Consumer Search and Choice Overload

Volker Nocke [1] Patrick Rey [2]

[1]University of Mannheim [2]TSE

CRETE 2022
Tinos, July 13, 2022

Motivation

- *Steering* is prevalent in online (and also present in offline) markets
  - Recommendations / rankings are a necessity because of large # of products
  - Evidence that retailers/platforms often divert search by consumers
  - Steering is increasingly scrutinized by agencies and policy-makers
    - EU: *Google shopping* (2017); US DOJ: *Google* complaint (2020 -)
    - EU: DMA and DSA; US: Congress
- Lack of (informative) recommendations can create "choice overload"
  - Can discourage consumers
  - Experimental evidence
    - Lyengar and Lepper (2000): Jam tasting booth in upscale grocery store
    - Boatwright and Nunes (2001): Experiment run by online grocery

Main Insights

- Issues
  - What drives consumers' decisions to start / keep searching?
  - Best positioning strategy?
  - Is there choice overload?

- Findings
  - Trade-off between *extensive* and *intensive* search margins
  - Divergence of interest
    - *perfect* positioning is best for consumers
    - *garbled* positioning may be more profitable

    ... despite apparently aligned interests
    - no self-preferencing
    - no bias in favor of specific products
    - no price discrimination motive
  - Choice and *garbling* overload

Related Literature

Setting

- Set $\mathcal{I}_N \equiv \{1, ..., N\}$ of available products
  - Each product i characterized by match probability μ_i (popularity)
  - $1 > \mu_1 > \cdots > \mu_N > 0$
- Supply
  - Monopolist chooses
    - ★ which products to offer ($\mathcal{I} \subseteq \mathcal{I}_N$)
    - ★ their prices $(p_i)_{i \in \mathcal{I}}$
    - ★ their positions (more on this later)
  - Constant and identical marginal costs (normalized to zero), no fixed costs
- Demand
  - Unit mass of consumers, each with unit demand
  - Consumers differ in
    - ★ search cost c, $\backsim$ c.d.f. $G(\cdot)$ over $\mathbb{R}_+$
    - ★ match-conditional valuation v, $\backsim$ c.d.f. $F(\cdot)$ over $\mathbb{R}_+$
  - c and v are i.d. across consumers but the *same* across products

Setting

- Positioning: firm assigns (possibly randomly) products to *distinguishable slots*
  - Supermarkets: aisles, shelves, ...
  - Online platforms: recommendations, rankings, ...
- Search: consumers choose which slots to inspect and in which order

  Having observed c, consumer decides whether to inspect a slot
  - If not, leaves the market
  - If yes
    - incurs c, observes the price and whether he has a match; if so, also learns v
    - consumer decides whether to purchase the product
      - if yes, pays the price and leaves the market
      - if not, decides whether to inspect a second slot
    - and so on

Setting

- Timing:

  Stage 1 Firm chooses size n of product line (publicly observed)

  Stage 2 Firm privately chooses $\mathcal{I}$ (s.t. $|\mathcal{I}| = n$), positions and prices; consumers sequentially decide which slots to inspect (if any)

- For this presentation: role of *positioning* (downplay role of prices)
  - Focus on monopoly pricing equilibria (*PBE* with *passive beliefs*)

$$p^m \equiv \arg\max_p p[1 - F(p)]$$

$$\pi^m \equiv p^m[1 - F(p^m)] \text{ (profit conditional on match)}$$

$$s^m \equiv \int_{p^m}^{\infty} [v - p^m] dF(v) \text{ (exp. CS conditional on match)}$$

  - Unique pure-strategy equilibrium features monopoly pricing
    intuition: no profitable price discrimination
  - [Alternative: free services (zero lower bound)]

Perfect Positioning

Consider pure-strategy continuation equilibria, given the assortment size n

Proposition (**perfect positioning**)

For any $n \in \mathcal{I}_N$, there exists a unique pure-strategy continuation equilibrium:

- *the firm offers the most popular products ($\mathcal{I} = \mathcal{I}_n$) at the monopoly price p^m*
- *consumers follow a myopic search pattern*
  - *inspect products by decreasing order of popularity*
  - *those with $c \leq \mu_k s^m$ inspect up to k products until finding a match*

Intuition: the interests of consumers and of the firm are *aligned*

- If the firm positions the more popular products in the first slots,
  ... consumers with search cost $c \leq \mu_k s^m$ inspect only the first k slots
- The firm thus has an incentive to position the popular products in first slots

Perfect Positioning

- Overall game also has a unique pure-strategy equilibrium

  The firm offers *all* available products

- Perfect information: resulting search pattern is efficient

  Conditional on monopoly pricing, consumer surplus is maximal

- Yet, can be profitable to *garble* the information provided to consumers
  - Discourages participation (*extensive* search margin)
  - But encourages active consumers to keep searching (*intensive* search margin)

No Positioning

- Suppose that the firm uniformly randomizes over its positioning strategy

  Let α_k denote the probability of a match on k^{th} inspection, conditional on having had *no match* on previous inspections

Lemma (increasing optimism)

- α_k *is strictly increasing in* k
- *active consumers keep searching until finding a match*

  Intuition: when an inspection does not yield a match, the consumer assigns a higher probability to the inspected product being one of the less popular ones

- It follows that the optimal search rule is *not* myopic

Remark (search addiction)

Some consumers who do not want to start searching may choose to keep searching if coerced to do a first search, the result of which turns out to be unsuccessful

No Positioning

The marginal consumer is $c_n^N \equiv \mathrm{M}_n s^m / \Gamma_n$, where M_n is the probability of a match and Γ_n is the expected number of inspections

Proposition (**no positioning**)

For any $n > 1$, there exist no positioning continuation equilibria:

- *the firm offers the n most popular products*
- *consumers with search cost below c_n^N inspect slots until finding a match*

Intuition

- If consumers keep searching until finding a match, the firm is:
  - indifferent about its positioning strategy
  - willing to randomize uniformly
- In the absence of positioning, consumers
  - are indifferent about their search sequences
  - any (random or deterministic) search sequence can be sustained in equilibrium
  - in particular, consumers can uniformly randomize or follow a given sequence

Choice Overload

Impact of an increase in the assortment size n (if no positioning for any n)?

Corollary (choice overload)

Absent positioning, expanding the product line reduces consumer participation

c_n^N is strictly decreasing in n

- Consumers who are barely willing to search are harmed
- Consumers with sufficiently low search costs however benefit from more choice

No Positioning vs. Perfect Positioning

- Perfect and no positioning are two (polar) forms of *steering*
  - Active consumers all inspect the slots in the same order
    some may however stop earlier than others
  - Perfect positioning: active consumers strictly prefer doing so
    - *extensive search:* maximizes participation
    - some *do* stop earlier than others
  - No positioning: active consumers are indifferent but willing to do so
    - *intensive search:* they all keep searching until finding a match
    - however, fewer start searching

- From now on: focus on steering equilibria
  - Characterize their general structure
  - Characterize worst and best equilibria

Steering Equilibria

Lemma (block structure)

For any n and any steering continuation equilibrium, there exists

$$\mathcal{N} = \{n_1, \ldots, n_{|\mathcal{N}|} = n\} \quad \textit{and} \quad \hat{\mathbf{c}} = (\hat{c}_1, \ldots, \hat{c}_{|\mathcal{N}|})$$

such that $\hat{c}_k > \hat{c}_{k+1}$ *and:*

- *the firm assigns products from* $\mathcal{B}_k^{\mathcal{N}} \equiv \{n_{k-1}+1, \ldots, n_k\}$ *to slots in* $\mathcal{B}_k^{\mathcal{N}}$
  and does so randomly whenever $|\mathcal{B}_k^{\mathcal{N}}| > 1$

$$\underbrace{1 \,,\, \ldots \,,\, n_1}_{\mathcal{B}_1} ,\, \underbrace{n_1+1 \,,\, \ldots \,,\, n_2}_{\mathcal{B}_2} ,\, \ldots ,\, \underbrace{n_{|\mathcal{N}|-1}+1 \,,\, \ldots \,,\, n}_{\mathcal{B}_{|\mathcal{N}|}}$$

- *active consumers inspect the slots in increasing order*
  those with $c \leq \hat{c}_k$ *inspect up to* n_k *slots until finding a match*

- Intuition
  - If some consumers stop early, then any product allocated to a later block must be less popular than any product allocated to an earlier block
  - Within blocks: *intensive search* $\rightarrow$ requires *garbled positioning*
  - Between blocks: some consumers drop out even if they had no match

Intensive Search

- The firm's profit is given by:

$$\Pi\left(\mathcal{N},\hat{\mathbf{c}}\right)\equiv\sum_{k=1}^{|\mathcal{N}|}\Pi_k^I\left(\mathcal{N},\hat{c}_k\right)$$

where, using $\lambda_i \equiv 1-\mu_i$, $\Lambda\left(\mathcal{I}\right)\equiv\prod_{i\in\mathcal{I}}\lambda_i$ and $\mathrm{M}\left(\mathcal{I}\right)\equiv 1-\Lambda\left(\mathcal{I}\right)$:

$$\Pi_k^I\left(\mathcal{N},\hat{c}_k\right)\equiv\Lambda\left(\mathcal{I}_{n_{k-1}}\right)G\left(\hat{c}_k\right)\mathrm{M}\left(\mathcal{B}_k^{\mathcal{N}}\right)\pi^m$$

- Additive separability
  - given $\mathcal{N}$, profit only depends on $\hat{\mathbf{c}}$
  - c_k (*participation*) depends only on (garbled) positioning used in block $\mathcal{B}_k^{\mathcal{N}}$
- We can thus focus on intensive search *within individual blocks*

Worst Intensive Search

Lemma (worst intensive search)

Among the equilibria with intensive search over $\mathcal{B}$ (s.t. $|\mathcal{B}| > 1$), the no positioning equilibrum minimizes firm's profit as well as consumer surplus

Intuition

- If no positioning, consumers are willing to randomize uniformly over their search sequences *and* keep searching until finding a match
- Conversely, if consumers do so, their expected surplus is independent of the firm's positioning strategy
- Hence, consumers can always secure the payoff obtained under no positioning
- Any directed intensive search equilibrium gives them higher expected surplus
- Hence, more consumers participate, which also increases the firm's profit

Best Intensive Search

- For a given block $\mathcal{B}$, what is the best intensive search equilibrium?
- Difficult problem
  - $|\mathcal{B}|!$ possible slot assignments $\rightarrow |\mathcal{B}|! - 1$ positioning probabilities
  - active consumers must keep searching until finding a match
  - they must prefer the targeted search sequence to any other

$\longrightarrow$ no less than $\sum_{k=1}^{|\mathcal{B}|} \frac{|\mathcal{B}|!}{(|\mathcal{B}|-k)!} - 1 \, (>> |\mathcal{B}|! - 1)$ incentive constraints

- In general, the optimal search sequence is moreover hard to characterize

  depending on the positioning strategy, finding no match in one slot may make the consumer more or less optimistic about finding a match at the next slot

Best Intensive Search

Yet, the best intensive search equilibria exhibit a rather simple structure

- Problem
  - Encouraging participation: position popular products in first slots
  - Inducing intensive search: position popular products in last slots
- Solution:
  - Assign popular products with "just" enough probability to induce marginal consumers to keep searching
  - Achieves the same profit as if the block were holding a single product with (geometric) average no-match probability $\bar{\lambda}(\mathcal{B}) \equiv \sqrt[|\mathcal{B}|]{\prod_{i \in \mathcal{B}} \lambda_i}$

Proposition (best intensive search equilibrium)

Among the equilibria with intensive search over $\mathcal{B}$, those maximizing the firm's profit as well as consumer surplus induce consumers with search cost $c \leq \hat{c}(\mathcal{B}) \equiv \left[1 - \bar{\lambda}(\mathcal{B})\right] s^m$ to participate.

Most Profitable Block Structure

- Consider "merging" two successive blocks $\mathcal{B}_1$ and $\mathcal{B}_2$ into $\mathcal{B} \equiv \mathcal{B}_1 \cup \mathcal{B}_2$

$$\underbrace{1, \ldots, n_1}_{\mathcal{B}_1}, \underbrace{n_1 + 1, \ldots, n_2}_{\mathcal{B}_2} \quad \rightarrow \quad \underbrace{1, \ldots, n_1, n_1 + 1, \ldots n_2}_{\mathcal{B}}$$

Corollary (garbling overload)

For any adjacent blocks $\mathcal{B}_1$ and $\mathcal{B}_2$: $\hat{c}(\mathcal{B}_1) > \hat{c}(\mathcal{B})$

- Merger creates a trade-off between extensive and intensive search margins

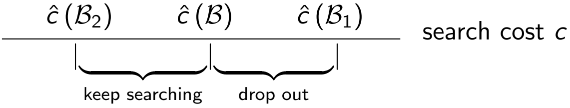

  - cost of discouraging participation

$$[G(\hat{c}(\mathcal{B}_1)) - G(\hat{c}(\mathcal{B}))]\,\mathrm{M}(\mathcal{B}_1)\,(> 0)$$

  - benefit of inducing active consumers to keep searching until finding a match

$$[G(\hat{c}(\mathcal{B})) - G(\hat{c}(\mathcal{B}_2))]\,\Lambda(\mathcal{B}_1)\,\mathrm{M}(\mathcal{B}_2)$$

Most Profitable Block Structure

- Drivers?
  - popularity of the various products
  - distribution of consumers' search costs

- Fix the distribution of search costs

  The following proposition shows that combining two blocks can be desirable only when their products are neither too different nor too similar:

Proposition (directed search: product popularity)

Combining $\mathcal{B}_1$ and $\mathcal{B}_2$ is not profitable if:

(*i*) $\bar{\lambda}(\mathcal{B}_1)$ *is close to* 0 *or* $\bar{\lambda}(\mathcal{B}_2)$ *is close to* 1

(*ii*) $\bar{\lambda}(\mathcal{B}_2)$ *and* $\bar{\lambda}(\mathcal{B}_1)$ *are close to each other*

Sketch of the proof

Most Profitable Block Structure

- Fix the distribution of search costs
- Suppose the search cost distribution has a constant hazard rate $\gamma > 0$

$$G(c) = 1 - \exp(-\gamma c)$$

Proposition (directed search: cost distribution)

$\exists\hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2)$ s.t. combining $\mathcal{B}_1$ and $\mathcal{B}_2$ is profitable if and only if $\gamma > \hat{\gamma}(\mathcal{B}_1, \mathcal{B}_2)$

- Intuition: $\nearrow \gamma \rightarrow \nearrow$ relative importance of intensive over extensive margin

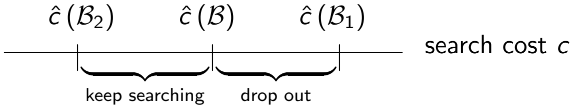

- Proceeding by induction, the most profitable equilibrium involves
  - complete garbling for γ sufficiently large
  - perfect positioning for γ sufficiently small

Extensions

- Disclosure
- Aggregate uncertainty
- Platforms
- From sales to clicks

Conclusion

- Garbling incentives
  - The firm can have an incentive to "garble" its positioning strategy
  - Even with interests apparently aligned with consumers'
    - per sale rather than per click $\rightarrow$ favour more popular products
    - same expected profit per sale
    - so price discrimination motive
    - no self-preferencing
  - More likely if
    - products are neither too popular nor too unpopular
    - products are not too similar either
    - distribution of search cost has high hazard rate

Conclusion

- Choice overload under directed search
  - From *choice overload* to *garbling overload*

    expanding the range of garbled positioning
    - ⋆ discourages participation
    - ⋆ fosters search intensity
  - Choice overload can still be present (always the case under random search)

    consider adding product 3 to block $\mathcal{B} = \{1, 2\}$
    - ⋆ if leads to $\{1, 2, 3\}$: choice overload

      reduces participation

      intensifies search
    - ⋆ if leads instead to $\{1\}, \{2, 3\}$: boosts participation

Related Literature

- Sequential search (Wolinsky *QJE* 1986)
  - multiproduct firms
    - Rhodes (*RES* 2014), Zhou (*AER* 2014), Rhodes et al. (*JPE* 2021)
    - within-firm sequential search
      - Hagiu & Jullien (*RAND* 2011): controlled search, per click
      - Petrikaité (*RAND* 2018): obfuscation
    - here: exogenous search costs, induced search, per sale, no price discrimination
  - platforms
    - positioning auctions: Chen & He (*EJ* 2011), Athey & Ellison (*QJE* 2011)
    - here: positioning by the platform
- Choice overload and contextual inference
  - Kamenica (*AER* 2008)
  - Villas-Boas (*MS* 2009), Kuksov & Villas-Boas (*MS* 2010), Ke, Shen & Villas-Boas (*MS* 2016)
  - here: directed search, garbling overload

Back

Best Intensive Search: Sketch of the Proof

- Let β_k be the probability of no match with the first k slots, for $k \in \mathcal{I}_{|\mathcal{B}|-1}$
  - inspecting up to k slots yields a match with probability $1 - \beta_k$
  - the expected number of inspections is $\hat{\Gamma}(\boldsymbol{\beta}) \equiv 1 + \beta_1 + \cdots + \beta_{k-1}$
  - the value can therefore be expressed as (with the convention $\beta_0 = 1$):

$$V_k(c) = (1 - \beta_k)\, v - \left(\sum_{i=1}^{k} \beta_{i-1} \right) c$$

  - the marginal consumer, $\hat{c}$, is such that $V_{|\mathcal{B}|}(\hat{c}) = 0$
  - he is willing to inspect all slots if, for $k \in \mathcal{I}_{|\mathcal{B}|-1}$, $V_{|\mathcal{B}|}(\hat{c}) \geq V_k(\hat{c})$, or:

$$[\beta_k - \Lambda(\mathcal{B})]\, s^m \geq \left(\sum_{i=k}^{|\mathcal{B}|-1} \beta_i \right) \hat{c} \qquad (IC_k)$$

- Consider the following relaxed problem
  - choose $\boldsymbol{\beta} = \left(\beta_0, \beta_1, \ldots, \beta_{|\mathcal{B}|} \right)$ so as to minimize $\hat{\Gamma}(\boldsymbol{\beta})$
  - subject to $\beta_0 = 1$, $\beta_{|\mathcal{B}|} = \Lambda(\mathcal{B})$, and $(IC_k)_{k \in \mathcal{I}_{|\mathcal{B}|}}$

Best Intensive Search: Sketch of the Proof

Lemma (lower bound on expected number of searches)

The expected number of searches in block $\mathcal{B}$ cannot be lower than

$$\underline{\Gamma}(\mathcal{B}) \equiv \frac{1-\Lambda(\mathcal{B})}{1-\overline{\lambda}(\mathcal{B})}$$

Furthermore, this lower bound can be achieved only if, for $k \in \mathcal{I}_{|\mathcal{B}|}$:

$$\beta_k = \beta_k^l(\mathcal{B}) \equiv \left(\bar{\lambda}(\mathcal{B})\right)^k$$

- Maximizing participation amounts to minimizing expected number of searches
  - requires minimizing no-match probabilities
  - incentive constraints (IC_k) are all binding
- The marginal consumer
  - is kept indifferent between inspecting another slot or not
  - is such that

$$\hat{c} = \frac{1-\beta_{|\mathcal{B}|}^l(\mathcal{B})}{\underline{\Gamma}(\mathcal{B})} s^m = \frac{1-\Lambda(\mathcal{B})}{\underline{\Gamma}(\mathcal{B})} s^m = [1-\overline{\lambda}(\mathcal{B})] s^m = \hat{c}(\mathcal{B})$$

Best Intensive Search: Sketch of the Proof

Lemma (implementation)

If the firm adopts the positioning strategy $\hat{\boldsymbol{\rho}}(\mathcal{B})$ then, for every $k \in \mathcal{I}_{|\mathcal{B}|}$:

- $\beta_k = \left(\bar{\lambda}(\mathcal{B})\right)^k$
- *inspecting the first k slots maximizes the match probability of k inspections*

- It follows that
  - active consumers inspect slots by decreasing order of popularity
  - consumers with $\hat{c}(\mathcal{B})$ are willing to participate

Back

Product Popularity: Sketch of the Proof

- If $\mathcal{B}_2$ offers highly unpopular products (i.e., $\bar{\lambda}(\mathcal{B}_2) \sim 1$):
  - the benefit of inducing active consumers to inspect these products is negligible
  - the cost of reduced participation is instead non-negligible

- Likewise, if $\mathcal{B}_1$ offers highly popular products (i.e., $\bar{\lambda}(\mathcal{B}_1) \sim 1$):
  - almost all active consumers obtain a match in $\mathcal{B}_1$
  - the benefit of inducing the few unlucky ones to keep searching is thus again negligible compared to the cost of reduced participation

Product Popularity: Sketch of the Proof

- Interestingly, combining the blocks is not profitable either
  ... when they offer similar products (i.e., $\bar{\lambda}(\mathcal{B}_1) \sim \bar{\lambda}(\mathcal{B}_2)$)
  - the participation thresholds are also similar: $\hat{c}(\mathcal{B}_1) \sim \hat{c}(\mathcal{B}) \sim \hat{c}(\mathcal{B}_2)$
  - the cost and benefit of combining the two blocks are both small
  - however, the per-consumer cost of reduced participation is given by the probability of having a match in the first block
  - the per-consumer benefit of intensive search is instead given by the probability of having a match only in the second block
  - for blocks of equal size, the latter is smaller and the participation thresholds are not only similar but moreover evenly spaced:
  - the argument extends to unequal sizes Back