Algorithms, learning, and pricing

Heski Bar-Isaac and Sandro Shelegia

April 2025

U Toronto, Competition Bureau, and CEPR; and UPF, BSE, and CEPR

Caveat: Which algorithm?

This talk will focus on ranking algorithms (think who is listed first on Walmart marketplace) not algorithms that determine price

These algorithms naturally affect on-platform prices (Castellini, Fletcher, Ormosi and Savani, 2023; Calvano, Calzolari, Denicolo and Pastorello, 2025; Johnson, Rhodes, Wildenbeest, 2023)

This talk will (primarily) focus on implications for off-platform prices and two allocative distortions:

- where (on which venue) do consumers buy? (a la "showrooming" Hagiu and Wright, 2024, Bar-Isaac and Shelegia, 2023)
- how platforms learn product quality and determine which products to feature so that consumers become aware of them?

Introduction

Policy context

Concern that platforms such as Amazon, Bookings, etc condition rankings on prices elsewhere (cf Hunold, Kesler, Laintenberger and Schlutter, 2018; Hunold, Laitenberger and Thebaudin, 2022)

Specifically, alleged that they do this to deter disintermediation/showrooming, hamper potential rivals etc.

This concern has led to legislation against such practices

- Article 6(5) of the DMA restricts platforms from using non-public data of competitors;
- UK's DMCC Bill prohibits anti-competitive ranking practices

And calls for "algorithm audits" to ensure that platforms are not introducing "bias"

Algorithms that do not condition on rival prices

Even if an algorithm does not condition on rival prices *directly*, it might do so *indirectly* through conditioning on sales

Conditioning on past sales is natural/necessary when new goods/varieties are introduced

We ask: How does a platform-optimal algorithm that does not condition on off-platform prices compare to one that does? in terms of

- (off-platform) prices
- what gets featured by the algorithm
- platform profits
- welfare

Sketch of the set up

Monopoly platform designs an algorithm to determine what is featured and charges a fee for every on-platform transaction

Consumers become aware of a product after it is featured and can buy on- or off-platform (so platform may have concerns about showrooming)

Seller chooses prices on and off-platform

Repeated twice so second-period algorithm can depend on first-period sales

Sketch of the set up II

Consumers vary in their preferences for buying on- or off-platform (in particular, some may prefer buying off-platform)

We start by supposing that the type of the seller is known

And then the (more natural?) case where the platform learns the product's popularity

The core idea and its relevance

Sellers try to manipulate ranking algorithms

Going back to early days of search there was a large advice industry in "search engine optimization" (SEO)

An important relevant search engine now is the ranking algorithm that for Amazon Booking etc first page of results etc. SEO has not disappeared...

The internet is rife with advice like *Your first option for generating sales is by driving both internal and external traffic to your Amazon listing."* (Bigcommerce.com)

One way to send more sales to the platform is to make own channel more expensive (complete diversion is equivalent to a prohibitively high price)... (and platform might not mind that)

Preview of findings

- Welfare can be higher when a platform charges higher fees: seller gets more aggressive in the direct channel and consumer allocation across platform and direct channel can improve
- Platform algorithm might distort the sales threshold (the level that a product must reach to be featured)
- Conditioning on prices directly may be better for welfare than not: a sales-based algorithm may try to replicate the price-based algorithm but does so in a more distortive way.

Model (baseline)

Baseline Model (Known type seller)

Monopoly platform M where consumers learn about a featured seller and can buy.

Seller S who can only become known to consumers via M. Once a consumer knows about the product, might also buy directly.

Consumers' willingness to pay for the (featured) seller's good on the platform is v.

- later we suppose that only some fraction of consumers (which the platform must estimate) are interested and uninterested consumers value the good at 0

Seller sets price $p^{platform}$ on the platform and must pay a proportionate fee f on all sales on the platform

Baeline Model cont.

Seller also has a direct channel where choose price p^{direct}

If a buyer buys from the direct channel, she suffers a disutility δ which is distributed according to $H(\delta)$ on $[\underline{\delta}, \bar{\delta}]$ with $\bar{\delta} > 0$ and $\underline{\delta} \in [-v, 0]$

When $\underline{\delta} < 0$ some consumers prefer to buy off-platform.

Consumer of type δ (who is aware of the product) buys from platform when $p^{platform} < p^{direct} + \delta$ or when *direct channel discount* $\Delta \equiv p^{platform} - p^{direct} < \delta$

Timing

There are two periods, seller and platform both put weight β on period 2

Before either period M chooses an algorithm to determines whether the product is featured or not

- if it cannot condition on off-platform prices, algorithm depends on
  - $p_1^{platform}$ in period 1
  - $p_2^{platform}$, $p_1^{platform}$, and S_1 corresponding to first-period platform sales in period 2
- if it can condition on off-platform prices, it depends on
  - $p_1^{platform}$ and p_1^{direct} in period 1
  - $p_2^{platform}$, p_2^{direct}, $p_1^{platform}$, p_1^{direct}, and S_1 in period 2

Costs

Seller sets prices each period and incurs constant marginal cost c on any units sold

In any period that the platform does not feature the seller it earns some alternative A

Analysis

Second-period pricing

It is immediate that the platform can “force” the on-platform price to be as high as it likes and so it is immediate that $p_2^{platform} = v$

Consider any featured seller and as above $\Delta_2 = v - p_2^{direct}$ as the discount on the direct channel

Those consumers with $p_2^{platform} \leq p^S_{direct} + \delta$ buy from the platform i.e. $1 - H(\Delta_2)$ buy on the platform and $H(\Delta_2)$ buy on the direct channel

Seller’s average profit per consumer is

$$\pi(\Delta_2) \equiv v(1-f) - c + H(\Delta)(v - \Delta_2).$$

Second-period pricing: characterization

Proposition

In period 2, seller sets Δ_2^* *as the unique* [1] *to* $\Delta = fv - \frac{H(\Delta)}{h(\Delta),}$ *if* $vf - \underline{\delta} - \frac{1}{h(\underline{\delta})} < 0$,

and $\Delta_2^* = \underline{\delta}$ *otherwise.*

In the latter case, all buy from platform

In the former, seller prices optimally given differentiation and higher margin on the direct channel (where it does not incur the platform fee)

[1] under some fairly innocuous conditions: strict log concavity of H(.) and $lim_{x \to 0} \frac{H(x)}{h(x)} = 0$ which hold for uniform, triangular, beta

Consumers can benefit from positive fees

Given setup (which makes on-platform pricing problem trivial) there is never any on-platform consumer surplus

Discount may be negative For example if $f = 0$ and $\underline{\delta} < 0$ (i.e. some consumers prefer buying from the direct channel)

Socially optimal discount is 0 so that consumers are optimally allocated

Optimal fee is positive when $\underline{\delta} < 0$, as a consequence of the two observations above. (cf Hagiu and Wright (forthcoming) who consider a related model where $\underline{\delta} = 0$)

Non-zero fee can raise platform profits, welfare, and consumer surplus relative to 0 fees!

Second-period pricing when algorithm conditions on direct-channel prices

Trivially the platform will only feature the product if $p_2^{platform} = v$ and $\Delta_2 \leq \underline{\delta}$ (so that no sales leak to direct channel)

Worse for consumers since the on-platform price is the same as in the no-conditioning case but there is no opportunity to buy off-platform (which consumers do when it makes them better off)

Welfare a little more subtle: misallocation to direct channel vs to platform channel?

First-period pricing

In the first period, the algorithm conditions only on $p_1^{platform}$ and trivially, as above "forces" $p_1^{platform} = v$

In addition to requiring $p_2^{platform} = v$, the algorithm also sets a threshold T level of sales to be featured in period 2. If it does so, then it earns $\pi(\Delta_2^*)$ as characterized above.

Thus the seller chooses Δ_1 to maximize

$$\pi(\Delta_1) \text{ if } 1 - H(\Delta_1) > T$$
$$\pi(\Delta_1) + \beta\pi(\Delta_2^*) \text{ otherwise}$$

Characterization of algorithm and first-period pricing

Proposition
In equilibrium, the seller sets the first period discount Δ_1^* *as the solution to* $\pi(\Delta) = (1-\beta)\pi(\Delta_2^*)$ *if* $\beta \geq 1 - \frac{\pi(\underline{\delta})}{\pi(\Delta_2^*)}$ *and sets* $\Delta_1^* \leq \underline{\delta}$ *otherwise (i.e. makes no direct channel sales).*

The platform optimally features the seller if first-period sales (given by $1 - H(\Delta_1)$*) are at least* $T^* = 1 - H(\Delta_1^*)$ *and does not feature the seller otherwise.*

Here, the seller's best outside option is to get the static profit-maximizing discount. As before this is Δ_2^* and allows it earn $\pi(\Delta_2^*)$ today but nothing in the next period, instead setting the anticipated discount gives $\pi(\Delta_1^*) + \beta\pi(\Delta_2^*)$.

First-period pricing when the algorithm can condition on prices

Again, the platform will only feature if $p_1^{platform} = v$ and $\Delta_1 \leq \underline{\delta}$ (so that no sales leak to direct channel)

Given that can condition on first-period prices on- and off-platform, platform does not need to condition on sales or gain by doing so

Again and just as for the second-period worse for consumers than sales-based algo since on-platform prices same as the no-conditioning case but no opportunity to buy off-platform

Seller of uncertain type

Seller of uncertain type

Platform has to learn seller's type in the first period

- realistic
- a natural rationale for conditioning on sales
- whereas in the baseline model, allocative efficiency (in equilibrium) is associated with the venue from which consumers buy; uncertain types potentially introduces an additional concern: is the product featured when it should be?
    - introduces new economic considerations and changes results
    - notably, a force through which welfare may be higher from an algorithm that conditions on prices

Extending the model for seller of uncertain type

With probability θ consumers find the seller's good appealing.

If they do not find the good appealing then they value it at 0 (both on and off platform)

If they find it appealing then, as in the baseline, they value it at v on the platform and $v - \delta$ on the direct channel

If the platform knows θ, analysis is identical to the baseline model (but for the factor θ that would appear in profits)

Initially θ is unknown and follows some distribution $G(\theta)$ on $[0, 1]$ with mean μ. Later we will suppose that $r(\theta) \equiv \theta^2 g(\theta)$ is quasi-concave on $[0, 1]$ and increasing on $[0, \mu]$

Second-period pricing

Identical to known type ... same results apply $p_2^{platform} = v$ and

Proposition
In period 2, when the algorithm cannot condition on the off-platform price, the seller sets Δ_2^* as the unique solution to $\Delta = fv - \frac{H(\Delta)}{h(\Delta)}$ if $vf - \underline{\delta} - \frac{1}{h(\underline{\delta})} < 0$ and sets $\Delta_2^* = \underline{\delta}$ otherwise.

When algorithm can condition on off-platform price platform, it will only feature if $p_2^{platform} = v$ and $p_2^{direct} \geq v - \underline{\delta}$ (so that no sales leak to direct channel)

Optimal algorithm: preliminaries

Lemma
It is optimal to use a threshold strategy; that is to set a threshold T such that sales in period 1 higher than T (and a second-period price of $p_2^{platform} = v$) are required to be featured in period 2.

Convenient assumptions

- outside option not too high so worth featuring good
- Sellet's profit is quasi-concave in Δ_1

Critical Types and Notation

Given an anticipated equilibrium discount Δ_1^*, it is convenient to define a marginal type that just clears the sales threshold. We write $\Omega^* \equiv \frac{T}{1-H(\Delta_1^*)}$ and the seller is featured in period 2 if

$$\theta \geq \frac{T}{1 - H(\Delta_1)} = \Omega^* \frac{1 - H(\Delta_1^*)}{1 - H(\Delta_1)}$$

Also convenient to introduce notation to reflect a type $\alpha \equiv \frac{A}{fv(1-H(\Delta_2^*))}$ corresponding to alternative A that the platform could earn and yields the same payoff

Optimal algorithm and first period pricing

Seller's expected profit is given by

$$\pi(\Delta_1)\mu + \beta\pi(\Delta_2^*) \int_{\frac{\Omega(1-H(\Delta_1^*))}{(1-H(\Delta_1))}}^{1} \theta g(\theta)\, d\theta$$

Proposition

When $\beta < \frac{fv\mu}{\pi(\Delta_2^*)\Omega^2 g(\Omega)}$ and $T \leq 1 - H(\Delta_2^*)$, the seller sets an off-platform price discount, Δ_1^*, that satisfies

$$\Delta_1^* = fv - \frac{H(\Delta_1^*)}{h(\Delta_1^*)} - \frac{\beta\pi(\Delta_2^*)\Omega^2 g(\Omega)}{\mu(1 - H(\Delta_1^*))}$$

This is decreasing in β and c and increasing in f and in v.
Otherwise, it foregoes the direct channel (sets $\Delta_1^* \leq \underline{\delta}$).

Comparing first- and second-period off platform discounts

$$\Delta_1^* = fv - \frac{H(\Delta_1^*)}{h(\Delta_1^*)} - \frac{\beta\pi(\Delta_2^*)\Omega^2 g(\Omega)}{\mu(1 - H(\Delta_1^*))}$$

This is a "career concerns" kind of effect and depends on density and value around the critical Ω threshold.

The new term is always negative, i.e. lower discount in the first period than in the second.

Δ_1^* is (weakly) larger than the discount for a known type (platform cannot distinguish off-platform price deviation from uncertainty about type and cannot extract rent as efficiently).

Δ_1^* is decreasing in β and c and increasing in v and f.

Comparison to “non-strategic” (no-commitment) platform

A myopic platform would feature a seller in period 2 if this earned more than the alternative A (recall our hitherto little-used notation for the outside option).

Equivalently, we can consider the (“neutral”) threshold type α

Characterization preliminaries

Derivative of optimal algorithm profit function

$$\Pi'_M(\Omega) = fv\left[\beta g(\Omega)(1 - H(\Delta_2^*))(\alpha - \Omega) - \mu h(\Delta_1^*)\frac{\partial \Delta_1^*}{\partial \Omega}\right]$$

Neutral/myopic/(?efficient) to set $\Omega = \alpha$

Here distort Ω in direction that that leads to more on-platform sales and immediate that $sign(\frac{\partial \Delta_1^*}{\partial \Omega}) = -sign(r'(\Omega))$

Given assumption that $r(\theta) \equiv \theta^2 g(\theta)$ is increasing on $[0, \mu]$, $\Omega^* < \alpha$

Characterizing the threshold

For β high enough

Proposition

There exist $\overline{\alpha}$ and $\underline{\alpha}$ (corresponding to $\overline{A}$ and $\underline{A}$) such that

1. *(Great platform alternative) If $\alpha \geq \overline{\alpha}$ then $\Omega^* = \alpha$ and $\Delta_1^* = \underline{\delta}$ (no distortion in allocation, no off-platform sales)*
2. *(Intermediate platform alternative) If $\alpha \in [\underline{\alpha}, \overline{\alpha})$ then $\Omega^* = \overline{\alpha}$ and $\Delta_1^* = \underline{\delta}$ (distorted allocation that induces no off-platform sales)*
3. *(Lousy platform alternative) If $\alpha < \underline{a}$ then $\Omega^* > \alpha$ and $\Delta_1^* > \underline{\delta}$.*
   (distorted allocation and some off-platform discounting)

Threshold type as a function of outside option

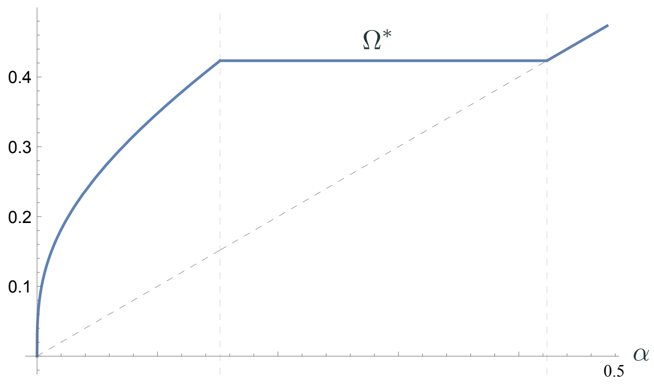

Figure 1: Optimal threshold Ω^* as a function of α with $f = 0.25$, $\beta = 2$, $G \sim U(0, 1)$ and $H \sim U(-\frac{1}{3}, 1)$

Discount as a function of outside option α

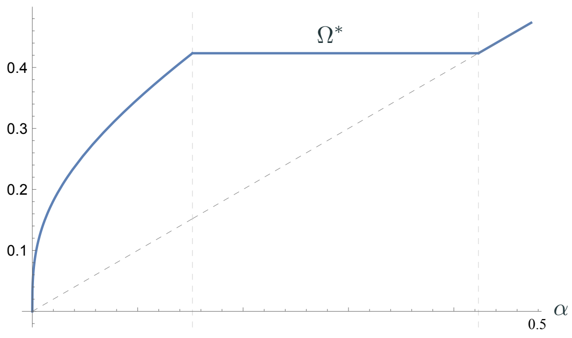

Figure 2: Optimal first and second period discounts with platform commitment (blue) at $f = 0.25$, $\beta = 2$, $G \sim U(0, 1)$ and $H \sim U(-\frac{1}{3}, 1)$

Comparison to algorithm that conditions on off-platform price

Again, an algorithm that conditions on prices, shuts down off-platform sales (i.e. will not feature in period 1 unless $\Delta_1 < \underline{\delta}$.

The algorithm that conditions on prices, allocates efficiently features second-period type if and only if θ (which is accurately estimated regardless of off-platform pricing since it is observed) is higher than α

Second-period allocative Distortions and First-Period Surplus relative to conditioning on prices

When α high enough then conditioning on prices and sales are equivalent in the first period and so only second-period trade-offs.

For α in an intermediate range then first-period discount is same as second-period but threshold to retain is affected allocation is worse, consumers are worse-off

For α low there is a distrortion to the threshold but there is a bigger off-platform discount benefit from lower off-platform prices

Welfare as a function of outside option for price-contingent algo (orange) and sales-contingent (blue)

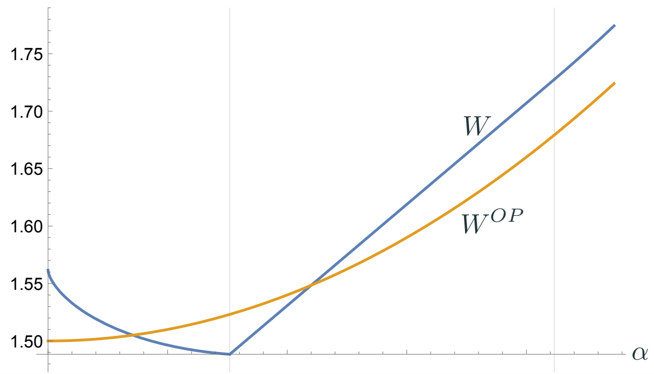

Figure 3: Welfare as a function of α at $f = 0.25$, $\beta = 2$, $G \sim U(0,1)$ and $H \sim U(-\frac{1}{3}, 1)$

Discussion and conclusion

Discussion

- Simple two-period model; With many periods (a la career concerns) distortions likely larger
- misallocation has a welfare implication but consumers never earn on-platform surplus since assume unit demand, if scope to earn some consumer surplus on platform then additional effects.
- competing platforms (fees "off-platform", discoverability and leakage in both directions?)
- endogenous fees?

Conclusions

- Contrast algorithms that rely on off-platform prices to those that do not
- In a one-shot case algorithms that rely on off-platform prices are worse for consumers, but welfare can go either way if some consumers prefer off-platform
- In a one-shot case where algorithms do not rely on off-platform prices, consumers may benefit from higher fees
- Even if cannot condition on off-platform prices directly, this happens indirectly when condition on sales
- When platforms learn about then sales-based algorithms may distort what gets featured
- Algorithmic audits are subtle since algorithms and on-platform outcomes impact on and from off-platform behavior

Haiku summary

Good algorithm?

Might rely on prices elsewhere.

So tough to audit.