Ranking algorithms, learning, and pricing

Heski Bar-Isaac and Sandro Shelegia

U Toronto, Hebrew University, CRESSE, and CEPR; and UPF, BSE, and CEPR

Economics of the DMA Workshop, Brussels, September 2025.

Supported by Shelegia's ERC project FAPoD (101044072)

Introduction

Ranking algorithms shape product visibility and sales.

Platforms design them strategically; regulators and policymakers are paying close attention.

Key concern: use of off-platform prices in rankings.

Another concern (predating e-commerce): Price Parity Clauses (PPCs).

New rules (DMA, DMCC) restrict what algorithms can do; some call for algorithmic audits.

Our Core Question

How do ranking algorithms operate in a learning environment where sellers set prices on- and off-platform?

How does an algorithm that ignores off-platform prices compare to one that uses them?

Do contractual restrictions on off-platform prices, such as PPCs, benefit consumers?

Preview of Findings

Even without conditioning on off-platform prices, algorithms can influence cross-channel behavior.

Sellers face "career concerns": raising off-platform prices boosts on-platform sales and promotion chances.

Platforms may distort promotion to exert even more from sellers.

Banning off-platform price use may lower welfare.

PPCs can improve both sales-channel allocation and promotion efficiency.

The Model

A monopoly platform and a single seller.

The seller can sell on the platform (and pay a fee f) or via an off-platform channel (at no fee).

The seller's type is $\theta \in [0, 1]$ (CDF G), which, when unknown to the platform, is also unknown to the seller.

The Model (2)

Consumers can only find the seller if the seller is promoted by the platform.

Consumers arrive in two periods (period 2 has a scale β) and have two independent types of heterogeneity:

- **Product taste**: probability θ that they find the product appealing, in which case WTP is 1, otherwise it is 0.
- **Channel preference**: Disutility $\delta \in [\underline{\delta}, 1]$ (CDF H, $\underline{\delta} \leq 0$) from purchasing off-platform so WTP there is $1-\delta$ if the product is appealing.

A consumer lives for one period and buys if the best WTP minus the price is positive.

Timing

- The platform commits to an algorithm that determines if the seller is featured in each period.
- In period $t = 1, 2$, the seller sets the on-platform price p_t and the off-platform price $\tilde{p}_t$.
- As per the algorithm, the seller is either featured or not.
- If the seller is featured, consumers observe p_t and $\tilde{p}_t$, decide whether to buy, and choose a channel. Otherwise, the seller has no sales and the platform earns fA.
- At the end of the period t, both observe platform sales s_t. Off-platform sales are not observed by the platform.

Types of Algorithms

Restricted Algorithm: Uses only on-platform information (p_1, p_2, s_1).

OPP Algorithm (Off-Platform Price): Conditions on on-platform information as well as off-platform prices $(\tilde{p}_1, \tilde{p}_2)$.

PPC Algorithm (Price Parity Clause): Excludes sellers whose off-platform price is lower than their on-platform price; in addition, it uses on-platform information.

Period 2 Equilibrium with Restricted Algorithm

G and H are assumed to be well-behaved, e.g. $G \sim U[0,1]$ and $H \sim U(\underline{\delta}, 1)$.

The algo induces on-platform price $p_2^* = 1$ and the seller sets off-platform discount $\Delta_2^* = 1 - \tilde{p}_2^*$ that solves

$$\Delta_2^* = f - \frac{H(\Delta_2^*)}{h(\Delta_2^*)}$$

.

The discount trades off losses from a lower price and savings on fees.

The 'discount' may actually be negative – more expensive to buy off-platform (requires $\underline{\delta} < 0$).

Period 2 Comparison with OPP and PPC

The restricted algorithm cannot affect Δ_2^* because no future remains.

OPP pushes all sales to the platform ($\Delta_2^* \leq \underline{\delta}$)

- Sales channel efficiency requires $\Delta_2^* = 0$, OPP achieves it only when $\underline{\delta} = 0$.

PPC either does not bind or equalizes price across channels ($\Delta_2^* = 0$), thus can only increase channel efficiency.

Period 1 Equilibrium with Known Seller Type

The platform commits to a threshold rule: it will feature the seller in period 2 if s_1 exceeds a threshold T.

If β is large (e.g. $\beta \geq 1$) the platform shuts down off-platform sales ($T^* = \theta$), whereas if β is small, the platform forces Δ_1 down as much as it can, still leaving some off-platform sales.

Simple insight: the restricted algorithm may replicate the OPP/PPC algorithm by using sales data as a mirror of off-platform price.

Period 1 Equilibrium with Unknown Seller Type

A threshold rule is still optimal.

The platform uses first-period on-platform sales $s_1 = \theta(1 - H(\Delta_1))$ to update its belief about the seller's type θ (no uncertainty in equilibrium).

We convert T and A into equivalent equilibrium seller types τ and α.

The platform's updating on θ introduces a 'career concerns' incentive for the seller:

- Downward deviation on Δ_1 increases s_1, raises the probability of surpassing T and being featured in period 2.

Due to career concerns, the seller sets a lower discount than in period 2, $\Delta_1^* < \Delta_2^*$.

Period 1 Equilibrium with Unknown Seller Type (2)

Naive platform would set $\tau = \alpha$ but rational platform may end up with $\tau > \alpha$ in order to reinforce career concerns.

Indeed, the platform's optimal algorithm is inefficient ($\tau^* > \alpha$) unless the outside option is high.

This *promotion inefficiency* is due to the use of imperfect tool to raise off-platform price.

Off-platform sales may be fully choked off, but things are not as bad as with the known seller type.

Algorithmic audit is going to be tricky - what's wrong with the algorithm that promotes sellers with high sales?

Equilibrium Algorithm

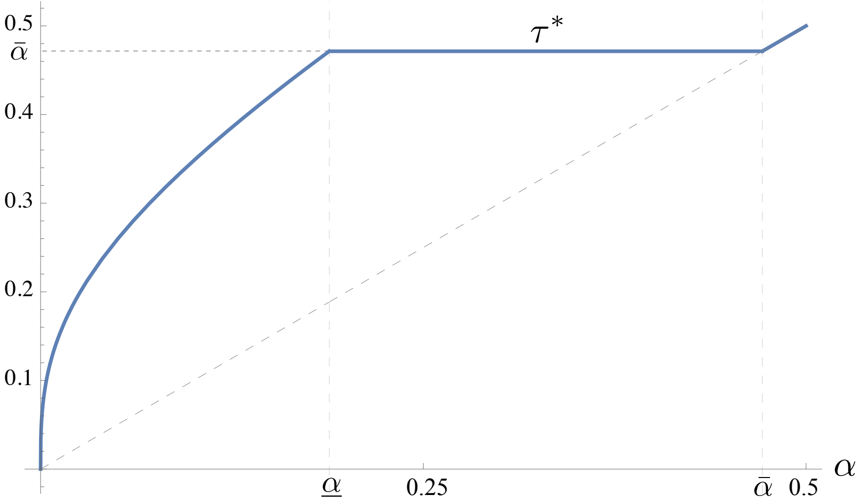

Equilibrium Discounts

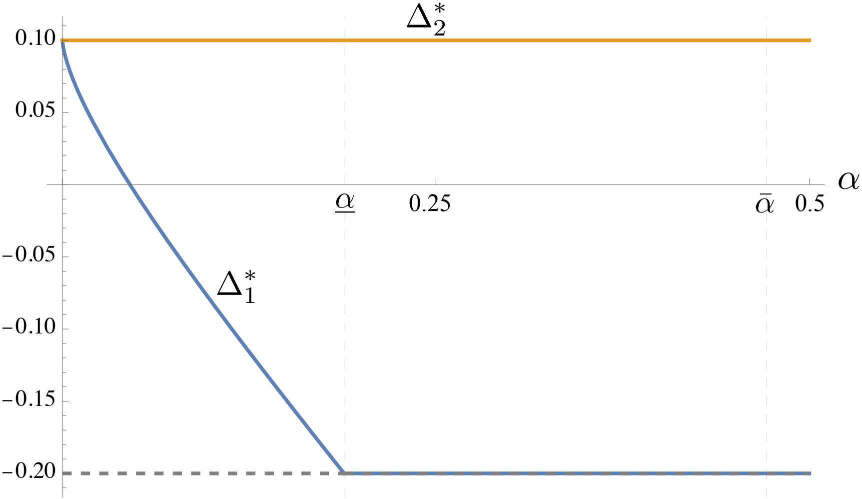

Welfare Analysis

Two key sources of inefficiency:

1. **Sales-channel misallocation**: Consumers may purchase through a lower-value (for them) channel if prices differ.
2. **Promotion inefficiency**: The platform may not feature a seller even when it would generate greater social value than the alternative.

Comparing Restricted and OPP Algorithms

Proposition 1

For $\underline{\delta} = 0$, the OPP algorithm maximizes welfare and therefore increases social welfare relative to the restricted algorithm. For $\underline{\delta} < 0$, the welfare associated with the OPP algorithm may be higher or lower.

If $\underline{\delta} = 0$ the direct channel is inefficient, OPP kills it and simultaneously eliminates promotion inefficiency. Otherwise, there's a tradeoff between promotion and sales channel efficiency.

Welfare with OPP

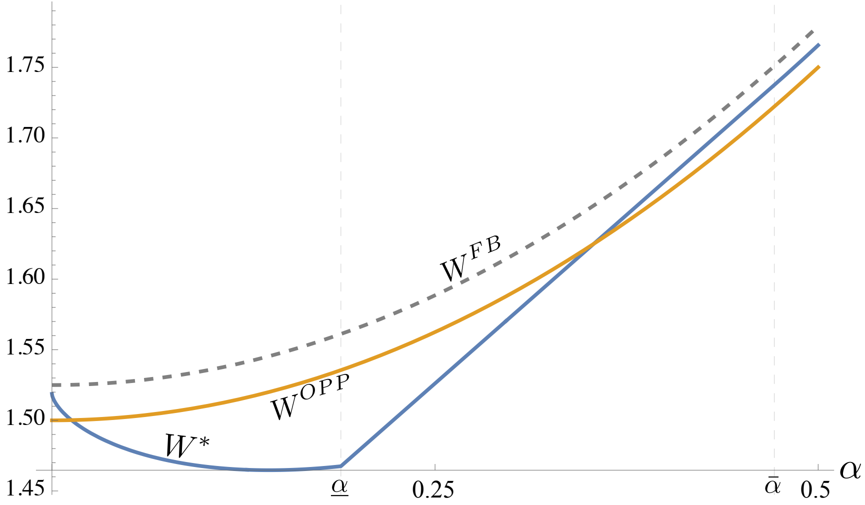

Comparing Restricted and PPC Algorithms

Proposition 2

If PPC binds in both periods, then $\tau^{PPC} = \alpha$ and the PPC algorithm maximizes welfare.

PPC naturally pushes prices together thus eliminating sales-channel inefficiency; if it binds, it also eliminates the promotion inefficiency because it allows the platform to promote efficiently.

Welfare with PPC

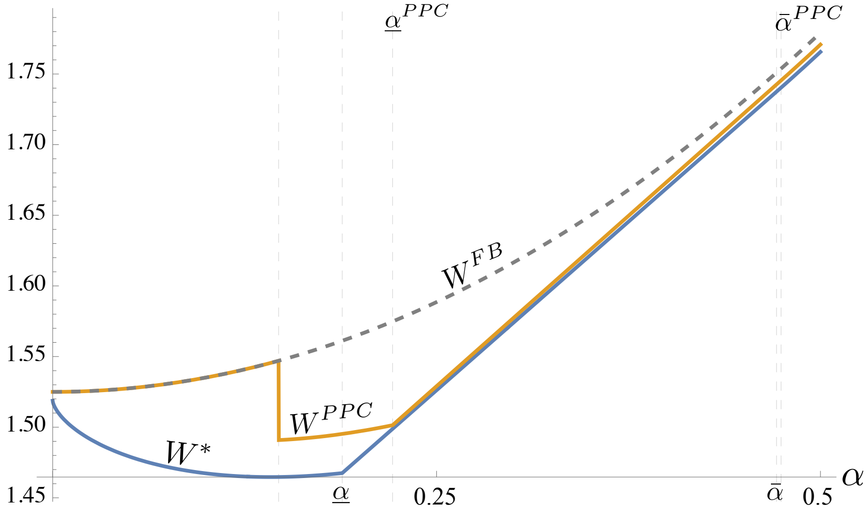

Conclusion

A dynamic model in which the ranking algorithm influences seller behavior across channels.

Even a *neutral/naive* algorithm induces the seller to increase off-platform prices via a novel career concerns effect.

Algorithmic design further reinforces career concerns and can go as far as replicating full PPCs.

Prohibiting platforms from using off-platform prices can reduce welfare via misallocation.

Regulation (e.g. algorithmic audits) may require detailed information on the broader market, suggesting limited value.

Extensions

The paper considers several extensions:

- Alternative fee structures
- Sellers who know their types
- Fees in the off-platform channel
- Repeat purchases
- Downward-sloping demand

Related Literature

Jiang et al. (2011), Choi et al. (2024), Madsen and Vellodi (2024) focus on "Sherlocking" whereby platforms copy successful sellers – sellers want to hide their success.

Related papers in the Holmstrom (1999) tradition: Casas-Arce (2010), Bar-Isaac and Deb (2014), Miklos-Thal and Ullrich (2015).

Akoz et al. (2021) and Drugov and Jeon (2025) study review manipulation.

Hagiu and Wright (2025), Wang and Wright (2020, 2024), Bar-Isaac and Shelegia (2023), Bergemann and Bonatti (2024), Bergemann et al. (2024) among many, study showrooming.

Boik and Corts (2016) and Johnson (2017) study the agency model and PPCs.

Ma et al. (2024) study contractual PPCs on Booking.com.