Steering Fallible Consumers

Paul Heidhues[1] Mats Köster[2] Botond Kőszegi[2]

[1] Düsseldorf Institute for Competition Economics (DICE)

[2] Central European University (CEU)

Regulating the Digital Economy Workshop, Toulouse

May 13, 2022

Motivation

General recognition that internet companies — like Google — **steer** consumers.

$\longrightarrow$ They direct a consumer's attention to products that they are likely to buy.

Conventional wisdom: in symmetric settings (e.g. no self-preferencing) with fixed prices, this kind of steering always benefits consumers.

Academic research: e.g. Varian (1996), Bergemann and Bonatti (2011), de Corniere (2016), Marotta et al. (2018), Teh and Wright (2021), Hidir and Vellodi (2021).

Recent policy reports: e.g. Furman et al. (2019), ACCC (2019), CMA (2020).

Advertising industry: "giving consumers the most relevant ads creates a win-win situation".

(Implicit) Assumption: consumers are rational, and steering is preference-based.

- But many consumers are **fallible:** they make mistakes in evaluating options.
  Same policy reports argue "behavioral effects" are important in digital markets.
- Arguably, "A/B testing" of how to market a given product based on mistakes.

Guiding website design using A/B tests (Kohavi et al. 2020).

Example: Bing ran A/B test shifting information from an ad's text to its title.

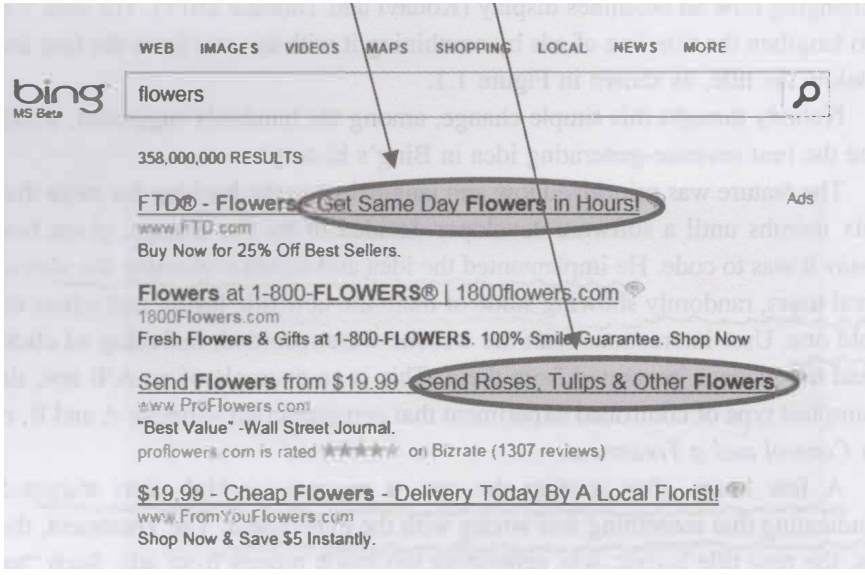

This subtle design change increased annual sales by more than $100 million.

- Presumably irrelevant for valuation (same product and information).
- A/B testing to figure out the "best" framing of a fixed product.

Research question and preview of results

Research question: what is the effect of steering when consumers are fallible?

Long answer. It depends on several factors:

- The **type** of information that is used to steer consumers.
- The **strength** of the technology in identifying attractive products.
- The **reasonability** of the consumer in buying and refraining from buying.

Likely empirically relevant, short answer:

Steering often lowers consumer welfare, in some cases severely so.

Basic setup: the products and the consumer

There are $I \geq 2$ products that can be sold at a fixed price of zero. (Relax later.)

Consumer observes a signal of the consumption value from buying product i:

$$w_i = \underbrace{v_i}_{\text{true value}} + \underbrace{m_i}_{\text{mistake}},$$

where all v_i and m_i are i.i.d. according to "nice" G and F. (Nice: twice differntiable, full support, monotone hazard rate approaching infinity, signals w_i satisfy MLRP)

Consumer sees a single product i and buys if and only if the perceived value

$$\underbrace{\tilde{v}_i}_{\text{perceived value}} = \tilde{v}(w_i) \geq 0,$$

where $\tilde{v}(\cdot)$ is a fixed increasing function with full support. Our fallible consumer

- evaluates products with noise;
- is strategically naive since $\tilde{v}_i$ does not condition on how i is selected; and
- may make additional valuation mistakes encoded in $\tilde{v}(\cdot)$.

Consumer naivete and microfoundations for mistakes

Consumer decides based on face value $\tilde{v}(w_i) = w_i$ instead of true value v_i.

- Random utility (Luce 1959, McFadden 1980): taste shocks or "errors of perception and optimization".

"Dark Patterns". Consumer mis-perceives either the quality or price of a product.

- With hidden prices the consumer believes the price to be $-m_i$ for $m_i > 0$.
- Creating a false sense of urgency, firms inflate quality perception by $m_i > 0$.

Projection bias. Consumer's valuation is colored by the current state of the world.

- Specialized companies help firms to "sync Facebook ads with weather".

Misinterpretation of available information.

- Consumer overweighs recent performance in finance, sellers using "models to persuade" (Schwartzstein & Sunderam 2021) in other contexts.

Basic setup: intermediary and steering

Trade occurs through an intermediary suggesting one product to the consumer.

Consumer buys product i if and only if it is suggested to her and $\tilde{v}_i \geq 0$.

Intermediary aims to maximize the probability of purchase, using an informative signal $s_i \in \mathbb{R}$ about either

- the consumer's true value v_i (preference-based steering);
- the consumer's mistake m_i (mistake-based steering);
- the consumer's perceived value $\tilde{v}_i$ (perceived-value-based steering).

Signal structure satisfies MLRP, is informative, and signals conditional on underlying variable are independent of each other and other variables.

Steering technologies and properties of suggested product

Preference-based steering: possible signal structures.

- Additive signal: observe, for any i, a signal $v_i + \epsilon_i$ with all ϵ_i being i.i.d..
- Perfect signal: $s_i = v_i$.
- Binary signal: observe, for any i, a signal of whether or not $v_i \geq v^c$.
- More generally: MLRP implies higher signal is "good news".

By Bayes' rule, the **distribution of value conditional on i being suggested** is:

$$\frac{g(v|i)}{g(v)} = \frac{\mathbb{P}[\text{product } i \text{ is recom.}|v_i = v]}{\mathbb{P}[\text{product } i \text{ is recom.}]} = I\mathbb{P}[\text{product } i \text{ is recom.}|v_i = v],$$

Since the intermediary suggests the product with the highest signal, i is more likely to be recommended when v_i is higher, and hence

$$g(\cdot|i) \text{ and } g(\cdot) \text{ satisfy MLRP,}$$

while the distribution of the mistake m_i remains unchanged.

Reasonability of the consumer's behavior absent systematic steering

Let $\tilde{v}(\bar{w}) = 0$; then for $v_i = v$ the consumer buys if and only if $m_i \geq \bar{w} - v$.

The consumer buys reasonably if the avg value of purchased products is positive:

$$\int_{-\infty}^{\infty} v \underbrace{\big(1 - F(\bar{w} - v)\big)}_{= \mathbb{P}[\text{buy} \,|\, v_i = v]} dG(v) > 0.$$

And she refrains reasonably if the avg value of rejected products is negative:

$$\int_{-\infty}^{\infty} vF(\bar{w} - v)\ dG(v) < 0.$$

Neither of the two implies the other, and any consumer satisfies at least one.

Within our setup, refraining reasonably and not buying reasonably often likely.

- Avg value of products advertised on the internet is likely to be negative.
- Firms have an incentive to induce mistaken purchases, but not restraint.

Preference-based steering: benchmarks

Benchmark: preference-based steering always benefits an infallible consumer.

Further Benchmark: Allowing for noisy signals

Proposition

A fully rational consumer always benefits from (i) preference-based steering and (ii) perceived value-based steering, as well as (iii) from perfect "mistake-based steering" for sufficiently high I.

- A rational consumer benefits from seeing (i) better products or (ii) products she is more likely to purchase.
- Rough intuition for (iii):
  - Suppose support of mistakes were bounded.
  - With many products, intermediary selects one very close to maximal mistake.
  - Consumer then subtracts this amount and her signal becomes more informative.
  - Our thin-tail assumption approximates a bounded support so that the variance of the signal becomes small as I grows. Hence (after heavily discounting) the signal becomes very informative.

Preference-based steering induces more and, conditional on buying, better purchases

Now consider our fallible and strategically naive consumers.

The net effect depends on the consumer's reasonability and the steering technology

Proposition

I. The following are equivalent conditions on F and G:

(a) Buying reasonably. *(b) Any preference-based steering raises welfare.*

II. Sufficiently strong preference-based steering raises welfare.

Sketch of proof: Second part follows almost immediately from the previous slide.

(a) $\Rightarrow$ (b): Extra purchases due to steering cannot be worse on average. Hence, if the consumer buys reasonably, then extra purchases are reasonable as well.

(b) $\Rightarrow$ (a): Construct a binary technology that lowers welfare of a consumer who doesn't buy reasonably. If cutoff v^c is sufficiently low, steering raises likelihood of buying "near-average" product, and, thus, induces more erroneous purchases. (This "weak stearing" lowers welfare.)

Mistake-based steering hurts consumers who refrain, but don't buy reasonably

Mistake-based steering yields extra purchases of bad & valuable products:

- Extra purchases are less valuable than products purchased absent steering.
- Extra purchases are more valuable than products rejected absent steering.

Proposition

I. The following are equivalent conditions on F and G:

(a) Not buying reasonably. (b) Any mistake-based steering lowers welfare.

II. The following are equivalent conditions on F and G:

(a) Not refraining reasonably. (b) Any mistake-based steering raises welfare.

III. If consumer refrains reasonably, strong mistake-based steering lowers welfare.

Perceived-value-based steering has similar effects as preference-based steering

Ass. imply: The conditional expectation $\mathbb{E}[v_i|\tilde{v}_i]$ is increasing in $\tilde{v}_i$.

Proposition

The following are equivalent conditions on F and G:

(a) Buying reasonably. (b) Any perceived-value-based steering raises welfare.

Steering weakly improves selection of products, and a consumer who buys reasonably benefits thereof.

A binary technology with $v^c < 0$ increases prob of purchase without affecting welfare conditional on purchase. Bad if consumer does not buy reasonably.

Taking a seller's optimal pricing response into account

With many products and (almost) perfect steering likely to choose a high parameter. Focus on this case. Suppose intermediary suggests optimal price.

Proposition (Welfare Effect of Perfect Steering with Many Products)

Fix any G and F.
I. If the selected v_i is sufficiently high, then preference-based steering benefits the consumer.
II. If the selected m_i is sufficiently high, then mistake-based steering hurts the consumer.
III. Suppose that $\mathbb{E}[\tilde{v}_i - v_i|\tilde{v}_i]$ is weakly increasing, and non-negative for sufficiently high $\tilde{v}_i$. If the selected $\tilde{v}_i$ is sufficiently high, then perceived-value-based steering hurts the consumer.

Intuition for I: high margin and fear of lost sales implies CS rent.
II and III: high price to exploit consumer mistake.

Taking stock: what have we learned so far?

The effect of steering depends on its type and strength, as well as reasonability.

Reasonability: for non-initiated search, consumers likely refrain reasonably (and plausibly do not buy reasonably).

Strength: with the cost of information going down, (absent regulatory changes) likely close to optimal (in near future).

	Fixed Prices	**Endogenous Prices**
Preference-based	consumer welfare ↑	consumer welfare ↑
Mistake-based	consumer welfare ↓	consumer welfare ↓
Perceived-value-based	consumer welfare ↑/↓	consumer welfare ↓

Type: a bit harder to figure out, but key for thinking about policy implications.

Data on consumer behavior won't help, but can look at **what intermediaries do**.

What type of steering is actually going on? — A/B testing

Intermediaries "run thousands of A/B tests every year" (Kohavi et al. 2020).

Consumers decide based on perceived values, so what do intermediaries learn?

- What products to suggest to which users, which holding the framing fixed amounts to perceived-value-based steering.
- Which framing yields most sales of a fixed product (mistake-based steering).
- Harder to imagine how A/B testing could lead to preference-based steering.

What type of steering is actually going on? — Machine learning

Not known or understood which patterns the algorithms use to predict behavior.

e.g. Wills and Tatar (2012), Datta et al. (2015), Bashir et al. (2019), and many more.

Still, hard to imagine preference-based steering (as trained based on A/B tests).

But extremely easy to imagine how it could be used for mistake-based steering.

Ex 1: Companies like WeatherAds help firms to condition their ads on weather.

Ex 2: Google steers users searching "high-return investment" to fraudulent ads.

Ex 3: Uber develops an algorithm that predicts likelihood of a user being drunk.

- Likely correlates with mistaken purchases, and might be used for steering.

Ex 4: Spotify uses speech recognition to steer users based on emotional state.

- Offers potential to exploit mistaken purchases as a result of projection bias.

Conclusion and policy discussion

- Thinking about steering in terms of a rational model is often misleading.
- Allowing for mistakes not enough, as steering typically not preference-based.
  - Aside: transparency on steering mechanisms highly desirable.
- With decreasing info costs, A/B test-based steering likely detrimental.
  - One alternative: steering based on stated interests.
  - Especially when self-initiated, such search can be beneficial if (endogenous) price effect not too big.
  - If a consumer engaging in self-initiated search looks for multiple products anyhow, then steering additionally reduces search costs.
  - When only considering a narrow range of products, mistakes more likely to be correlated across these $\rightarrow$ resembling preference-based steering.