Taming Tech Giants' Algorithms: An analysis of consumers' choices on Amazon *

Michela Boldrini† Francesco Clavorà Braulin‡

November 14, 2025

Click here for the most recent version

Abstract

Online marketplaces rely on opaque ranking algorithms that steer consumers' choices. We study how transparency about such algorithms affects consumer behavior on Amazon's marketplace, focusing on the Buy Box—the feature that preselects a default seller for most products. We conduct an online (framed) field experiment with Italian Amazon users, varying the amount of information provided about the Buy Box to study resulting changes in consumer search and deal selection. Informing consumers increases their willingness to inspect alternatives and reduces reliance on the default deal, without increasing search time or effort. At the end of the experiment, most participants hold critical views of the Buy Box mechanism and are willing to pay a fee to opt out of it in favor of a customized selection service. Our results highlight how a limited understanding of platform algorithms reinforces consumer inertia and demonstrate that simple information interventions can improve welfare without altering platform design.

Keywords: Consumer Search; Algorithmic transparency; Information Provision; Field Experiment; E-Commerce. **JEL Codes:** C93, D12, D83, L81.

*The authors gratefully acknowledge financial support from the Diligentia Foundation for Empirical Research (Köln, Germany), the Einaudi Institute for Economics and Finance (Rome, Italy), and Lear (Rome, Italy). We are grateful to Maria Bigoni, Emanuele Tarantino, Michele Fioretti, Rafael Jimènez-Duràn, Katrin Gödker, Muxin Li, Leon Musolff, Irene Bertschek and Dominik Rehse for the discussions, and to the audiences at the University of Bologna, Lear Competition Festival, EAYE Workshop, EARIE Conference, ESA European Meeting, TUM Workshop on Field Experiments in Economics and Business, University of Vienna, University of Insubria, Joint Research Center of the European Commission, and the ZEW Conference on the Economics of ICT for their comments.

†Bocconi University & IGIER. Email: michela.boldrini@unibocconi.it

‡Italian Competition Authority. Email: francesco.clavorabraulin@agcm.it. The views and opinions expressed are strictly those of the author and do not reflect in any way those of the Institution to which he is affiliated. This work was primarily carried out while the author was affiliated with ZEW Mannheim, during a postdoctoral position.

1 Introduction

How does algorithmic transparency affect consumers' purchasing decisions on digital marketplaces? Online platforms such as Amazon increasingly rely on ranking algorithms that determine which sellers and products users see first, thereby shaping consumer behavior and market outcomes (Athey and Ellison, 2011). Yet, little is known about whether consumers understand these algorithmic mechanisms or how greater transparency might influence their reliance on them. This paper's key contribution is to provide causal evidence on how information about algorithmic preselection affects consumers' search and purchasing behavior.

Using a *framed* field experiment with real Italian Amazon users, we test whether increasing transparency about the functioning of Amazon's Buy Box—the algorithmic feature that automatically selects a default seller for most products—alters consumers' search patterns and deal selection choices. The Buy Box represents an ideal setting to study how consumers interact with algorithmic intermediation. When multiple sellers list the same product, Amazon's proprietary algorithm selects one as the default seller, and its offer is displayed most prominently on the product page and automatically preselected for purchase. Consumers cannot customize how the default selection is operated, nor filter or personalize results on offers by other sellers, which are relegated to a secondary webpage. Hence, consumers can essentially either rely on the algorithmic selection or opt out, engaging in an extended search process. This design feature has been estimated to channel the vast majority of sales on Amazon, both in the EU and the US. Regulators have expressed concerns that such opaque mechanisms may reinforce consumer inertia and enable the platform to privilege its own retail business and/or ancillary product services.

The design of our experimental intervention is guided by a theoretical framework that formalizes the consumer's search (and selection) problem in the presence of algorithmic pre-selection. What we model is a very specific moment of consumers' experience on the Amazon marketplace, where they usually face a series of nested decisions: the process starts with the identification of their product of interest, often after a generic query on the website's homepage or a search engine, and ends with the ultimate choice of the deal that specifies purchasing terms. We focus on this final step, in which the consumer has already identified her product of interest and is left with choosing her preferred deal. This choice determines which seller will provide her with the desired product, as well as what price, shipping, and other conditions will be implemented upon purchasing. In particular, we are interested in understanding what drives the consumer's propensity to select the Buy Box deal over other alternatives. Accordingly, we model consumers' purchasing choice process as a series of binary choices, in which consumers compare each of the other deals offered by competing sellers with the Buy Box deal. In this context, information about the Buy-Box deal is free and

immediately available, while acquiring information on the deals offered by other sellers requires some effort, such that opting for the Buy Box deal can save a considerable amount of time. We model the consumer's propensity to rely on the Buy Box as a function of: (1) her individual search cost, and (2) her trust in the platform and in the Buy Box algorithm, which affects her perception of the quality of the Buy Box pre-selection, hence of its time-saving service.

Our framed field experiment involves real Italian Amazon users, who are engaged, at the core of the experimental procedure, in an e-shopping task on the real Amazon.it marketplace. After a few questions about their demographics, e-shopping habits, and experience with the Amazon e-commerce platform, participants are randomly assigned to one of our three information treatments (between-subjects design). All three conditions provide some generic information on the Amazon.it platform, yet each contains a differential amount of information on the Buy Box mechanism: in the Baseline treatment, no information on the Buy Box is provided; in the Partial Information treatment, we inform participants that the Buy Box contains information only on the one specific deal pre-selected by Amazon's algorithm, and about the possibility to inspect alternatives, which are displayed (one-click-away) on the secondary page containing all available listings; in the the Full Information treatment, on top of the latter information, participants also learn that Amazon's Buy Box algorithm is not public, and what are supposed to be the main selection criteria according to the tech giant's self-declarations.

After the information intervention, all participants proceed to the e-shopping task. They access the Amazon.it e-commerce platform through a parallel browser page governed by us, which guides them through a series of pre-selected steps and allows us to track and scrape in real-time some relevant information on what they see and do on the pages they visit. At the core of the task, they land on a product page and are asked to add to their cart what they consider the best deal for the product, given an exogenously set budget. They know that a randomly selected subset of all participants will receive the product in-kind, according to all terms and conditions embedded in the selected deal, together with a cash transfer equal to the amount of the budget (eventually) saved. After the e-shopping task, we debrief all participants on what the Buy Box is and elicit self-reported measures of their literacy and reliance on it before the experiment, together with their current evaluation and willingness to rely on it in the future. Through a follow-up survey, disseminated a few weeks after the main data collection, we gather information about participants' behavior and (updated) beliefs about the Buy Box, together with an incentivized measure of their willingness to pay for its pre-selection service.

We have four key findings. First, without any information intervention, most consumers rely on the default option given by the Buy Box, rarely inspecting alternative

offers, even when cheaper options are available; this result aligns with regulatory findings on the relevance of the Buy Box in other European countries based on sales data, and is coherent with participants' low self-reported literacy of its functioning. Second, providing transparent information about how the algorithm selects the default deal (Full Information treatment) substantially reduces this reliance: consumers become more likely to inspect alternative deals and to opt out of the default option. Thanks to our ability to characterize users' behavior by observing both their search and selection choices, we learn that most of this effect is driven by the induced change in their search behavior: conditional on inspecting all available options, the Boy Box deal is selected only half the times; the probability to select the Buy Box deal is also significantly lower among those who inspect a restricted set of alternatives, who divert their choices towards, on average, cheaper deals. Third, consumers who opt out of the Buy Box deal do not spend more time or effort to finalize their deal selection. Fourth, after participating in the experiment and being debriefed on what the Buy Box is, the majority of participants hold a negative view of it, irrespective of treatment assignment. Most argue that it fails to always give prominence to what they believe to be the best deal available. This evidence is not driven by experimenter demand effects and persists over time. In the follow-up survey, most participants still hold a negative view of the Buy Box and, on average, would be willing to pay a fee of approximately 4% of the product price for a customized deal selection tailored to their preferences, different from the default one operated by Amazon's Buy Box algorithm.

The novel contribution of these results is threefold. First, we quantify the role of Amazon's Buy Box in steering consumers' decisions, relying on firsthand data collected through a rigorous experimental investigation in the field, which allows us to separately identify consumers' search and deal selection choices. This confirms regulatory claims by the European Commission (EC, 2022) and the US Federal Trade Commission (FTC, 2023). It also complements the recent evidence obtained via structural estimation from sales data (Lee and Musolff, 2025), which is conditional on a set of assumptions on consumers' sophistication, and subject to data constraints in the identification of deals actually shown to consumers at each point in time (and their ranking). Second, we test causally how consumers' behavior, and hence the Buy Box's steering potential, is affected by consumers' informational endowment on how such algorithmic preselection is operated. Different from the literature studying default effects (Madrian and Shea, 2001; Johnson and Goldstein, 2003), we do not manipulate consumers' choice setting: platform design characteristics, determining how the default option is selected and shown to consumers, as well as opting-out procedures, are held fixed. By changing only the amount of information consumers have on *how the default option is determined*, we explicitly focus on how (the lack of) such information shapes consumers' evaluation and willingness to rely on algorithmic recommendations,

underscoring the role of information asymmetries between platforms and users, and their welfare consequences. Third, we isolate through which channels the information provision affects consumers' deal selection, distinguishing between direct and indirect effects, mediated by changes in their willingness to engage in a more thoughtful search of alternative non-default options.

The results of this paper highlight that even among experienced Amazon users, there is little awareness about the functioning of the Buy Box, and that a simple information provision intervention can affect consumers' search and selection behavior, with positive welfare consequences. Thanks to the novelty of our experimental design and the richness of the field data collected, we can isolate and quantify the effect of information provision, and identify through which channels information provision interventions operate, advancing our knowledge of consumer behavior in algorithmically shaped digital environments. The relevance of our findings goes beyond purely scientific interest and signals room for non-invasive regulatory interventions to discipline how platforms should provide more transparent information on how their ranking algorithms work.

Related Literature. This paper speaks to several strands of the literature. First, we relate to the broad literature on consumer search from economics and marketing (Honka and Chintagunta, 2017; Santos et al., 2012; Jolivet and Turon, 2019; Abaluck et al., 2020), which investigates the importance of consumers' preferences and information acquisition or processing costs. We build on the findings from this literature to derive the simple theoretical framework that guides the design of our experimental investigation.

Second, we contribute to the literature examining how search results' prominence on platforms can influence consumer choices, welfare, and e-commerce sales, both in the absence and in the presence of personalized rankings or recommendations (Armstrong et al., 2009; Ghose et al., 2014; Armstrong, 2017; Ursu, 2018; Derakhshan et al., 2022; Heidhues et al., 2023; Donnelly et al., 2024; Compiani et al., 2024). Different from these works, we are not interested in quantifying the role and the consequences of the algorithmically-driven mechanisms that influence results' prominence. Instead, we investigate consumers' interaction with and knowledge of such mechanisms, exploring the causal effect of providing more transparent information about their functioning.

Third, we relate to the growing body of work from the industrial organization literature that focuses on the case of Amazon to study price competition and consumer search on online marketplaces (Crawford et al., 2022; Tai Lam, 2023; Gutiérrez, 2021), including research on self-preferencing practices and steering (Zhu and Liu, 2018; Farronato et al., 2023; Chen and Tsai, 2023; Waldfogel, 2024; Li, 2025; Dubé, 2024; Farronato et al., 2025; de Corniere et al., 2025; Hartzell and Haupt, 2025) and on the

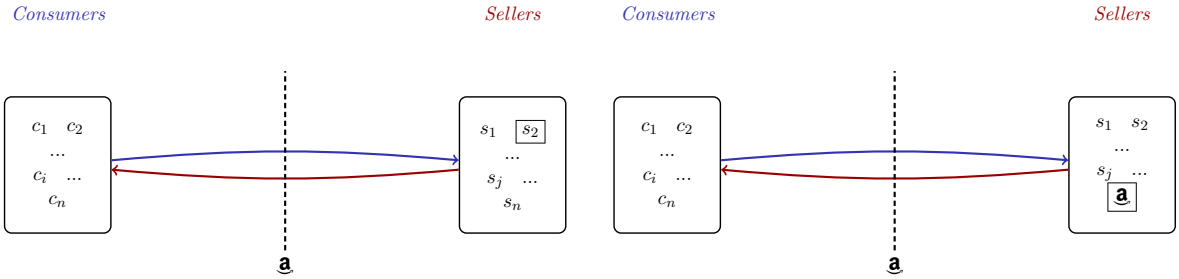
effects of the Buy Box (Raval, 2022; Hunold et al., 2022; Lee and Musolff, 2025). We complement this literature, by studying the role of the Buy Box on Amazon's marketplace, with a new consumer-centered perspective. Most closely related to our work is the paper by Lee and Musolff (2025), where they estimate a structural model using Amazon sales data to capture the effects of Amazon's Buy Box on consumer choices, seller pricing, and entry. On the supply side, their findings highlight that the Buy Box selection is very sensitive to competitive pricing, yet skewed towards Amazon's offers. On the demand side, their model distinguishes between sophisticated consumers, who always compare all available options before finalizing their choice, and naive ones, who always opt for the Buy Box option without even inspecting alternatives: their structural estimates are consistent with 95% of consumers belonging to the latter category, leaving substantial opportunities for demand steering. Different from Lee and Musolff (2025), we rely on field data on consumers' complete set of options, search behavior, and deal choices to study how consumers interact with the Buy Box. Our data allows us to derive consumers' categorization ex-post, conditional on observed behavior, and to quantify the causal effects of providing (more) information on the Buy Box functioning on consumers' choices, exploiting the (induced) exogenous variation in consumers' literacy driven by our treatments.

Lastly, we contribute, also methodologically, to the empirical literature relying on field data collection to study consumers' behavior on e-commerce platforms. Some of these papers implement experimental designs in purely field settings to isolate the causal effect of providing some specific pieces of information on sellers – e.g. quality – (Jin and Kato, 2006; Resnick et al., 2006; Cui et al., 2019) on platforms like eBay or Amazon. Other papers exploit natural experiments or collaborations with platform owners (Dinerstein et al., 2018; Fang et al., 2024) to investigate the effect of platform design changes, shaping consumers' experience, on purchasing choices and market outcomes. Most recently, and similarly to our empirical approach, some papers scraped data on consumers' behavior over time, through browser extension tools that were programmed by the researchers and deliberately downloaded and installed by recruited platform users on their devices (Allcott et al., 2025; Farronato et al., 2025). The closest to our setting is the paper by Farronato et al. (2025): their browser extension manipulates Amazon brands' visibility among search results to study the effect of their removal on consumers' welfare, looking at changes in users' organic browsing patterns and choices in incentivized tasks with some pre-selected products. Different from them, we focus on a distinct feature of Amazon's platform (the Buy Box) and don't manipulate Amazon's platform design across treatment conditions. Instead, we only vary the content of the information provided to users, with the ultimate objective to causally isolate whether informing users about the Buy Box functioning affects their search and deal selection choices.

2 Amazon's Buy Box

Many digital platforms operate as two-sided markets, facilitating the matching between sellers and buyers. The Amazon case is particularly interesting, both from a scientific and policy point of view, because of the dual role the tech giant plays on its platform, where it both controls the rules of the game and participates as a player: Amazon shapes buyers and sellers' interactions, through its platform's design choices, and operates as a seller itself, leveraging its gatekeeper power and the information collected on both sides of the market.

On Amazon, just like on other online platforms, many platform features are governed by ranking algorithms, which determine, for example, whether and how some products and sellers are shown to consumers after a query. In this paper, we focus on a very specific algorithmic-driven feature of Amazon's marketplace: the Buy Box.



(a) "Pure" marketplace. (b) Amazon's dual role.

Figure 1: The marketplace. Panel (a): Amazon matches consumers and sellers, one of which is prominently featured (boxed seller 2 in the example). Panel (b): Amazon is also a seller on its own marketplace, possibly self-selected for the Buy Box.

The Buy Box represents an appealing shortcut for customers. After having searched and identified their product of interest, they can directly proceed to checkout by clicking on the "add-to-cart" button contained in the box located in the upper-right corner of the product page.[1] By clicking, they automatically agree to purchase from the "default" seller placed in the Buy Box. When multiple sellers list the same product, Amazon's proprietary algorithm determines who wins this prominent position, while all other sellers are relegated to a secondary webpage. Consumers cannot filter results, personalize their search, or customize the default selection.

The Buy Box algorithm represents an extreme case of ranking algorithms, where the selection process operated by the algorithm leads to a binary winning-losing outcome. Since the rules behind its functioning are neither observable nor decodable, fears of matching inefficiencies arise should the algorithm fail to always give prominence to the best deals, hence leading consumers to second-order choices. Especially if consumers lack appropriate information on how this selection is operated, which could lead them

[1] See Figure 11 in Appendix C for a concrete visual example.

to underestimate the value of searching for other-than-default deals, by overestimating the algorithm's ability to identify the best deal for them.

The criticality of this design choice was also recognized by the European Commission (2020), which has examined "whether the criteria that Amazon sets to select the winner of the Buy Box lead to preferential treatment of Amazon's retail business and/or ancillary product services".[2]

A few empirical papers, mostly from the computer science literature, have sought to unravel the factors that drive the selection of the Buy Box winner, governed by Amazon's proprietary algorithm (Chen et al., 2016b; Chen, 2017; Duch-Brown and Gomez-Losada, 2019; Gomez-Losada et al., 2022). These studies rely on a reverse engineering approach. They model the probability of being selected as the Buy Box winner as a function of available deals' characteristics, usually obtained through data scraping, and treat it as a prediction problem. Results indicate that sellers relying on sophisticated algorithmic pricing are more likely to secure this prominent position, and that competitive pricing and customer feedback/ratings are the most relevant predictors. However, the combination of these factors cannot fully explain the Buy Box winning probability, substantiating the concerns related to a lack of transparency.

Available evidence suggests that the Buy Box plays a crucial role in buyers' and sellers' interactions, channeling the vast majority of purchases, despite some geographical variation. In Europe, the European Commission (EC, 2022) claims that "based on 2017-2019 sales data provided by Amazon for its German and French websites, sales made directly through the Featured Offer accounted for [70-80] to [80-90]% of the overall sales". In the US, this figure is reported to be even bigger, with the Federal Trade Commission (FTC, 2023) arguing that "nearly 98% of Amazon sales are made through the Buy Box". The use of ranking algorithms to steer consumers towards a specific deal in the presence of homogeneous products is not rare, especially in the context of price comparison websites or "web aggregators".[3] However, the case of Amazon and its Buy Box is undoubtedly the most interesting because of the vast and unparalleled volume of sales.[4]

[2]European Commission. "Antitrust: Commission sends Statement of Objections to Amazon for the use of non-public independent seller data and opens second investigation into its e-commerce business practices" *Press release (Brussels, 10 November 2020)*. On July 7, 2022, Amazon sent a commitment proposal to the European Commission concerning marketplace seller data and fair access to Buy Box and Prime. These commitments do not apply to Italy, where remedies regarding the Italian marketplace were already required by the Italian Competition Authority in 2021, *Press release (Brussels, 14 July 2022,* - see also `https://ec.europa.eu/commission/presscorner/detail/en/ip_22_4522`).

[3]For example, the price comparison websites for accommodation services Trivago adopts a fairly similar platform design choice, where one deal is given special prominence and no full information on all, or at least a consistent number of alternative deals is shown on the same page (see Figures 12-14 in Appendix C for concrete visual examples on Trivago and other aggregator websites).

[4]The Amazon.com marketplace net sales revenues in 2023 totaled nearly 138.943 million U.S. dollars, excluding income from subscription services and provisions for third-party sales and fulfillment (Annual report - ECDB analysis, 2023).

3 Theoretical framework

To study how consumers interact with Amazon's Buy Box and react to information disclosed about the Buy Box functioning, we model consumers' purchasing choice environment in a stylized way based on Jolivet and Turon (2019). In reality, Amazon users make a series of nested decisions before they finalize their purchase on the platform: the process starts from the identification of their product/variety of interest among all alternative options displayed after a product query is launched on the website's homepage, and ends with the final choice over the preferred deal, which determines what seller, price, shipping terms, and other conditions will be implemented after the purchase. In this paper, we deliberately focus only on consumers' final "best deal" selection problem, which comes after they have identified their desired product. Since we are interested in understanding what drives consumers' propensity to select the Buy Box deal, we model the "best deal" selection process of consumers as a binary choice in which consumers compare the other deals offered by competing sellers with the Buy Box deal. In this context, acquiring information on the deals offered by other sellers is costly and opting for the Buy Box deal saves a considerable amount of time.

Model. Consider n sellers, indexed by $i \in \{1, ..., n\}$, listing a homogeneous product on the Amazon marketplace. Each seller posts a price p_i and is endowed with a quality index q_i which summarizes its records in terms of reviews, percentage of positive feedback and stars. We assume a very simple form for the consumers' utility function: when buying the product from seller i at price p_i, consumers get a net utility $u_i = q_i – p_i$, which is the same for all consumers. Since the Buy Box algorithm places a unique seller in a prominent position, the Buy Box seller on the first page is $i = 1$ by convention and the utility consumers get from purchasing the product from the Buy Box seller is $u_1 = q_1 – p_1$.

One-dimensional heterogeneity of consumers. In the simplest form of the model, we assume that when consumers explore the second page showing all other sellers listing the product, they incur in a consumer specific search cost θ (i.e. effort, time). The parameter θ is drawn from the interval $[\underline{\theta}, \bar{\theta}]$ according to a density function $f(\theta)$ with associated cumulative distribution function $F(\theta)$. The mass of consumers is normalized to one without loss of generality. The type θ^*, which is indifferent between purchasing the product from the Buy Box seller and searching for the best possible deal considering

also all the other offers on the second page, satisfies:

$$q_1 - p_1 = -\theta + E[\max\{q_i - p_i\}] \tag{1}$$

where $\max\{q_i – p_i\}$ is a compact way of writing $\max\{q_1 – p_1, ..., q_n – p_n, 0\}$. Solving for θ^* we find that all types above the threshold would purchase from the Buy Box. The market shares of interest depend on the distribution of types θ, so that the share of consumers buying from the Buy Box seller is equal to $\phi_{BB} = 1 - F(\theta^*)$.

$$\theta > E[\max\{q_i - p_i\}] - (q_1 - p_1) = \theta^* \tag{2}$$

Two dimensional heterogeneity of consumers. Allowing consumers to have different beliefs on how the Buy Box algorithm is effective in selecting the best deal for them or, more in general, to have different levels of "trust" in the Buy Box, we introduce a second layer of heterogeneity. Consumers have therefore a two-dimensional type (θ, γ), where γ identifies consumers' trust in the platform and, more precisely, measures the belief that the Buy Box algorithm proposes the best and unbiased purchasing advice for them. The utility consumers get from purchasing from the Buy Box seller now writes $u_1 = \gamma q_1 – p_1$.[5] Solving for θ^*, we get:

$$\theta(\gamma) > E[\max\{q_i - p_i\}] - (\gamma q_1 - p_1) = \theta^* \tag{3}$$

where $\gamma \in [0, +\infty]$ and $\partial\theta(\gamma)/\partial\gamma < 0$. Given qualities and prices, lower trust ($\gamma \leq 1$) is associated with a greater share of consumers opting out of the Buy Box ($\phi_{BB} = 1 - F(\theta(\gamma)^*)$.

4 Experimental Design

Our online *framed field* experiment combines the following main features. First, we collect individual-level data on the search and deal selection behavior of real (Italian) Amazon users. Second, we observe consumers' behavior on the real `Amazon.it` marketplace, without any manipulation to the current design of the platform and its Buy Box feature, hence holding fixed incentives for sellers competing for prominence. Third, this data is collected in a controlled environment, that is, we don't observe consumers' everyday behavior on the e-commerce platform, but within our experimental procedure. This primarily allows us to manipulate the amount of information disclosed to consumers about the functioning of the Buy Box, before they engage in any e-shopping activity, and to control what type of products and budgets participants face across the different treatment conditions. In addition, this also permits us to elicit additional

[5] Whenever $\gamma = 1$ for all consumers the model boils down to the one dimensional case.

self-reported information about consumers' characteristics and habits, and to measure their perceptions. Lastly, we complement this evidence with a follow-up questionnaire designed to assess whether the information provision effects last over time.

The experiment is programmed in oTree (Chen et al., 2016a) and run entirely online, through a series of asynchronous sessions, where participants always follow the same path of actions in the same order, irrespective of treatment assignment.
Figure 2 presents the experimental flow: the experiment is composed of three mandatory parts, followed by a fourth non-compulsory part and lastly a short follow-up questionnaire, distributed a few weeks after the main experimental data collection.

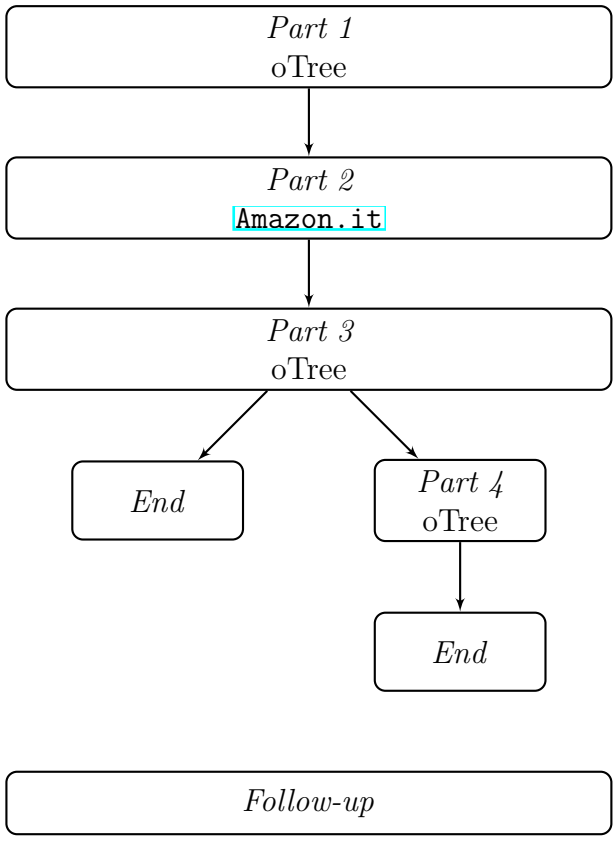


Figure 2: Experimental flow.

In *Part 1*, participants are asked to self-report a few personal and demographic information, and to answer a series of questions concerning their usual online activities (e.g. e-shopping on other platforms/websites, use of social media, etc.), including their habits and experience on the Amazon marketplace.

In *Part 2*, we first elicit prior beliefs about some characteristics of Amazon's marketplace; for participants in other than the Baseline condition, we also measure awareness of how the e-commerce works on the sellers' side.[6] After the elicitation is over, we

[6]Specifically, we investigate how often participants in the Partial [Full] Information treatment believe that best-selling products are listed by multiple sellers [Amazon itself as a seller].

provide treatment-specific information about the topics they have been asked about, plus some additional information on either Amazon's marketplace in general, to participants in the Baseline, or on the features and functioning of the Buy Box, to the others. Lastly, once the information provision is over, participants engage in the core activity of the experiment: the e-shopping task on the real Amazon marketplace. After receiving the instructions, participants are redirected to Amazon's webpage through our in-house web-tracking software, which allows us to keep track of their actions on the platform. Upon launch, the software prompts the opening of a small window containing an interactive recap of the instructions, and a (new) browser page, where participants navigate a pre-determined set of webpages within Amazon's e-commerce platform:[7] after logging in with their personal Amazon credentials, participants land on the product page, and are asked to select what they believe to be the best deal for the product, given an exogenously set budget. The task is over once the preferred deal is added to their personal cart.

In *Part 3*, participants are asked to answer a last set of questions, mostly concerning their search and shopping habits on Amazon's marketplace, and their evaluation and satisfaction with the current platform design and the Buy Box. In addition, they also answer a few questions aimed at verifying their recollection of the information provided before the e-shopping task, to check for inattention and manipulation effectiveness.

Part 4 is non-mandatory: after Part 3 is over, participants can choose to either leave or continue to the last part of the experiment. Before deciding, they are informed about the activity they would engage in and of the possibility of accruing some additional small earnings based on their choices. If they opt in, participants are presented a standard counting-zero task (Abeler et al., 2011) and given up to 10 additional minutes to play, under a small-stakes piece-rate incentives system. They receive 20 Euro cents for each correctly solved table, and after the first three minutes of play, they are again given the choice to leave or continue playing for up to an additional seven minutes.[8]

Follow-up. After a few weeks, the same participants are invited to answer a short follow-up questionnaire and paid a small fixed fee upon completion. The questionnaire asks about their experience on Amazon's marketplace since their participation in the main experiment, and about their (updated) perception of the Buy Box. To

[7] We developed the web-tracking and monitoring tool in partnership with the external company Deltologic. Participants are provided with instructions on how to download and launch the software before taking part in the e-shopping activity and reminded to eliminate it from their devices at the very end of the experiment. The software captures and stores only a limited subset of all the information displayed on participants' parallel browser screens, which is used exclusively for scientific purposes, as clearly stated in the Informed Consent form. No sensitive personal information is retained, log-in credentials are not recorded, and the tracking activity stops as the e-shopping activity is over, when the parallel browser automatically closes. The software only works on the desktop version of the website. Participants can proceed with the experiment only if they access the provided link through a desktop or laptop computer and have downloaded the latest version of Chrome.

[8] These earnings add up to the fixed participation fee and the eventual e-shopping task rewards.

measure information retention and possible interferences, participants are also asked recall questions on the information provided during the main experiment, and whether they individually searched and acquired additional information on the Buy Box.

They are also presented with an incentivized willingness to pay (WTP) task, which allows us to measure how much participants would be willing to pay for a personalized deal selection service, as opposed to the default deal selection provided by Amazon's Buy Box algorithm. First, we elicit participants' preferences on how they would customize the Buy Box algorithm if they could. Second, as we did for the main experiment, we preselect a few retail Electronics products and randomly assign participants to one of them. Third, we show participants their assigned product and inform them that a randomly selected sample will receive the product in-kind as an additional reward for their participation.[9] At this point, we ask them to decide, should they be selected for the reward, if they would prefer us to select for them the Buy Box winning deal or the deal that would best match their interest, to be identified through a customized procedure tailored to their stated preferences. Such customization service comes at a small cost: we rely on an iterative multiple price list elicitation à la Becker-DeGroot-Marschak for different levels of cost; in the first question, this cost is equal to 5% of the product price, and subsequent cost options increase or decrease the cost stepwise, until the WTP is placed in intervals that go from (-∞, 0] to [10, ∞), with 1% increases.

Sample. We recruit Amazon users partnering with the external recruiting company Bilendi. We implement a stratified sampling strategy over age groups, gender and geographical location, to have a higher degree of heterogeneity and collect data from a sample that does not markedly differ from the actual Amazon users' population in Italy over these dimensions.[10] As the main eligibility criterion, we restrict our sample to active users, defined as users having an Amazon account, either Prime or standard, at least three months old.[11]

Information Provision. Our main treatment dimension relates to the amount of information provided about Amazon's platform design features and, specifically, about

[9] Only 1 in every 25 participants in the Follow-Up is randomly selected to receive the assigned product as an additional in-kind reward.

[10] Statistics on Amazon users' population in Italy: Gender distribution = Women (48%), Men (52%); Age brackets = 18-24 (9%), 25-34 (18%), 35-44 (18%), 45-54 (21%), 55+(34%); Geo-location = North (43%), Center (20%), South (37%).
Source: PrimeDay Shoppers (YouGov x Statista 2022):
https://www.statista.com/statistics/1047043/amazon-primeday-shoppers-by-gender-in-italy/
https://www.statista.com/statistics/1047068/amazon-primeday-shoppers-by-gender-in-italy/
https://www.statista.com/statistics/1047073/amazon-prime-day-shoppers-by-macro-area-in-italy/

[11] In order to proceed to the e-shopping task, participants are required to log in on Amazon.it with their credentials and as a first step they land on the page containing their membership details: this allows us to check the existence of their account, and to collect information on whether they have a Standard or Prime account and, eventually, their membership cost.

the Buy Box mechanism. As shown in Table 1, moving from the baseline scenario, in which only unrelated information about Amazon's marketplace numbers and facts is provided, we have two other treatment conditions, where we progressively provide participants with an increasing amount of information about Amazon's platform design features. More in detail, in the *Baseline* treatment (T0) we only provide general information on the characteristics and relevance of the Amazon marketplace in Italy, with no mention of the role or functioning of the Buy Box mechanism. In the *partial information* treatment (T1), in addition to little information on the Amazon marketplace in Italy, we provide some information about Amazon's platform design features and the role and functioning of the Buy Box: Participants find out about the possibility to learn, at the cost of an additional click, about other offers; other offers include those made by all the other sellers listing the product, which may better suit participants' specific needs and tastes, since the Buy Box winner does not necessarily correspond to the deal with the lowest price or with the highest rankings (Chen et al., 2016b). Lastly, in the *full information* treatment (T2), we provide participants with the same information provided in T1, but we additionally mention policy concerns arising from the lack of transparency in the mechanism behind the Buy Box and potential conflicts of interest of Amazon, which acts on its own platform as a competing seller itself.
Participants are evenly randomized across the three treatment conditions, and each participant is exposed to one information condition only (between-subjects design).

T0 **Baseline**	**T1** **Partial information**	**T2** **Full information**
- Relevance of Amazon.it - Facts & figures on Amazon.it - Info on other e-commerce platforms in Italy	- Relevance of Amazon.it - Info on AMZ platform design & Buy Box functioning	- Relevance of Amazon.it - Info on AMZ platform design & Buy Box functioning - Lack of Buy Box transparency

Table 1: Summary of information provision contents across treatments.

Products. To increase control, we preselect which products our participants face during the e-shopping task. Each participant is exposed to one product, and the same products are used across the three treatment conditions. To reduce noise and increase non-monetary incentives, we restrict our empirical exercise to Electronics best-sellers and inform our participants during the recruitment that by taking part in the experiment, they might have the chance to receive a retail electronics product as an additional in-kind reward. We focus on electronics products because they are easily resalable, which makes the products generally appealing to a more heterogeneous pool of individuals, and because they are among the most popular items on the Amazon

marketplace.[12] To select target products, in the week preceding the data collection, we monitor the Amazon Italy best-selling products in the retail Electronics category relying on the Keepa.com website.[13] We acknowledge that the intensity of our information treatments may vary along product-specific dimensions that can affect how consumers value the selection operated by the current Buy Box and the opportunity cost of checking outside options. For this reason, we select our products, ensuring some variation over two main dimensions: price and number of competing sellers listing the product.

We identify products satisfying our heterogeneity conditions after implementing some additional "stability" criteria that help us select products whose relevance on the platform and characterization over the two dimensions of interest is stable over time and, hence, is likely not to be subject to abrupt changes over a time window in which data collection sessions are run. To this aim, we restrict our attention to relatively popular products (with at least 100 reviews), with a number of competing sellers $n \geq 3$ and whose presence in the best-sellers' group is stable over the week.

Incentives. Participants are paid a fixed fee for completing all mandatory parts. Upon invitation, participants are listed a few exemplary retail electronics products and informed that they might be eligible to receive one of such products as an additional in-kind reward for their participation. During the experiment, when they receive the instructions for the e-shopping task, they learn that the exact product they might receive is the one they will be shown during the task. Approximately 4% of participants are randomly selected to receive the product: if selected, participants receive the product according to their personal deal selection,[14] plus a cash transfer equal to the amount of budget saved (if any).[15]

Main outcomes of interest. Our main interest lies in consumers' search and purchase behavior on the platform. To analyze how these two dimensions of behavior are affected by our treatment manipulations, we will rely on the data collected during the incentivized e-shopping task (Part 2), from the moment participants landed on the product page through the parallel browser, until they finalized their preferred deal

[12] During Amazon Prime Days 2017 in Italy, 57% of users purchased at least one item in the retail Electronics category, which includes consumer and other electronics, smartphones and tablets, smart home devices, photo and video cameras, and televisions (Statista). More generally, while Shoes, Clothing and Electronics represent the top three categories of interest for shoppers *searching* for products online in Italy, consumers use digital channels to mainly *purchase* products belonging to the Electronics category. In 2021, 81% of Italian shoppers searched for shoes, while 72% looked for electronics. However, only 36% of consumers ultimately purchased shoes online, whereas 44% of users bought an electronic product (Source: https://www.statista.com/global-consumer-survey). Chen et al. (2016b) confirms that this category is the most popular one, also on the US Amazon marketplace.

[13] For more details on products' selection see Appendix A.

[14] We implement participants' choices based on the indication provided through the best deal selection (seller identity, price, product condition, shipping terms and times, etc.).

[15] In the experiment, participants have no outside option: they have to select their preferred deal among those available and cannot choose to refrain from selecting one to save the whole budget.

selection, which we use as a proxy for their purchase behavior. In addition, we will also analyze: participants' self-reported measures of evaluation of the Buy Box, prior literacy on the functioning of the algorithm, past reliance on its default selection, and willingness to rely on it in the future. These measures are elicited at the end of the e-shopping task (Part 3). For some of these measures, we will also analyze how participants' responses evolve over time, looking at self-reported measures collected through the follow-up survey.

Main hypotheses. Through the experimental data, we can observe consumers' behavior and test the effect of our information treatments, which are designed to manipulate both consumers' perceived search cost and reliance on the selection operated by the Buy Box algorithm.
Based on previous evidence and on how we expect our information provision will affect consumers, we aim to test the following main hypotheses:

H0: In the *"Baseline"* most consumers (at least 50%) go through the Buy Box for their purchases. We expect to replicate the percentage of Buy Box users reported by previous studies: e.g. Chen et al. (2016b) report that more than 80% of sales on Amazon.com go through the Buy Box.

H1: *"Partial Information"* provision in T1 increases consumers' willingness to engage in the search for out-of Buy Box deals, and/or reduces Buy Box oriented purchase choices, compared to Baseline (T1 vs. T0). We expect that the information provided, increasing consumers' knowledge on the wider menu of sellers' options available just one-click-away and on the possibility to find a better deal, would reduce consumers' marginal cost of search.

H2: *"Full Information"* provision in T2 increases consumers' willingness to engage in the search for out-of Buy Box deals, and/or reduces Buy Box oriented purchase choices, compared to Baseline (T2 vs. T0). We expect that the information provided would negatively affect consumers' trust in the platform and, therefore, the perceived utility from the Buy Box deal.

H3: *"Full Information"* effect is stronger than *"Partial Information"* effect. We expect T2 to lead to stronger treatment effects than T1 on the outcomes of interest: T2 vs. T0 $>$ T1 vs. T0.

Heterogeneous Treatment Effects. The main heterogeneity analyses of interest cover the following moderators: (i) product price/durability; (ii) product market's level of competition, measured through the number of active sellers - other than the Buy Box seller - listing the product on Amazon; (iii) participants' individual-specific search costs, proxied by willingness to engage in the low-pay time-consuming non mandatory effort

task; (iv) participants' individual-specific ex-ante satisfaction with Amazon recommendations, proxied by self-reported information on personal experience and e-shopping habits on the platform.

With respect to the first two moderators, we expect treatment effects to be *stronger* for products characterized by: (i) *higher* price/durability ($\uparrow p$), holding fixed the level of competition; assuming consumers' search effort is positively correlated to the intrinsic value of the product, we expect consumers' incentives to inspect other purchasing options beyond the Buy Box to be higher, the higher the price/durability of the product; (ii) *lower* level of competition ($\downarrow n$), holding fixed the price/durability of the product; assuming consumers' search effort is increasing in the number of alternative options, outside of the Buy Box, to be evaluated, we expect consumers' opportunity cost of engaging in search to be lower, as the task becomes faster and easier, requiring a lower cognitive burden.

With respect to the last two moderators, we expect treatment effects to be *stronger*: (iii) the *lower* individual search costs ($\downarrow sc$), holding constant other product-specific factors and trust; assuming consumers' search effort is negatively correlated to individual search cost, we expect consumers' incentives to engage in the search for other, possibly better-fitting, purchasing options beyond the Buy Box to be higher when their cost of time is lower; (iv) the *less extreme* individual ex-ante reliance on Amazon ($\downarrow t$), holding fixed the other product-specific factors and individual search cost; assuming consumers' responsiveness to our information provision is negatively correlated to the distance between individuals' ex-ante trust priors and 'induced' trust levels (after information-led updating) we expect consumers' incentives to engage in search and/or to opt-out of the Buy Box offer to be lower when they are polarized on extreme ex-ante levels of trust, especially if extremely positive (which would require a notable updating effort).

Pre-registration and implementation. The study was approved by the Bocconi University Ethics Committee on May 9, 2023 (SA000607). We pre-registered on the Open Science Framework Registry the main experiment on May 10, 2023 (doi:x2tzu), and the follow-up questionnaire on August 17, 2023 (doi:3me7v). The main data collection was administered over different waves, which took place between May 11, 2023 and December 29, 2023. Participants were equally divided across the three information treatments. The final sample counts observations from N=363 Amazon users.[16]

[16]See Appendix B for details on the data cleaning procedures. We pre-registered to collect observations on the behavior of N = 750 Amazon users, equally divided across information treatments. We further planned to implement, within each information condition, equal exposure to different types of products, to ensure we had a balanced number of products with low/high price and low/high number of total competing deals across treatment conditions. Due to unforeseeable technical issues, causing our scraping routine to fail over some steps of the e-shopping task and challenging a safe retrieval of

5 Results

Recruited Amazon users. The socio-demographic characteristics of recruited participants broadly align with those of the target population: almost equally balanced over gender, a median age of 40 with central age groups (25-54) being the most largely represented, and most participants residing in the Northern and Southern regions.[17] The majority of participants are employed at the time of the data collection, and have at least a high-school degree.

All recruited participants are active Amazon users, with some experience with the e-commerce platform, and electronic products' purchasing in particular. They all report having had their Amazon account for at least one year, roughly 2/3 are Amazon Prime users, and most of those have had their Prime membership active for years. In terms of frequency of use, roughly the entire sample reports having used Amazon's e-commerce services at least once over the previous three months, with more than half of the sample declaring an average of more than 2 orders per month. In terms of most frequently bought product categories, the modal answer is by large that of Electronics, followed first by the broad category of home, gardening and pet products, and second by music, books, movies and videogames. Once they identified their product of interest on Amazon, roughly 1 out of 5 participants immediately checks out the product: the majority adds the product to the their cart (54%) or Wishlist (16%) to finalize the purchase later; among non instantaneous buyers, most wait at least a few days before they eventually purchase the product of interest, but less than a week.

Zooming out of Amazon, the vast majority of participants also relies on other platforms and channels when searching and/or purchasing products online, with eBay and Subito.it (one of the main Italian platforms for second-hand products) emerging as the second most popular. When searching for products online, most (84%) cross-check different websites and platforms before finalizing the purchase and tend to invest more time searching when the product is expensive, they believe a more accurate search could get them a better deal, and they are buying the product for the very first time; when asked to rate what deal characteristics they weigh the most during the search, the vast majority ranks the deal price first, followed by previous users' reviews.

---

the information, we had to perform data cleaning operations resulting in a drop of the final count of observations from N=776 users to N=363.

For consistency, we show results on all outcomes of interest on the same sample (N=363), including all users for which we can safely retrieve information on Buy Box winning deals and all their actions over the different steps of the e-shopping procedure. As robustness, in Sub-section D.1 of the the Appendix, we replicate the analyses concerning participants' search behavior also on a larger sample, including all observations for which at least data on search decisions and Buy Box winning deal(s) is available (504 participants). Results are qualitatively unchanged.

[17] Gender distribution: Male: 47%, Female: 53%; Age brackets = 18-24 (6%), 25-34 (27%), 35-44 (29%), 45-54 (24%), 55+(14%); Geo-location = North (55%), Center (11%), South (34%).

Satisfaction with Amazon. We measure participants' satisfaction with Amazon services before they are exposed to the information provision and e-shopping task. We assess satisfaction on a scale from 1 "Not at all" to 10 "Completely" with different aspects and services provided by the e-commerce platform. On average, they all report high evaluations for products' quality (8.5), Amazon's services and customer care (8.9), and the amount and quality of information provided on the deal before a purchase is finalized (8.3) — including price, shipping, and return terms. To build a proxy for participants' satisfaction with Amazon's algorithmic recommendations while avoiding to explicitly mention the Buy Box to avoid priming, we measure participants' satisfaction with two other forms of algorithmic recommendation: badges signaling specific features when in the stage of scouting which specific product or variety to purchase after a generic query, and frequently bought together advices shown when browsing the product page of a related product, at the bottom. While satisfaction with the first type of recommendations is still quite high (8), participants appear less satisfied with Amazon's recommendations on frequently bought together products, reporting on average lower and more heterogeneous evaluations (6.7).

Information provision: Attention and Recall. Before being exposed to the information provision, respondents are asked to report their prior beliefs on some characteristics of Amazon's marketplace in Italy. After the e-shopping task is over, we measure participants' recall of the information provided: for all recall questions, we measure a positive shift towards correct answers compared to the answers provided to the same questions during the priors' elicitation phase, suggesting participants actively read and paid attention to the information provided.[18] We further check participants' attention immediately after the information provision, testing — for participants exposed to the Partial and Full Information treatment — their retention and comprehension of the information provided about the functioning of the Buy Box. In both treatments, the modal answer is the correct one, with a higher share of correct answers in the Full Information treatment.[19]

[18]The percentage of correct answers to the first recall (multiple-choice question) on the amount of Amazon.it revenues from the e-commerce, which is common across all treatment conditions, is 88%, compared to a correct answers' rate of 35% when the same question was asked before the information intervention. Similarly, the share of correct answers to the additional multiple-choice recall question asked in the Baseline — concerning the percentage of Italians engaging in online shopping activities — is equal to 71%, compared to an initial correct answers' rate of 45%. In the Partial and Full information treatments, the additional recall question we ask is a numerical-input question concerning, respectively, the relevance of multiple sellers and Amazon itself among the Electronics best-selling products' listings: in both cases, the mean and median answers get closer to the true values in the recall questions, compared to initial answers. In general, participants tend to overestimate both the share of products listed by multiple sellers, and for which Amazon operates itself active as a seller.

[19]Note that when answering this (unanticipated) attention check question, participants didn't have the chance to go back to the information provision sheet.

Search Environment. The characteristics of the search environment vary with the products participants are exposed to, which are equally randomized across treatment conditions.

Since our focus is on how consumers interact with the Buy Box, we only select products displaying the Buy Box feature. In some cases, more than one deal is shown inside the Buy Box, with the additional deal being displayed in a smaller shaded area at the bottom of the Buy Box, just below the main winning deal.[20] We denote the main deal shown in the Buy Box — the Buy Box winner — as BB1, and the second-most prominent deal shown at the bottom of the Buy Box as BB2. Given our selection of products, roughly half the participants saw two deals in the Buy Box while engaging in the e-shopping task, as shown in Table 2, and in almost all cases the second deal (BB2) corresponded to a used deal.[21]

In addition to the Buy Box, most product pages shown to our participants also contained an additional feature: scrolling down on the right-end side of the page, a preview of a small selection of up to three alternative deals was provided.[22] This preview is contained in a box denoted "Compare Offers on Amazon" (COA), positioned slightly below the Buy Box. This feature was available, in our sample, in roughly 82% of the cases and the number of deals shown in the preview varied between 1 (4%), 2 (24%) and 3 (54%).[23]

All participants, during the e-shopping task, are given a budget tailored to the price of the product they had been assigned. Budgets range between 25 and 165 Euros, with an average (median) budget of 68 (60) Euros. We define a product as expensive if its associated budget is above the median budget. In our sample, opting out of the Buy Box winning deal could lead to cheaper deals in the vast majority of cases (83%), with possible budget savings of around 6.6 Euros on average. If we restrict our attention to New deals only, the chances of encountering cheaper deals by opting out of the Buy Box are still substantial but decrease to approximately 36%, with an average amount of budget savings around 1.3 Euros.[24]

The number of competing deals listed for each product ranged between 2 and 35, with

[20] See Figure 15 in Appendix C for a concrete visual example.

[21] Note that the main Buy Box winner (BB1) always corresponds to a new deal in our sample, which is also generally the case for all products for which at least one new deal is listed on the platform.

[22] See Figure 16 in Appendix C for a concrete visual example of how the product page looks like when this additional feature is present.

[23] As a result of the randomization across treatment conditions, the distribution of products, in terms of number of deals shown both in the Buy Box and COA preview section, is balanced.

[24] The distribution product budgets and cases in which savings are possible, considering all available deals and restricting to new deals only, are overall balanced across treatment conditions; the only exception is a significantly lower (at 1% level) probability to get out-BB1 savings in the Full Information treatment compared to the Baseline, when considering the full set of available deals (both new and used condition).

Table 2: Descriptive Statistics: Search Environment

	Overall	**Baseline**	**Partial Information**	**Full Information**
BB1 only	172	53	57	62
BB1 & BB2	191	51	68	72
— *BB2 is used deal*	96%	94%	100%	93%
n. COA deals = 0	64	15	24	25
n. COA deals = 1	15	4	4	7
n. COA deals = 2	88	28	28	32
n. COA deals = 3	196	57	69	70
Budget* (€)	68 (60)	72 (60)	62 (45)	70 (60)
	[25; 165]	[25; 165]	[30; 165]	[25; 165]
Out-BB1 savings (% Yes; €)	83%, 6.6	88%, 6.5	87%, 6.8	77%, 6.5
— *Subset of new deals only*	36%, 1.3	41%, 1.5	34%, 1.3	34%, 1.2
Number of deals*	16 (15)	17 (20)	16 (15)	16 (12)
	[2; 35]	[2; 35]	[2; 35]	[2; 35]
Used deals available (% YES)	76%	77%	80%	72%
— *Share of used deals over total*	24%	25%	26%	21%
Observations	363	104	125	134

Notes. * Mean values; median values in parentheses; minimum and maximum values in squared brackets. † for each product, only all possible deals in the budget are considered.

an average (median) number of total deals per product equal to 16 (15). We consider a product to have a high number of competing deals if the overall number of available deals is above the median. In most of the cases (76%), used deals are also available and they sum up to roughly 1/4 of all listed deals for the product.[25]

5.1 Literacy of the Buy Box.

After the e-shopping task, we ask participants to self-report their literacy on the functioning of the Buy Box, and their reliance on it for past purchases. To ensure all participants have the same shared understanding of what the Buy Box is, irrespective of treatment assignment, we debrief all with some neutral information at the end of the e-shopping task. The content of this information is equivalent to what participants in the Partial Information treatment are exposed to before the task. All questions explicitly ask participants to report their experience *prior* to their participation in our study, and no significant treatment differences are detected.

[25]The distribution of cases in which used deals are available, and the share of used deals over the total, is balanced across treatment conditions.

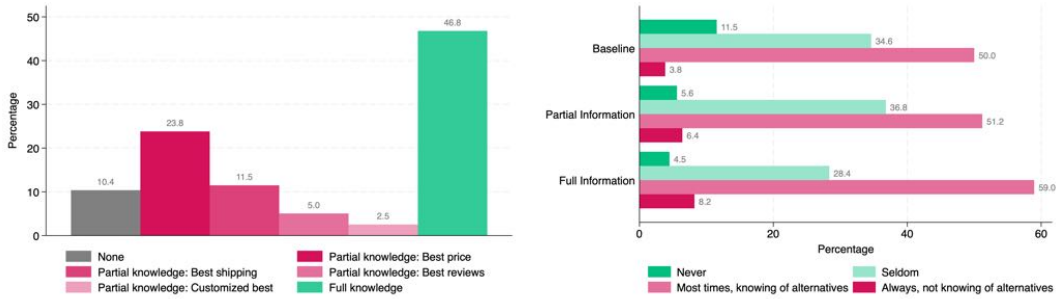
Table 3: Descriptive Statistics: Buy Box Literacy and Prior Use

	Overall	**Baseline**	**Partial Info**	**Full Info**	Test
No Literacy	10%	9%	13%	9%	0.39
Partial Literacy	44%	47%	42%	43%	0.53
Full Literacy	46%	44%	46%	48%	0.23
Observations	363	104	125	134	

Notes. Test values refer to Chi-square statistics from a non-parametric Kruskal-Wallis test, whose null hypothesis is that samples from different treatments are from the same population.

As shown in Table 3, only 10% of the participants report having no prior knowledge of the Buy Box functioning (algorithmic selection criteria) and the possibility to inspect alternative deals before finalizing their purchase. Almost half of the sample claims full prior knowledge, and the remaining, almost equally sizeable, part reports somewhat partial knowledge. Among those, the majority believed the default deal highlighted by the platform as the Buy Box winner always corresponded to the one offering the best price (24%) or shipping terms (12%), as detailed in Figure 3 (Left panel).[26] In terms of past reliance on the Buy Box, overall more than half of participants declare they have predominantly relied on the Buy Box for their past purchases on Amazon, and in most cases knowing of alternatives (see Figure 3, Right panel).[27]

Figure 3: Buy Box Literacy and Prior Use



Notes. Pooled sample of participants across all treatment conditions.

5.2 Search behavior

To analyze search behavior and evaluate the impact of information provision treatments, we rely on different metrics about participants' search patterns collected during

[26] Looking at the distribution of answers across treatment conditions no major significant differences arise.

[27] Looking at the distribution of answers across treatment conditions no major significant differences arise, except for a higher share of participants answering 'Never' in the Baseline treatment compared to the Partial and Full information treatment — significant at the 10% and 5% level, respectively.

the e-shopping task. Specifically, we primarily look at the choice to inspect deals outside the main product page by accessing the separate window showing the complete list of other deals available. In addition, to proxy search effort, we analyze how much time our participants spent finalizing their best deal selection and the number of scrolls.

Table 4: Search behavior across treatment conditions.

	Probability to inspect the secondary page		Time spent to finalize		N. of scrolls on product page	
Partial Information	0.037	0.045	-2.485	-6.603	-0.113	-0.060
	(0.045)	(0.043)	(18.993)	(19.645)	(0.720)	(0.707)
Full Information	0.183***	0.170***	39.347	38.716	1.037	0.999
	(0.051)	(0.049)	(29.969)	(28.044)	(0.788)	(0.752)
Constant	0.115***	0.125	132.317***	61.072	5.769***	4.682
	(0.031)	(0.199)	(13.852)	(116.403)	(0.502)	(3.042)
Observations	363	363	363	363	363	363
R2	0.041	0.154	0.008	0.055	0.008	0.051
Controls	No	Yes	No	Yes	No	Yes

Notes. One observation is one participant, we include participants assigned to all treatment conditions, pooled. (1-2) Linear Probability Model estimates; the dependent variable is a binary indicator equal to 1 when the participant decides to inspect the secondary page listing alternative deals other than the subset of those shown in the main product page. (3-6): OLS estimates; the dependent variables are continuous measures of (3-4) the amount of time spent, in seconds, and (5-6) the number of scrolls performed on the product page, before reaching the cart stage. Controls include: product characteristics (indicators for expensive products and products with a high number of deals available, and the total number of deals shown on the main product page); Buy Box deal characteristics (indicators for Amazon's presence, absence of instant shipping, fulfillment by Amazon, and a continuous measure of percentage price difference wrt the cheapest deal available); individual characteristics (indicators for Prime status, fast purchasing style, low satisfaction with Amazon recommendations, high Buy Box literacy, low cost of time, plus standard socio-demographic traits — including gender, income, age, education and employment status). Robust standard errors in parentheses. Significance levels: $^{*}p < 0.10$,$^{**}p < 0.05$,$^{***}p < 0.01$.

Baseline. Looking at incentivized choice data from Baseline treatment, in which no information about the Buy Box functioning is provided, we observe that almost 90% of participants do not inspect the secondary page with additional deals. On average, participants spend slightly more than two minutes to complete the core part of the e-shopping task, scrolling over the product page approximately 6 times.

Information provision causal effect. Comparing incentivized choice data from all three treatment conditions, we observe that, overall, the share of participants who choose to extend their search and inspect the secondary page with the other deals is

only slightly (and not significantly) higher than in the Baseline under the Partial Information treatment. Instead, when Full Information is provided, this share increases substantially to 30% (see Table 11).

Following our pre-registered hypotheses, we further test whether the effect of our information provision is heterogeneous over four main dimensions: (i) the product price; (ii) the number of competing deals available for the same product; (iii) individual search cost, and (iv) ex-ante satisfaction with Amazon's recommendations. The former two dimensions relate to product characteristics, which we have exogenously varied, ensuring the same products are shown across the three treatment conditions. As predicted, the treatment effect is stronger for expensive products, and products for which a lower number of competing deals is displayed (see Figure 4, Panel A).[28] The other two dimensions relate to individual-specific traits. To measure individual search costs, we rely on participants' incentivized choices in the counting zero task presented in the last part of the experiment.[29] Participants' cost of time is proxied by their willingness to engage in (and stick to) this tedious and low-paid task: we consider them to have a 'Low cost of time' if they decide to engage in the counting-zero task for more than three minutes. We gauge participants' ex-ante position concerning Amazon's recommendations through their self-reported answers, elicited before the e-shopping task and the information provision: we consider them to have a low appreciation of Amazon's recommendations ('Low AMZrec liking') if their rating of Amazon's recommendations through badges and Frequently Bought Together advices is below the median (equal to 7.5, on a 1-10 scale). As predicted, the treatment effect is stronger for participants with a low cost of time, and low ex-ante satisfaction with Amazon recommendations (see Figure 4, Panel B).

The other metrics we rely on to evaluate search behavior reveal no further stark, nor economically relevant, treatment differences: similarly to the Baseline, on average, the time spent on the product page ranges between 2-3 minutes, and the number of scrolls between 5 and 6, in all treatments, as shown in Table 11. Interestingly, this suggests that the change in participants' propensity to inspect a wider set of deals, triggered by our information intervention, does not come at the cost of more time and effort.

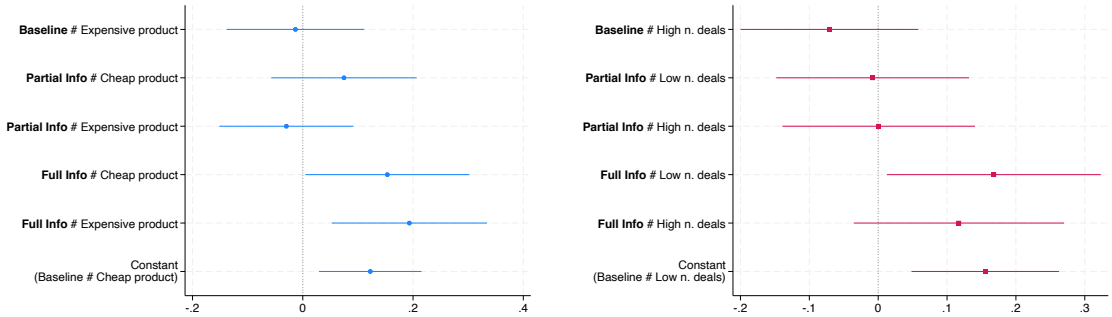
All results on treatment effects are qualitatively unchanged if we control for product and individual characteristics.

Search decision. Introducing product features and individual characteristics as controls in the analysis allows us to draw some exploratory conclusions on their role in

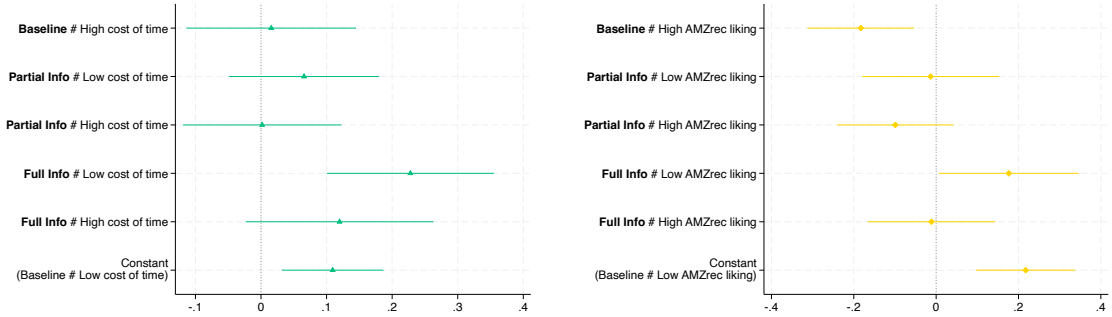
[28] We define products as 'expensive' if their price is above the median price (60 Euros) and with a 'high number of deals' if the number of available deals for the product is above the median (15). See Table 2 for more details on the distribution.

[29] See Section 4 for more details on the design of the task.

Figure 4: Heterogeneous treatment effects on the propensity to inspect the secondary page.



Panel A. Heterogeneity by product characteristics.



Panel B. Heterogeneity by individual characteristics.

Notes. Linear Probability Model estimated coefficients (and constant term). The dependent variable is a binary indicator equal to 1 when the participant decides to inspect the secondary page listing alternative deals not visible on the main product page.

participants' search decisions. Results reported in Table 12 show that, in this peculiar search context, where Amazon users land once they have already identified what they are interested in purchasing, product characteristics don't exert any significant effect. On the contrary, individual characteristics do play a significant role.[30] The decision to extend the search to inspect the other deals not shown on the main product page is strongly and positively associated with (ex-ante) low satisfaction with Amazon's recommendations, and high literacy on the Buy Box functioning. It also emerges that the probability of searching for other deals decreases with participants' age and active employment status. No other significant patterns emerge, except for a reasonable negative association between the amount of time needed to finalize a purchase and participants' age and self-reported tendency to immediately proceed to check out for products of interest.

To analyze search behavior and evaluate the impact of information provision treatments, we rely on different metrics about participants' search patterns collected during

[30] As robustness, in Sub-section D.1 of the Appendix, we replicate this analysis on the extended sample of 504 participants, including all observations for which at least data on search decisions and Buy Box winning deal(s) is available. Results are qualitatively unchanged.

the e-shopping task. Specifically, we primarily look at the choice to inspect deals outside the main product page by accessing the separate window showing the complete list of other deals available. In addition, to proxy search effort, we analyze how much time our participants spent finalizing their best deal selection and the number of scrolls.

5.3 Deal selection

To analyze deal selection behavior and evaluate if information provision interventions ultimately affect participants' welfare by altering their intentions on which deals to purchase, we rely on data describing participants' selection decisions from different angles.

Our primary outcome of interest is how frequently Buy Box winning deals are selected, across different treatments. Given the features of the search environment our participants have been exposed to, we can distinguish three main cases, that is whether participants opted for: the Buy Box winning deal (BB1), one of the deals shown on the main product page but different from BB1 — either the second-most prominent deal shown at the bottom of the Buy Box (BB2) or one of the deals highlighted in the "Compare Offers on Amazon" (COA) preview, when these features are available — or one of the other deals that can be inspected only by accessing the secondary page containing all the listings (OTHER).

Combining this information with participants' search decisions gives us a more detailed picture of their behavior. On the one hand, we can measure how often participants choose the Buy Box winner after inspecting the full list of deals compared to when they only consider the limited set visible on the main product page. On the other hand, we can observe which alternative deals participants turn to when they opt out of the Buy Box winner, conditional on whether they restricted their search to the main page or extended it to the full list. By characterizing participants' course of action on the platform, we can distinguish the direct effects of information provision on deal selection from the indirect ones, which operate through an increased propensity to extend the search to deals outside the main product page.

Baseline. Table 5 reveals that in the absence of any information provision, more than 90% of the participants opt for the Buy Box winning deal (BB1), and most of those who opt out of this option pick one of the alternative deals listed in the secondary page (OTHER).

Conditioning on whether participants inspected the secondary page or not, we observe that when they don't, almost the totality of participants selects the Buy Box winner (almost 98%, see Figure 5)with other deals on the main page playing a minor and insignificant role. When participants choose, instead, to inspect the secondary page, the probability BB1 is selected drops to approximately 50%, while the probabil-

ity that one of the deals shown in the secondary page is selected significantly increases (41.7%).

Information provision causal effect. Comparing the incentivized choice data from all three treatment conditions, we observe that - overall - the share of participants who choose to select the BB1 deal decreases significantly, with respect to the Baseline, only under the Full Information treatment. At the same time, the selection of deals displayed on the main product page other than the Buy Box winner increases; this leads to both a significant and sizeable increase in the probability that one of the COA deals is selected, and a significant yet smaller increase in the probability that the BB2 deal is selected (see Table 5).

The analysis of participants' search behavior (see Subsection 5.2) has highlighted that our Full Information treatment has a strong first-order effect on participants' propensity to inspect the secondary page listing other deals available. It becomes then crucial to understand whether the information provision intervention affects participants' deal selection behavior solely through this channel, that is by increasing the width of participants' consideration set, or not. In principle, other than this *indirect* effect on selection mediated by a change in extended search propensity, the information provision could also alter the way participants select deals once they have determined the size of their consideration set.

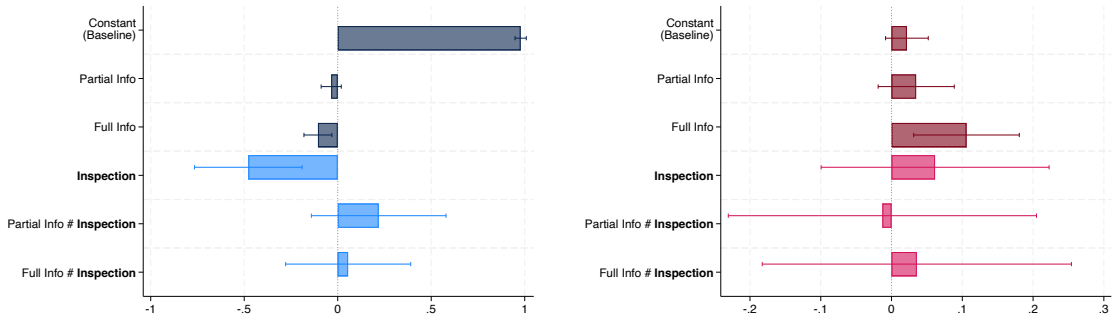
Interestingly, we observe that within the set of participants who choose to inspect the full list of deals available, deal selection choices are not statistically different across treatment conditions. On the contrary, when we look at participants who restrict their search only to the set of deals visible on the main product page, we observe that exposure to the Full Information treatment triggers a *direct* negative effect on the probability the Buy Box winner is selected, compensated by a corresponding positive effect on the propensity to select one of the other deals displayed on the main product page different from the Buy Box winner — e.g. BB2 or COA deal (see Figure 5). Under the Full Information treatment, the latter two direct effects sum up to the main (indirect) effect mediated by the increase in participants' willingness to extend their search, resulting in the sizeable reduction in BB1 selection trends observed overall.

Table 5: Deal selection across treatment conditions.

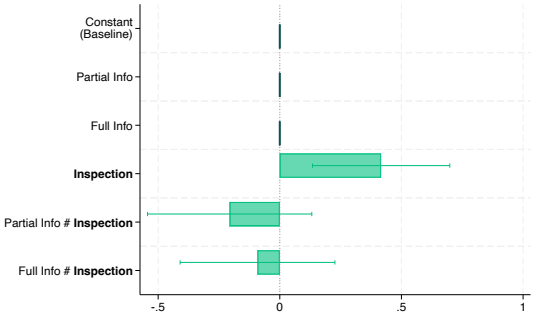
	BB1 deal	Not-BB1 deal on main page			OTHER deal on secondary page
		Overall	BB2 deal	COA deal	
Partial Information	-0.019	0.040	0.009	0.031	-0.016
	(0.037)	(0.030)	(0.009)	(0.029)	(0.026)
Full Information	-0.177***	0.139***	0.041**	0.098***	0.049
	(0.046)	(0.039)	(0.018)	(0.035)	(0.033)
Constant (Baseline)	0.923***	0.032*	0.000	0.032*	0.048**
	(0.026)	(0.018)	(0.000)	(0.018)	(0.021)
Observations	363	329	329	329	363
R2	0.051	0.039	0.017	0.023	0.014

Notes. One observation is one participant, we include participants assigned to all treatment conditions. Linear Probability Model estimates; the dependent variable is a binary indicator equal to 1 when the participant selects (1) the BB1 deal, (2) one of the deals visible on the main product page, other than BB1, (3) the BB2 deal, (4) one of the COA deals, (5) one of the OTHER deals visible only by inspecting the secondary page. Robust standard errors in parentheses. Models (2-4) are estimated on the restricted sample of N=329 participants for whom alternative deals other than BB1 where visible on the main product page. Significance levels: $^{*}p < 0.10$,** $p < 0.05$,*** $p < 0.01$.

Figure 5: Deal selection across treatment conditions, by secondary page inspection.



Panel A. Selection: BB1 deal (Left) and Not-BB1 deal on the main page (Right).



Panel B. Selection: OTHER deal on the secondary page.

Notes. Linear Probability Model estimated coefficients (and constant term). The dependent variable is a binary indicator equal to 1 when the participant decides to select the BB1 deal (Panel A - Left), one of the deals visible on the main product page other than BB1 (Panel A - Right), one of the OTHER deals visible only by inspecting the secondary page (Panel B).

Table 6: Selected deal characteristics across treatment conditions.

	Used Condition	Price	Shipping Time	Shipping Cost	Amazon is seller	Fulfilled by Amazon
Panel A. Overall (N=363)						
Partial Information	0.007	-0.151	0.447	0.053	-0.044	0.008
	(0.015)	(0.380)	(1.455)	(0.079)	(0.034)	(0.035)
Full Information	0.037	-0.810*	1.457	0.032	-0.060	-0.001
	(0.022)	(0.460)	(3.473)	(0.095)	(0.036)	(0.040)
Panel B. Subset: secondary page not inspected (N=292)						
Partial Information	0.007	-0.387	0.480	0.000	-0.044	0.004
	(0.010)	(0.325)	(1.974)	(0.000)	(0.036)	(0.042)
Full Information	0.043**	-1.517**	-1.120	0.010	-0.053	0.009
	(0.020)	(0.701)	(2.979)	(0.010)	(0.043)	(0.059)
Panel C. Subset: secondary page inspected (N=71)						
Partial Information	0.059	3.221	-0.233	0.213	-0.096	0.069
	(0.109)	(2.641)	(9.980)	(0.921)	(0.181)	(0.195)
Full Information	0.015	2.819	9.787	-0.083	-0.177	-0.010
	(0.125)	(1.862)	(13.884)	(0.670)	(0.175)	(0.143)

Notes. One observation is one participant, we include participants assigned to all treatment conditions. Column 1, 5-6: Linear Probability Model estimates; the dependent variable is a binary indicator equal to 1 when the participant selects and adds to cart (1) a used deal, (5) a deal sold by Amazon as a seller, and (6) a deal fulfilled by Amazon. Columns 2-4: OLS estimates; the dependent variable is (2) a continuous measure of the selected deal's price, in Euros; (3) the average estimated shipping time in hours, and (4) and the shipping cost in Euros. To control for product-specific effects and heterogeneity, product indicators are included in all regressions and standard errors, reported in parentheses, are clustered at the product level. Significance levels: $^{*}p < 0.10$,$^{**}p < 0.05$,$^{***}p < 0.01$.

Selected deal characteristics. Overall, selected deal characteristics reveal no strongly significant differences across treatment conditions, as shown in Table 6. However, coherently with previous results on deal selection behavior, we observe different patterns within the two subsets of participants who extend their search beyond the main product page and not. Among participants whose search is restricted to listings on the main product page only, we observe a significant treatment effect under the Full Information condition: participants tend to select Used products more often and, on average, cheaper deals, compared to the Baseline.[31]

This evidence is further corroborated by reduced form estimates of participants' deal selection choice, reported in Table 10 in Appendix D. Leveraging our ability to

[31] We do not look at whether treatment conditions affect the ratings of sellers associated to selected deals because this information is only accessible in the secondary page, irrespective of whether a deal is visible on the main product page or not.

observe the characteristics of all deals inspected by our participants, we estimate a conditional logit model where, for each participant, the probability of selecting a given deal from her observed consideration set is a function of deals' attributes. These include price and positioning at the top of the Buy Box, along with shipping time and cost, information about the product condition, and whether fulfillment by Amazon is available. Estimated coefficients capture how an attribute changes the probability a deal is selected relative to others in the same set. Results on the interaction between treatment indicators and the main features of interest - price and prominent positioning - confirm that our information intervention reduces the weight of being located at the top of the Buy Box, and increases the sensitivity to price. In particular, we observe that participants who extend their search beyond the main product page exert, on average, higher price sensitivity, yet the estimated weight of price is not statistically different across treatments. On the contrary, those who restrict their search to main product page only, exhibit lower sensitivity to price on average, yet significant treatment differences emerge.

5.4 Consumers' evaluation of the Buy Box

At the end of the e-shopping task, we elicit consumers' evaluation of the Buy Box through some (unincentivized) self-reported questions.[32] Table 7 shows that, overall, the majority of participants reveal a negative view of the Buy Box: despite some descriptive differences, we observe no significant treatment effects and the modal reason why participants express a negative view is the criticism on how the Buy Box algorithm fails to always give prominence to — what they believe to be — the best deal. Roughly half of the participants report they will search for deals out of the Buy Box more often before finalizing their purchase in the future and, again, no strong statistical differences emerge across treatments.

When asked about what features of the Buy Box algorithm they would improve, almost half of the participants report they would prefer a Buy Box that clearly indicates the reason a specific deal has been picked as the Buy Box winner (e.g. best price, reviews, etc.), in line with the major criticism expressed on its current design; second to this, roughly one-third of participants signal they would like the other deals, not in the Buy Box, to be granted more visibility. Only 1 out of 10 thinks the current Buy Box design needs no improvement.

To measure the scope for experimenter demand effects, we also ask participants

[32]Note that before answering these post-experimental questions, to make sure all participants knew what we refer to as the Buy Box, we provide all participants with some information about its functioning, which is content-equivalent to the information provided to participants in the Partial Information treatment before the e-shopping task. As a result, what we measure here is the overall result of participants' experience and reaction to information provision, which took place either only after or both before and after the e-shopping task.

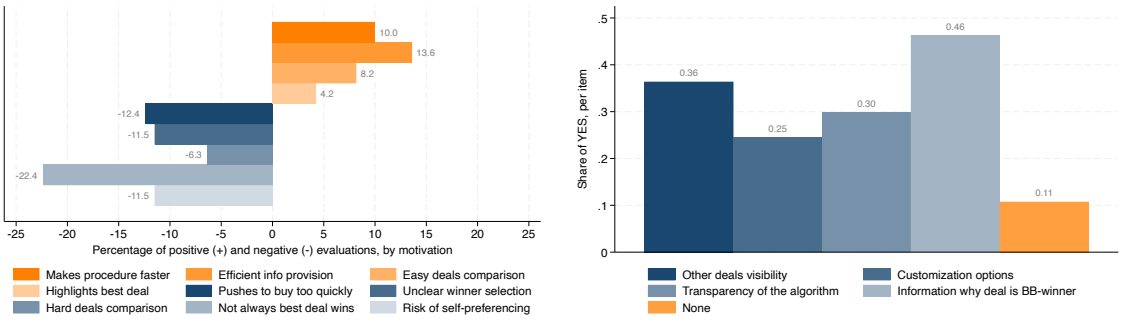
what view they think the researchers conducting the study hold on Amazon. The majority believe the researchers have a neutral view of Amazon (45%), with the second-most popular answer being that the researchers hold a 'positive' view (33%). Only less than 5% perceive that the researchers have a 'negative' view of Amazon.

Table 7: Self-reported evaluations after the e-shopping task

	Overall	**Baseline**	**Partial Info**	**Full Info**	Test
Negative view on the Buy-Box	0.63 [0.48]	0.70 [0.46]	0.62 [0.49]	0.60 [0.49]	1.80
Will search more out of Buy-Box	0.53 [0.50]	0.46 [0.50]	0.58 [0.49]	0.54 [0.50]	2.57
Researchers' view of Amazon					
- Positive	33%	32%	34%	32%	0.07
- Negative	4%	3%	5%	5%	0.10
- Ambiguous	18%	19%	18%	18%	0.03
- Neutral	45%	46%	43%	45%	0.15
Observations	363	104	125	134	

Notes. Mean values; standard deviation in squared brackets. Test values refer to Chi-square statistics from a non-parametric Kruskal-Wallis test, whose null hypothesis is that samples from different treatments are from the same population.

Figure 6: Self-reported evaluation of the Buy Box after the e-shopping task



Panel a. *Panel b.*

Notes. Panel a: Reasons why the Buy Box algorithm is positively or negatively evaluated; Panel b: Features of the Buy Box algorithm consumers would improve.

5.5 Follow-Up

Through the follow-up questionnaire, disseminated to the same participants a few weeks after the main experiment, we collect information about participants' behavior and updated beliefs on the platform. A total of 266 participants complete the Follow-Up study.

Follow-up results show that the effects of our information provision intervention are

not short-lived. The vast majority of participants still hold a negative evaluation of the Buy Box, with no statistically significant differences across treatments; a within-subject test on how participants' perceptions changed over time reveals that, especially for participants exposed to the Full Information treatment, perceptions about the Buy Box eventually became even more negative. Consistently, a relevant share of participants declare to have searched for deals out of the Buy Box more than before, after their participation in the main study. This is true across all treatment conditions, despite participants' actual efforts to inspect out of Buy Box deals being lower than their intentions (as declared at the end of the main experiment).
We also asked participants if they found the information we provided about the Buy Box and the functioning of its algorithm useful: out of a 1-10 scale, the average evaluation provided by our participants was 7.86.

To further evaluate participants' updated perception of the Buy Box, we asked them whether they found that the Buy Box winning deal corresponded to the best deal available, thinking of the searches they engaged in after their participation in the main experiment. On average, roughly 40% of participants state that the Buy Box winning deal was not the best deal available.

When exposed to the incentivized willingness to pay task (WTP),[33] participants report being willing to pay, on average, an amount of around 4% of the product price for a customized deal selection, as opposed to the one operated by the default Amazon's Buy Box algorithm (Table 8).

We also asked participants if, after taking part in the main experiment, they searched for additional information regarding the Buy Box, or more generally, the Amazon e-commerce platform. To check if this could have potentially interfered with our information provision, we also asked them to state to what extent the newly acquired information aligned with what we provided. Roughly 47% of the participants declared they searched for additional information. Among those, most report that the additional information found was either completely or at least partially coherent with the information we provided.

[33] See Section 4 for more details on the design of the task.

Table 8: Self-reported evaluations of the Buy Box in the Follow-Up

	Overall	**Baseline**	**Partial Info**	**Full Info**
Negative view on the Buy-Box	0.74 [0.74]	0.76 [0.43]	0.69 [0.46]	0.77 [0.43]
Main Experiment	0.64 [0.48]	0.69 [0.47]	0.61 [0.49]	0.63 [0.48]
Test on difference [Main vs. FU]	0.008 (224)	0.467 (51)	0.144 (83)	0.028 (90)
Searched more out of Buy-Box	0.37 [0.48]	0.26 [0.44]	0.47 [0.50]	0.36 [0.48]
Main Experiment	0.50 [0.50]	0.36 [0.36]	0.55 [0.50]	0.54 [0.50]
Test on difference [Main vs. FU]	0.008 (240)	0.108 (58)	0.106 (88)	0.000 (94)
Buy-Box information useful [1;10]	7.86 [1.55]	7.75 [1.56]	7.91 [1.56]	7.88 [1.54]
Buy-Box not best deal	0.42 [0.49]	0.43 [0.50]	0.46 [0.50]	0.37 [0.48]
WTP for customized selection	4.22 [3.72]	4.48 [3.78]	3.96 [3.69]	4.24 [3.75]
Observations	266	68	94	104

Notes. Mean values; standard deviation in squared brackets. Tests on difference [Main vs. FU]: we report the p-values of the non-parametric Wilcoxon signed-rank test between participants' answers in the Follow-up and main experiment; the test sample corresponds to the subset of participants who report an answer different from 'I don't know' in both occasions, whose size is reported in parentheses. Statistics from the "Main Experiment", as a reference, are obtained from the same restricted sample used for the within-test. WTP for customized selection: participants' willingness to pay for a customized preference-based deal selection vs. Buy-Box winner default; values are expressed as percentages of the product price.

6 Conclusion

In this paper, we focus on consumers and investigate how they interact with a specific algorithm-driven feature that shapes the design of Amazon's marketplace: the Buy Box.

Our objective is to respond to a fast-growing regulatory demand to investigate the functioning of algorithms and shed light on human-algorithm interaction.

We find that more than half of the recruited Amazon users lack full knowledge of the Buy Box mechanism and hold a negative view of it, primarily criticizing the absence of transparent indications of why a given deal is assigned to such a prominent position. Our novel granular data on users' search and deal selection on the platform allow us to validate this self-reported information, and to obtain a more detailed picture of consumers' behavior on the platform, so far only proxied via sales data. Coherently with self-reported information, and to some extent with regulatory findings, we observe that — in the absence of any information provision — more than 90% of users opt for the Buy Box winning deal. Most of these users rely on this prominent deal without even extending their search to the secondary page listing all deals available, and pay very little attention to the few alternative deals displayed on the main page, different from the Buy Box winner. Among the few who conduct a more attentive and extended search, the probability of selecting the Buy Box winner drops to 50%.
Consistent with this evidence of low Buy Box literacy and tendency to search attentively, we find that providing information about the Buy Box functioning changes both search and deal selection decisions. First and foremost, providing information boosts the propensity to search, causing users to inspect a larger set of deals before making their ultimate selection choice. Similarly to when no additional information is provided, those who inspect a larger set of options tend to rely on the Buy Box less, which leads to a reduced tendency to select the Buy Box winner. Second, providing information reduces the tendency to select the Buy Box winner also among those who restrict their search only to the deals displayed on the main product page, possibly through an increased salience mechanism. As a result of this latter effect, used and cheaper deals are selected more often, with no collateral negative effects on users' welfare in terms of increased search efforts, measured in terms of time needed to finalize the selection. Interestingly, our rich individual-level data reveal that, in this peculiar search environment — reached by consumers once they have already identified their product of interest — individual characteristics play a crucial role. Factors like pre-existing literacy of the Buy Box and satisfaction with Amazon's recommendations largely outperform search environment and product characteristics as predictors for users' decisions.

Through our follow-up, we can verify that the effect of informing consumers has lasting effects on their search behavior and attitude towards this peculiar platform

design feature.

The novel contribution of these results is to shed light on whether and how transparent information provision can affect consumers' behavior in digital settings where platform design is governed by algorithms. Through a rigoruous experimental investigation conducted in the field, we are able to collect rich data on actual consumers' behavior on the platform, holding fixed the current platform design structure and algorithm functioning. This allows us to both isolate the causal effect of information provision on consumers' reliance on algorithmic recommendations and to identify the channels through which our information intervention impacts their behavior.

The relevance of our findings goes beyond purely scientific interest and contributes to informing policy-makers on how to design a more effective regulation of algorithms and digital platforms, to tackle the issue of consumer protection. Consumer protection is a first-order policy concern, and our findings can help understand how to implement measures to foster transparency and consumer awareness in online marketplaces. Our results show that thoughtful yet soft interventions like the one we rely on in the experiment can prove effective and long-lasting, on top of being less intrusive and more viable for the regulator than other remedies such as platforms' design changes, which might also have direct and immediate effect on sellers' strategies and competition - possibly leading to unexpected and unanticipated market outcomes.

References

Abaluck, J., Compiani, G., and Zhang, F. (2020). A method to estimate discrete choice models that is robust to consumer search. Technical report, National Bureau of Economic Research Cambridge, MA.

Abeler, J., Falk, A., Goette, L., and Huffman, D. (2011). Reference points and effort provision. *American Economic Review*, 101(2):470–92.

Allcott, H., Castillo, J. C., Gentzkow, M., Musolff, L., and Salz, T. (2025). Sources of market power in web search: Evidence from a field experiment. Technical report, National Bureau of Economic Research.

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association*, 15(5):989–1024.

Armstrong, M., Vickers, J., and Zhou, J. (2009). Prominence and consumer search. *The RAND Journal of Economics*, 40(2):209–233.

Athey, S. and Ellison, G. (2011). Position auctions with consumer search. *The Quarterly Journal of Economics*, 126(3):1213–1270.

Chen, D. L., Schonger, M., and Wickens, C. (2016a). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9(C):88–97.

Chen, L. (2017). *Measuring the algorithms in online marketplaces*. PhD thesis, Northeastern University.

Chen, L., Mislove, A., and Wilson, C. (2016b). An empirical analysis of algorithmic pricing on Amazon marketplace. In *International World Wide Web Conference Committee (IW3C2)*.

Chen, N. and Tsai, H.-T. (2023). Steering via algorithmic recommendations. *The RAND Journal of Economics*.

Compiani, G., Lewis, G., Peng, S., and Wang, P. (2024). Online search and optimal product rankings: An empirical framework. *Marketing Science*, 43(3):615–636.

Crawford, G., Courthoud, M., Seibel, R., and Zuzek, S. (2022). Amazon entry on amazon marketplace. *CEPR Press Discussion Paper No. 17531*.

Cui, R., Zhang, D. J., and Bassamboo, A. (2019). Learning from inventory availability information: Evidence from field experiments on amazon. *Management Science*, 65(3):1216–1235.

de Corniere, A., Jerath, K., and Taylor, G. (2025). Fulfilled by amazon: Seller-side tying of platform services. *Working Paper*.

Derakhshan, M., Golrezaei, N., Manshadi, V., and Mirrokni, V. (2022). Product ranking on online platforms. *Management Science*, 68(6):4024–4041.

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7):1820–1859.

Donnelly, R., Kanodia, A., and Morozov, I. (2024). Welfare effects of personalized rankings. *Marketing Science*, 43(1):92–113.

Dubé, J.-P. (2024). Amazon private brands: Self-preferencing versus traditional retailing. *Antitrust Law Journal*, 86(1):115–148.

Duch-Brown, N. and Gomez-Losada, A. (2019). Competing for the amazon's buybox: A machine learning approach. In *W. Abramowicz and R. Corchuelo: BIS 2019 Workshops*.

EC (2022). Commission decision of 20.12.2022 relating to cases at.40462 and at.40703. Technical report, European Commission. Available at: https://ec.europa.eu/competition/antitrust/cases1/202310/AT_40703_8990760_1533_5.pdf. Accessed: 17/08/25.

Fang, L., Chen, Y., Farronato, C., Yuan, Z., and Wang, Y. (2024). Platform information provision and consumer search: A field experiment. Technical report, National Bureau of Economic Research.

Farronato, C., Fradkin, A., and MacKay, A. (2023). Self-preferencing at amazon: evidence from search rankings. In *AEA Papers and Proceedings*, volume 113, pages 239–243. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Farronato, C., Fradkin, A., and MacKay, A. (2025). Vertical integration and consumer choice: Evidence from a field experiment. Technical report, National Bureau of Economic Research.

FTC (2023). Complaint case no. 2:23-cv-01495-jhc. Technical report, Federal Trade Commission and State Attorneys General. Available at: https://www.ftc.gov/system/files/ftc_gov/pdf/1910134amazonecommercecomplaintrevisedredactions.pdf. Accessed: 17/08/25.

Ghose, A., Ipeirotis, P. G., and Li, B. (2014). Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science*, 60(7):1632–1654.

Gomez-Losada, A., Asencio-Cortes, G., and Duch-Brown, N. (2022). Automatic eligibility of sellers in an online marketplace: a case study of Amazon algorithm. *Information*, 13(44).

Gutiérrez, G. (2021). The welfare consequences of regulating Amazon.

Hartzell, O. and Haupt, A. (2025). Platform preferencing and price competition i: Evidence from amazon. *Available at SSRN 5126918*.

Heidhues, P., Köster, M., and Kőszegi, B. (2023). Steering fallible consumers. *The Economic Journal*, 133(652):1430–1465.

Honka, E. and Chintagunta, P. (2017). Simultaneous or sequential? search strategies in the us auto insurance industry. *Marketing Science*, 36(1):21–42.

Hunold, M., Laitenberger, U., and Thébaudin, G. (2022). Bye-box: An analysis of non-promotion on the amazon marketplace.

Jin, G. and Kato, A. (2006). Price, quality and reputation: Evidence from an online field experiment. *The RAND Journal of Economics*, 37(4):983–1004.

Johnson, E. J. and Goldstein, D. (2003). Do defaults save lives?

Jolivet, G. and Turon, H. (2019). Consumer search costs and preferences on the Internet. *Review of Economic Studies*, 86:1258–1300.

Lee, K. H. and Musolff, L. (2025). Two-sided markets shaped by platform-guided search.'.

Li, M. (2025). Gatekeepers and self-preferencing: Incentives and welfare trade-offs in two-sided markets. *CEPR Discussion Paper No. 20668*.

Madrian, B. C. and Shea, D. F. (2001). The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly journal of economics*, 116(4):1149–1187.

Raval, D. (2022). Steering in one click: Platform self-preferencing in the amazon buy box.

Resnick, P., Zeckhauser, R., Swanson, J., and Lockwood, K. (2006). The value of reputation on eBay: A controlled experiment. *Experimental Economics*, (9):97–101.

Santos, B. D. l., Hortaçsu, A., and Wildenbeest, M. R. (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American economic review*, 102(6):2955–2980.

Tai Lam, H. (2023). Platform search design and market power.

Ursu, R. M. (2018). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, 37(4):530–552.

Waldfogel, J. (2024). Amazon self-preferencing in the shadow of the digital markets act. Technical report, National Bureau of Economic Research.

Zhu, L. and Liu, Q. (2018). Competing with complementors: An empirical look at Amazon.com. *Strategic Management Journal*, 39:2618–2642.

Appendix

A Product selection

We download data about the 10000 best selling products from the "Electronic" category. This is the maximum number of products that the Keepa.com interface allows to download for a specific category. We replicate the download procedure for several days, so to have one observation per calendar day, and we create an (unbalanced) panel dataset with the most popular products. Given that we explicitly focus on stable products, we drop listings that exit (and eventually re-enter) from the market due to stock-outs, so to create a balanced panel dataset from which we extrapolate the products of interest for the experiment.

First, we exclude products priced below a value that is twice the basic fixed participation fee (i.e. $p \geq 20$€): too cheap items could induce participants to divert their attention from the main task, and we want the price to be a salient part of the experiment. Second, we keep only relatively popular products, namely those with at least 100 reviews. Thirdly, to avoid design changes in the product pages to be scraped, we focus on markets with at least three active competitors. Suppose instead that only two sellers are active. If so, the platform could show the preferred seller in the Buy Box, but at the same time the second one – being alone - would not be truly relegated to a less prominent secondary page. That seller would be displayed just below the Buy Box, making the comparison quite straightforward for the consumer. As we are interested in situations in which other options are not available at a glance in the consumers' consideration set, we consider at least $n \geq 3$. Only then, we split the resulting set of products along the dimensions of interest (i.e. low vs high price and low vs high number of competing sellers). Finally, from each one of the four resulting subsets of products, we randomly extract one product among the most popular ones.

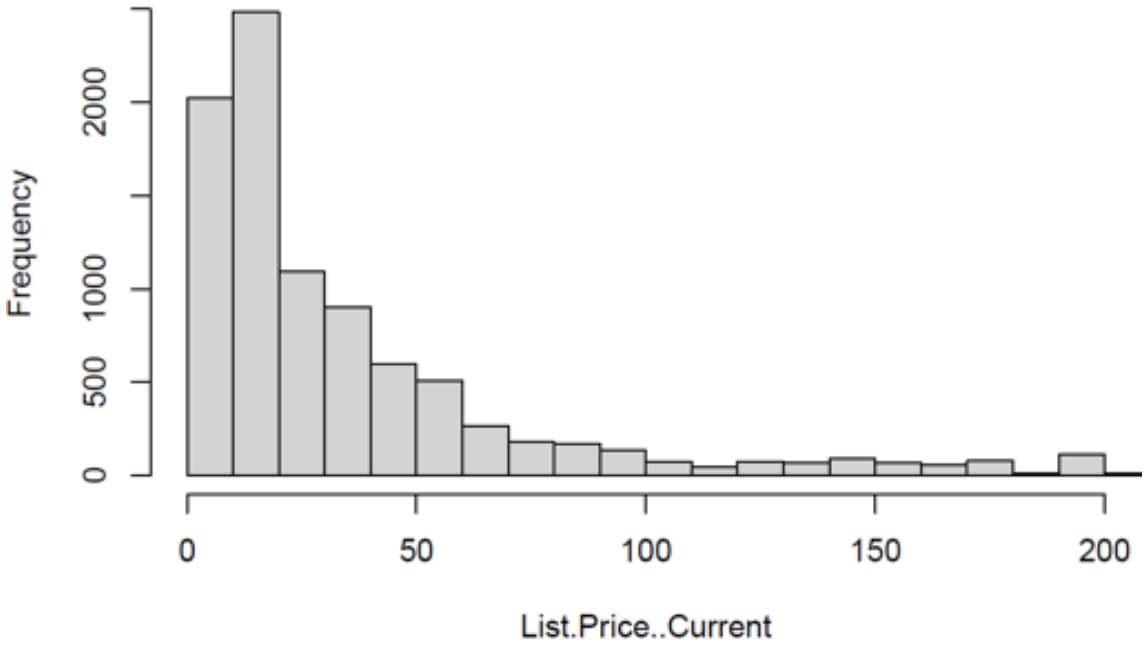


Figure 7: Price of the best selling items in the Electronics category (*Source:* `Keepa.com`).

B Data cleaning

Step 0. After the end of data collection procedures, the raw dataset counts 10656 observations (776 participants).[34] This dataset includes observations referred to all participants who reached the final stage of the e-shopping task and for which the quality check of the data collection algorithm does not fail.

Step 1: Buy Box Information.

The data. The identification of the Buy Box winning deal(s) happens when participants load the product page [TabB]. Information on winning deals' characteristics (including seller identity,[35] deal price, condition, and sub-condition) plus the ranking, when multiple winning deals are present. As a backup measure, the same information is recorded, with a short time delay, after participants terminate the e-shopping task [TabE]. Through SP (Selling Partner) API, the information on all available deals' characteristics is stored, and the deal(s) identified as Buy Box winner(s) are flagged. The identification is based only on the seller identity, the condition, and the sub-condition of the deal(s): all deals in the list with matching characteristics are marked as Buy

[34] The number of total observations in the dataset corresponds to the product between the number of participants and the number of entries per each participant; each entry (row in the dataset) corresponds to a deal available for the product shown, hence the number of entries per participant may vary across products and time. The initial dataset counts 10767 observations (781 participants), from which we remove 111 observations (5 participants) for which the quality check of the data collection algorithm fails.

[35] We refer to the alphanumeric unique identifier of each seller (ASIN) and not its respective name, as displayed on-screen to users.

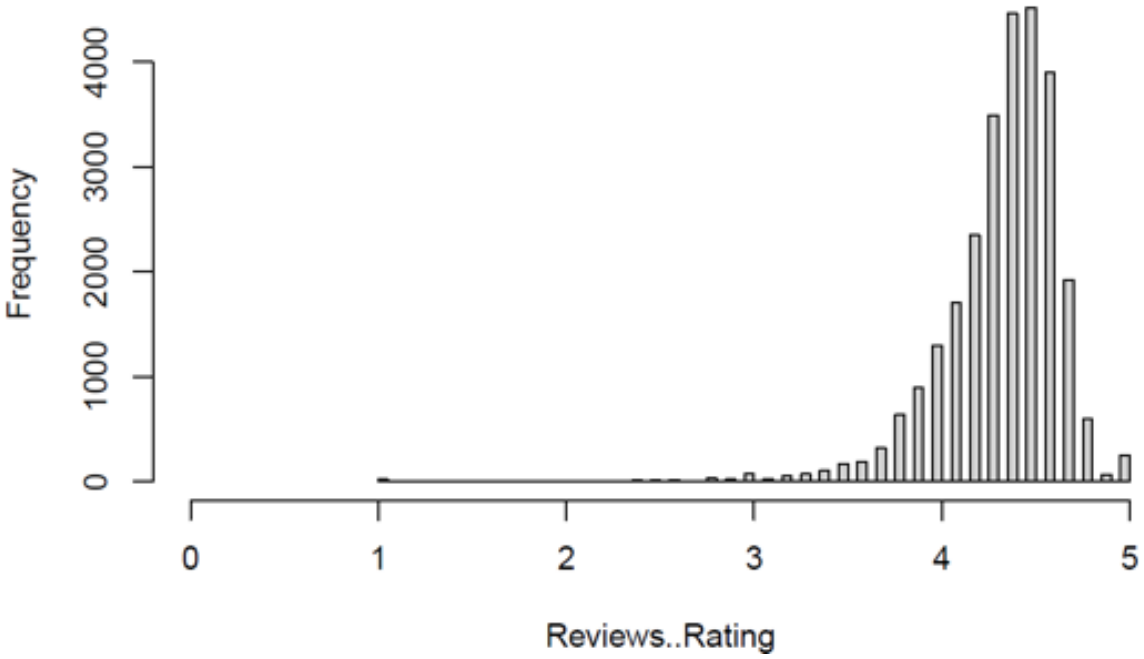


Figure 8: Ratings (i.e. number of stars: 1-5) of the best selling items in the Electronics category (*Source:* Keepa.com).

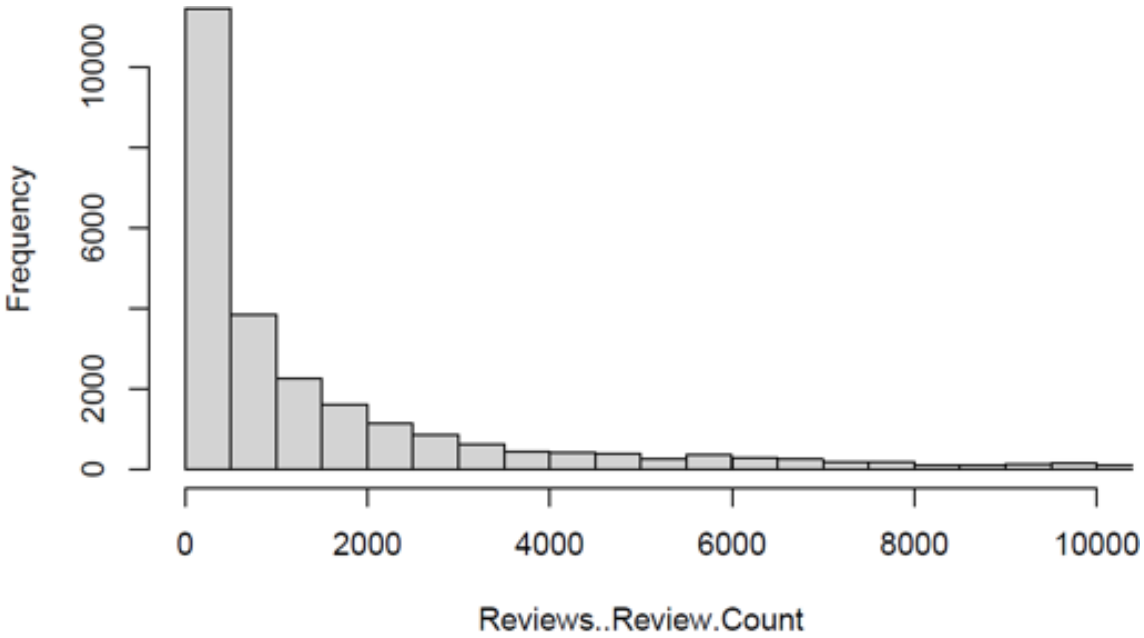


Figure 9: Number of reviews left by consumers (*Source:* Keepa.com).

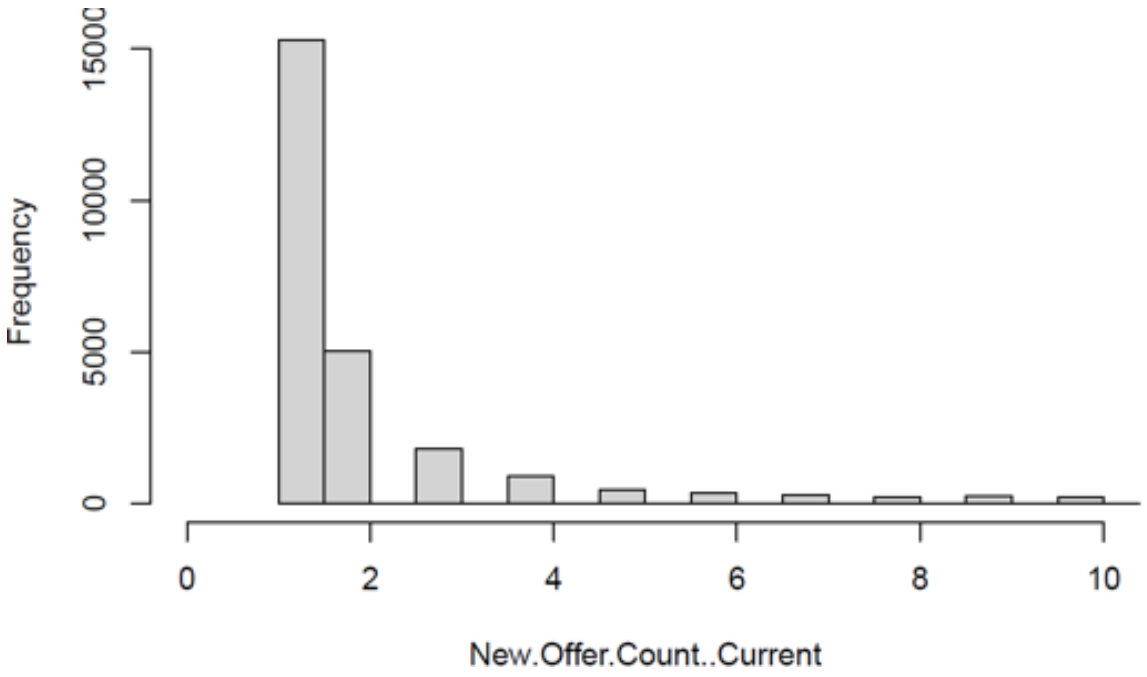


Figure 10: Number of active sellers offering a new product (*Source:* `Keepa.com`).

Box winner(s),[36] and no ranking information is provided when multiple winning deals are present.

Data cleaning. We retain all observations for which the identification of Buy Box winning deal(s) is successful through the primary information source [TabB].[37] This includes 9715 observations [666 participants] for which complete information on both Buy Box winning deals is obtained from the primary source, and an additional 63 observations [13 participants] for which all information is taken from the primary source except the deal price, which is retrieved from the backup source. Furthermore, we exclude observations for which just one deal with "Used" condition is identified as the only Buy Box winner, despite the availability of "New" deals in the list, which signals that the identification of Buy Box winning deals partially failed: this results in dropping 1007 observations (79 participants).[38] To ensure we can fully retrieve all the additional information concerning the Buy Box winning deals (e.g. shipping terms, reputation, size, etc.), which is only contained in the list encompassing all available deals [TabE], we check whether winning deals are included in such list, adjusting for

[36] Before imputation, some data cleaning operations are needed to reduce the number of deals flagged as Buy Box winners in the presence of duplicates in terms of seller identity and deals' condition and sub-condition. Buy Box winning deals' ranking is imputed based on deals' condition and price: "New(er)" and cheaper deals are ranked first.

[37] We exclude: 84 observations (3 participants) for which no price information for the first winning deal is available from the primary source and no univocal imputation can be performed through the backup source; 694 observations (90 participants) for which no information is available from the primary source; 100 observations (4 participants) for which no information on the Buy Box winning deals is available from either source. After this stage, the dataset counts a total of 9778 observations (679 participants).

[38] The dataset now counts 8771 observations (600 participants).

minor price differences in (due to the fact that the information contained in this list is stored with some delay with respect to the moment the product page information is loaded and stored).[39] We drop 569 + 80 observations (45 + 7 participants) for whom this check fails, signaling the retrieval of information on all deals available is incomplete. As a final quality check, we drop 540 observations (51 participants) for which the information on the amount of time spent on the product page is not recorded.

After all these operations, the dataset counts a total of 7662 observations (504 participants).

Step 2: Selected deal Information.

The data. The identification of which deal participants select happens at two different points in time. First, when the participants click on the "Add to cart" button immediately before leaving the main product page [TabB]. Second, at the very last step of the e-shopping procedure, when participants are asked to load their "Cart" information page before terminating the activity [TabG]. In the first case, information on the seller identity, as well as the price, condition, and sub-condition of the selected deal, is stored.[40] In the second case, only the selected deal price information and seller identity information is stored.[41] We primarily rely on the information stored from the product page. This information is more complete and more closely reflects what is seen by participants when they choose. We rely on the information stored subsequently as a secondary source whenever the same information is not available from the primary source. As a backup measure, just like for the Buy Box winning deals, selected deals are also identified, with a short time delay, after participants terminate the e-shopping

[39] This imputation involves: 941 (942) observations (43 [40] participants) for which no duplicates the first [second] Buy Box winning deal are detected in the list in terms of deal's condition, sub-condition and seller identity yet minor price differences are present, hence the price recorded at the moment the product page is loaded is imputed; 69 [83] observations (6 [3] participants) for which at least one duplicate of the first [second] Buy Box winning deal is detected in the list in terms of deal's condition, sub-condition and seller identity, hence the price recorded in the moment the product page is loaded is imputed to the corresponding deal in the list with the smallest price difference.

[40] Since the identification of some of these characteristics sometimes fails for various technical reasons, related – for example – to page visualization (e.g. the presence of a "Prime" badge often prevents retrieving information on the identity of the seller in charge of the selected deal), we maximize the number of observations [participants] for which a full identification of the selected deal is possible by retrieving, if needed, the partial information that is missing from the list of all available deals. This imputation is feasible only when, based on the partial information available, it is possible to univocally pin down which of all deals listed corresponds to the selected deal. Through this procedure, we retrieve full information on the selected deal for 52 observations (6 participants) for which the seller identity information was missing and on 106 observations (6 participants) for which the deal price information was missing. As a result, full information on participants' selected deal is available for 3605 observations (208 participants).

[41] Also in this case, we leverage the list of all available deals to retrieve missing information on participants' selected deal. In particular, we impute information on seller identity for 3723 observations (223 participants). As a result, we have complete information on selected deals' price and seller identity for 5230 observations (327 participants). Out of these, we retrieve full information – including also the selected deal condition and sub-condition – for 5123 observations (323 participants).

task [TabE], through a flag that identifies the selected deal among all deals available stored in the list.

Data cleaning. . We retain all observations for which we can identify which deal participants added to cart relying on the information stored either through the primary or secondary information source. Specifically, we identify the selected deal based on the information available from the primary source [TabB] for 3605 observations (208 participants). Furthermore, we identify the selected deal based on the information from the secondary source [TabG] for additional 2295 observations (157 participants).[42] In this last case, only two characteristics of the selected deal are stored - price and seller identity - and the information regarding the deal condition and sub-condition is retrieved from the list including all available deals' characteristics [TabE].[43] To ensure this imputation is reliable, we include only observations for which the identification from the secondary source [TabG], in terms of seller identity and price deal, matches the information stored also in the back-up source [TabE]. As for Buy Box winning deals, we check whether all deals identified as added to cart are contained in the list of all deals available for each product and participant [TabE], adjusting for small price differences (due to the fact that the information contained in this list is stored with some delay with respect to the moment the product and cart pages are loaded and the

[42] We restrict this imputation to observations for which the information extracted from the secondary source coincides with the information stored in the back-up source.

[43] We perform a small imputation to address a minor anomaly emerged during the testing phase and regarding sellers listing multiple deals, with different conditions (and possibly sub-conditions), for the same product. When a deal listed by one of these sellers is added to cart, the secondary source [TabG] correctly identifies the identity of the seller but might fail to distinguish which of the multiple deals listed by this seller is actually selected by the participant, imputing automatically the selection to the best new offer listed by the identified seller. This anomaly, in a few cases, lead to obvious misclassifications, which we correct for. This happens when the seller added to cart lists the second Buy Box winning deal but not the first one, which has a lower price than the best new offer listed by the seller added to cart, and participants have not inspected the other deals listed on the additional page: the best new offer listed by the seller added to cart is marked is mistakenly added to cart. We correct for this misclassification by marking the second Buy Box winning deal as added to cart; in our data, this correction affects 33 observations (2 participants). We do not perform additional corrections, yet other cases might also lead to biased added-to-cart deal identification. First, cases when the seller added to cart lists the second Buy Box winning deal but not the first, and participants inspected the other deals listed on the additional page: the best new offer listed by the seller added to cart is identified as the selected deal, yet participants could have opted for the second Buy Box winning deal; in our data, this situation affects 62 observations (4 participants) and could possibly lead to an over-estimation of the frequency non Buy Box winning deals are selected. Second, cases when the seller added to cart lists both deals shown in the Buy Box, and – irrespective of whether other deals listed on the additional page are inspected - the best new offer listed by the seller added to cart is identified as the selected deal, yet participants could have opted for the second Buy Box winning deal; in our data, this situation affects 66 observations (2 participants), and could possible inflate the count of times the first Buy Box winner is selected, contrary to the previous case. The last, and least, problematic case refers to when the seller added to cart lists the first Buy Box winning deal but not the second one (when present), and participants inspected the other deals listed on the additional page: the first Buy Box winner is identified as the selected deal, yet participants could have, in principle, also opted for the second Buy Box winning deal; in our data, this situation affects 58 observations (4 participants), and could possibly lead to an over-estimation of the frequency Buy Box winning deals are selected.

information on deals added to cart is saved). We drop 50 observations (2 participants) for which this check fails, signaling the list of all deals available is incomplete.

After all these operations, the final dataset counts 5850 observations [363 participants] in total.

C Additional Figures

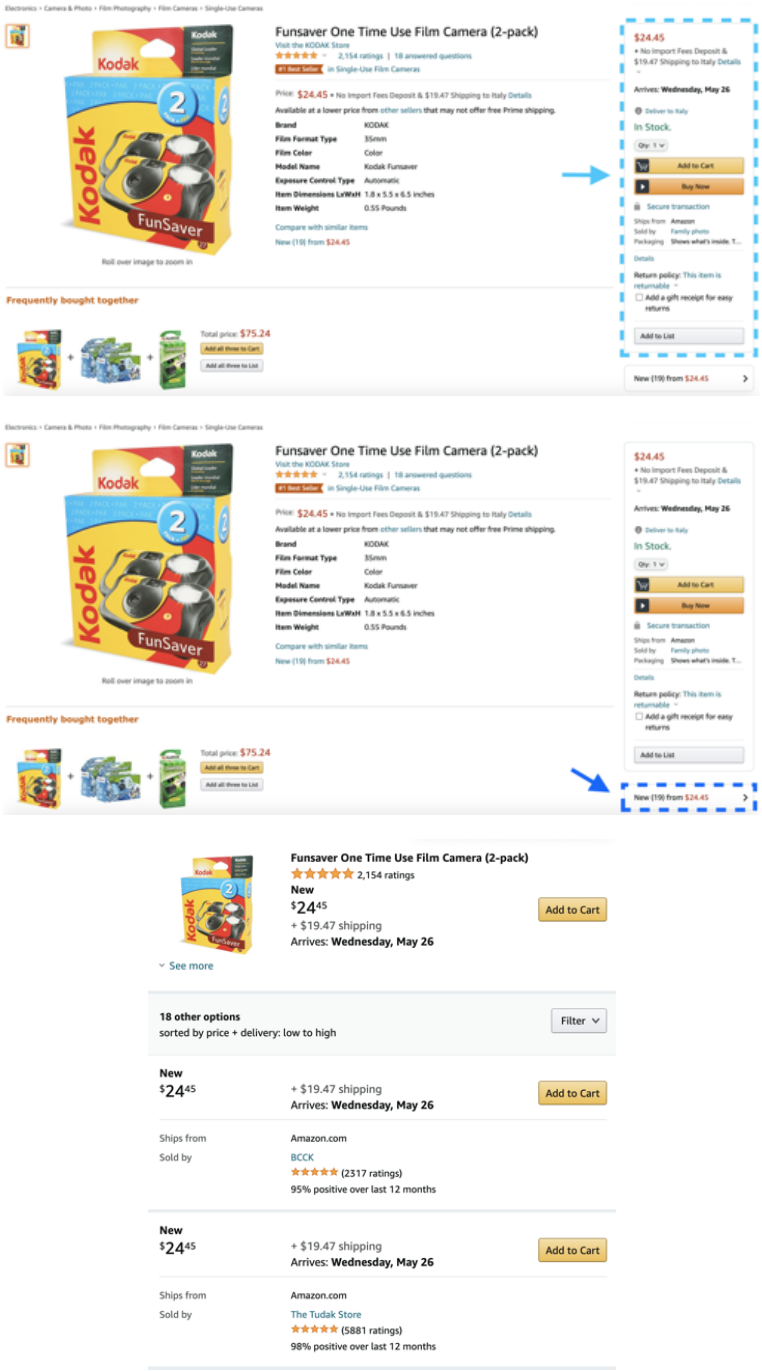


Figure 11: Location of the Buy Box (light blue) and button to access the secondary page with alternative deals (dark blue) on the main product page: Amazon (Aug 24, 2022).

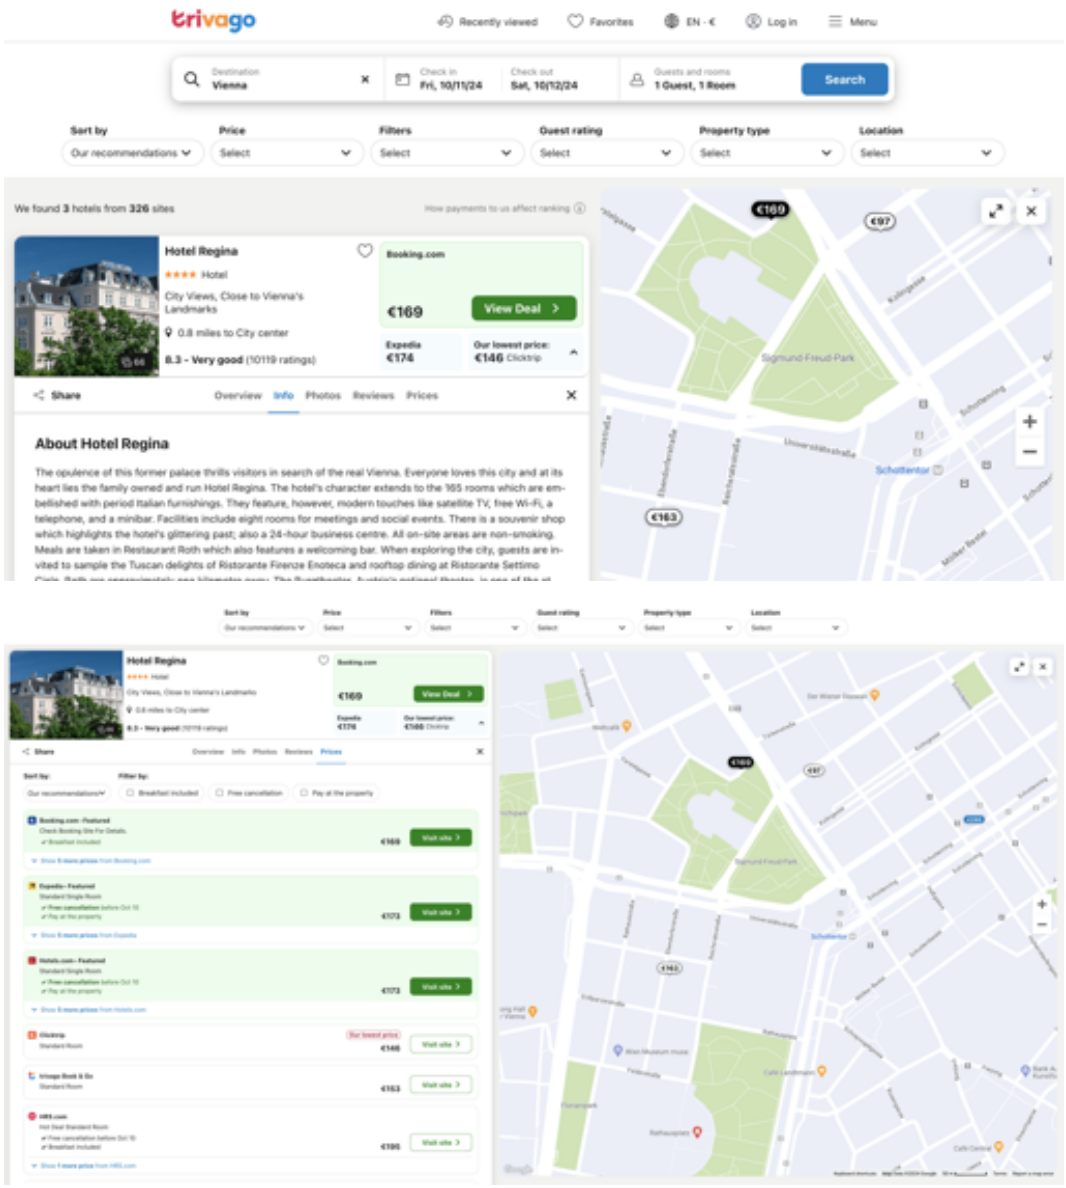


Figure 12: Other examples of 'extreme' forms of ranking algorithms: Trivago (Oct 10, 2024).

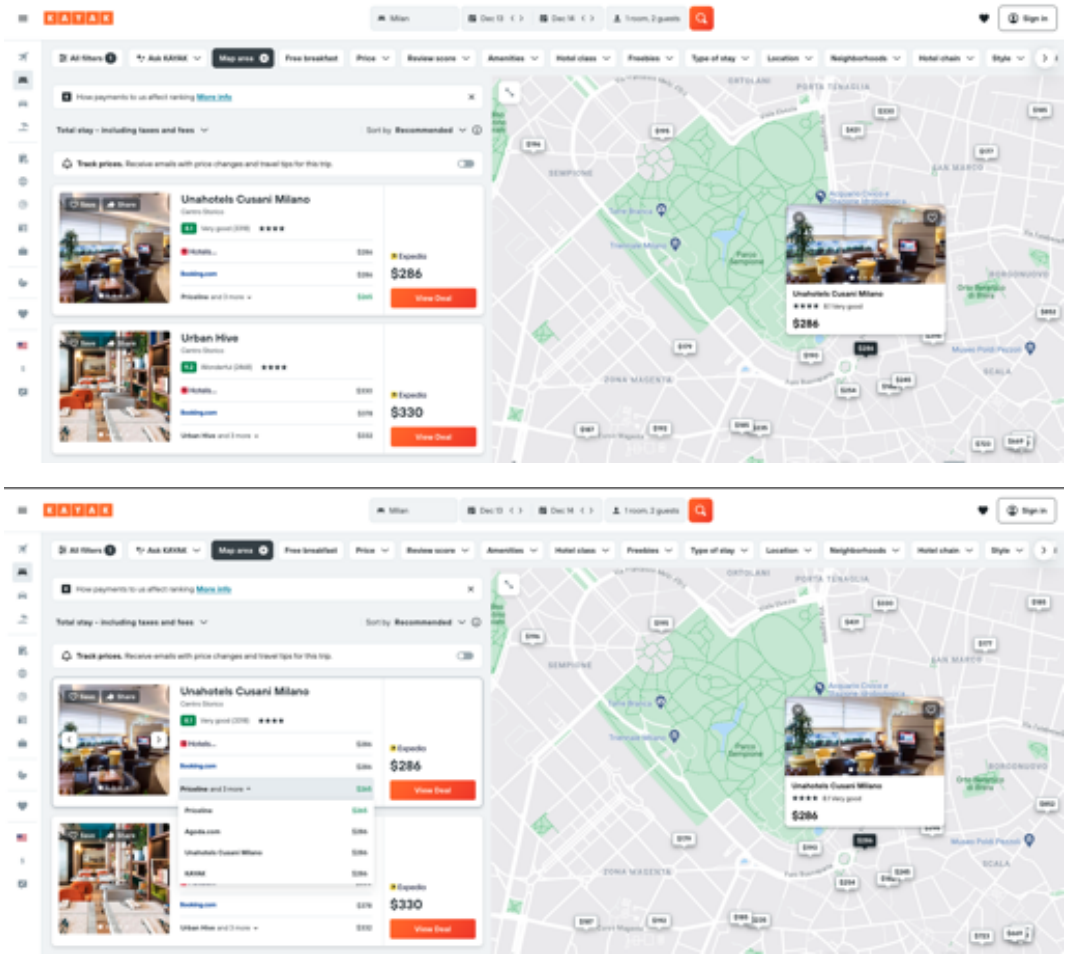


Figure 13: Other examples of 'extreme' forms of ranking algorithms: Kayak (Dec 12, 2024).

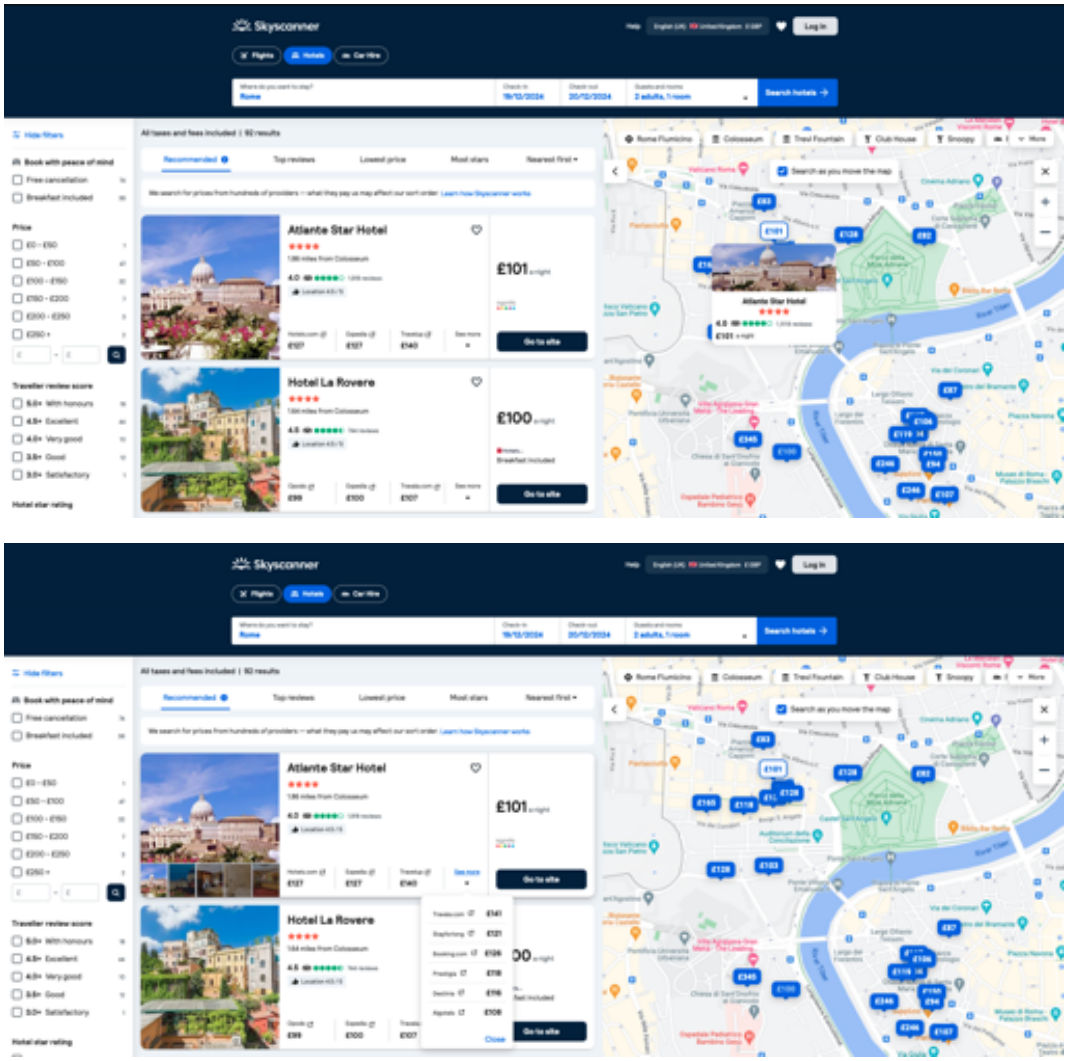


Figure 14: Other examples of 'extreme' forms of ranking algorithms: Skyscanner (Dec 12, 2024).

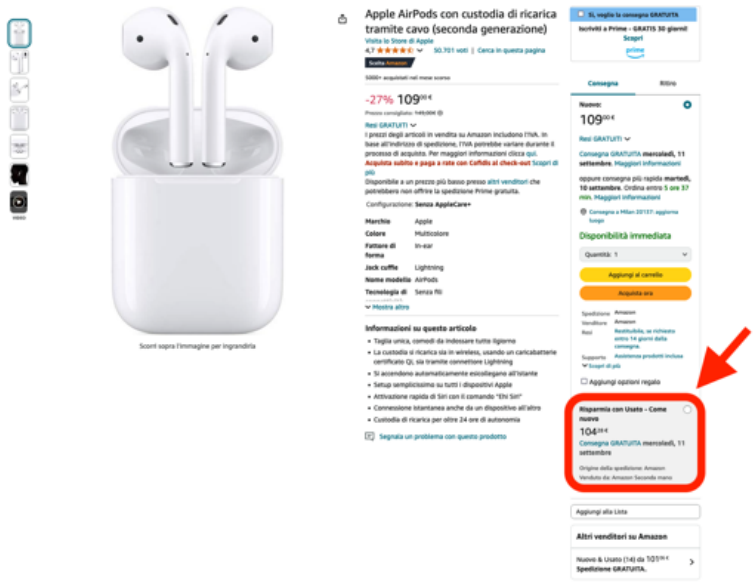


Figure 15: Positioning of the 'second' Buy Box winning deal (BB2) on the product page.

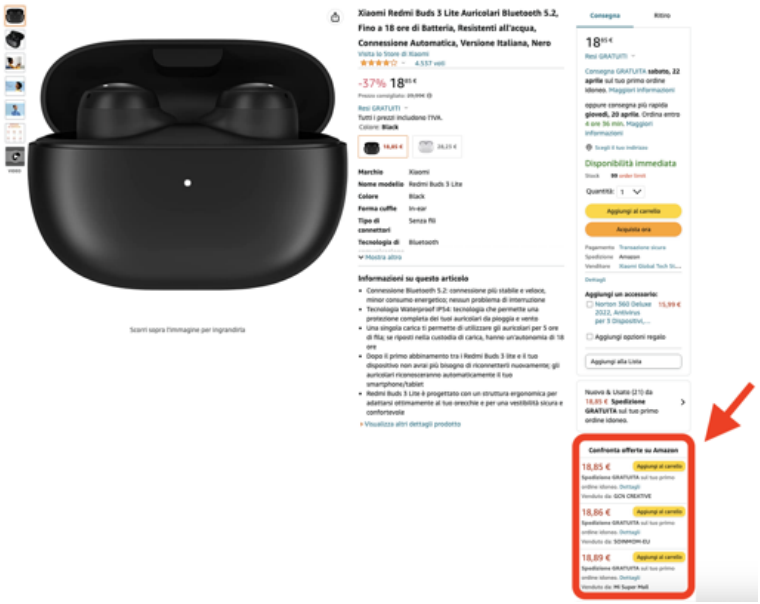
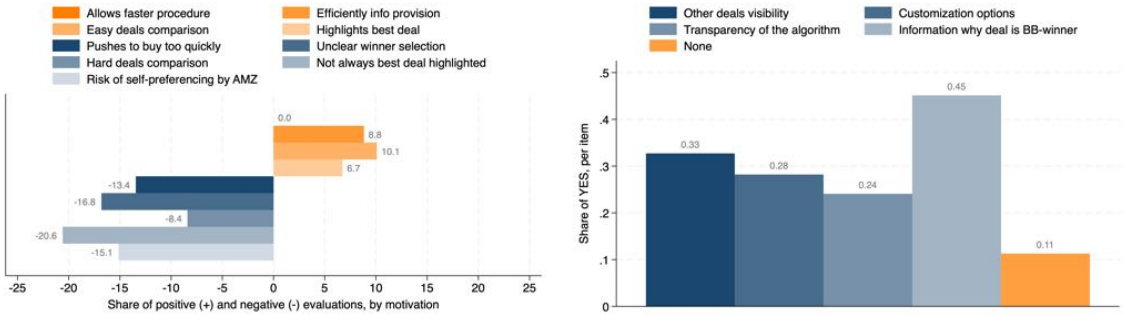


Figure 16: Positioning of the 'Compare Offers on Amazon' (COA) box on the product page.

Figure 17: Self-reported evaluation of the Buy Box in the Follow-Up: drivers



Panel a. *Panel b.*

Notes. Panel a: Reasons why the Buy Box algorithm is positively or negatively evaluated; Panel b: Features of the Buy Box algorithm consumers would improve.

D Additional Tables

Table 9: Search behavior across treatment conditions — Control variables estimates.

	Probability to inspect the secondary page	Time spent to finalize	N. of scrolls on product page
N. deals main page	-0.003	7.869	0.350
	(0.022)	(7.593)	(0.267)
Expensive product (Y/N)	-0.039	-4.427	-0.673
	(0.047)	(24.600)	(0.758)
High n. deals (Y/N)	0.034	21.041	0.197
	(0.055)	(22.767)	(0.800)
Buy Box: Contains AMZ (Y/N)	0.061	45.183	0.929
	(0.065)	(30.510)	(0.949)
Buy Box: No instant shipping (Y/N)	0.043	-62.940	-1.731
	(0.115)	(89.267)	(1.920)
Buy Box: Fulfilled by AMZ (Y/N)	-0.051	-49.933	-1.528
	(0.114)	(91.190)	(1.987)
Buy Box: Price gain wrt cheapest	0.294	10.258	-1.874
	(0.225)	(102.260)	(3.029)
Prime user (Y/N)	0.054	-28.938	-0.257
	(0.043)	(29.699)	(0.713)
Instant buyer (Y/N)	-0.039	-42.055**	-0.356
	(0.045)	(18.631)	(0.722)
Low AMZrec liking (Y/N)	0.150***	-15.246	0.401
	(0.044)	(20.191)	(0.677)
High BB literacy (Y/N)	0.120***	-11.994	0.333
	(0.042)	(21.059)	(0.670)
Low cost of time (Y/N)	0.038	23.357	0.878
	(0.041)	(18.559)	(0.619)
Female	-0.026	34.019	1.101*
	(0.042)	(21.725)	(0.645)
Income quantile	-0.000*	-0.018	0.001
	(0.000)	(0.027)	(0.001)
Age	-0.004**	2.582***	0.034
	(0.002)	(0.788)	(0.028)
Education	0.024	-2.034	0.006
	(0.030)	(16.847)	(0.422)
Employed (Y/N)	-0.094**	-16.799	-1.437
	(0.047)	(24.009)	(0.908)

Notes. Treatment effects are estimated but not reported. Controls include: product characteristics (indicators for expensive products and products with a high number of deals available, and the total number of deals shown on the main product page); Buy Box deal characteristics (indicators for Amazon's presence, absence of instant shipping, fulfillment by Amazon, and a continuous measure of percentage price difference wrt the cheapest deal available); individual characteristics (indicators for Prime status, fast purchasing style, low satisfaction with Amazon recommendations, high Buy Box literacy, low cost of time, plus standard socio-demographic traits — including gender, income, age, education and employment status).

Table 10: Deal selection choice estimation

	Overall	Consideration set:	
		Main page deals	Main + Other deals
BB1 · Baseline	2.939***	6.028***	1.233
	(0.416)	(0.871)	(0.800)
BB1 · Partial Information	2.626***	2.945***	1.896***
	(0.356)	(0.494)	(0.592)
BB1 · Full Information	1.590***	2.242***	0.744**
	(0.253)	(0.398)	(0.370)
Price · Baseline	-0.079**	0.220***	-0.086*
	(0.039)	(0.054)	(0.050)
Price · Partial Information	-0.180***	-0.153***	-0.235***
	(0.045)	(0.054)	(0.086)
Price · Full Information	-0.185***	-0.204***	-0.185***
	(0.045)	(0.056)	(0.058)
Shipping Time	0.002	0.003	0.001
	(0.002)	(0.010)	(0.002)
Shipping Cost	-0.070	-0.121	-0.129
	(0.070)	(0.168)	(0.082)
Sold by Amazon	0.693	1.366*	0.334
	(0.478)	(0.744)	(0.613)
Used	-5.715***	-5.407**	-6.279***
	(1.839)	(2.375)	(2.314)
Used · Condition	0.789**	0.561	0.982**
	(0.360)	(0.512)	(0.423)
Fulfilled by Amazon	0.399	0.173	0.447
	(0.446)	(0.752)	(0.544)
Observations	1947	993	954
Pseudo-R2	0.631	0.778	0.430
Log-Likelihood	-188.719	-75.225	-98.049
Price elasticity	-10.231***	-6.546***	-13.501***

Notes. One observation is a deal inspected by a participant. We include only observations on deals actually inspected by participants, pooling observations from all treatment conditions. We restrict the sample to participants whose consideration set, defined by the sum of all inspected deals, is greater than one (N=309). Conditional Logit estimates; the dependent variable is a binary indicator equal to 1 for the deal selected by the participant, and 0 for all other deals the participant had in her consideration set. Deal attributes include: a binary indicator equal to 1 if the deal was the Buy Box winner, a continuous measure of price, a continuous measure of average shipping time and cost, binary indicators for whether the deal is sold or fulfilled by Amazon, a binary indicator equal to 1 if the product condition associated to the deal is used, and a categorical variable measuring the condition quality (1=Acceptable, 2=Good, 3=Very Good, 4=Like New). Significance levels: $^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$.

D.1 Robustness: Extended dataset on search behavior

Table 11: Search behavior across treatment conditions.

	Probability to inspect the secondary page		Time spent to finalize		N. of scrolls on product page	
Partial Information	0.076**	0.065*	-20.394	-18.780	-0.447	-0.417
	(0.038)	(0.035)	(17.206)	(17.492)	(0.627)	(0.635)
Full Information	0.162***	0.152***	20.054	26.042	-0.115	-0.045
	(0.041)	(0.039)	(26.030)	(24.297)	(0.654)	(0.636)
Constant	0.097***	0.193	147.065***	137.368	6.279***	5.690**
	(0.024)	(0.162)	(13.591)	(103.144)	(0.449)	(2.716)
Observations	504	504	504	504	504	504
R2	0.029	0.135	0.006	0.063	0.001	0.029
Controls	No	Yes	No	Yes	No	Yes

Notes. One observation is one participant, we include participants assigned to all treatment conditions, pooled. (1-2) Linear Probability Model estimates; the dependent variable is a binary indicator equal to 1 when the participant decides to inspect the secondary page listing alternative deals other than the subset of those shown in the main product page. (3-6): OLS estimates; the dependent variables are continuous measures of (3-4) the amount of time spent, in seconds, and (5-6) the number of scrolls performed on the product page, before reaching the cart stage. Controls include: product characteristics (indicators for expensive products and products with a high number of deals available, and the total number of deals shown on the main product page); Buy Box deal characteristics (indicators for Amazon's presence, absence of instant shipping, fulfillment by Amazon, and a continuous measure of percentage price difference wrt the cheapest deal available); individual characteristics (indicators for Prime status, fast purchasing style, low satisfaction with Amazon recommendations, high Buy Box literacy, low cost of time, plus standard socio-demographic traits — including gender, income, age, education and employment status). Robust standard errors in parentheses. Significance levels: $^{*}p < 0.10$,$^{**}p < 0.05$,$^{***}p < 0.01$.

Table 12: Search behavior across treatment conditions — Control variables estimates.

	Probability to inspect the secondary page	Time spent to finalize	N. of scrolls on product page
Expensive product	-0.038	-13.645	-0.411
	(0.039)	(21.830)	(0.654)
High n. deals	0.029	-4.813	-0.169
	(0.045)	(20.703)	(0.683)
N. deals main page	-0.002	-1.707	0.162
	(0.017)	(7.973)	(0.232)
Buy Box: Contains AMZ	0.032	61.428**	0.816
	(0.050)	(28.708)	(0.815)
Buy Box: No instant shipping	0.015	-64.325	-1.704
	(0.095)	(82.508)	(1.797)
Buy Box: Fulfilled by AMZ	-0.060	-55.253	-1.207
	(0.094)	(83.698)	(1.825)
Buy Box: Price gain wrt cheapest	0.284*	-19.165	-1.650
	(0.166)	(86.082)	(2.340)
Prime user	0.018	-24.770	-0.496
	(0.037)	(24.217)	(0.582)
Instant buyer	-0.026	-32.485*	-0.273
	(0.038)	(19.136)	(0.607)
Low AMZrec liking	0.125***	-12.750	0.607
	(0.035)	(18.262)	(0.555)
High BB literacy	0.129***	-11.008	0.268
	(0.035)	(18.725)	(0.554)
Low cost of time	0.020	15.194	0.153
	(0.034)	(17.493)	(0.517)
Female	-0.023	2.055	0.629
	(0.035)	(19.289)	(0.534)
Income	-0.000***	0.011	0.001
	(0.000)	(0.030)	(0.001)
Age	-0.005***	2.727***	0.035
	(0.001)	(0.749)	(0.024)
Education	0.021	5.168	0.111
	(0.025)	(14.970)	(0.347)
Employed	-0.060	-47.892*	-0.921
	(0.039)	(25.535)	(0.720)

Notes. Treatment effects are estimated but not reported. Controls include: product characteristics (indicators for expensive products and products with a high number of deals available, and the total number of deals shown on the main product page); Buy Box deal characteristics (indicators for Amazon's presence, absence of instant shipping, fulfillment by Amazon, and a continuous measure of percentage price difference wrt the cheapest deal available); individual characteristics (indicators for Prime status, fast purchasing style, low satisfaction with Amazon recommendations, high Buy Box literacy, low cost of time, plus standard socio-demographic traits — including gender, income, age, education and employment status).