breaking rule is given by*

$$\pi_i(\theta,\omega) = \mathbf{1}_{E_i(z)}(\theta,\omega) \quad (F\otimes\gamma_\alpha)\text{-almost everywhere}, \tag{8}$$

where $E_i(z) \subseteq \Theta\times\Omega$ denotes the set of type profile and customer pairs (θ,ω) for which merchant i attains the highest CTR-weighted score $\omega_i\overline{\phi}_{\eta,i}(\theta_i;z_i)$, that is,

$$E_i(z) := \Big\{(\theta,\omega)\in\Theta\times\Omega :\ \omega_i\overline{\phi}_{\eta,i}(\theta_i;z_i) > \omega_j\overline{\phi}_{\eta,j}(\theta_j;z_j) \text{ for all } j\in N,\ j\neq i\Big\}.$$

Moreover, if the critical worst-off type $\hat{\theta}_i(z_i)$ is interior, then z_i solves

$$a_i = \int_\Omega \omega_i\, \mathbb{P}_{F_{-i}}\Big(\omega_i z_i > \omega_j\overline{\phi}_{\eta,j}(\theta_j;z_j) \text{ for all } j\neq i\Big)\, d\gamma_\alpha(\omega). \tag{9}$$

Fix type profile θ in Θ, and let $E_i^\theta(z) := E_i(z)\cap(\{\theta\}\times\Omega)$ be the set of customers at which merchant i of type θ_i has the highest CTR-weighted score. Assumption 1 implies that, whenever there exists at least one merchant with nonzero score $\overline{\phi}_{\eta,j}(\theta_j;z_j)\neq 0$, the collection of sets $\big\{E_i^\theta(z)\big\}_{i\in N}$ partition Ω up to a γ_α-null set. We can therefore ignore ties and completely characterize targeted ads according to (8). Thus determining the optimal mechanism reduces to finding ironing parameters $z = (z_i)_{i\in N}$ which are solutions to a system of at most $|N|$ equations given by (9).

3.3 Proof Sketch of the Main Results

The proofs of Theorem 1 and Proposition 1 have three main steps. After applying standard Myersonian techniques, we write (IC) in terms of expected clicks Q_i and a critical type θ'. At the optimum, we determine a critical worst-off type, and ensure that (IR) is satisfied as long as the worst-off type's participation constraint is satisfied. Second, we formulate the designer's problem as a max-min problem akin to Loertscher and Wasser (2019) and establish a saddle point consisting of the worst-off types and targeted ads. Third, we apply ironing techniques of Myerson (1981) to characterize the saddle point.

Given (interim) payoffs $U_i(\theta_i)$ defined by

$$U_i(\theta_i) := \theta_i(Q_i(\theta_i) - a_i) - T_i(\theta_i),$$

by the envelope theorem, (IC) holds if and only if

$$U_i(\theta_i) = U_i\left(\theta_i'\right) + \int_{\theta_i'}^{\theta_i} (Q_i(y) - a_i)\, dy, \quad \forall \theta_i, \theta_i' \in \Theta_i,$$

$$Q_i(\theta_i) \text{ is nondecreasing in } \theta_i. \tag{M}$$

Hence, for any critical type θ_i', transfers are pinned down by

$$T_i(\theta_i) = \theta_i(Q_i(\theta_i) - a_i) - \int_{\theta_i'}^{\theta_i} (Q_i(y) - a_i)\, dy - U_i\left(\theta_i'\right). \tag{10}$$

For a critical vector $\theta' \in \Theta$, define the corresponding *virtual type* by

$$\phi_{\eta,i}\left(\theta_i; \theta_i'\right) = \begin{cases} \phi_{\eta,i}^S(\theta_i) & \text{if } \theta_i < \theta_i', \\ \phi_{\eta,i}^B(\theta_i) & \text{if } \theta_i \geq \theta_i', \end{cases}$$

and the corresponding *virtual objective* by

$$\psi_\eta\left(x, \theta'\right) = \sum_{i \in N} \mathbb{E}_{F_i}\left[(Q_i(\theta_i) - a_i)\phi_{\eta,i}\left(\theta_i; \theta_i'\right)\right].$$

Substituting (10) into the designer's objective function and integrating by parts gives

$$\Psi_\eta(x, t) = \psi_\eta\left(x, \theta'\right) - \eta_r \sum_{i \in N} U_i\left(\theta_i'\right), \quad \forall \theta' \in \Theta. \tag{11}$$

The difficulty is that a merchant may be either a net seller or a net buyer, so the worst-off type is not known ex-ante. Following Loertscher and Wasser (2019), for a fixed targeted ads x the set of worst-off types, defined as the product set of types θ_i with least payoffs $U_i(\theta_i)$, is given by

$$\widehat{\Theta}(x) = \operatorname*{argmin}_{\theta' \in \Theta} \psi_\eta\left(x, \theta'\right)$$

and therefore the designer's problem can be formulated as a max-min

$$\max_{x \in \mathscr{X}} \min_{\theta' \in \Theta} \psi_\eta\left(x, \theta'\right)$$

where $\mathscr{X}$ denotes the space of targeted ads x that satisfy (Fsb) and (M).

Then, showing that $\mathscr{X}$ is compact, and applying Fan's minimax theorem,[15] we show that there exists a saddle point $\left(x^*, \hat{\theta}\right)$ such that $\hat{\theta} \in \widehat{\Theta}(x^*)$ is a worst-off type and that

$$\max_{x \in \mathscr{X}} \psi_\eta\left(x, \hat{\theta}\right) = \psi_\eta\left(x^*, \hat{\theta}\right) = \min_{\theta' \in \Theta} \psi_\eta\left(x^*, \theta'\right).$$

Therefore, fixing a worst-off type $\hat{\theta}$, the optimal targeted ads x solves

$$\max_{x \in \mathscr{X}} \psi_\eta\left(x, \hat{\theta}\right). \tag{12}$$

[15] Fan's minimax theorem extends von Neumann's result from bilinear functions on finite-dimensional simplices to concave-convexlike functions on more general (Hausdorff) spaces; see Fan (1953, Theorem 1).

However, when $\hat{\theta}_i$ is interior, the virtual type function $\phi_{\eta,i}\big(\theta_i;\hat{\theta}_i\big)$ is generally non-monotone. Therefore, pointwise maximization of (12) need not satisfy the monotonicity condition (M).

To ensure monotonicity of expected clicks $Q_i(\theta_i)$, we apply the ironing technique of Myerson (1981). Let $z_i = z_i\big(\hat{\theta}_i\big)$ be the ironing parameter satisfying

$$\mathbb{E}\big[\overline{\phi}_{\eta,i}(\theta_i; z_i)\big] = \mathbb{E}\big[\phi_{\eta,i}\big(\theta_i;\hat{\theta}_i\big)\big] = \eta_w + \eta_v \mathbb{E}[\theta_i] + \eta_r \hat{\theta}_i.$$

For either CTR environment, the problem in (12) is then equivalent to maximizing

$$\mathbb{E}\left[\sum_{i\in N}\Big(\int_\Omega \omega_i\, x_i(\theta; d\omega) - a_i\Big)\overline{\phi}_{\eta,i}(\theta_i; z_i)\right],$$

subject to the same feasibility constraints. Because the scoring functions $\overline{\phi}_{\eta,i}(\theta_i; z_i)$ are nondecreasing, any pointwise maximizer generates nondecreasing expected clicks. When Ω is finite, the integral separates across CTR realizations. Fixing $\omega \in \Omega$ and a type profile $\theta \in \Theta$ gives the pointwise problem

$$\max_{\{x_i^\omega(\theta)\}_{i\in N}} \sum_{i\in N} x_i^\omega(\theta)\, \omega_i \overline{\phi}_{\eta,i}(\theta_i; z_i) \quad \text{subject to} \quad \sum_{i\in N} x_i^\omega(\theta) \le \gamma_\alpha(\omega), \quad x_i^\omega(\theta) \ge 0.$$

Thus, all mass $\gamma_\alpha(\omega)$ is assigned to merchants with the highest nonnegative CTR-weighted ironed score. This is exactly the scoring mechanism in (5), with $x_i^\omega(\theta) = \gamma_\alpha(\omega)\pi_i^\omega(\theta)$ and customer-specific tie-breaking rule π^ω chosen so that $\hat{\theta}$ is a worst-off type. Transfers are recovered from (10) evaluated at $\theta' = \hat{\theta}$.

The continuous-CTR characterization follows from the same pointwise argument, with the measure-theoretic details deferred to the appendix. Lastly, if the critical worst-off type $\hat{\theta}_i$ is interior, then $Q_i\big(\hat{\theta}_i\big) = a_i$. Substituting (8) into $Q_i\big(\hat{\theta}_i\big)$ and applying Tonelli's theorem to interchange the order of integration, then gives (9).

4 Limit Market

In this section, we directly formulate the continuum-of-merchants market (henceforth, *limit market*). The limit market preserves the key economic trade-offs from the finite model while eliminating market-level uncertainty, thereby yielding a tractable indirect implementation of the optimal mechanism. We characterize the optimal mechanism in Sec. 4.2, study efficiency in Sec. 4.3, and provide proof sketches in Sec. 4.4. In the Online Appendix we show that the finite market converges to this limit market along a suitably defined sequence of finite markets.

4.1 Model

There is a unit mass of merchants partitioned into finitely many observable tiers indexed by $k \in K$. Tier k has mass $\beta_k > 0$, with $\sum_k \beta_k = 1$. Each merchant in tier k has a private type $\theta \in \Theta_k$ drawn from regular distribution F_k supported on $\Theta_k = \left[\underline{\theta}_k, \overline{\theta}_k\right] \subset \mathbb{R}_+$. Let $\Theta := \prod_{k \in K} \Theta_k$ denote the type space.

There is also a unit mass of customers, initially distributed across merchants' platform-built audiences. We normalize initial audience contributions proportionally to merchant mass: any group of merchants with mass $\beta \in [0,1]$ collectively contributes audience mass β.[16] Thus, tier k contributes audience mass β_k, and $\sum_{k \in K} \beta_k = 1$ implies that the pooled audience has unit mass. By advertising to her own audience, a tier-k merchant generates $\mu_k \in (0,1)$ expected clicks, giving a type θ merchant an outside option of $\mu_k \theta$.

In the limit market, we impose supply capacity constraints to rule out cases in which a single merchant serves an arbitrarily large number of customers relative to her market size. We normalize this capacity to one, meaning that each merchant can meet at most one unit of demand, measured by the expected clicks generated by her ads. By participating in the platform, a tier-k merchant can therefore increase her expected clicks from μ_k to at most one, achieving a maximum potential targeting gain of $1 - \mu_k$.

Each customer in the pooled audience is represented by *tier-specific CTRs* (henceforth, *CSRs*)

$$\omega = (\omega_k)_{k \in K} \in [0,1]^K$$

where ω_k is the probability that the customer clicks on an ad from a tier-k merchant. CSRs ω are observable and distributed according to $G \in \Delta\left([0,1]^K\right)$. We assume that the pooled audience weakly outperforms each tier's initial audience in terms of expected clicks:

$$\mathbb{E}_G[\omega_k] \geq \mu_k \quad \text{for all } k \text{ in } K. \tag{13}$$

Mechanism Design. As in the finite market, we restrict attention to direct mechanisms. Let $\mathcal{M}_+\left([0,1]^K\right)$ denote the space of finite nonnegative Borel measures on $[0,1]^K$. A direct mechanism (x,t) consists of a tier-specific targeted ads allocation $x = (x_k)_{k \in K}$ (henceforth, targeted ads) and transfers $t = (t_k)_{k \in K}$ such that $x_k : \Theta_k \to \mathcal{M}_+\left([0,1]^K\right)$ and $t_k : \Theta_k \to \mathbb{R}$ are Borel measurable for every $k \in K$. For each $\theta \in \Theta_k$ and Borel set $B \subseteq [0,1]^K$, $x_k(\theta; B)$ is

[16] If audience contributions vary within tier k, with type θ merchants contributing $a_k(\theta) > 0$, redefine the type distribution by

$$d\widetilde{F}_k(\theta) := \frac{a_k(\theta)}{\int_{\Theta_k} a_k(y)\, dF_k(y)} dF_k(\theta).$$

the per-merchant mass of customers with CSR vectors in B assigned to a representative type θ merchant in tier k. The direct mechanism must satisfy the usual incentive-compatibility and individual-rationality constraints:

$$\theta q_k(\theta) - t_k(\theta) \geq \theta q_k\left(\tilde{\theta}\right) - t_k\left(\tilde{\theta}\right) \qquad \text{for all } \theta, \tilde{\theta} \in \Theta_k \text{ and } k \text{ in } K \tag{IC$_k$}$$

$$\theta q_k(\theta) - t_k(\theta) \geq \mu_k \theta \qquad \text{for all } \theta \in \Theta_k \text{ and } k \text{ in } K. \tag{IR$_k$}$$

where $q_k(\theta)$ denotes the *expected clicks* assigned to a type θ merchant in tier k, defined by

$$q_k(\theta) := \int_{[0,1]^K} \omega_k \, x_k(\theta; d\omega), \tag{14}$$

which are subject to the capacity constraints

$$q_k(\theta) \leq 1 \quad \text{for every } \theta \in \Theta_k \text{ and } k \in K. \tag{C$_k$}$$

Lastly, because there is no market-level uncertainty, targeted ads x are subject to aggregate feasibility constraints

$$\sum_{k \in K} \beta_k \int_{\Theta_k} x_k(\theta; B) \, dF_k(\theta) \leq G(B) \quad \text{for every Borel set } B \subseteq [0,1]^K. \tag{G-Fsb}$$

The left-hand side aggregates per-merchant assignments across the mass β_k of each tier, while the right-hand side is the mass of customers in the pooled audience with CSR vectors in B. Unlike (Fsb), (G-Fsb) does not require each type-level allocation measure $x_k(\theta; \cdot)$ to be absolutely continuous with respect to G for every type.

We show next that (G-Fsb) admits a simpler reduced-form characterization using a family of linear inequalities in expected clicks $q = (q_k)_{k \in K}$.

Definition 2 (Implementable Expected Clicks)**.** *An expected clicks allocation $q = (q_k)_{k \in K}$, where $q_k : \Theta_k \to \mathbb{R}_+$ for every $k \in K$, is implementable if there exist targeted ads $x = (x_k)_{k \in K}$ that satisfy* (G-Fsb) *and yield q as in* (14)*.*

Proposition 2 (Reduced-Form Feasibility)**.** *An expected clicks allocation $q = (q_k)_{k \in K}$, where $q_k : \Theta_k \to \mathbb{R}_+$ for every $k \in K$, is implementable if and only if*

$$\sum_{k \in K} \lambda_k \, \beta_k \, \mathbb{E}_{F_k}[q_k(\theta)] \leq \mathbb{E}_G\left[\max_{k \in K}\{\lambda_k \omega_k\}\right] \quad \textit{for every } \lambda \in \mathbb{R}^K_+. \tag{R-Fsb}$$

We interpret each λ_k as the shadow price per expected click in tier k (henceforth, shadow price). The right-hand side of (R-Fsb) therefore gives the largest shadow value that the pooled audience can generate at the given shadow prices λ. Thus, (R-Fsb) requires the value generated from allocating ads across all tiers to remain below this bound for every possible vector of shadow prices. Lastly, note that the expected clicks $q = (q_k)_{k \in K}$ that ensure the outside option $q_k \equiv \mu_k$ for all k in K are implementable since they satisfy

(R-Fsb): for every $\lambda \in \mathbb{R}^K_+$,

$$\sum_{k\in K} \lambda_k \beta_k \mu_k \leq \sum_{k\in K} \lambda_k \beta_k \mathbb{E}_G[\omega_k] \leq \max_{k\in K}\{\lambda_k \mathbb{E}_G[\omega_k]\} \leq \mathbb{E}_G\left[\max_{k\in K}\{\lambda_k \omega_k\}\right],$$

where the first inequality uses (13), the second uses $\sum_{k\in K} \beta_k = 1$, and the third applies Jensen's inequality since the maximum is a convex function.

Designer's Problem. For given expected clicks $q = (q_k)_{k\in K}$, define merchant surplus and customer engagement for type θ in tier k by

$$v_k(\theta) := \theta(q_k(\theta) - \mu_k) \quad \text{and} \quad w_k(\theta) := q_k(\theta) - \mu_k,$$

respectively. Given welfare weights $\eta = (\eta_v, \eta_w, \eta_r) \in \mathbb{R}^3_+$ satisfying $\eta_v + \eta_w + \eta_r = 1$ and $\eta_r > 0$, the designer's reduced-form problem is

$$\begin{aligned} \max_{q,t} \quad & \sum_{k=1}^{K} \beta_k \mathbb{E}_{F_k}[\eta_v v_k(\theta) + \eta_w w_k(\theta) + \eta_r t_k(\theta)] \\ \text{subject to} \quad & \text{(R-Fsb) and } q_k(\theta) \geq 0 \text{ for every } \theta \in \Theta_k \text{ and } k \in K, \\ & (\mathrm{C}_k), (\mathrm{IC}_k), \text{ and } (\mathrm{IR}_k) \text{ for every } k \in K. \end{aligned}$$

4.2 Optimal Mechanism

The designer solves a separate screening problem within each tier, while ensuring that the allocation across tiers is feasible.

First, we impose a tie-free condition on G, analogous to the condition imposed on continuous CTRs in the finite market. The condition simplifies the characterization of the optimal mechanism, although the results hold for more general distributions.

Assumption 2 (Tie-Free CSRs)**.** *We say that G is tie-free if*

$$\operatorname*{argmax}_{k\in K}\{\lambda_k \omega_k\}$$

is a singleton G-almost everywhere for every $\lambda \in \mathbb{R}^K_+ \setminus \{\mathbf{0}\}$.

The tie-free CSR condition therefore gives a unique winning tier for almost every CSRs ω. If the condition fails, the same within-tier optimal market structure holds, but the designer may need to split a positive mass of tied CSRs ω across tiers.

Next, for each tier k, define the weighted virtual cost and value functions by

$$\phi^S_{\eta,k}(\theta) := \eta_w + \eta_v \theta + \eta_r\left(\theta + \frac{F_k(\theta)}{f_k(\theta)}\right) \quad \text{and} \quad \phi^B_{\eta,k}(\theta) := \eta_w + \eta_v \theta + \eta_r\left(\theta - \frac{1 - F_k(\theta)}{f_k(\theta)}\right),$$

respectively. Thus, $\mathrm{MC}_k(\theta) = \phi^S_{\eta,k}(\theta)$ is the marginal cost of acquiring expected clicks from type θ, while $\mathrm{MR}_k(\theta) = \phi^B_{\eta,k}(\theta)$ is the marginal revenue from assigning expected clicks to type θ. For each $j = B, S$ and k in K, we use the inverse convention that $\left(\phi^j_{\eta,k}\right)^{-1}(z) = \underline{\theta}_k$ whenever $z < \phi^j_{\eta,k}(\underline{\theta}_k)$ and $\left(\phi^j_{\eta,k}\right)^{-1}(z) = \overline{\theta}_k$ whenever $z > \phi^j_{\eta,k}\left(\overline{\theta}_k\right)$.

We state next one more condition on the model primitives that eliminates the optimality of the degenerate market structure in which no merchant is treated as a net seller. Let $\underline{p}^B_k$ be the cutoff type for which the marginal revenue is zero, that is,

$$\underline{p}^B_k := (\mathrm{MR}_k)^{-1}(0).$$

Assumption 3 (Pooled-Audience Scarcity)**.** *The tier-specific type distributions $(F_k)_{k\in K}$ and the pooled-audience distribution G jointly satisfy*

$$\sum_{k\in K} \beta_k \left[1 - (1-\mu_k) F_k\left(\underline{p}^B_k\right)\right] > \mathbb{E}_G\left[\max_{k\in K} \omega_k\right].$$

The left-hand side of Assumption 3 gives total demand generated by the exchange and buying markets with cutoff types $\underline{\theta}_k$ and $(\mathrm{MR}_k)^{-1}(0)$, respectively. Its right-hand side is the largest total number of expected clicks that the pooled audience can generate under G. Thus, the condition makes the pooled audience scarce relative to demand: the tier-specific screening problems cannot all be supported when expected clicks have zero shadow prices ($\lambda = \mathbf{0}$).

We are now ready to state our main theorem.

Theorem 2 (Optimal (Indirect) Mechanism)**.** *Suppose Assumptions 2 and 3 hold. An optimal mechanism exists and can be implemented through three separate markets within each tier $k \in K$. In particular, there exist per-click posted prices $\underline{\theta}_k \le p^S_k \le p^B_k \le \overline{\theta}_k$ such that*

i. (selling market) types $\theta < p^S_k$ receive no targeted ads and are compensated by the lump-sum transfer $t_k(\theta) = -\mu_k p^S_k$;

ii. (exchange market) types θ in $\left[p^S_k, p^B_k\right]$ receive μ_k expected clicks and zero transfers;

iii. (buying market) types $\theta > p^B_k$ receive targeted ads that yield one unit of expected clicks and pay the lump-sum transfer $t_k(\theta) = (1-\mu_k) p^B_k$;

where the prices are uniquely pinned down by strictly positive shadow prices $\lambda = (\lambda_k)_{k\in K}$ that lie in $\prod_{k\in K}\left(0, \phi^S_{\eta,k}\left(\overline{\theta}_k\right)\right]$ and satisfy

$$p^j_k = \left(\phi^j_{\eta,k}\right)^{-1}(\lambda_k) \quad \text{for all } j = B, S \text{ and } k \text{ in } K, \qquad (\mathrm{MR}_k{=}\mathrm{MC}_k)$$

where the shadow price λ_k is unique whenever at least one of p_k^S and p_k^B is interior. The shadow prices are such that demand and supply clears in each tier: for every $k \in K$,

$$\beta_k\Big[1 - F_k\Big(p_k^B\Big) + \mu_k\Big(F_k\Big(p_k^B\Big) - F_k\Big(p_k^S\Big)\Big)\Big] = \mathbb{E}_G\left[\omega_k \mathbf{1}\left\{k = \operatorname*{argmax}_{j\in K} \lambda_j\omega_j\right\}\right]. \qquad (\mathrm{D}_k{=}\mathrm{S}_k)$$

Remark 1. *The LHS of ($\mathrm{D}_k{=}\mathrm{S}_k$) is demand in tier k, $\beta_k \mathbb{E}_{F_k}\left[q_k^\lambda(\theta)\right]$, where we let q_k^λ denote the expected clicks allocated accordingly in the exchange and buying markets with posted prices determined according to shadow prices λ_k. Then multiplying the market-clearing conditions in ($\mathrm{D}_k{=}\mathrm{S}_k$) by the corresponding optimal shadow prices and summing over k gives binding (R-Fsb) at the optimal λ, that is,*

$$\sum_{k\in K} \lambda_k\beta_k \mathbb{E}_{F_k}\left[q_k^\lambda(\theta)\right] = \mathbb{E}_G\left[\max_{k\in K}\{\lambda_k\omega_k\}\right].$$

Remark 2. *Let $\overline{\lambda}_k := \phi_{\eta,k}^S\left(\overline{\theta}_k\right)$. Whenever $\lambda_k \geq \overline{\lambda}_k$, both cutoff types equal $\overline{\theta}_k$, so tier k demands no expected clicks. Raising λ_k further can only increase the tier's priority in the allocation of the pooled audience, even though its demand remains zero. Such an increase is therefore unnecessary: lowering λ_k to $\overline{\lambda}_k$ leaves tier-k demand unchanged and weakly redirects the pooled audience toward tiers with positive demand. Hence, an optimal shadow price vector can always be selected with $\lambda_k \leq \overline{\lambda}_k$.*

The characterization in Theorem 2 allows a tier to consist entirely of sellers when its equilibrium demand and supply are both zero. We next impose a richness condition ensuring that every tier receives positive supply at every strictly positive shadow price vector and show that market clearing then rules out such all-seller tiers.

Assumption 4 (Tier-wise CSR Richness)**.** *For every $\lambda \in \mathbb{R}_{++}^K$ and every $k \in K$, tier k has a positive probability of a sufficiently strong comparative CSR advantage:*

$$G\left(\left\{\omega \in [0,1]^K : \lambda_k\omega_k > \max_{j\neq k} \lambda_j\omega_j\right\}\right) > 0.$$

Assumption 4 holds whenever G admits a positive density on $(0,1)^K$. Now, by Theorem 2 and Assumption 4, every tier receives strictly positive supply in the optimal mechanism, so market clearing requires its demand to be strictly positive. Then clearly $p_k^S < \overline{\theta}_k$. This observation immediately yields the following corollary.

Corollary 1 (No All-Seller Tiers)**.** *Suppose in addition that Assumption 4 holds. Then every optimal shadow price vector satisfies*

$$0 < \lambda_k^* < \phi_{\eta,k}^S\left(\overline{\theta}_k\right) \quad \text{for every } k \in K.$$

Consequently,

$$p_k^S < \overline{\theta}_k \quad \textit{for every } k \in K,$$

so every tier contains an exchange market, a buying market, or both; in particular, no tier consists only of sellers.

Discussion of Results. Theorem 2 has two main takeaways. The first is the role of λ_k, an endogenous *shadow price* per unit of expected clicks in tier k. The second is the optimal three-market composition within each tier.

First, the shadow price λ_k plays two related roles. Within tier k, it determines the selling and buying cutoff types. Whenever both cutoffs are interior, the marginal cost of acquiring expected clicks from the selling market and the marginal revenue from allocating them to the buying market both equal λ_k. More generally, the cutoff types satisfy (MR_k=MC_k), including the relevant boundary cases. Across tiers, the shadow prices $\lambda = (\lambda_k)_{k\in K}$ couple the tier-specific screening problems through the market-clearing conditions in (D_k=S_k). Except on a measure-zero set of ties, the designer assigns each CSR vector $\omega = (\omega_k)_{k\in K}$ to the tier that maximizes its shadow value, $\lambda_k\omega_k$.

Holding the shadow prices of the other tiers fixed, an increase in λ_k increases the supply of expected clicks to tier k. At the same time, it raises both cutoff types and thereby reduces the tier's demand for expected clicks. In equilibrium, the strictly positive shadow prices $\lambda = (\lambda_k)_{k\in K}$ equate demand and supply in every tier.

Although the shadow price determines the locations of the cutoff types, the three-market structure itself emerges from the optimal screening problem within each tier. Low types give up their platform-built audiences and receive compensation, intermediate types exchange them for targeted ads generating the same expected clicks, and high types purchase the additional $1 - \mu_k$ expected clicks up to capacity. Intermediate types are pooled: they receive their outside option expected clicks, pay no transfer, and earn zero surplus relative to their outside option. By contrast, merchants in the selling and buying markets earn positive rents. Note that, intermediate types need not receive the same customers. The designer may substitute its platform-built audience with a different group of customers that generates the same expected clicks.

Taken together, these results provide an indirect implementation of the optimal mechanism. Posted prices induce merchants to self-select into the three markets, while shadow prices acts as a visible hand that decentralizes the allocation of targeted ads across tiers and ensure that demand equals supply in every tier.

Discussion of Tie-Free CSRs. The tie-free CSR assumption makes the allocation of the pooled audience across tiers uniquely determined by the shadow price vector: almost every customer is assigned to the unique tier maximizing $\lambda_k \omega_k$. It is not needed for the fixed-shadow screening problem within each tier, and hence not for the resulting selling, exchange, and buying markets. When ties occur with positive probability, the same cutoff and transfer formulas continue to apply, but the pooled audience may be divided among shadow value maximizing tiers to clear their respective markets.

4.3 Merchant and Customer Efficiency

We now study when merchant- and customer-efficient platform design is possible without an outside subsidy, that is, when it can be implemented by a mechanism satisfying the no-deficit constraint:

$$\sum_{k \in K} \beta_k \mathbb{E}_{F_k}[t_k(\theta)] \geq 0. \tag{ND}$$

For merchant and customer efficiency, the designer puts zero weight $\eta_r = 0$ on revenues and nonnegative weights $(\eta_v, \eta_w) \geq \mathbf{0}$ on merchant surplus and customer engagement. We call the resulting platform design *merchant- and customer-efficient*; when $\eta_v = 1$, so the designer places all weight on merchant surplus, we call it *merchant-efficient*. Thus, expected clicks $q = (q_k)_{k \in K}$ are merchant- and customer-efficient if they maximize

$$\sum_{k \in K} \beta_k \mathbb{E}_{F_k}[(\eta_w + \eta_v \theta)(q_k(\theta) - \mu_k)] \tag{15}$$

subject to (R-Fsb) and (C_k).

Note that $\underline{p}_k^B = \underline{\theta}_k$ when $\eta_r = 0$. Hence, the pooled-audience scarcity condition in Assumption 3 reduces to

$$\mathbb{E}_G\left[\max_{k \in K} \omega_k\right] < 1. \tag{ES}$$

Condition (ES) means that the pooled audience is scarce relative to the allocation that provides one expected clicks to every merchant, exhausting everyone's supply capacity.

Proposition 3 (Merchant and Customer Efficiency without an Outside Subsidy)**.** *Suppose $\eta_v > 0$, and both Assumption 2 and (ES) hold. Then merchant- and customer-efficient expected clicks $q^E = \left(q_k^E\right)_{k \in K}$ are implementable by a mechanism satisfying (IC_k), (IR_k), and (ND) if and only if*

$$\sum_{k \in K} \beta_k p_k^E \left[1 - \mu_k - F_k\left(p_k^E\right)\right] \geq 0, \tag{16}$$

where the corresponding cutoff types $p^E = (p_k^E)_{k\in K}$ *are determined by*

$$p_k^E = \min\left\{\overline{\theta}_k,\ \max\left\{\underline{\theta}_k,\ \frac{\lambda_k^E - \eta_w}{\eta_v}\right\}\right\} \quad \textit{for every } k \in K, \tag{17}$$

for some shadow prices $\lambda^E = \left(\lambda_k^E\right)_{k\in K} \in \mathbb{R}^K_{++}$ *that satisfy market clearing conditions*

$$\beta_k\left[1 - F_k\left(p_k^E\right)\right] = \mathbb{E}_G\left[\omega_k \mathbf{1}\left\{k = \operatorname*{argmax}_{j\in K} \lambda_j^E \omega_j\right\}\right] \quad \textit{for every } k \in K. \tag{D_k^E=S_k^E}$$

The merchant- and customer-efficient cutoff vector p^E *is uniquely pinned down, and the shadow price* λ_k^E *is unique whenever* p_k^E *is interior.*

In tier k, high types above p_k^E pay for the additional $1-\mu_k$ expected clicks they receive, whereas low types below p_k^E are compensated for giving up μ_k expected clicks. Hence, tier k makes the net contribution

$$\beta_k p_k^E\left[1 - F_k\left(p_k^E\right) - \mu_k\right]$$

to the designer's budget. Summing these across tiers gives (16). Importantly, the no-deficit condition need not hold tier by tier: deficits in some tiers may be financed by surpluses in others, provided that total expected transfers are nonnegative.

Next we show that merchant-efficient platform design is always possible without an outside subsidy.

Proposition 4 (Merchant-Efficient Platform Design)**.** *Suppose Assumption 2 and* (ES) *hold, and let* $\eta_v = 1$*. Then merchant-efficient expected clicks* $q^E = \left(q_k^E\right)_{k\in K}$ *are implementable by a mechanism satisfying* (IC_k)*,* (IR_k)*, and* (ND)*, and hence without an outside subsidy.*

The result contrasts with the efficient partnership dissolution benchmark of Cramton et al. (1987). In their environment, sufficiently symmetric partnerships can always be dissolved efficiently, whereas efficient dissolution may fail when the initial ownership shares are sufficiently asymmetric. Here, by contrast, the masses β_k of merchants and the outside option expected clicks μ_k need not be sufficiently symmetric for merchant efficiency to be possible. Conceptually, partnership dissolution reallocates a fixed asset among finitely many agents, whereas our model reallocates a continuum of customers across a continuum of merchants. Accordingly, our analysis is ex-post at the market level unlike the interim incentive considerations that arise in Cramton et al. (1987).

4.4 Proof Sketch of the Main Results

We now present the proof sketches of optimal mechanism and efficiency results, and provide the remaining details in Appendix B.

4.4.1 Proof Sketch of Optimal Mechanism

The proof of Theorem 2 has three main steps. First, we represent aggregate feasibility constraints by a finite-dimensional demand set and then derive the reduced-form feasibility constraints. Second, we introduce a pseudo-Lagrangian and apply Fan's minimax theorem to obtain strong duality, after which we solve the resulting problem for every fixed shadow price vector λ. Third, we minimize the reduced value over a compact set of shadow prices and use the first-order conditions to establish optimal market clearing shadow prices. We state only the main definitions and arguments here and leave the remaining details to Appendix B.

Reduced-Form Feasibility. Let $d_k(q)$ denote total expected clicks demand in tier k (henceforth, demand), defined by

$$d_k(q) := \beta_k \mathbb{E}_{F_k}[q_k(\theta)].$$

Then define the feasible demand set $\mathcal{D}(G) \subset \mathbb{R}^K_+$ by

$$\mathcal{D}(G) := \left\{ \begin{array}{c} d \in \mathbb{R}^K_+ : \text{there exists a } G\text{-measurable } m : [0,1]^K \to \mathbb{R}^K_+ \text{ such that} \\ \sum_{k \in K} m_k(\omega) \leq 1 \quad G\text{-a.e.}, \\ d_k \leq \int_{[0,1]^K} \omega_k m_k(\omega)\, dG(\omega) \quad \text{for every } k \in K. \end{array} \right\} \tag{18}$$

Here, $m_k(\omega)$ is the tier-specific share of customers with CSR vector ω assigned to tier k. Our first step is showing that implementability is equivalent to requiring feasibility in $\mathcal{D}(G)$:

Lemma 1. *The expected clicks allocation $q = (q_k)_{k \in K}$ is implementable if and only if $d(q) \in \mathcal{D}(G)$.*

Henceforth, we work with the feasible demand set $\mathcal{D}(G)$, whose key properties are summarized next.

Lemma 2. *The set $\mathcal{D}(G)$ is nonempty, convex, compact, and downward closed, that is, if $d \in \mathcal{D}(G)$ and $0 \leq \hat{d} \leq d$ then $\hat{d} \in \mathcal{D}(G)$.*

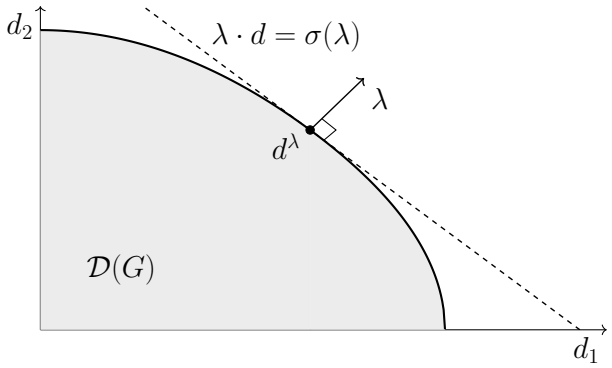


Figure 4: Feasible demand set $\mathcal{D}(G)$ for $K = 2$.

Now let $\sigma : \mathbb{R}^K_+ \to \mathbb{R}_+$ denote the support function of $\mathcal{D}(G)$ restricted on $\mathbb{R}^K_+$, defined by

$$\sigma(\lambda) := \max_{d \in \mathcal{D}(G)} \lambda \cdot d \quad \text{for every } \lambda \in \mathbb{R}^K_+.$$

Lemma 3. *For every $\lambda \in \mathbb{R}^K_+ \setminus \{\mathbf{0}\}$, the support function $\sigma(\lambda)$ satisfies*

$$\sigma(\lambda) = \mathbb{E}_G\Big[\max_{k \in K}\{\lambda_k \omega_k\}\Big],$$

which is attained by the demand vector $d^\lambda \in \mathcal{D}(G)$ defined by $d^\lambda_k := \mathbb{E}_G\Big[\omega_k m^\lambda_k(\omega)\Big]$, with tie-breaking rule chosen such that

$$m^\lambda_k(\omega) := \mathbf{1}\left\{k = \min \operatorname*{argmax}_{j \in K} \lambda_j \omega_j\right\} \quad G\text{-almost everywhere.}$$

Fig. 4 illustrates the feasible demand set when there are two tiers and $\omega_1, \omega_2 \overset{\text{i.i.d.}}{\sim} \text{Uniform}[0,1]$. The shaded region represents $\mathcal{D}(G)$, which is a nonempty, convex, compact, and downward closed, and the solid curve traces its upper frontier. In direction $\lambda = (1,1)$, the support function attains the value $\sigma(1,1) = 2/3$ at $d^\lambda = (1/3, 1/3)$; accordingly, the dashed line $(1,1) \cdot d = 2/3$ supports $\mathcal{D}(G)$ at d^λ. The corresponding support-function and frontier calculations are provided in Appendix C.3.

Using Lemmas 1 to 3 we then prove Proposition 2.

Pseudo-Lagrangian Formulation. Let $\mathcal{Q}^0_k$ be the set of nondecreasing measurable functions $q_k : \Theta_k \to [0,1]$, and set

$$\mathcal{Q}^0 := \prod_{k \in K} \mathcal{Q}^0_k.$$

Let $\mathcal{Q} \subseteq \mathcal{Q}^0$ denote the subset satisfying (R-Fsb). For a vector of critical types $\theta' = (\theta'_k)_{k\in K} \in \Theta$, define the tier-$k$ virtual type by

$$\phi_{\eta,k}(\theta;\theta'_k) := \begin{cases} \phi^S_{\eta,k}(\theta), & \theta < \theta'_k, \\ \phi^B_{\eta,k}(\theta), & \theta \geq \theta'_k, \end{cases}$$

and the virtual objective by

$$\psi_\eta(q,\theta') := \sum_{k\in K} \beta_k \mathbb{E}_{F_k}[(q_k(\theta) - \mu_k)\phi_{\eta,k}(\theta;\theta'_k)]. \tag{19}$$

Thus, similarly as in the finite market, the designer's problem reduces to

$$\sup_{q\in\mathcal{Q}} \inf_{\theta'\in\Theta} \psi_\eta(q,\theta'). \tag{20}$$

Now fix $\lambda \in \mathbb{R}^K_+$ and define the feasibility slack from (R-Fsb) by

$$\Delta(q,\lambda) := \sigma(\lambda) - \sum_{k\in K} \lambda_k \beta_k \mathbb{E}_{F_k}[q_k(\theta)]$$

and the *pseudo-Lagrangian*

$$\mathcal{L}(q,\theta',\lambda) := \psi_\eta(q,\theta') + \Delta(q,\lambda). \tag{21}$$

We refer to (21) as a pseudo-Lagrangian because the Lagrangian of the family of constraints in (R-Fsb) assigns multipliers according to some nonnegative measure on $\mathbb{R}^K_+$. However, the pseudo-Lagrangian provides an upper bound on the virtual objective in (19), which we show that it is sufficient to obtain a pseudo-duality structure that solves (20).

Lemma 4. *The designer's problem satisfies*

$$\sup_{q\in\mathcal{Q}} \inf_{\theta'\in\Theta} \psi_\eta(q,\theta') = \sup_{q\in\mathcal{Q}^0} \inf_{(\theta',\lambda)\in\Theta\times\mathbb{R}^K_+} \mathcal{L}(q,\theta',\lambda) = \inf_{(\theta',\lambda)\in\Theta\times\mathbb{R}^K_+} \sup_{q\in\mathcal{Q}^0} \mathcal{L}(q,\theta',\lambda). \tag{22}$$

We next evaluate the dual problem by fixing the shadow price vector.

Constructing Saddle-Points. For each tier $k \in K$, define the lowest and highest feasible ironing values by

$$\underline{z}_k := \left(\phi^B_{\eta,k}(\underline{\theta}_k)\right)_+ \quad \text{and} \quad \overline{z}_k := \phi^S_{\eta,k}\left(\overline{\theta}_k\right),$$

and, for every $\lambda_k \geq 0$, let

$$z_k(\lambda_k) := \min\{\overline{z}_k, \max\{\lambda_k, \underline{z}_k\}\}.$$

Define as in (7) the corresponding critical type

$$\hat{\theta}_k(\lambda_k) := \mathbb{E}_{F_k}\left[\overline{\phi}_{\eta,k}(\theta; z_k(\lambda_k)) - \eta_w - \eta_v\theta\right]/\eta_r \ \in \left[p^S_k(\lambda_k), p^B_k(\lambda_k)\right], \tag{23}$$

where, using the inverse convention introduced before, the cutoff types are defined by

$$p_k^j(\lambda_k) := \left(\phi_{\eta,k}^j\right)^{-1}(\lambda_k), \quad j \in \{S, B\}.$$

Lastly, construct expected clicks by

$$q_k^\lambda(\theta) := \begin{cases} 0, & \theta < p_k^S(\lambda_k), \\ \mu_k, & p_k^S(\lambda_k) \le \theta \le p_k^B(\lambda_k), \\ 1, & \theta > p_k^B(\lambda_k). \end{cases} \tag{24}$$

Now define the value function $\ell : \mathbb{R}_+^K \to \mathbb{R}$ by

$$\ell(\lambda) := \mathcal{L}\left(q^\lambda, \hat{\theta}(\lambda), \lambda\right) \quad \text{for all } \lambda \in \mathbb{R}_+^K. \tag{25}$$

Lemma 5 (Fixed-Shadow Saddle Point)**.** *For every $\lambda \in \mathbb{R}_+^K$, the pair $\left(q^\lambda, \hat{\theta}(\lambda)\right)$ is a saddle point of $\mathcal{L}(\cdot, \cdot, \lambda)$ on $\mathcal{Q}^0 \times \Theta$. Consequently,*

$$\ell(\lambda) = \sup_{q \in \mathcal{Q}^0} \inf_{\theta' \in \Theta} \mathcal{L}(q, \theta', \lambda) = \inf_{\theta' \in \Theta} \sup_{q \in \mathcal{Q}^0} \mathcal{L}(q, \theta', \lambda). \tag{26}$$

Combining Lemmas 4 and 5, the value of the designer's problem is

$$\sup_{q \in \mathcal{Q}} \inf_{\theta' \in \Theta} \psi_\eta(q, \theta') = \inf_{\lambda \in \mathbb{R}_+^K} \left(\inf_{\theta' \in \Theta} \sup_{q \in \mathcal{Q}^0} \mathcal{L}(q, \theta', \lambda) \right) = \inf_{\lambda \in \mathbb{R}_+^K} \ell(\lambda). \tag{27}$$

It remains to establish existence and characterize the minimizing shadow price vector λ^*, and then to verify that the corresponding expected clicks q^{λ^*} attains this value.

Existence and Characterization of Optimal Shadow Prices.

Lemma 6 (Compact Shadow-Price Domain)**.** *Let*

$$\overline{\lambda}_k := \phi_{\eta,k}^S\left(\overline{\theta}_k\right) \quad \textit{and} \quad \Lambda := \prod_{k \in K} \left[0, \overline{\lambda}_k\right].$$

Then

$$\inf_{\lambda \in \mathbb{R}_+^K} \ell(\lambda) = \inf_{\lambda \in \Lambda} \ell(\lambda)$$

and a minimizer $\lambda^* \in \Lambda$ exists.

By (27), any minimizer λ^* satisfies

$$\sup_{q \in \mathcal{Q}} \inf_{\theta' \in \Theta} \psi_\eta(q, \theta') = \ell(\lambda^*). \tag{28}$$

This identifies the optimal value. We next characterize the minimizing shadow price vector and verify that the corresponding constructed expected clicks attains this value.

Lemma 7 (Optimal Shadow Prices)**.** *Suppose Assumptions 2 and 3 hold. There exists a shadow price vector $\lambda^* \in \mathbb{R}^K_+$ that minimizes ℓ, where $\lambda^* \in \mathbb{R}^K_{++}$ and satisfies the market-clearing conditions in* (D_k=S_k)*.*

As shown in the proof of Lemma 7, ℓ is convex, every minimizer is strictly positive, and the fixed-shadow saddle-point envelope theorem gives

$$\frac{\partial \ell(\lambda)}{\partial \lambda_k} = \frac{\partial \sigma(\lambda)}{\partial \lambda_k} - \beta_k \mathbb{E}_{F_k}\left[q_k^\lambda(\theta)\right] \tag{29}$$

for every $\lambda \in \mathbb{R}^K_{++}$. Hence, at any minimizer, the first-order conditions set (29) equal to zero. By Danskin's theorem and the tie-free CSR condition,

$$\frac{\partial \sigma(\lambda^*)}{\partial \lambda_k} = \mathbb{E}_G\left[\omega_k \mathbf{1}\left\{k = \operatorname*{argmax}_{j \in K} \lambda_j^* \omega_j\right\}\right]. \tag{30}$$

Substituting the constructed expected clicks (24) into (29) and using (30) gives the market-clearing conditions in (D_k=S_k).

Transfers are then recovered from the envelope theorem by setting rent equal to zero throughout the exchange market.

4.4.2 Proof of Efficiency Results

The key step is to characterize the solution to the merchant- and customer-efficiency problem in (15).

Lemma 8 (Merchant- and Customer-Efficient Cutoff Characterization)**.** *Suppose $\eta_v > 0$, and both Assumption 2 and* (ES) *hold. There exists $\lambda^E \in \mathbb{R}^K_{++}$ such that every merchant- and customer-efficient expected clicks q^E is given, for every $k \in K$, by*

$$q_k^E(\theta) = \begin{cases} 1, & \theta \geq p_k^E, \\ 0, & \theta < p_k^E, \end{cases} \tag{31}$$

where the cutoff types $p^E = \left(p_k^E\right)_{k \in K}$ are given by (17) and satisfy the market-clearing conditions in (D_k^E=S_k^E)*.*

Given these cutoffs, the envelope theorem with individual rationality binding at p_k^E yields the maximal implementable transfers, whose total gives (16) and proves Proposition 3. The complete argument is provided in Appendix B.2.

5 Applications and Extensions

This section applies the model to other economic settings and then discusses natural extensions of the baseline model.

5.1 Applications

The baseline model is broad in scope and applies to other economic settings in which multiple agents trade multiple heterogeneous goods, commonly referred to as *combinatorial exchanges*. This section demonstrates applications to community time banks and greenhouse gas (GHG) credit markets.

5.1.1 Community Time Banks

Community time banks are reciprocal exchange systems in which members exchange service time Andersson et al. (2021), typically valuing services one-for-one by hours of work rather than money.[17] Members earn time credits by providing services, such as tutoring, repairs, or caregiving, and spend those credits to receive services from others. Traditional time banks generally do not introduce money into the exchange.

Consider instead a hybrid system that allows monetary transfers alongside time credits. Some members may prefer to buy or sell services with money, while others may still value accumulating credits for future reciprocal exchange. The finite-market model is well suited to a static version of this setting.

Let N be the set of community members and let $\mathcal{K}$ be the set of services. Member $i \in N$ can provide services $K_{S,i} \subseteq \mathcal{K}$ and demands services $K_{D,i} \subseteq \mathcal{K}$, where $K_{S,i} \cap K_{D,i} = \emptyset$. Member i can supply up to $e_i^k \geq 0$ hours of service $k \in K_{S,i}$. Normalize total available hours so that

$$\sum_{i \in N} \sum_{k \in K_{S,i}} e_i^k = 1.$$

Define the analogue of a CTR realization for a service hour supplied by i as $\omega(i,k) = (\omega_j(i,k))_{j \in N}$, where

$$\omega_j(i,k) := \begin{cases} 1, & \text{if } j = i \text{ or } k \in K_{D,j}, \\ 0, & \text{otherwise.} \end{cases}$$

Let $\Omega := \{\omega(i,k) : i \in N,\ k \in K_{S,i}\}$. Member i's share of the total endowment and condi-

[17] Prominent platforms include hOurworld, TimeRepublik, and TimeBanks.

tional distribution of service hours are

$$\alpha_i := \sum_{k \in K_{S,i}} e_i^k \quad \text{and} \quad \gamma_i(B) := \frac{1}{\alpha_i} \sum_{k \in K_{S,i}} e_i^k \mathbf{1}\{\omega(i,k) \in B\} \quad \text{for every } B \subseteq \Omega,$$

where we restrict attention to members with $\alpha_i > 0$. The pooled endowment has distribution $\gamma_\alpha = \sum_{i \in N} \alpha_i \gamma_i$. The owner coordinate $\omega_i(i,k) = 1$ records member i's outside option: if the hour is not reassigned, member i avoids the private opportunity cost of providing it. For $j \neq i$, the coordinate $\omega_j(i,k) = 1$ records that member j can use service k. Thus, θ_j is a common private value per usable service hour and, for hours initially supplied by j, its private opportunity cost of provision. An hour with payoff vector ω therefore gives member j value $\theta_j \omega_j$. This is exactly the finite market environment with binary payoffs. The same scoring mechanism determines how service hours should be reassigned, while transfers allow the mechanism to compensate members who contribute more time than they receive. A limitation of the application is that the static model abstracts from the accumulation and use of time credits over time.

5.1.2 Greenhouse Gas Credit Markets

A greenhouse gas (GHG) credit is a tradable instrument that represents a claim to avoided emissions or enhanced removal of greenhouse gases from the atmosphere.[18] Article 6 of the 2015 Paris Agreement allows countries to cooperate in meeting their emissions targets, including through voluntary market mechanisms that transfer verified emissions reductions across parties.

Consider a market in which firms buy and sell rights corresponding with different GHG reductions. Let Ω be the set of credit types, and let a credit $\omega \in \Omega$ have a vector $\omega = (\omega_i)_{i \in N}$, where ω_i measures the compliance value of that credit in firm i's industry or jurisdiction. After normalizing the total mass of credits to one, firm i initially controls a portfolio share $\alpha_i > 0$ whose credit types are distributed according to $\gamma_i \in \Delta(\Omega)$. If θ_i is firm i's private marginal compliance cost, then the value of assigning credit ω to firm i is $\theta_i \omega_i$. Older, emissions-intensive technologies may have a larger θ_i than cleaner technologies because an additional unit of compliance is more costly to achieve internally.

This makes the firm's initial credit portfolio its outside option. The designer can reallocate credits across firms and use transfers to compensate firms that give up credits they already hold. This setting helps study how credit quality, industry needs, and private abatement costs affect optimal trading.

[18] https://en.wikipedia.org/wiki/Carbon_offsets_and_credits

5.2 Revisiting Model Assumptions

The baseline model abstracts from several features of real-world advertising markets to isolate the poaching risks created by pooling audiences. We next revisit some of the main assumptions and discuss how relaxing them would potentially affect the analysis.

5.2.1 Limits of Pooling Audiences

"Can I exclude specific customers from my audiences?" is another question addressed on Shopify's FAQ page.[19] Shopify does not currently give merchants this flexibility: "...you can't add or exclude customers to a specific audience." In the model, participating merchants contribute their full platform-built audiences, and the designer observes the CTRs of each customer. Selective withholding of audiences is therefore not part of the merchant's choice.

The analysis would change if merchants knew more about their customers than the designer or could conceal part of their audiences. A merchant might then share customers selectively in order to influence the composition of the pool or the competition it faces. The designer would anticipate this behavior, making audience disclosure itself part of the mechanism. Formally, this connects the problem to mechanism design with limited commitment, pioneered by Bester and Strausz (2001) and studied more recently by Doval and Skreta (2022). Such an extension is beyond the scope of this paper and is left for future research.

5.2.2 Sequential Targeted Advertising

The model assigns at most one ad to each customer. If the platform can observe in near real time that a customer did not click the first ad, it could retarget the customer with the next-ranked ad. Because the ranking combines CTR and merchant type through the score $\omega_i h_i(\theta_i)$, the next-ranked merchant may have a higher CTR even when it has a lower type. A non-click may therefore leave value for a later impression. The process could continue until a purchase occurs or the platform exhausts its options.

At the same time, later impressions raise questions about customer experience. The first ad should not be evaluated only by its immediate click probability because it may also create brand awareness (Percy and Rossiter, 1992). Moreover, the effectiveness of later impressions may be lower because repeated exposure may become intrusive and generate

[19] https://help.shopify.com/en/manual/promoting-marketing/shopify-audiences/faq#can-i-exclude-specific-customers-from-my-audiences

ad fatigue. The model abstracts from these learning and behavioral effects. Incorporating them, and studying how they reshape targeted advertising, remains a promising direction for future research.

5.2.3 Data Complementarities

In our model, customers have unit demand and the platform assigns at most one merchant to each customer. In practice, the value of pooled data may also come from complementary products. Shopify emphasizes, for example, that its audience algorithms may draw insights from customers who previously purchased complementary products.[20] A merchant selling swimsuits may therefore be interested in customers who recently purchased sandals, sunglasses, or sunscreen.

Such complementarities may create value from showing several ads or product recommendations. Extending the model to allow complementary products and customers with multi-unit demand across products is left for future research.

6 Conclusions

We study the optimal design of targeted advertising platforms that pool merchants' platform-built audiences. Pooling improves customer–merchant matching but exposes merchants to customer poaching. The central design problem is therefore to improve targeting while preserving merchants' incentives to contribute their audiences.

With finitely many merchants, the optimal mechanism assigns targeted ads using a single score for each merchant. This score combines the virtual cost of selling customer access with the virtual value of buying it and is ironed for intermediate types. With continuously distributed CTRs, ties disappear almost everywhere, and the optimal mechanism is characterized by a system of equations in the ironing parameters.

In the limit market, the optimal indirect mechanism segments each merchant tier into selling, exchange, and buying markets. Low types sell their audience contributions, intermediate types exchange them while preserving their expected clicks, and high types buy additional expected clicks. An endogenous vector of shadow prices determines the market cutoffs, assigns customers across tiers, and clears expected clicks demand and supply. Finally, we characterize when merchant- and customer-efficient design requires no outside subsidy and show that merchant-efficient platform design can always be achieved.

[20] https://www.shopify.com/partners/blog/shopify-audiences

References

Andersson, T., Á. Cseh, L. Ehlers, and A. Erlanson (2021): "Organizing time exchanges: Lessons from matching markets," *American Economic Journal: Microeconomics*, 13, 338–373.

Babaioff, M., P. Briest, and P. Krysta (2008): "On the approximability of combinatorial exchange problems," in *International Symposium on Algorithmic Game Theory*, Springer, 83–94.

Bergemann, D. and A. Bonatti (2011): "Targeting in advertising markets: Implications for offline versus online media," *The RAND Journal of Economics*, 42, 417–443.

——— (2015): "Selling cookies," *American Economic Journal: Microeconomics*, 7, 259–294.

——— (2024): "Data, competition, and digital platforms," *American Economic Review*, 114, 2553–2595.

Bergemann, D., A. Bonatti, and T. Gan (2022): "The economics of social data," *The RAND Journal of Economics*, 53, 263–296.

Bergemann, D., A. Bonatti, and N. Wu (2024): "How Do Digital Advertising Auctions Impact Product Prices?" *Review of Economic Studies*, rdae087.

Bester, H. and R. Strausz (2001): "Contracting with imperfect commitment and the revelation principle: the single agent case," *Econometrica*, 69, 1077–1098.

Bulow, J. and J. Roberts (1989): "The simple economics of optimal auctions," *Journal of political economy*, 97, 1060–1090.

Craig, C. S., B. Sternthal, and C. Leavitt (1976): "Advertising wearout: An experimental analysis," *Journal of Marketing Research*, 13, 365–372.

Cramton, P., R. Gibbons, and P. Klemperer (1987): "Dissolving a partnership efficiently," *Econometrica: Journal of the Econometric Society*, 615–632.

Debreu, G. and H. Scarf (1963): "A limit theorem on the core of an economy," *International Economic Review*, 4, 235–246.

Doval, L. and V. Skreta (2022): "Mechanism design with limited commitment," *Econometrica*, 90, 1463–1500.

Edelman, B., M. Ostrovsky, and M. Schwarz (2007): "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords," *American economic review*, 97, 242–259.

Eliaz, K. and R. Spiegler (2020): "Incentive-compatible advertising on nonretail platforms," *The RAND Journal of Economics*, 51, 323–345.

Elliott, M., A. Galeotti, A. Koh, and W. Li (2022): "Matching and information design in marketplaces," *Available at SSRN 4283968*.

Fan, K. (1953): "Minimax theorems," *Proceedings of the National Academy of Sciences*, 39, 42–47.

Farboodi, M. and L. Veldkamp (2023): "Data and markets," *Annual Review of Economics*, 15, 23–40.

Galperti, S., A. Levkun, and J. Perego (2024): "The value of data records," *Review of Economic Studies*, 91, 1007–1038.

Galperti, S. and J. Perego (2022): "Competitive markets for personal data," *Available at SSRN 4309966*.

Gomes, R. and A. Pavan (2016): "Many-to-many matching and price discrimination," *Theoretical Economics*, 11, 1005–1052.

——— (2024): "Price customization and targeting in matching markets," *The RAND Journal of Economics*, 55, 230–265.

Loertscher, S. and C. Wasser (2019): "Optimal structure and dissolution of partnerships," *Theoretical Economics*, 14, 1063–1114.

Myerson, R. B. (1981): "Optimal auction design," *Mathematics of operations research*, 6, 58–73.

Myerson, R. B. and M. A. Satterthwaite (1983): "Efficient mechanisms for bilateral trading," *Journal of economic theory*, 29, 265–281.

Percy, L. and J. R. Rossiter (1992): "A model of brand awareness and brand attitude advertising strategies," *Psychology & Marketing*, 9, 263–274.

Pigou, A. C. (1920): *The economics of welfare*, London: Macmillan.

Robinson, J. (1933): *The economics of imperfect competition*, London: Macmillan.

Shapley, L. and H. Scarf (1974): "On cores and indivisibility," *Journal of mathematical economics*, 1, 23–37.

Varian, H. R. (2007): "Position auctions," *international Journal of industrial Organization*, 25, 1163–1178.

Yang, F. (2025): "Nested bundling," *American Economic Review*, 115, 2970–3013.

Appendix

A Proofs for the Finite Market

A.1 Completing the Proof of Theorem 1

Definition 3 (Worst-Off Types). *For every $i \in N$, let $\widehat{\Theta}_i(x,t)$ denote the set of worst-off types θ_i induced by mechanism (x,t), i.e.*

$$\widehat{\Theta}_i(x,t) := \operatorname*{argmin}_{\theta_i' \in \Theta_i} U_i\left(\theta_i'\right),$$

and define $\widehat{\Theta}(x,t) := \prod_{i \in N} \widehat{\Theta}_i(x,t)$.

Characterizing Worst-Off Types à la Cramton et al. (1987). By envelope theorem, $U_i'(\theta_i) = Q_i(\theta_i) - a_i$ almost everywhere, implying that $U_i(\theta_i)$ is convex in θ_i by (M) and absolute continuity of U_i.[21] Then following Cramton et al. (1987) one characterizes the set of worst-off types $\widehat{\Theta}_i(x,t)$ in terms of expected clicks $Q_i(\theta_i)$ alone.

Eliminating (IR). Fix a mechanism (x,t) satisfying (IC). By definition, it is sufficient that (IR) holds at every $\hat{\theta} \in \widehat{\Theta}(x,t)$, which by optimality must satisfy $U_i\left(\hat{\theta}_i\right) = 0$. Now note that the left-hand side of (11) is constant in $\theta' \in \Theta$. By definition, for any arbitrary $\hat{\theta} \in \widehat{\Theta}(x,t)$ and $\theta' \in \Theta \setminus \widehat{\Theta}(x,t)$ we have $U_i\left(\theta'\right) \geq U_i\left(\hat{\theta}_i\right)$ for all $i \in N$, with at least one inequality being strict. It then follows that

$$\psi_\eta\left(x, \theta'\right) > \psi_\eta\left(x, \hat{\theta}\right).$$

Since Θ is compact, following Loertscher and Wasser (2019) the worst-off types admit the following characterization as a function of the targeted ads x alone:

$$\widehat{\Theta}(x) = \operatorname*{argmin}_{\theta' \in \Theta} \psi_\eta\left(x, \theta'\right).$$

Using (11) it then follows that for any $\theta' = \hat{\theta}(x) \in \widehat{\Theta}(x)$

$$\Psi_\eta(x,t) = \psi_\eta\left(x, \hat{\theta}(x)\right) = \min_{\theta' \in \Theta} \psi_\eta\left(x, \theta'\right).$$

Therefore, the designer's problem is equivalently given by

$$\sup_{x \in \mathcal{X}} \min_{\theta' \in \Theta} \psi_\eta\left(x, \theta'\right).$$

[21] If a function h is absolutely continuous and h' is increasing *almost everywhere* then h is convex.

Existence. Suppose Ω is finite. Endow $L^\infty(F)$, the space of essentially bounded measurable functions on the type space Θ, with the weak* topology, and endow finite products of $L^\infty(F)$ with the corresponding product weak* topology. By (Fsb), for every $\omega \in \Omega$ and $\theta \in \Theta$,

$$x_i^\omega(\theta) \geq 0 \quad \text{and} \quad \sum_{i \in N} x_i^\omega(\theta) \leq \gamma_\alpha(\omega).$$

Hence each x_i^ω belongs to the unit ball of $L^\infty(F)$, so $\mathscr{X}$ is contained in a finite product of such unit balls. Feasibility and interim monotonicity are weak*-closed linear constraints, so $\mathscr{X}$ is convex and weak* closed. By the Banach–Alaoglu theorem, this finite product of unit balls is weak* compact. Therefore, $\mathscr{X}$ is weak* compact.

Next, observe first that $\psi_\eta\left(\cdot, \theta'\right) : \mathscr{X} \to \mathbb{R}$ is a weak*-continuous linear functional for every $\theta' \in \Theta$. On the other hand, $\psi_\eta(x, \cdot) : \Theta \to \mathbb{R}$ is continuous at every $x \in \mathscr{X}$ since we can write $\psi_\eta\left(x, \theta'\right)$ as

$$\psi_\eta\left(x, \theta'\right) = \sum_{i \in N} \left(\int_{\underline{\theta}_i}^{\theta_i'} (Q_i(\theta_i) - a_i) \phi_{\eta,i}^S(\theta_i) dF_i(\theta_i) + \int_{\theta_i'}^{\overline{\theta}_i} (Q_i(\theta_i) - a_i) \phi_{\eta,i}^B(\theta_i) dF_i(\theta_i) \right).$$

Since Q_i, $\phi_{\eta,i}^S$ and $\phi_{\eta,i}^B$ are measurable functions, it follows that (see Royden, 4th Ed., Chapter 5, Theorem 14) $\psi_\eta(x, \cdot)$ is absolutely continuous and has partial derivatives given by

$$\frac{\partial \psi_\eta\left(x, \theta'\right)}{\partial \theta_i'} = \eta_r \left[Q_i\left(\theta_i'\right) - a_i \right] \quad \text{and} \quad \frac{\partial^2 \psi_\eta\left(x, \theta'\right)}{\partial \theta_j' \partial \theta_i'} = 0 \quad dF\text{-a.e.}$$

Then, $\partial \psi_\eta(x, \cdot)/\partial \theta'$ is increasing dF-almost everywhere by (M), hence $\psi_\eta(x, \cdot)$ is convex at every $x \in \mathscr{X}$. Lastly, $\psi_\eta\left(\cdot, \theta'\right)$ is continuous at each $\theta' \in \Theta$. Hence, applying Fan's minimax theorem, we get strong duality

$$\max_{x \in \mathscr{X}} \min_{\theta' \in \Theta} \psi_\eta\left(x, \theta'\right) = \min_{\theta' \in \Theta} \max_{x \in \mathscr{X}} \psi_\eta(x, \theta'),$$

which implies the existence of a saddle point.

A.2 Completing the Proof of Proposition 1

Existence. For all targeted ads x, (Fsb) implies *absolute continuity* of the measures: $x_i(\theta; \cdot) \ll \gamma_\alpha$ for every $i \in N$ and $\theta \in \Theta$. Define the non-negative residual measure

$$x_0(\theta; B) := \gamma_\alpha(B) - \sum_{i \in N} x_i(\theta; B)$$

The measurable Radon–Nikodym theorem gives jointly measurable Radon–Nikodym densities $\rho_i(\theta, \omega)$ for every $i \in N \cup \{0\}$. Moreover, since $\sum_{i \in N \cup \{0\}} x_i(\theta; \cdot) = \gamma_\alpha$, uniqueness of

Radon–Nikodym densities implies that, for every $\theta \in \Theta$ and Borel set $B \subseteq \Omega$,

$$x_i(\theta; B) = \int_B \rho_i(\theta, \omega)\, d\gamma_\alpha(\omega), \quad \rho_i(\theta, \omega) \geq 0, \quad \sum_{i \in N} \rho_i(\theta, \omega) \leq 1 \quad \gamma_\alpha\text{-a.e.} \tag{32}$$

We may therefore identify targeted ads in the continuous-CTR environment with

$$\mathcal{R} := \left\{ \rho \in (L^\infty(F \otimes \gamma_\alpha))^{|N|} : \rho_i \geq 0 \text{ and } \sum_{i \in N} \rho_i \leq 1 \quad (F \otimes \gamma_\alpha)\text{-a.e.} \right\}.$$

The set $\mathcal{R}$ is convex, contained in the unit ball, and weak* closed because its nonnegativity and feasibility restrictions are weak* closed. The Banach–Alaoglu theorem therefore implies that $\mathcal{R}$ is weak* compact. We may then apply the same minimax argument as in the proof of Theorem 1 to show existence a saddle point.

Partitioning Ω into Winning Sets. If $z \neq \mathbf{0}$, there exists a merchant i with $z_i > 0$. Because $\phi^S_{\eta,i}$ is nonnegative and strictly increasing, the score $\overline{\phi}_{\eta,i}(\theta_i; z_i)$ is strictly positive for F_i-almost every θ_i. Hence, the vector of merchant scores is nonzero at F-almost every type profile. By Assumption 1,

$$\operatorname*{argmax}_{j \in N} \left\{ \omega_j \overline{\phi}_{\eta,j}(\theta_j; z_j) \right\}$$

is a singleton γ_α-almost everywhere. Therefore, for F-almost every type profile, the regions $\left\{ E^\theta_i(z) \right\}_{i \in N}$ partition Ω up to a γ_α-null set.

Pointwise Maximizers. By (32), feasibility is equivalent to

$$\sum_{i \in N} \rho_i(\theta, \omega) \leq 1 \quad \gamma_\alpha\text{-a.e. } \omega \in \Omega, \quad \text{for every } \theta \in \Theta.$$

Now, in the continuous-CTR environment, the problem in (12) is equivalent to maximizing

$$\mathbb{E}\left[\sum_{i \in N} \left(\int_\Omega \omega_i \rho_i(\theta, \omega) d\gamma_\alpha(\omega) - a_i \right) \overline{\phi}_{\eta,i}(\theta_i; z_i) \right],$$

subject to feasibility constraints.

Let x be a pointwise maximizer of the ironed virtual objective. Fix $i \in N$ and consider the winning region $E^\theta_i(z)$. On this region, merchant i has the unique highest strictly positive CTR-weighted ironed score. Hence assigning positive mass to any merchant $j \neq i$ cannot be optimal, so

$$\rho_j(\theta, \omega) = 0 \quad \text{for all } j \neq i, \quad \gamma_\alpha\text{-a.e. on } E^\theta_i(z).$$

Moreover, feasibility binds on $E^\theta_i(z)$: otherwise, unassigned mass with a strictly positive winning score could be assigned to merchant i, strictly increasing the ironed virtual objec-

tive. Therefore,

$$\rho_i(\theta,\omega) = \mathbf{1}_{E_i(z)}(\theta,\omega) \quad (F\otimes\gamma_\alpha)\text{-a.e..}$$

Define therefore the tie-breaking rule of the scoring mechanism by $\pi_i(\theta,\omega) = \rho_i(\theta,\omega)$ for every $i \in N$ and

$$\pi_0(\theta,\omega) := 1 - \sum_{i\in N} \rho_i(\theta,\omega).$$

Finally, suppose that the critical worst-off type $\hat{\theta}_i$ is interior. By (7), it lies in the ironing region, so $\overline{\phi}_{\eta,i}\big(\hat{\theta}_i; z_i\big) = z_i$. Using the optimal allocation above,

$$\begin{aligned} Q_i\big(\hat{\theta}_i\big) &= \mathbb{E}_{F_{-i}}\Big[\int_\Omega \omega_i\, \mathbf{1}_{E_i(z)}\Big(\big(\hat{\theta}_i,\theta_{-i}\big),\omega\Big) d\gamma_\alpha(\omega)\Big] \\ &= \int_\Omega \omega_i\, \mathbb{P}_{F_{-i}}\Big(\omega_i z_i > \omega_j \overline{\phi}_{\eta,j}(\theta_j; z_j) \text{ for all } j\neq i\Big) d\gamma_\alpha(\omega), \end{aligned}$$

where the second equality follows by Tonelli's theorem and the independence of CTRs and types. By the characterization of worst-off types à la Cramton et al. (1987), we have $Q_i\big(\hat{\theta}_i\big) = a_i$. This gives

$$a_i = \int_\Omega \omega_i\, \mathbb{P}_{F_{-i}}\Big(\omega_i z_i > \omega_j \overline{\phi}_{\eta,j}(\theta_j; z_j) \text{ for all } j\neq i\Big) d\gamma_\alpha(\omega).$$

B Proofs for the Limit Market

B.1 Proof of Theorem 2

Proof of Lemma 1. Suppose first that $q = (q_k)_{k\in K}$ is implemented by targeted ads $x = (x_k)_{k\in K}$ satisfying (G-Fsb). For each $k \in K$, define the total tier measure

$$\overline{x}_k(B) := \beta_k \int_{\Theta_k} x_k(\theta; B)\, dF_k(\theta) \quad \text{for every Borel set } B \subseteq [0,1]^K.$$

By (G-Fsb),

$$\sum_{k\in K} \overline{x}_k(B) \le G(B) \quad \text{for every Borel set } B \subseteq [0,1]^K.$$

Since the measures $\overline{x}_k$ are nonnegative, it follows that $\overline{x}_k \le G$ and hence $\overline{x}_k \ll G$ for every $k \in K$. The Radon–Nikodym theorem therefore gives measurable tier-specific shares

$$m_k := \frac{d\overline{x}_k}{dG}.$$

Moreover, for every Borel set B,

$$\int_B \sum_{k\in K} m_k(\omega)\, dG(\omega) = \sum_{k\in K} \overline{x}_k(B) \le G(B),$$

which implies

$$\sum_{k\in K} m_k(\omega) \le 1 \quad G\text{-a.e.}$$

Finally, Tonelli's theorem gives

$$\begin{aligned} d_k(q) &= \beta_k \int_{\Theta_k} \int_{[0,1]^K} \omega_k \, x_k(\theta; d\omega)\, dF_k(\theta) \\ &= \int_{[0,1]^K} \omega_k \, d\overline{x}_k(\omega) \\ &= \int_{[0,1]^K} \omega_k m_k(\omega)\, dG(\omega). \end{aligned}$$

Hence, $d(q) \in \mathcal{D}(G)$.

Conversely, suppose $d(q) \in \mathcal{D}(G)$, so there exists a vector of tier-specific shares m satisfying (18). For each $k \in K$, define

$$\gamma_k := \begin{cases} \dfrac{d_k(q)}{\int_{[0,1]^K} \omega_k m_k(\omega)\, dG(\omega)}, & \displaystyle\int_{[0,1]^K} \omega_k m_k(\omega)\, dG(\omega) > 0, \\ 0, & \text{otherwise}, \end{cases}$$

and set

$$\widetilde{m}_k(\omega) := \gamma_k m_k(\omega).$$

Then $\gamma_k \in [0,1]$,

$$\sum_{k\in K} \widetilde{m}_k(\omega) \le 1 \quad G\text{-a.e.},$$

and

$$\int_{[0,1]^K} \omega_k \widetilde{m}_k(\omega)\, dG(\omega) = d_k(q).$$

If $d_k(q) > 0$, define

$$\rho_k(\theta, \omega) := q_k(\theta) \frac{\widetilde{m}_k(\omega)}{d_k(q)};$$

if $d_k(q) = 0$, set $\rho_k(\theta, \omega) = 0$. These densities are measurable. Define targeted ads by

$$x_k(\theta; B) := \int_B \rho_k(\theta, \omega)\, dG(\omega) \quad \text{for every Borel set } B \subseteq [0,1]^K. \tag{33}$$

For $d_k(q) > 0$,

$$\beta_k \int_{\Theta_k} \rho_k(\theta, \omega)\, dF_k(\theta) = \widetilde{m}_k(\omega).$$

The same equality holds when $d_k(q) = 0$. Therefore, for every Borel set B,

$$\begin{aligned} \sum_{k\in K} \beta_k \int_{\Theta_k} x_k(\theta; B)\, dF_k(\theta) &= \int_B \sum_{k\in K} \widetilde{m}_k(\omega)\, dG(\omega) \\ &\le G(B), \end{aligned}$$

so x satisfies (G-Fsb). Finally,

$$\int_{[0,1]^K} \omega_k \, x_k(\theta; d\omega) = \frac{q_k(\theta)}{d_k(q)} \int_{[0,1]^K} \omega_k \widetilde{m}_k(\omega) \, dG(\omega) = q_k(\theta),$$

where the zero demand case is immediate. Thus, x implements q. □

Proof of Lemma 2. The zero vector belongs to $\mathcal{D}(G)$, so the set is nonempty. Convexity follows by taking convex combinations of the tier-specific shares that implement any two demand vectors. Downward closedness follows because, if $d \in \mathcal{D}(G)$ and $0 \leq d' \leq d$ coordinate-wise, the same tier-specific shares that implement $d \in \mathcal{D}(G)$ also implement $d' \in \mathcal{D}(G)$.

It remains to prove compactness. Let

$$\mathcal{C}(G) := \left\{ m \in (L^\infty(G))^K : m_k(\omega) \geq 0, \quad \sum_{k \in K} m_k(\omega) \leq 1 \quad G\text{-a.e.} \right\}.$$

The set $\mathcal{C}(G)$ is contained in the unit ball of $(L^\infty(G))^K$ and is weak* closed because its constraints are weak* closed. It is therefore weak* compact by the Banach–Alaoglu theorem. Define the continuous linear map $A : \mathcal{C}(G) \to \mathbb{R}^K_+$ by

$$A_k(m) := \int_{[0,1]^K} \omega_k m_k(\omega) \, dG(\omega).$$

Then $A(\mathcal{C}(G))$ is compact, and

$$\mathcal{D}(G) = \left\{ a \odot b : a \in A(\mathcal{C}(G)), \ b \in [0,1]^K \right\},$$

where $\odot$ denotes coordinatewise multiplication. Hence, $\mathcal{D}(G)$ is the continuous image of the compact set $A(\mathcal{C}(G)) \times [0,1]^K$ and is therefore compact. □

Proof of Lemma 3. Fix $\lambda \in \mathbb{R}^K_+$. For any $d \in \mathcal{D}(G)$, let m be a vector of tier-specific shares satisfying (18). Then

$$\begin{aligned} \lambda \cdot d &\leq \int_{[0,1]^K} \sum_{k \in K} \lambda_k \omega_k m_k(\omega) \, dG(\omega) \\ &\leq \int_{[0,1]^K} \max_{k \in K} \{\lambda_k \omega_k\} \sum_{k \in K} m_k(\omega) \, dG(\omega) \\ &\leq \mathbb{E}_G \left[\max_{k \in K} \{\lambda_k \omega_k\} \right]. \end{aligned}$$

Taking the maximum over $d \in \mathcal{D}(G)$ gives

$$\sigma(\lambda) \leq \mathbb{E}_G \left[\max_{k \in K} \{\lambda_k \omega_k\} \right].$$

For the reverse inequality, fix any $\lambda \in \mathbb{R}^K_+ \setminus \{\mathbf{0}\}$ and define $\kappa^\lambda : [0,1]^K \to K$ by

$$\kappa^\lambda(\omega) := \min \operatorname*{argmax}_{k \in K} \{\lambda_k \omega_k\}.$$

Because K is finite, κ^λ is a well-defined measurable selection from the maximizing tiers. Set the tier-specific shares m_k^λ by

$$m_k^\lambda(\omega) := \mathbf{1}\Big\{\kappa^\lambda(\omega) = k\Big\}.$$

These shares sum to one G-almost everywhere. Therefore taking $d^\lambda \in \mathcal{D}(G)$ defined by

$$d_k^\lambda := \int_{[0,1]^K} \omega_k m_k^\lambda(\omega)\, dG(\omega)$$

attains the value on the support since

$$\lambda \cdot d^\lambda = \int_{[0,1]^K} \sum_{k\in K} \lambda_k\omega_k \mathbf{1}\Big\{\kappa^\lambda(\omega) = k\Big\}\, dG(\omega) = \mathbb{E}_G\Big[\max_{k\in K}\{\lambda_k\omega_k\}\Big].$$

If $\lambda = \mathbf{0}$, then $\sigma(\mathbf{0}) = 0$. □

Proof of Proposition 2. If q is implementable, then Lemma 1 gives $d(q) \in \mathcal{D}(G)$. Therefore, for every $\lambda \in \mathbb{R}^K_+$,

$$\sum_{k\in K} \lambda_k\beta_k\mathbb{E}_{F_k}[q_k(\theta)] = \lambda \cdot d(q) \le \sigma(\lambda) = \mathbb{E}_G\Big[\max_{k\in K}\{\lambda_k\omega_k\}\Big],$$

where the final equality follows from Lemma 3. Hence, (R-Fsb) is necessary.

Conversely, suppose q satisfies (R-Fsb), but $d(q) \notin \mathcal{D}(G)$. Since $\mathcal{D}(G)$ is nonempty, convex, and compact, strict separation gives a vector $a \in \mathbb{R}^K$ such that

$$a \cdot d(q) > \max_{d\in\mathcal{D}(G)} a \cdot d.$$

Let a_+ denote the coordinatewise positive part of a. Because $d(q) \ge 0$,

$$a_+ \cdot d(q) \ge a \cdot d(q).$$

Moreover, downward closedness of $\mathcal{D}(G)$ implies

$$\max_{d\in\mathcal{D}(G)} a \cdot d = \max_{d\in\mathcal{D}(G)} a_+ \cdot d.$$

Indeed, for any $d \in \mathcal{D}(G)$, setting to zero the coordinates for which $a_k < 0$ gives another vector in $\mathcal{D}(G)$ and weakly increases the value of the objective. It follows that

$$a_+ \cdot d(q) > \sigma(a_+),$$

contradicting (R-Fsb) at $\lambda = a_+$. Hence, $d(q) \in \mathcal{D}(G)$, and Lemma 1 implies that q is implementable. □

Proof of Lemma 4. Fix $q \in \mathcal{Q}^0$. If q satisfies (R-Fsb), then $\Delta(q,\lambda) \ge 0$ for every $\lambda \in \mathbb{R}^K_+$, while $\Delta(q,0) = 0$. Hence,

$$\inf_{\lambda\in\mathbb{R}^K_+} \Delta(q,\lambda) = 0.$$

If q violates (R-Fsb), there exists $\bar{\lambda} \in \mathbb{R}^K_+$ such that $\Delta\big(q, \bar{\lambda}\big) < 0$. Homogeneity of $\Delta\big(q, \bar{\lambda}\big)$ in $\bar{\lambda}$ implies

$$\Delta\big(q, c\bar{\lambda}\big) = c\,\Delta\big(q, \bar{\lambda}\big) \longrightarrow -\infty \quad \text{as } c \to \infty.$$

Therefore,

$$\inf_{(\theta', \lambda) \in \Theta \times \mathbb{R}^K_+} \mathcal{L}(q, \theta', \lambda) = \begin{cases} \inf\limits_{\theta' \in \Theta} \psi_\eta(q, \theta'), & q \in \mathcal{Q}, \\ -\infty, & q \notin \mathcal{Q}. \end{cases}$$

Taking the supremum over $q \in \mathcal{Q}^0$ gives the first equality in (22).

For the second equality, endow $\mathcal{Q}^0$ with the weak* topology inherited from $\prod_{k \in K} L^\infty(F_k)$. The set $\mathcal{Q}^0$ is convex and weak* compact: it is contained in the unit ball, and the bounds and monotonicity restrictions defining it are weak* closed. The set $\Theta \times \mathbb{R}^K_+$ is convex. For every (θ', λ), the map $q \mapsto \mathcal{L}(q, \theta', \lambda)$ is affine and weak* continuous, and jointly continuous in (q, θ', λ).

Conversely, fix $q \in \mathcal{Q}^0$. The map $\lambda \mapsto \Delta(q, \lambda)$ is convex and continuous because σ is a support function and the remaining term is linear in λ. Moreover,

$$\frac{\partial \psi_\eta(q, \theta')}{\partial \theta'_k} = \eta_r \beta_k (q_k(\theta'_k) - \mu_k) \quad \text{for almost every } \theta'_k.$$

Since q_k is nondecreasing, this derivative is nondecreasing in θ'_k. Hence, $\theta' \mapsto \psi_\eta(q, \theta')$ is convex and continuous. Because the dependence on θ' and λ is additively separable, $(\theta', \lambda) \mapsto \mathcal{L}(q, \theta', \lambda)$ is jointly convex and continuous. Fan's second minimax theorem (Fan (1953, Theorem 2)) therefore gives

$$\max_{q \in \mathcal{Q}^0} \inf_{(\theta', \lambda) \in \Theta \times \mathbb{R}^K_+} \mathcal{L}(q, \theta', \lambda) = \inf_{(\theta', \lambda) \in \Theta \times \mathbb{R}^K_+} \max_{q \in \mathcal{Q}^0} \mathcal{L}(q, \theta', \lambda),$$

which proves the second equality in (22). □

Proof of Lemma 5. Fix $\lambda \in \mathbb{R}^K_+$. Since $\Delta(q, \lambda)$ does not depend on θ', minimizing the pseudo-Lagrangian over θ' is identical to minimizing the virtual objective $\psi_\eta(q, \theta')$. Thus, the worst-off-type characterization used in the finite market applies tier by tier without modification. By the construction of $\hat{\theta}(\lambda)$ and q^λ,

$$\hat{\theta}(\lambda) \in \operatorname*{argmin}_{\theta' \in \Theta} \mathcal{L}\big(q^\lambda, \theta', \lambda\big). \tag{34}$$

Next hold λ and $\hat{\theta}(\lambda)$ fixed. After omitting terms that do not depend on q, maximizing the pseudo-Lagrangian reduces to maximizing

$$\sum_{k \in K} \beta_k \mathbb{E}_{F_k}\Big[q_k(\theta)\Big(\phi_{\eta,k}\Big(\theta; \hat{\theta}_k(\lambda_k)\Big) - \lambda_k\Big)\Big].$$

The same ironing argument as in the finite market replaces each virtual type by $\overline{\phi}_{\eta,k}(\cdot; z_k(\lambda_k))$. By the construction of $z_k(\lambda_k)$, the expected clicks q^λ is a pointwise maximizer of the ironed objective and is constant on the ironing interval, so the original and ironed objectives agree. Hence,

$$q^\lambda \in \operatorname*{argmax}_{q \in \mathcal{Q}^0} \mathcal{L}\Big(q, \hat{\theta}(\lambda), \lambda\Big). \tag{35}$$

Combining (34) and (35) gives

$$\mathcal{L}\Big(q, \hat{\theta}(\lambda), \lambda\Big) \leq \mathcal{L}\Big(q^\lambda, \hat{\theta}(\lambda), \lambda\Big) \leq \mathcal{L}\Big(q^\lambda, \theta', \lambda\Big) \tag{36}$$

for every $q \in \mathcal{Q}^0$ and $\theta' \in \Theta$. The equalities in (26) follow immediately by the saddle point property.

□

Proof of Lemma 6. We prove each claim separately.

Restriction on the Compact Domain Λ. For each tier k, define the tier-specific reduced value

$$L_k(\lambda_k) := \max_{q_k \in \mathcal{Q}_k^0} \min_{\theta'_k \in \Theta_k} \beta_k \mathbb{E}_{F_k}[(q_k(\theta) - \mu_k)\phi_{\eta,k}(\theta; \theta'_k) - \lambda_k q_k(\theta)].$$

Because the feasible sets and the objective separate across tiers,

$$\ell(\lambda) = \sigma(\lambda) + \sum_{k \in K} L_k(\lambda_k). \tag{37}$$

Fix a tier k. If $\lambda_k \geq \overline{\lambda}_k = \phi^S_{\eta,k}\Big(\overline{\theta}_k\Big)$, then the inverse convention gives

$$p^S_k(\lambda_k) = p^B_k(\lambda_k) = \overline{\theta}_k.$$

For every critical type θ'_k and every $\theta \in \Theta_k$,

$$\phi_{\eta,k}(\theta; \theta'_k) \leq \phi^S_{\eta,k}\Big(\overline{\theta}_k\Big) = \overline{\lambda}_k \leq \lambda_k.$$

Hence, for every θ'_k, the maximization over q_k is attained by $q_k \equiv 0$. The remaining tier objective is independent of λ_k, and its derivative with respect to θ'_k is $-\eta_r \beta_k \mu_k < 0$, so it is minimized at $\theta'_k = \overline{\theta}_k$. It follows that

$$L_k(\lambda_k) = L_k\Big(\overline{\lambda}_k\Big) \quad \text{for every } \lambda_k \geq \overline{\lambda}_k.$$

Now take any $\lambda \in \mathbb{R}^K_+$ and define

$$\widetilde{\lambda}_k := \min\Big\{\lambda_k, \overline{\lambda}_k\Big\}.$$

Since each CSR coordinate is nonnegative, the support function σ is nondecreasing in every coordinate, implying $\sigma\Big(\widetilde{\lambda}\Big) \leq \sigma(\lambda)$, and hence $\ell\Big(\widetilde{\lambda}\Big) \leq \ell(\lambda)$.

Since $\widetilde{\lambda} \in \Lambda$, every shadow price vector outside Λ can be replaced by one in Λ without increasing the reduced value. Therefore,

$$\inf_{\lambda \in \mathbb{R}^K_+} \ell(\lambda) = \inf_{\lambda \in \Lambda} \ell(\lambda).$$

Existence. Endow $\mathcal{Q}^0$ with the weak* topology. The sets $\mathcal{Q}^0$ and Θ are compact, and the pseudo-Lagrangian is jointly continuous in (q, θ', λ), where continuity in q is understood in the weak* topology. Applying Berge's maximum theorem first to the minimization over Θ and then to the maximization over $\mathcal{Q}^0$ shows that ℓ is continuous. By the preceding analysis, it is enough to minimize ℓ over the compact set Λ. Then Weierstrass theorem implies the existence of $\lambda^* \in \Lambda$ that minimizes ℓ. □

Proof of Lemma 7. We first establish convexity of the reduced value function. Since the feasibility slack term is independent of θ', the definition of ℓ gives

$$\ell(\lambda) = \sigma(\lambda) + \sup_{q \in \mathcal{Q}^0} \left\{ \inf_{\theta' \in \Theta} \psi_\eta(q, \theta') - \sum_{k \in K} \lambda_k \beta_k \mathbb{E}_{F_k}[q_k(\theta)] \right\}.$$

For every fixed $q \in \mathcal{Q}^0$, the expression inside braces is affine in λ. Its pointwise supremum over q is therefore convex. Since σ is a support function and hence convex, ℓ is convex on $\mathbb{R}^K_+$.

Now define the expected clicks demand generated by the fixed-λ construction by

$$d_k(\lambda_k) := \beta_k \mathbb{E}_{F_k}\left[q_k^\lambda(\theta)\right] = \beta_k \Big[1 - F_k\Big(p_k^B(\lambda_k)\Big) + \mu_k \Big(F_k\Big(p_k^B(\lambda_k)\Big) - F_k\Big(p_k^S(\lambda_k)\Big)\Big)\Big]. \tag{38}$$

The continuity of F_k and of the inverse cutoff types implies that d_k is continuous. Moreover, for a function $f : \mathbb{R}^K_+ \to \mathbb{R}$ and a feasible direction $h \in \mathbb{R}^K$ at $\lambda \in \mathbb{R}^K_+$, define the right directional derivative

$$D_h^+ f(\lambda) := \lim_{t \to 0^+} \frac{f(\lambda + th) - f(\lambda)}{t},$$

whenever the limit exists, where h is feasible at λ if $\lambda + th \in \mathbb{R}^K_+$ for all sufficiently small $t > 0$.

Let λ^* minimize ℓ, which exists by Lemma 6. We first show that $\lambda^* \neq \mathbf{0}$. By the positive homogeneity of σ,

$$D_{\mathbf{1}}^+ \sigma(\mathbf{0}) = \lim_{t \to 0^+} \frac{\sigma(t\mathbf{1}) - \sigma(\mathbf{0})}{t} = \sigma(\mathbf{1}) = \mathbb{E}_G\left[\max_{k \in K} \omega_k\right].$$

Since λ_k enters the tier-k reduced objective only through $-\lambda_k \beta_k \mathbb{E}_{F_k}[q_k(\theta)]$, the one-sided envelope theorem applied tier by tier gives

$$D_{\mathbf{1}}^+ \ell(\mathbf{0}) = \mathbb{E}_G\left[\max_{k \in K} \omega_k\right] - \sum_{k \in K} d_k(0).$$

At zero shadow prices, $p_k^S(0) = \underline{\theta}_k$ and $p_k^B(0) = \underline{p}_k^B$, and hence

$$d_k(0) = \beta_k\Big[1 - (1-\mu_k)F_k\Big(\underline{p}_k^B\Big)\Big].$$

It follows from Assumption 3 that

$$D_{\mathbf{1}}^+\ell(\mathbf{0}) < 0.$$

Consequently, $\ell(t\mathbf{1}) < \ell(\mathbf{0})$ for all sufficiently small $t > 0$, so $\mathbf{0}$ cannot minimize ℓ. Thus, $\lambda^* \neq \mathbf{0}$.

For any $\lambda \neq \mathbf{0}$ and any feasible direction h, the envelope theorem gives

$$D_h^+\ell(\lambda) = D_h^+\sigma(\lambda) - \sum_{k\in K} h_k d_k(\lambda_k) \tag{39}$$

Next define

$$m_\omega(\lambda) := \max_{j\in K} \lambda_j\omega_j,$$

so that $\sigma(\lambda) = \mathbb{E}_G[m_\omega(\lambda)]$. By Assumption 2, the maximizer is unique for G-almost every ω, hence Danskin's theorem gives

$$\lim_{t\to 0^+} \frac{m_\omega(\lambda+th) - m_\omega(\lambda)}{t} = \sum_{k\in K} h_k\omega_k \mathbf{1}\left\{k = \operatorname*{argmax}_{j\in K} \lambda_j\omega_j\right\}$$

for G-almost every ω. Moreover,

$$\left|\frac{m_\omega(\lambda+th) - m_\omega(\lambda)}{t}\right| \leq \max_{k\in K}\{|h_k|\omega_k\} \leq \|h\|_\infty.$$

The dominated convergence theorem (DCT) therefore implies

$$D_h^+\sigma(\lambda) = \sum_{k\in K} h_k \mathbb{E}_G\left[\omega_k \mathbf{1}\left\{k = \operatorname*{argmax}_{j\in K} \lambda_j\omega_j\right\}\right]. \tag{40}$$

In particular, if $\lambda_k > 0$, then the partial derivative exists and equals

$$\frac{\partial\sigma(\lambda)}{\partial\lambda_k} = \mathbb{E}_G\left[\omega_k \mathbf{1}\left\{k = \operatorname*{argmax}_{j\in K} \lambda_j\omega_j\right\}\right]. \tag{41}$$

Now Lemma 6 implies that λ^* also minimizes ℓ over $\mathbb{R}_+^K$. Thus, the only relevant constraints are the non-negativity constraints. Since ℓ is convex, the necessary and sufficient optimality conditions for $\lambda^* \in \mathbb{R}_+^K$ for every $k \in K$ are given by

$$\begin{cases} \dfrac{\partial\ell(\lambda^*)}{\partial\lambda_k} = 0 & \text{if } \lambda_k^* > 0, \\[2ex] D_{e_k}^+\ell(\lambda^*) \geq 0 & \text{if } \lambda_k^* = 0. \end{cases} \tag{42}$$

Suppose that $\lambda_k^* = 0$ for some $k \in K$. Then by setting $h = e_k$ in (40), we claim that

$$D_{e_k}^+\sigma(\lambda^*) = 0.$$

Indeed, since $\lambda_k^* = 0$ while $\lambda^* \neq \mathbf{0}$, tier k can maximize $\lambda_j^* \omega_j$ only as part of a tie. By Assumption 2, this event is G-null, proving the claim.

On the other hand,

$$d_k(0) = \beta_k \Big[1 - (1 - \mu_k) F_k\Big(\underline{p}_k^B\Big)\Big] \geq \beta_k \mu_k > 0.$$

Setting $h = e_k$ in (39) therefore gives

$$D_{e_k}^+ \ell(\lambda^*) = 0 - d_k(0) < 0,$$

contradicting the second case in (42). Therefore, $\lambda^* \in \mathbb{R}_{++}^K$.

Since every coordinate of λ^* is strictly positive, the first case in (42) applies to every $k \in K$, and hence

$$\nabla \ell(\lambda^*) = \mathbf{0}.$$

Thus, by the envelope theorem,

$$0 = \frac{\partial \ell(\lambda^*)}{\partial \lambda_k} = \frac{\partial \sigma(\lambda^*)}{\partial \lambda_k} - d_k(\lambda_k^*) \quad \text{for every } k \in K. \tag{43}$$

By (38) and (41), these identities imply the market-clearing conditions in (D_k=S_k). □

Completing the Proof of Theorem 2. Let λ^* minimize ℓ, and set

$$q^* := q^{\lambda^*}, \quad \hat{\theta} := \hat{\theta}(\lambda^*).$$

The market-clearing conditions, (24), and Lemma 3 imply

$$d(q^*) = d^{\lambda^*} \in \mathcal{D}(G).$$

Hence, Lemma 1 implies that q^* is implementable and therefore belongs to $\mathcal{Q}$. Moreover, since d^{λ^*} attains the support function at λ^*,

$$\Delta(q^*, \lambda^*) = \sigma(\lambda^*) - \lambda^* \cdot d(q^*) = 0.$$

By Lemma 5, $\hat{\theta}$ minimizes $\mathcal{L}(q^*, \theta', \lambda^*)$ over $\theta' \in \Theta$. Since $\Delta(q^*, \lambda^*)$ is independent of θ', $\hat{\theta}$ also minimizes $\psi_\eta(q^*, \theta')$. Therefore,

$$\inf_{\theta' \in \Theta} \psi_\eta(q^*, \theta') = \psi_\eta\Big(q^*, \hat{\theta}\Big) = \mathcal{L}\Big(q^*, \hat{\theta}, \lambda^*\Big) = \ell(\lambda^*).$$

Together with (28), this shows that q^* solves the designer's problem.

Finally, applying the envelope theorem and assigning zero rents to types on the exchange

interval, we get

$$u_k^*(\theta) = \begin{cases} \mu_k p_k^S, & \theta < p_k^S, \\ \mu_k \theta, & p_k^S \le \theta \le p_k^B, \\ \theta - (1-\mu_k) p_k^B, & \theta > p_k^B. \end{cases}$$

Therefore,

$$t_k^*(\theta) = \begin{cases} -\mu_k p_k^S, & \theta < p_k^S, \\ 0, & p_k^S \le \theta \le p_k^B, \\ (1-\mu_k) p_k^B, & \theta > p_k^B. \end{cases} \tag{44}$$

The allocation in (24), the transfers in (44), and the market-clearing conditions in (D_k=S_k) establish Theorem 2.

Lastly, we establish uniqueness. Let λ and $\widetilde{\lambda}$ be any two market-clearing shadow price vectors. Convexity of σ and market clearing imply

$$0 \le \left(\lambda - \widetilde{\lambda}\right) \cdot \left(\nabla\sigma(\lambda) - \nabla\sigma\left(\widetilde{\lambda}\right)\right) = \sum_{k\in K} \left(\lambda_k - \widetilde{\lambda}_k\right)\left(d_k(\lambda_k) - d_k\left(\widetilde{\lambda}_k\right)\right) \le 0.$$

Hence, every term in the sum is zero. Since each cutoff is nondecreasing in its shadow price, $f_k > 0$, and $\mu_k \in (0,1)$, equality implies that λ_k and $\widetilde{\lambda}_k$ induce the same p_k^S and p_k^B. Thus, all market-clearing shadow-price vectors induce the same posted prices. If at least one of these prices is interior, strict monotonicity of the corresponding virtual function then implies $\lambda_k = \widetilde{\lambda}_k$.

B.2 Proof of Proposition 3

Proof of Lemma 8. Let $g_k(\theta) := \eta_w + \eta_v \theta$. Up to the constant term $-\sum_k \beta_k \mu_k \mathbb{E}_{F_k}[g_k(\theta)]$, the objective in (15) is

$$\psi_\eta^E(q) := \sum_{k\in K} \beta_k \mathbb{E}_{F_k}[g_k(\theta) q_k(\theta)].$$

For $\lambda \in \mathbb{R}_+^K$, define

$$\ell^E(\lambda) := \sigma(\lambda) + \sum_{k\in K} \beta_k \mathbb{E}_{F_k}\left[(g_k(\theta) - \lambda_k)_+\right].$$

If q is feasible and $d_k(q) := \beta_k \mathbb{E}_{F_k}[q_k(\theta)]$, then $d(q) \in \mathcal{D}(G)$ by Lemma 1. Therefore, for every $\lambda \in \mathbb{R}_+^K$, clearly

$$\begin{aligned} \psi_\eta^E(q) &= \lambda \cdot d(q) + \sum_{k\in K} \beta_k \mathbb{E}_{F_k}[(g_k(\theta) - \lambda_k) q_k(\theta)] \\ &\le \sigma(\lambda) + \sum_{k\in K} \beta_k \mathbb{E}_{F_k}\left[(g_k(\theta) - \lambda_k)_+\right] = \ell^E(\lambda). \end{aligned} \tag{45}$$

The function ℓ^E is convex and continuous. Moreover, if $\lambda_k > g_k\big(\overline{\theta}_k\big)$, reducing λ_k to $g_k\big(\overline{\theta}_k\big)$ leaves the tier-k term unchanged and weakly decreases σ. Hence, there exists a minimizer λ^E of ℓ^E on the compact set

$$\prod_{k\in K}\Big[0, g_k\big(\overline{\theta}_k\big)\Big].$$

Since $\eta_v > 0$, the function g_k is strictly increasing. Because F_k is atomless, $g_k(\theta)$ has no atom at any λ_k, so the tier-k term in ℓ^E is differentiable at every λ_k. Applying DCT, its derivative is $-d_k^E(\lambda_k)$, where

$$d_k^E(\lambda_k) := \beta_k \mathbb{P}_{F_k}\{g_k(\theta) > \lambda_k\} = \beta_k\Big[1 - F_k\big(p_k^E\big)\Big]$$

with cutoff types p_k^E given by (17).

We now apply the argument used to characterize the minimizing shadow price vector in the proof of Lemma 7. The one-sided envelope theorem and the directional derivative of the support function give, for every feasible direction h at λ,

$$D_h^+\ell^E(\lambda) = D_h^+\sigma(\lambda) - \sum_{k\in K} h_k d_k^E(\lambda_k).$$

At $\lambda = \mathbf{0}$, positive homogeneity of σ and $d_k^E(0) = \beta_k$ imply

$$D_{\mathbf{1}}^+\ell^E(\mathbf{0}) = \sigma(\mathbf{1}) - \sum_{k\in K}\beta_k = \mathbb{E}_G\Big[\max_{k\in K}\omega_k\Big] - 1 < 0,$$

where the inequality follows from (ES). Thus, $\lambda^E \neq \mathbf{0}$.

Next, suppose that $\lambda_k^E = 0$ for some k. Since $\lambda^E \neq \mathbf{0}$, the tie-free CSR condition and the support derivative formula established in the proof of Lemma 7 give

$$D_{e_k}^+\sigma\big(\lambda^E\big) = 0.$$

Since $d_k^E(0) = \beta_k > 0$, it follows that

$$D_{e_k}^+\ell^E\big(\lambda^E\big) = -\beta_k < 0,$$

contradicting the optimality of λ^E. Therefore, $\lambda^E \in \mathbb{R}^K_{++}$.

The minimizer is interior, so by convexity and the envelope theorem, the necessary and sufficient conditions for optimal λ^E are

$$0 = \frac{\partial \ell^E\big(\lambda^E\big)}{\partial \lambda_k} = \frac{\partial \sigma\big(\lambda^E\big)}{\partial \lambda_k} - d_k^E\big(\lambda_k^E\big) \quad \text{for every } k \in K.$$

By the Danskin formula in (41), these conditions are exactly the market-clearing identities in (D_k^E=S_k^E).

□

By Lemma 8, it remains only to characterize the transfers that implement the efficient cutoff allocation and determine when they satisfy no deficit.

Following Cramton et al. (1987) and the characterization of worst-off types, the expected clicks q_k^E imply that $\theta = p_k^E$ is a worst-off type. Hence, setting $u_k\big(p_k^E\big) = \mu_k p_k^E$ makes (IR_k) bind at the worst-off type, which extracts the largest possible transfers. The envelope theorem then determines the maximal payments that can be extracted from high types and the minimum compensation to low types:

$$t_k^E(\theta) = \begin{cases} -\mu_k p_k^E, & \theta < p_k^E, \\ (1-\mu_k) p_k^E, & \theta \geq p_k^E. \end{cases} \tag{46}$$

The maximal expected transfer collected from tier k is therefore

$$\begin{aligned} \beta_k \mathbb{E}_{F_k}\big[t_k^E(\theta)\big] &= \beta_k\Big[(1-\mu_k) p_k^E\Big[1 - F_k\big(p_k^E\big)\Big] - \mu_k p_k^E F_k\big(p_k^E\big)\Big] \\ &= \beta_k p_k^E\Big[1 - \mu_k - F_k\big(p_k^E\big)\Big]. \end{aligned}$$

Summing over tiers gives the left-hand side of (16). Because individual rationality binds at the worst-off type, these transfers maximize the designer's expected revenues among all incentive-compatible and individually rational transfers implementing q^E. Consequently, a no-deficit implementation exists if and only if (16) holds.

Lastly, let λ^E and $\widetilde{\lambda}^E$ be any two market-clearing efficient shadow price vectors. Convexity of σ and market clearing imply

$$0 \leq \big(\lambda^E - \widetilde{\lambda}^E\big) \cdot \Big(\nabla\sigma\big(\lambda^E\big) - \nabla\sigma\big(\widetilde{\lambda}^E\big)\Big) = \sum_{k \in K} \big(\lambda_k^E - \widetilde{\lambda}_k^E\big)\Big(d_k^E\big(\lambda_k^E\big) - d_k^E\big(\widetilde{\lambda}_k^E\big)\Big) \leq 0.$$

Thus, every term is zero. Because $d_k^E(\lambda_k) = \beta_k\Big[1 - F_k\big(p_k^E(\lambda_k)\big)\Big]$, strict positivity of f_k implies that both shadow price vectors induce the same cutoff p_k^E. If this cutoff is interior, then $p_k^E = \big(\lambda_k^E - \eta_w\big)/\eta_v$ and $\eta_v > 0$ imply $\lambda_k^E = \widetilde{\lambda}_k^E$, proving the uniqueness claims.

B.3 Proof of Proposition 4

Let

$$d_k^E := \beta_k\Big[1 - F_k\big(p_k^E\big)\Big] \quad \text{and} \quad b_k := \beta_k \mu_k$$

denote the efficient demand and the total outside option demand in tier k, respectively. By Lemma 8 and the market-clearing conditions in (D_k^E=S_k^E), $d^E \in \mathcal{D}(G)$ and

$$\lambda^E \cdot d^E = \sigma\big(\lambda^E\big).$$

The expected clicks allocation that provides every merchant with her outside option is implementable, so $b \in \mathcal{D}(G)$. Hence, the definition of the support function gives

$$\lambda^E \cdot \left(d^E - b\right) = \sigma\left(\lambda^E\right) - \lambda^E \cdot b \geq 0. \tag{47}$$

Because $\eta_v = 1$, the efficient cutoff type in tier k is

$$p_k^E = \min\left\{\overline{\theta}_k, \max\left\{\underline{\theta}_k, \lambda_k^E\right\}\right\}.$$

We claim that

$$\left(p_k^E - \lambda_k^E\right)\left(d_k^E - b_k\right) \geq 0 \quad \text{for every } k \in K. \tag{48}$$

If $p_k^E \in \left(\underline{\theta}_k, \overline{\theta}_k\right)$, then $p_k^E = \lambda_k^E$, so the product is zero. If $p_k^E = \underline{\theta}_k$, then $p_k^E \geq \lambda_k^E$ and

$$d_k^E = \beta_k > \beta_k \mu_k = b_k.$$

Finally, if $p_k^E = \overline{\theta}_k$, then $p_k^E \leq \lambda_k^E$ and

$$d_k^E = 0 < \beta_k \mu_k = b_k.$$

Thus, (48) holds in every case.

The maximal expected transfers generated by the efficient allocation therefore satisfy

$$\begin{aligned}\sum_{k \in K} \beta_k p_k^E \left[1 - F_k\left(p_k^E\right) - \mu_k\right] &= p^E \cdot \left(d^E - b\right) \\ &= \lambda^E \cdot \left(d^E - b\right) + \left(p^E - \lambda^E\right) \cdot \left(d^E - b\right) \\ &\geq 0,\end{aligned}$$

where the inequality follows from (47) and (48). Hence, the transfers in (46) satisfy (ND), and the efficient allocation can be implemented without an outside subsidy.

Online Appendix

C Examples: Derivations

C.1 Derivations for the Heterogeneous-Audience Example

We now apply (9) to the example in Sec. 2.1.2. Let $L := 2\varepsilon - 1$, so ω_1 and ω_2 are independent and uniform on $[L, 1]$. The resulting CTR distribution is absolutely continuous and thus satisfies Assumption 1; consequently, ties occur only on a null set and no nontrivial tie-breaking rule is needed. For a type $\theta \sim \text{Uniform}[0, 1]$, define the distribution function of the ironed virtual type by

$$\Gamma(r; z) := \mathbb{P}\Big\{\overline{\phi}(\theta; z) < r\Big\} = \begin{cases} 0, & r \le 0, \\ r/2, & 0 < r \le z, \\ (1+r)/2, & z < r \le 1, \\ 1, & r > \max\{z, 1\}, \end{cases}$$

where the third case applies only when $z < 1$. Conditional on a customer having CTRs (ω_1, ω_2), merchant 1's critical type wins whenever $z_1\omega_1 > \omega_2\overline{\phi}_2(\theta_2; z_2)$, an event with probability $\Gamma(z_1\omega_1/\omega_2; z_2)$; the analogous expression applies to merchant 2. Ironing parameters (z_1, z_2) for which each critical type receive her outside option expected clicks exists, thus (z_1, z_2) satisfy

$$\begin{aligned} \alpha_1\varepsilon &= \frac{1}{(1-L)^2}\int_L^1\int_L^1 x\,\Gamma\Big(\frac{z_1\omega_1}{\omega_2}; z_2\Big)\,d\omega_2\,d\omega_1, \\ \alpha_2\varepsilon &= \frac{1}{(1-L)^2}\int_L^1\int_L^1 y\,\Gamma\Big(\frac{z_2\omega_2}{\omega_1}; z_1\Big)\,d\omega_1\,d\omega_2. \end{aligned} \tag{49}$$

Solving numerically this system for each ε gives the curves in Fig. 2.

C.2 Derivations for the Shopify Product-Tier Calibration

We verify the assumptions, cutoffs, and market-clearing conditions required by Theorem 2 for the low-, medium-, and high-tier products in Sec. 2.2. Fix any nonzero $\lambda \in \mathbb{R}^3_+$. Among the coordinates with $\lambda_k > 0$, the shadow values $\lambda_k\omega_k$ are independent and atomless and therefore do not tie with positive probability. They are also strictly positive almost surely, so coordinates with zero weight cannot maximize. Thus, Assumption 2 holds.

Pooled-audience scarcity also holds. Given $\eta_r = 1$ and uniform type distribution on

$\left[\underline{\theta}_k, \overline{\theta}_k\right]$, the virtual cost and value functions are

$$\phi_k^S(\theta) = 2\theta - \underline{\theta}_k \quad \text{and} \quad \phi_k^B(\theta) = 2\theta - \overline{\theta}_k.$$

At a zero shadow price, the buying cutoffs are 1/4, 3/8, and 1/2 in tiers L, M, and H, respectively. Hence, the left-hand side of Assumption 3 equals

$$\frac{1}{6}\left(1 - \frac{3}{4}\frac{1}{2}\right) + \frac{1}{6}\left(1 - \frac{5}{8}\frac{1}{4}\right) + \frac{2}{3} = \frac{175}{192}.$$

Moreover, ω_{L} can never be maximal because $\omega_{\mathrm{L}} \leq 1/3 < 1/2 \leq \omega_{\mathrm{H}}$. Therefore,

$$\mathbb{E}_G\left[\max_k \omega_k\right] = \mathbb{E}_G[\max\{\omega_{\mathrm{M}}, \omega_{\mathrm{H}}\}] = \frac{3}{4} + \frac{1}{351} = \frac{1057}{1404} < \frac{175}{192}.$$

Thus, pooled-audience scarcity holds strictly.

Consider shadow prices in the region

$$\frac{1}{2} < \lambda_{\mathrm{L}} < 1, \quad \frac{1}{4} < \lambda_{\mathrm{M}} < \frac{3}{4}, \quad 0 < \lambda_{\mathrm{H}} < \frac{1}{2}.$$

Inverting the virtual cost and value functions and applying the inverse convention gives

$$\left(p_{\mathrm{L}}^S, p_{\mathrm{L}}^B\right) = \left(\frac{\lambda_{\mathrm{L}}}{2}, \frac{1}{2}\right), \quad \left(p_{\mathrm{M}}^S, p_{\mathrm{M}}^B\right) = \left(\frac{\lambda_{\mathrm{M}} + 1/4}{2}, \frac{\lambda_{\mathrm{M}} + 3/4}{2}\right), \quad \left(p_{\mathrm{H}}^S, p_{\mathrm{H}}^B\right) = \left(\frac{1}{2}, \frac{\lambda_{\mathrm{H}} + 1}{2}\right).$$

The induced demands are therefore

$$\begin{aligned} d_{\mathrm{L}}(\lambda) &= \frac{1}{24}(1 - \lambda_{\mathrm{L}}), \\ d_{\mathrm{M}}(\lambda) &= \frac{1}{6}\left(\frac{15}{16} - \lambda_{\mathrm{M}}\right), \\ d_{\mathrm{H}}(\lambda) &= \frac{2}{3}\left(1 - \frac{3}{10}\lambda_{\mathrm{H}}\right). \end{aligned} \tag{50}$$

It remains to equate these demands with expected clicks supply from the pooled audience. Let G_k be the CDF of ω_k, and let

$$(\underline{\omega}_{\mathrm{L}}, \underline{\omega}_{\mathrm{M}}, \underline{\omega}_{\mathrm{H}}) = \left(\frac{1}{4}, \frac{1}{8}, \frac{1}{2}\right), \quad (\overline{\omega}_{\mathrm{L}}, \overline{\omega}_{\mathrm{M}}, \overline{\omega}_{\mathrm{H}}) = \left(\frac{1}{3}, \frac{2}{3}, 1\right).$$

For each tier,

$$G_k(x) = \min\left\{1, \max\left\{0, \frac{x - \underline{\omega}_k}{\overline{\omega}_k - \underline{\omega}_k}\right\}\right\}.$$

Independence implies that the expected clicks supply to tier k is

$$d_k^\lambda = \frac{1}{\overline{\omega}_k - \underline{\omega}_k} \int_{\underline{\omega}_k}^{\overline{\omega}_k} \omega_k \prod_{j \neq k} G_j\left(\frac{\lambda_k}{\lambda_j}\omega_k\right) d\omega_k, \quad k \in \{\mathrm{L}, \mathrm{M}, \mathrm{H}\}. \tag{51}$$

The product in the integrand is the conditional probability that tier k's weighted CSR exceeds those of both other tiers.

Solving $d_k^\lambda = d_k(\lambda)$ for every k gives

$$\lambda^* \approx (0.52133, 0.35176, 0.27707).$$

This solution lies in the region used to derive (50), and both sides of market clearing equal

$$d^* \approx (0.01994, 0.09762, 0.61125).$$

The resulting cutoff pairs are

$$\left(p_{\mathrm{L}}^S, p_{\mathrm{L}}^B\right) \approx (0.26067, 0.5), \quad \left(p_{\mathrm{M}}^S, p_{\mathrm{M}}^B\right) \approx (0.30088, 0.55088), \quad \left(p_{\mathrm{H}}^S, p_{\mathrm{H}}^B\right) \approx (0.5, 0.63853).$$

C.3 Derivations for Uniform Support-Function Example

Suppose there are two tiers and

$$\omega_1, \omega_2 \overset{\text{i.i.d.}}{\sim} \text{Uniform}[0, 1].$$

The support function of the feasible demand set is

$$\sigma(\lambda) = \mathbb{E}_G[\max\{\lambda_1\omega_1, \lambda_2\omega_2\}].$$

By positive homogeneity, for $\lambda_2 > 0$ we may normalize $\lambda_2 = 1$ and write $r := \lambda_1/\lambda_2$. A customer is then assigned to tier 1 when $r\omega_1 \geq \omega_2$ and to tier 2 otherwise. The maximizing demand vector is

$$d(r) = (d_1(r), d_2(r)),$$

where

$$d_1(r) = \mathbb{E}_G[\omega_1 \mathbf{1}\{r\omega_1 \geq \omega_2\}] \quad \text{and} \quad d_2(r) = \mathbb{E}_G[\omega_2 \mathbf{1}\{\omega_2 > r\omega_1\}].$$

Suppose first that $0 \leq r \leq 1$. Direct integration gives

$$\begin{aligned} d_1(r) &= \int_0^1 \int_0^{r\omega_1} \omega_1 \, d\omega_2 \, d\omega_1 = \frac{r}{3}, \\ d_2(r) &= \int_0^1 \int_{r\omega_1}^1 \omega_2 \, d\omega_2 \, d\omega_1 = \frac{1}{2} - \frac{r^2}{6}. \end{aligned}$$

It follows that

$$\sigma(r, 1) = r d_1(r) + d_2(r) = \frac{1}{2} + \frac{r^2}{6}.$$

Now suppose that $r \geq 1$. Splitting the integral for tier 1 at $\omega_1 = 1/r$ gives

$$\begin{aligned} d_1(r) &= \int_0^{1/r} \int_0^{r\omega_1} \omega_1 \, d\omega_2 \, d\omega_1 + \int_{1/r}^1 \int_0^1 \omega_1 \, d\omega_2 \, d\omega_1 \\ &= \frac{1}{2} - \frac{1}{6r^2}, \\ d_2(r) &= \int_0^1 \int_0^{\omega_2/r} \omega_2 \, d\omega_1 \, d\omega_2 = \frac{1}{3r}. \end{aligned}$$

Hence,

$$\sigma(r,1) = r d_1(r) + d_2(r) = \frac{r}{2} + \frac{1}{6r}.$$

Combining the two cases gives

$$\sigma(r,1) = \begin{cases} \dfrac{1}{2} + \dfrac{r^2}{6}, & 0 \le r \le 1, \\ \dfrac{r}{2} + \dfrac{1}{6r}, & r \ge 1, \end{cases} \tag{52}$$

and the associated support-maximizing demands are

$$d(r) = \begin{cases} \left(\dfrac{r}{3}, \dfrac{1}{2} - \dfrac{r^2}{6}\right), & 0 \le r \le 1, \\ \left(\dfrac{1}{2} - \dfrac{1}{6r^2}, \dfrac{1}{3r}\right), & r \ge 1. \end{cases} \tag{53}$$

The event $r\omega_1 = \omega_2$ has probability zero, so this demand vector is uniquely determined for every $r \ge 0$.

Eliminating r from (53) yields the upper frontier of $\mathcal{D}(G)$:

$$d_2 = \begin{cases} \dfrac{1}{2} - \dfrac{3}{2} d_1^2, & 0 \le d_1 \le \frac{1}{3}, \\ \sqrt{\dfrac{1 - 2d_1}{3}}, & \frac{1}{3} \le d_1 \le \frac{1}{2}. \end{cases} \tag{54}$$

Thus, the frontier runs from $(0, 1/2)$ through $(1/3, 1/3)$ to $(1/2, 0)$. In particular, setting $r = 1$ gives

$$d(1) = \left(\frac{1}{3}, \frac{1}{3}\right) \quad \text{and} \quad \sigma(1,1) = \frac{2}{3},$$

which produces the supporting line $d_1 + d_2 = 2/3$ shown in Fig. 4.