Data, Competition, and Digital Platforms

Dirk Bergemann[1] Alessandro Bonatti[2]

[1]Yale University

[2]MIT Sloan

Federal Trade Commission Microeconomics Conference

November 2022

Introduction

Digital platforms: **information** gatekeepers and **competition** managers.

- Surplus **creation** from matching consumers and products thanks to data from past and concurrent transactions.
- Concerns over surplus **extraction** from inducing seller market power.

Dual gatekeeper position under recent regulatory scrutiny:

> *One cannot exclude the possibility that a dominant platform could have incentives to sell "monopoly positions" to sellers by showing buyers alternatives which do not meet their needs.* *Crémer et al. (2019)*

European Unions Digital Markets Act went into effect November 1st, 2022.

Personalization and Its Limits

Personalized (sponsored) content drives both value creation and extraction.

- Retail platforms: eBay, Wayfair, Booking, Orbitz, Amazon...
- Advertising platforms (including display networks): Google, Meta, Microsoft, Twitter, Tiktok, Youtube, Criteo...

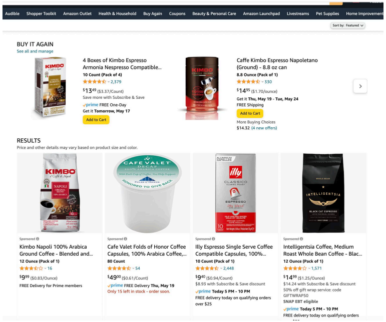
Audible Shopper Toolkit Amazon Outlet Health & Household Buy Again Coupons Beauty & Personal Care Amazon Launchpad Livestreams Pet Supplies Home Improvement
Sort by: Featured
BUY IT AGAIN
See all and manage
4 Boxes of Kimbo Espresso Armonia Nespresso Compatible...
10 Count (Pack of 4)
2,379
$13.49 ($3.37/Count)
Save more with Subscribe & Save
prime FREE One-Day
Get it Tomorrow, May 17
Add to Cart
Caffe Kimbo Espresso Napoletano (Ground) - 8.8 oz can
8.8 Ounce (Pack of 1)
330
$14.95 ($1.70/ounce)
Get it Thu, May 19 - Tue, May 24
FREE Shipping
Add to Cart
More Buying Choices
$14.32 (4 new offers)
RESULTS
Price and other details may vary based on product size and color.
Sponsored
Kimbo Napoli 100% Arabica Ground Coffee - Blended and...
12 Ounce (Pack of 1)
16
$9.99 ($0.83/Ounce)
FREE Delivery for Prime members
Sponsored
Cafe Valet Folds of Honor Coffee Capsules, 100% Arabica Coffee,...
80 Count
54
$49.00 ($0.61/Count)
prime FREE Delivery Thu, May 19
Only 15 left in stock - order soon.
Sponsored
Illy Espresso Single Serve Coffee Compatible Capsules, 100%...
10 Count (Pack of 1)
2,448
$9.40 ($0.94/Count)
$8.93 with Subscribe & Save discount
prime Today 5 PM - 10 PM
FREE delivery today on qualifying orders over $25
Sponsored
Intelligentsia Coffee, Medium Roast Whole Bean Coffee - Blac...
12 Ounce (Pack of 1)
1,571
$14.99 ($1.25/Ounce)
$14.24 with Subscribe & Save discount
50% off gift wrap service: code GIFTWRAP50
SNAP EBT eligible
prime Today 5 PM - 10 PM
FREE delivery today on qualifying orders


Instagram
Wharton Executive Education's Chief Techn...
Sponsored
Wharton
Aresty Institute of Executive Education
University of Pennsylvania
9 to 12 Months | Online + On Campus
CHIEF TECHNOLOGY OFFICER PROGRAM
Champion technological advances

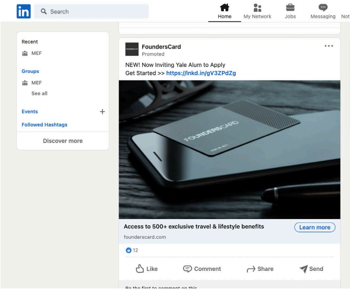
Search
Home
My Network
Jobs
Messaging
Recent
MEF
Groups
MEF
See all
Events
Followed Hashtags
Discover more
FoundersCard
Promoted
NEW! Now Inviting Yale Alum to Apply
Get Started >> https://lnkd.in/gV3ZPdZg
Access to 500+ exclusive travel & lifestyle benefits
Learn more
founderscard.com
12
Like
Comment
Share
Send


今日
买一送三

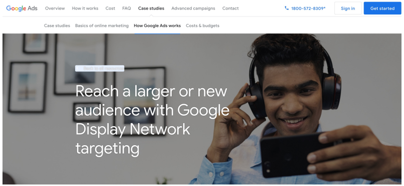
Google Ads
Overview
How it works
Cost
FAQ
Case studies
Advanced campaigns
Contact
1800-572-8309*
Sign in
Get started
Case studies
Basics of online marketing
How Google Ads works
Costs & budgets
Reach a larger or new audience with Google Display Network targeting

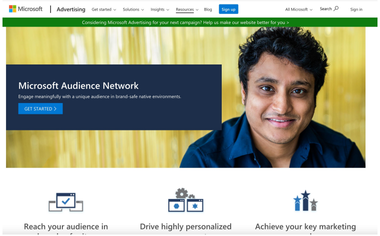
Microsoft
Advertising
Get started
Solutions
Insights
Resources
Blog
Sign up
All Microsoft
Search
Sign in
Considering Microsoft Advertising for your next campaign? Help us make our website better for you >
Microsoft Audience Network
Engage meaningfully with a unique audience in brand-safe native environments.
GET STARTED >
Reach your audience in
Drive highly personalized
Achieve your key marketing

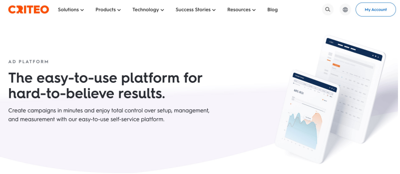
CRITEO
Solutions
Products
Technology
Success Stories
Resources
Blog
My Account
AD PLATFORM
The easy-to-use platform for hard-to-believe results.
Create campaigns in minutes and enjoy total control over setup, management, and measurement with our easy-to-use self-service platform.

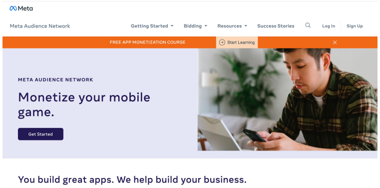
Meta
Meta Audience Network
Getting Started
Bidding
Resources
Success Stories
Log In
Sign Up
FREE APP MONETIZATION COURSE
Start Learning
META AUDIENCE NETWORK
Monetize your mobile game.
Get Started
You build great apps. We help build your business.

Personalization and Its Limits

Personalized (sponsored) content drives both value creation and extraction.

- Retail platforms: eBay, Wayfair, Booking, Orbitz, Amazon...
- Advertising platforms (including display networks): Google, Meta, Microsoft, Twitter, Tiktok, Youtube, Criteo...

Ubiquitous features of digital markets:

1. Platform leverages its informational advantage.
2. Product steering through personalized ads and offers.
3. Auction-based mechanisms to rank personalized content.

Managed Campaigns

17:49 LTE

support.google.com

Google Ads Help

About automated bidding

Automated bidding takes the heavy lifting and guesswork out of setting bids to meet your performance goals. Unlike Manual CPC bidding, there's no need to manually update bids for specific ad groups or keywords. Google Ads automatically sets bids for your ads based on that ad's likelihood to result in a click or conversion that helps you achieve a specific goal for your business.

Different types of automated bidding strategies can help you increase clicks, visibility and conversions. Automated bid strategies learn as they go, using information about a bid's performance to inform future bids. Learn how to determine a bid strategy based on your goals

This article describes different business goals and the automated bid strategy that best achieves each goal.

Note: If you'd like to automate your bidding specifically for a Shopping campaign, read About automated bidding for Shopping

17:50 LTE

support.google.com

Google Ads Help

About automated bidding for Shopping campaigns

Automated bid strategies for Shopping help you optimize your advertising spend. Using advanced machine learning, they monitor your campaign's performance and set a bid in every auction to help you achieve your goals.

This article explains the different automated bid strategies that are available for Shopping campaigns and how to choose the right one for you.

Benefits

- **You can focus on high level goals and allow Smart Bidding to set the right bid for you.** Whether you're trying to drive more visitors to your site or more revenue to your business, automated bidding allows you to start concentrating on overall performance of the campaign and less on how to set the perfect bid for each product group.
- **Your campaign's historical performance and future goals are always taken into**

Managed Campaigns

Example - Search Term Report

Search Engine Land SEO PPC Focuses SMX Webinars Intelligence Reports White Papers About

Search Engine Land » Google » Google Ads » Google's search terms move will make millions in ad spend invisible to advertisers

Google's search terms move will make millions in ad spend invisible to advertisers

The change removes visibility into more than 20% of search terms, one agency finds.

Ginny Marvin on September 3, 2020 at 3:58 pm

This morning, I negated a word that cost a campaign more than $3 for the one click it received in a brand campaign last week. I didn't add the whole query, just one irrelevant word that triggered a brand keyword. Going forward, I might not ever see that type word or know if it showed up across multiple low-volume queries.

As we reported yesterday, Google has notified advertisers the search terms report will "only include terms that were searched by a significant number of users." It has given no details about what "significant" means. The company told us the reason for the change is "to maintain our standards of privacy and strengthen our protections around user data."

Unsurprisingly, the move has angered advertisers.

https://searchengineland.com/googles-search-terms-move-will-make-millions-in-ad-spend-invisible-to-advertisers-340182

Example (1/5) - Drivers of optimizations are more automated

Before

- Structure granular
- Targeting Selection
- Manual Bid
- Query mining
- Ad Testing

Current

- Structure aggregated
- Broad targeting Selection
- Automated Bid
- Query mining limited visibility
- Ad Testing on autopilot

Questions

- How does the precision of the platform's data affect the creation and distribution of surplus, both on and off digital platforms?
- Do platforms with market power provide high-quality matches?
- Do platforms transfer market power "downstream" to sellers?
- Can we quantify the effect on prices both on and off the platform?
- What is the role of the platform's revenue model?
- How do different modes of data governance affect the creation and distribution of surplus, both on and off digital platforms?

  [DG:= mechanisms for collecting and transferring consumer data.]

Today

A model of digital platforms with four key features:

- Heterogeneity of buyer preferences and product characteristics.
- Personalization of sponsored content (ads, listings, offers).
- Product steering by managed campaigns but no personalized prices.
- Sellers with parallel sales channels (on- and off-platform).

Related Literature

- Information gatekeepers: Baye and Morgan (2001)...

  Armstrong and Zhou (2011), Rayo and Segal (2010), Inderst and Ottaviani (2012), de Corniere and de Nijs (2016).

→ Our model: multiple gates, heterogeneous products.

- Showrooming, steering, and multiple sales channels: Wang and Wright (2020), Miklos-Thal and Shaffer (2021), Bar-Isaac and Shelegia (2020), Idem (2021), Teh and Wright (2022), Lee (2022).

→ Our model: sponsored content, auction-like mechanism, no merchant fees.

- Platform self-preferencing: Anderson and Bedre-Defolie (2021), Hagiu et al. (2020), Gutierrez (2021), Kang and Muir (2021), Lam (2021), Lee and Musolff (2021), Raval (2022)...

Example

Single Seller

Single seller (Mussa and Rosen, 1978) and a unit mass of consumers.

Binary consumer type $\theta \in \{\theta_L, \theta_H\}$ with distribution $f(\theta)$.

Consumer type θ has valuation

$$\theta \cdot q$$

for a product of quality q.

Seller produces goods of quality q at cost $c(q) = q^2/2$.

$1 - \lambda$ consumers buy directly from the seller.

$\lambda \in [0, 1]$ consumers visit a monopolist platform that runs ads.

Managed Campaign and Posted Prices

Platform knows each consumer's type and offers seller a **managed campaign**.

Platform charges a fixed upfront fee t (e.g., campaign budget).

Seller "uploads" personalized offers $(q(\theta), p(\theta))$ for on-platform consumers.

Platform shows each consumer θ the relevant (profit-maximizing) offer.

Seller posts menu of quality, price pairs $(\hat{q}(\theta), \hat{p}(\theta))$ for off-platform consumers.

The on-platform offers may differ from the posted products and prices.

Summary

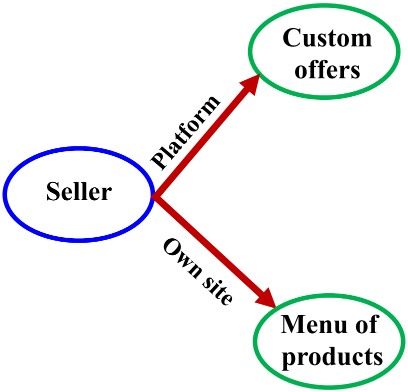

Summary

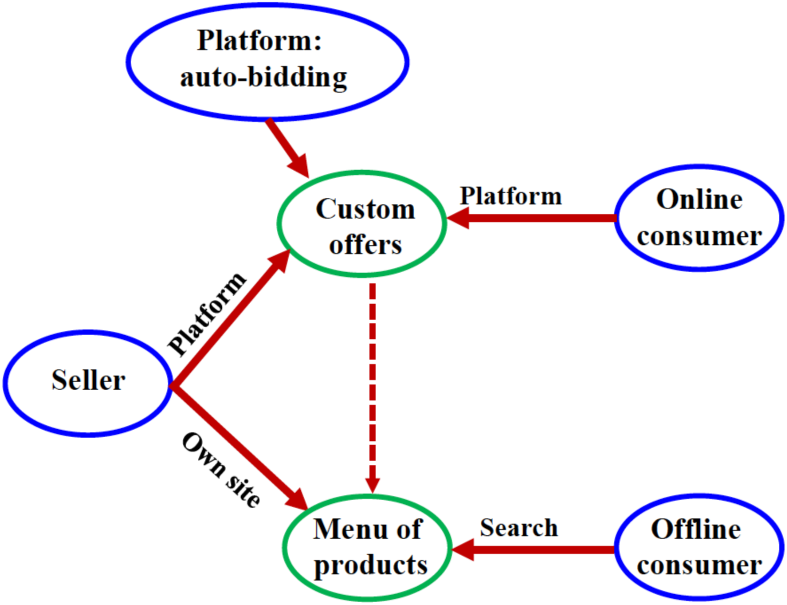

Platform:
auto-bidding
Custom
offers
Platform
Online
consumer
Platform
Seller
Own site
Menu of
products
Search
Offline
consumer

Seller's Problem

Off-platform, seller must screen consumer types as in Mussa and Rosen (1978).

- Low type θ_L obtains zero rents, high type θ_H obtains $\hat{U}(\theta_H) \geq 0$.

On-platform consumers see one product only: ideally, seller would offer

- socially efficient quality $q^*(\theta) = \arg\max \theta q - q^2/2 = \theta$,
- and charge consumer's full wtp, $p(\theta) = \theta q^*(\theta) = \theta^2$.

"Showrooming" constraint (on-platform consumers can buy directly from seller):

$$U(\theta) := \theta q(\theta) - p(\theta) \geq \theta \hat{q}(\theta) - \hat{p}(\theta) =: \hat{U}(\theta) \quad \forall \theta,$$

in addition to individual rationality and incentive compatibility off-platform.

$\Rightarrow$ Trade under symmetric information; limited ability to price discriminate.

Optimal Menus

Proposition (Single Seller, Binary Types)

The seller offers the efficient quality levels on-platform to each buyer type, $q(\theta) = \theta$ for all θ, and showrooming binds, $U(\theta) = \hat{U}(\theta)$ for all θ.

The optimal off-platform menu of products is given by

$$\begin{aligned} \hat{q}\left(\theta_L\right) &= \max\left\{0, \theta_L - \frac{f\left(\theta_H\right)}{f\left(\theta_L\right)}\left(\theta_H - \theta_L\right)\left(1 + \frac{\lambda}{1-\lambda}\right)\right\}, \\ \hat{q}\left(\theta_H\right) &= \theta_H. \end{aligned}$$

Without an off-platform store, the seller could charge each type their full wtp.

"Double" opportunity cost of serving the low type off-platform:
rents to high types off platform $\Rightarrow$ rents to high types on the platform.

Optimal Menus

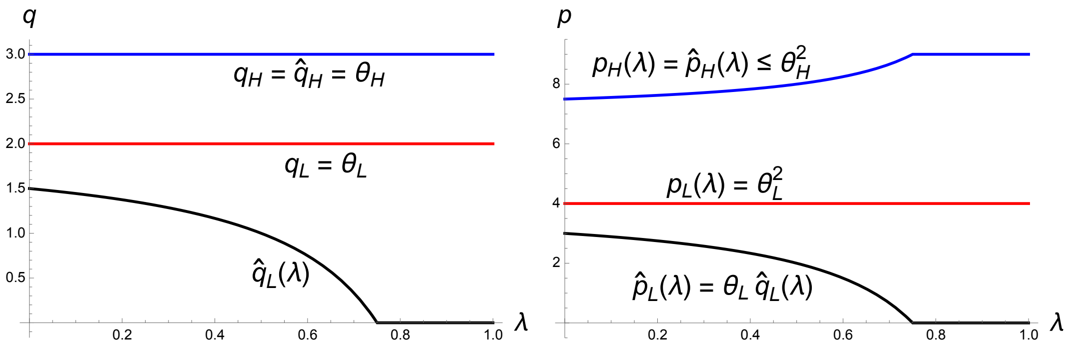

Taking Stock

Compared to Mussa and Rosen (1978) solution (i.e., $\lambda = 0$):

$\rightarrow$ lower quality $\hat{q}(\theta_L)$ and higher prices $\hat{p}(\theta_H)$ off the platform.

On platform: efficient quality; no rent for low types; positive rent for high types.

No need for MFN clauses: fixed fee $\Rightarrow$ lowest prices are on the platform.

What about the platform's fee t^*?

- Seller can guarantee the Mussa and Rosen (1978) profits off-platform.
- The platform's fee t^* extracts all the seller's extra profits: on-platform profits *minus* the losses from the distortions in the off-platform menu.
- This means on-platform profits $> t^*$. Campaign "delivers" $ROI > 0$.

Too Easy or Too Hard?

The example is “too easy:”

- Single seller vs. competing sellers.
- Two types vs. arbitrary type distributions.
- Symmetric information (consumers and platform).

It is also “too hard:”

- Linear pricing vs. nonlinear (quality) pricing?
- Personalized pricing vs. product steering.

Full model: can the platform create a “local monopolies” environment by managing the advertising campaigns of competing multiproduct sellers?

Full Model

Setup

J sellers and a unit mass of consumers with type $\theta = (\theta_1, \dots, \theta_j, \dots, \theta_J)$.

Consumer θ has value

$$\theta_j \cdot q_j$$

for a product of quality q_j by firm j.

Sellers offer vertically differentiated products with cost $c(q_j) = q_j^2/2$.

$\lambda \in [0, 1]$ consumers visit a monopolist platform that runs ads.

$1 - \lambda$ consumers buy directly from sellers.

Information Structure

Consumers' valuations θ_j with distribution F, i.i.d. across j.

The platform observes $\theta \in \mathbb{R}^J$ perfectly.

Every consumer observes a noisy signal s about θ.

Posterior mean $m_j = \mathbb{E}[\theta_j \mid s]$ with distribution G.

F is a mean-preserving spread of G. Assume same support.

On Platform: Managed Campaigns

Platform offers a **single** advertising slot per consumer.

Consumer type $\theta \sim$ *targeting category:* ads condition on her type.

Formally, the platform:

1. Charges a fixed fee t to participating sellers (e.g., campaign budget).
2. Solicits personalized ads from each seller j—functions $q_j(\theta)$ and $p_j(\theta)$.
3. Specifies which j gets the slot for consumer θ and what ad is shown.
4. Reveals to the consumer her θ_j for the selected product q_j.

Proposition (Optimal Mechanism)

The platform shows the seller j and the product (q_j, p_j) that maximize total surplus among all sellers that participate in the mechanism.

Off-Platform: Search

Off the platform, matching through consumer search:

- Consumer with expectations m faces search costs $\sigma > 0$ (first search free) as in Diamond (1971).
- Sellers j elicit consumers' wtp m_j through **menus** $(\hat{q}_j(m), \hat{p}_j(m))$ as in Mussa and Rosen (1978).
- Not an inspection good: learning θ_j requires the platform's data.

Timing

1. Platform announces managed campaign mechanism $\mathcal{M}$ and fee t.
2. Sellers simultaneously choose whether to participate in $\mathcal{M}$, their on-platform products and prices (q_j, p_j), and their off-platform menus $(\hat{q}_j, \hat{p}_j)$.
3. Type θ is realized, and a (seller, ad) pair is selected to be shown.
4. Consumer learns θ_j, buys on platform, or searches off-platform.

Equilibrium Search Patterns

Symmetric Equilibrium

Off platform, the Diamond (1971) paradox:

- $1-\lambda$ off-platform consumers with beliefs m face search costs $\sigma > 0$;
- they expect symmetric menus and visit $\hat{j} = \arg\max_j m_j$ only.

If platform has an informational advantage ($F \succ_{mps} G$):

- λ on-platform consumers infer $\theta_{j^*} = \max_j \theta_j$ (cannot detect deviations);
- they expect symmetric menus off-platform, both on and off path.

Proposition (Consideration Sets)

If $F \succ G$, every online consumer θ only compares the displayed seller j's online offer $(p_j(\theta), q_j(\theta))$ and its offline menu $(\hat{p}_j(\cdot), \hat{q}_j(\cdot))$.

...or in fewer words...

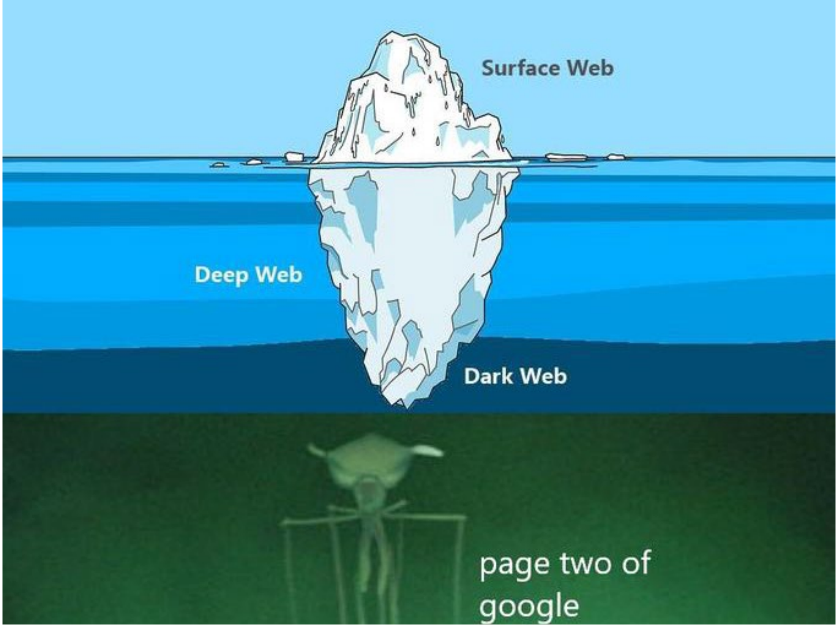

Equilibrium Search Patterns: Example

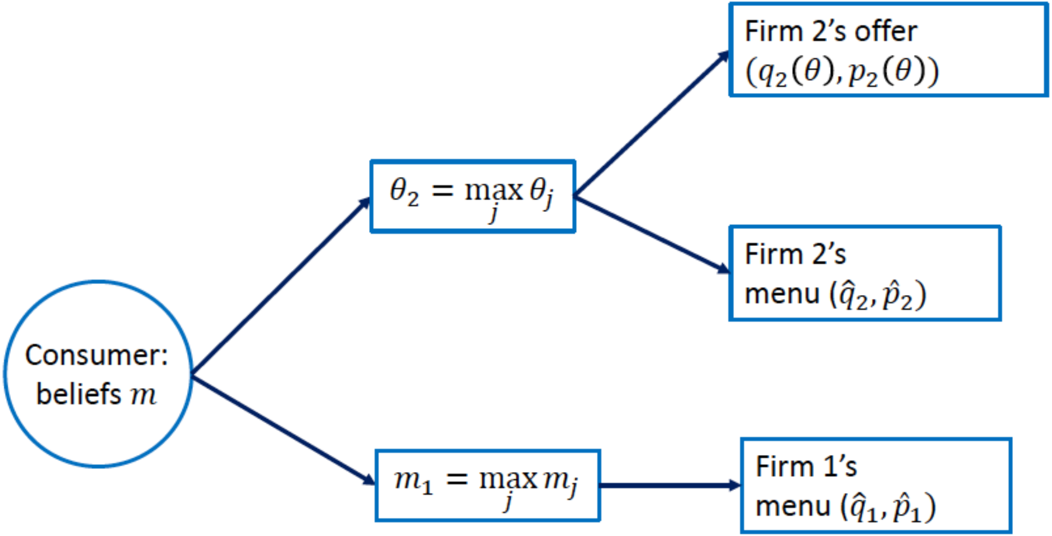

Interpretations

With a better-informed platform, equivalent interpretation:

- each brand has $(1 - \lambda)/J$ loyal (imperfectly informed) customers already shopping off-platform;
- the remaining λ consumers are not currently shoppers—they do not recognize any brands without the platform's data;
- these consumers can be turned into shoppers by informative advertising.

This result requires an (arbitrarily small) informational advantage:

- Without advantage vs. sellers: platform cannot make money.
- Without advantage vs. buyers, platform does not control outside options—consumers' beliefs determine where they search off platform.

Summary

A model of digital platforms where:

- Platforms monetize data through **managed campaigns**.
- Different information structures can be compared.
- On- and off-platform markets interact.

Superior information on the platform improves **match quality**:

- Consumers find their favorite sellers.
- Sellers offer consumers efficiently "tailored" products.

Information also introduces the potential for **surplus extraction**:

- Endogenously local monopolies.
- No price discrimination, but **product steering**.

Conclusions

Digital platforms monetize superior information about consumer preferences by auctioning access to the consumers' attention.

Product design and price decisions interact with modes of data governance (e.g., with the rules by which a platform shares its data).

Consumer surplus on- and off-platform is driven by information rents off platform and by the availability of organic information.

The growth of a platform's database (e.g., more consumers or better data) reduces each consumer's outside option and leads to higher prices.

Mitigating factors: fully **informed consumers**; free **organic content** (i.e., public off-platform prices); and **privacy protection** (e.g., cohort-based ads).

Looking Ahead: Privacy & Competition

Managed campaigns are privacy-preserving: sellers only learn ROI.

Sellers do not learn how much they "bid" for each buyer type.

Only the platform holds the data.

- Reduced risk of **leakages** and **spillovers**.
- A single firm uses the information at a time (in the model).

Yet, managed advertising campaigns restrict **competition** by design.

What about other mechanisms? (Bergemann, Bonatti, and Wu, 2023):
"Manual bidding" auctions (which are less private) also yield lower profits to platform, and higher prices to consumers, relative to auto-bidding.

Backup Slides

Equilibrium Product Lines

Matching and Product Steering

On platform, sellers can extract surplus through product steering.

"Showrooming constraint" for seller shown to θ:

$$U(\theta) \triangleq \theta\ q_j^*(\theta) - p_j(\theta) \geq \max_m \left[\theta\ q_j^0(m) - p_j^0(m)\right] \triangleq U^0(\theta).$$

Incentive-compatible menus off platform $\Rightarrow$ on-platform consumer compares

$$(q_j^*(\theta), p_j(\theta)) \text{ and } (q_j^0(\theta), p_j^0(\theta)).$$

On platform, clearly optimal to offer efficient quality $q^*(\theta) = \theta$.

On platform, surplus extraction limited by $U^0(\theta)$.

Seller j's Problem

Consider offline menu (q_j^0, U_j^0). Seller j's profits on online type θ_j:

$$\pi(\theta, U_j^0) = \theta_j^2/2 - U^0(\theta).$$

Seller's choice of menu off-platform:

$$\max_{q^0, U^0} (1-\lambda) \int_0^1 \left(\theta_j q^0(\theta_j) - q^0(\theta_j)^2/2 - U^0(\theta_j)\right) G^{J-1}(\theta_j) \mathrm{d}G(\theta_j)$$
$$+ \lambda \int_0^1 \left(\theta_j^2/2 - U^0(\theta_j)\right) F^{J-1}(\theta_j) \mathrm{d}F(\theta_j).$$

Equilibrium Menus

Proposition (Symmetric Equilibrium Menus)

The (unique, symmetric) equilibrium quality levels are given by

$$q^*(\theta) = \theta,$$

$$q^0(\theta) = \max\left\{0\ ,\ \underbrace{\theta - \frac{1-G^J(\theta)}{JG^{J-1}(\theta)g(\theta)}}_{\textit{MR quality}} - \frac{\lambda}{1-\lambda}\frac{1-F^J(\theta)}{JG^{J-1}(\theta)g(\theta)}\right\}$$

Furthermore,

$$U^*(\theta) = U^0(\theta) = \int_0^\theta q^0(m)dm.$$

Equilibrium Properties

Platform's data matches consumer to favorite brand.

Participating sellers invest in efficient quality (product customization).

Off platform, inefficient matching based on insufficient information, and inefficient quality under asymmetric information.

Opportunity cost of off-platform sales: positive rents on platform,

$$U^*(\theta) = U^0(\theta) > 0 \text{ iff } q^0(\theta) > 0.$$

Offline q^0 further distorted downward, more so the larger the platform size λ.

Quality Provision

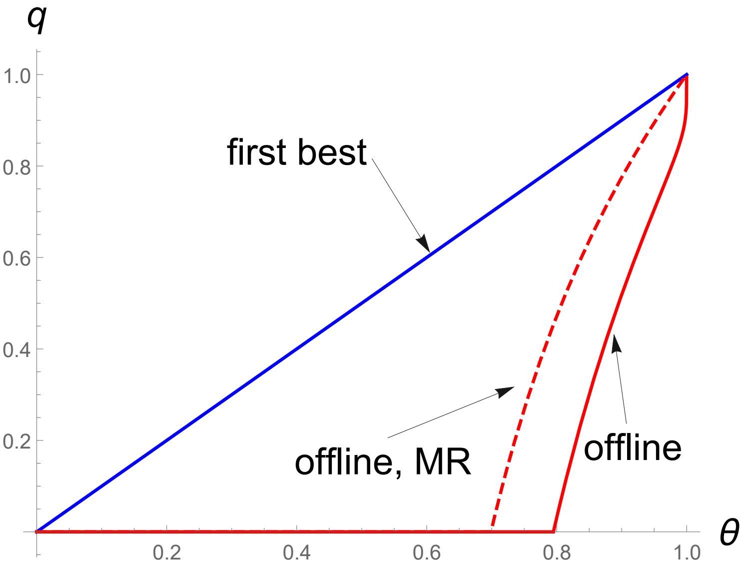

$$\lambda = 1/2, J = 5, G(m) = m, F(\theta) = Beta(\theta, 1/4, 1/4)$$

Payments

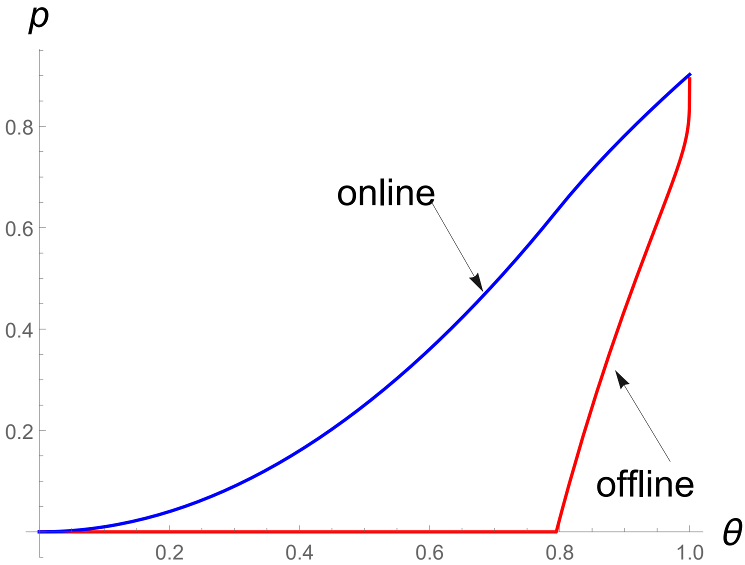

$\lambda = 1/2, J = 5, G(m) = m, F(\theta) = Beta(\theta, 1/4, 1/4)$

Nonlinear Tariffs

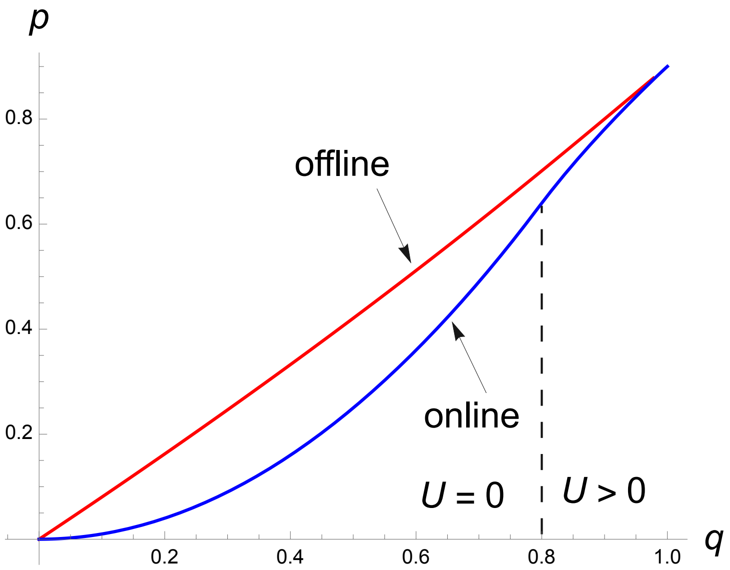

Every offline product is sold at a lower price online

Equilibrium and Consumer Value

$U^0(\theta) = U(\theta)$ for all θ.

But $F \succ_{mps} G$ and U convex $\Rightarrow \mathbb{E}_{F^J} U > \mathbb{E}_{G^J} U$.

This has several implications:

1. Ex ante, an individual consumer prefers to be on the platform.
2. Ex ante (holding prices fixed), an individual consumer wants the platform to disclose their information to sellers.
3. All consumers are worse off because of the platform.