Submitted to *Manufacturing & Service Operations Management*
manuscript

Algorithmic Assortment Curation: An Empirical Study of Buybox in Online Marketplaces

Santiago Gallino
The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, sgallino@wharton.upenn.edu

Nil Karacaoglu
Fisher College of Business, The Ohio State University, Columbus, OH 43210, karacaoglu.2@osu.edu

Antonio Moreno
Harvard Business School, Harvard University, Boston, MA 02163, amoreno@hbs.edu

Problem Definition: Online marketplaces have revolutionized online sales by creating platforms that connect millions of buyers and sellers. While the presence of numerous third-party sellers attracts customers, it also results in a proliferation of listings for each product, making it difficult for customers to choose between the available options. To address this issue, online marketplaces employ algorithmic tools to curate and present different product listings to customers. Although tools that assist customers in choosing between *different products*, such as recommender systems and reviews, have been extensively studied, there is limited evidence regarding tools that help customers choose between different *listings of the same product*. This paper focuses on the buybox algorithm, an algorithmic tool that prominently presents one option as the default choice to customers.

Methodology/Results: We assess the influence of the buybox on marketplace dynamics by examining its staggered introduction within a major product category in a leading online marketplace. Our results show that the implementation of buybox increases the number of orders and enhances the efficiency of the customer journey. This is evidenced by an increase in conversion rates and a more pronounced buybox effect on the mobile channel, where search frictions are higher compared to the desktop channel. The introduction of buybox simplifies the process of posting new products on the marketplace, potentially reducing friction for sellers. We find supporting evidence for this hypothesis, as the number of sellers offering a product increases after the introduction of buybox.

Managerial Implications: Our analysis reveals that a buybox is an effective tool for reducing search frictions and stimulating competition among sellers. Customers benefit from lower prices and higher average quality levels when competition in a buybox is intense. However, the marketplace becomes more concentrated following the introduction of the buybox, representing an unintended consequence that platforms and vendors should manage. Our study contributes to the growing literature on algorithms in platforms by examining how algorithmic curation affects marketplace participants and overall marketplace dynamics.

Keywords: marketplace operations, algorithmic curation, online marketplaces, retail operations, empirical operations, buybox

1. Introduction

The rise of online marketplaces, with global sales exceeding \$3.48 trillion in 2023, (Ghavami 2024), has transformed how we shop, connect, and sell. The success of leading marketplaces like Amazon, Alibaba, MercadoLibre, and Rakuten stems from their ability to attract vast customer bases and provide third-party sellers with access to millions of customers that were previously out of reach. For instance, Amazon boasts over two billion monthly visits and more than three million third-party sellers in North America alone (Kaziukenas 2021).

Marketplaces attract large numbers of third-party sellers offering a vast range of products, often with multiple sellers' listings for the same product. While these extensive assortments and numerous options for each SKU attract customers, the abundance of alternatives can lead to choice overload (Iyengar and Lepper 2000) and search frictions (Brynjolfsson and Smith 2000), making it difficult for buyers and sellers to connect. Moreover, unlike traditional retailers, marketplace operators often lack control over sellers' product selection and pricing. These factors amplify the need for marketplaces to curate their assortments in a customer-friendly manner, ensuring the best experiences and choices while enabling sellers to list their products uniformly. Algorithms are essential for performing this function at scale. However, the haphazard application of these algorithms can be problematic. Unlike traditional retail settings, where retailers primarily consider customers' responses to new algorithms or policy changes, marketplaces must account for the responses of all participants to assess whether the algorithms generate the desired effects. Therefore, it is crucial to study how these algorithms affect marketplace dynamics and examine the responses of the various parties involved.

Our paper examines the widely-used buybox algorithm, a “featured-offer" algorithm that allocates a dedicated product page for each item. It selects the seller with the best offering based on an algorithm and prominently presents this seller as the default option on the product page. While other sellers offering the same product remain accessible, they are available at a higher search cost for customers. Figure 1 showcases a product page and buybox from an Amazon template. Many marketplaces, from global platforms like eBay and Walmart to local sites like Hepsiburada and Trendyol in Turkey, have adopted this algorithmic tool to curate their assortments and provide customers with the best choices.

The implementation of a buybox algorithm brings significant changes to the marketplace for both customers and sellers. Sellers benefit from a standardized product catalog and product pages, simplifying the process of posting new products because they do not have to create individual product description pages. This can reduce friction for sellers. Additionally, a buybox algorithm influences the way products are displayed to customers, potentially affecting seller behavior. Concurrently, customers' search processes are considerably impacted as the algorithm curates and prominently showcases default options for each product.

We explore the impact of introducing the buybox algorithmic curation on customers, sellers, and marketplace dynamics by collaborating with a leading marketplace that introduced the buybox sequentially across categories and products. The marketplace provided data for a major electronics subcategory, which began implementing the buybox in September 2019. The marketplace introduced the buybox for different products in this subcategory in a staggered manner. When at least one seller participates in an available buybox, the buybox for that product becomes active. Consequently, each product's buybox became active and visible to customers at different times. We leverage this variation in the timing of buybox activation for different products to examine how its introduction affects the marketplace.

Our extensive dataset encompasses all listings, associated sellers, orders, purchase channels, visits, and the buybox rollout timeline within this subcategory. This comprehensive data enables us to examine the buybox's effects on customers, sellers, and marketplace dynamics in detail. To identify the impact of the buybox on the marketplace, we employ a generalized differences-in-differences (DiD) design that leverages the timing of buybox activation for each product. Our empirical analysis provides novel evidence on the consequences of buybox algorithmic curation, offering valuable insights for both academics and practitioners.

Our findings reveal that the adoption of the buybox leads to a significant increase in marketplace conversion rates (around 23%, $p < 0.01$) and transaction volumes (i.e., the number of orders; around 78%, $p < 0.001$). We conduct extensive analyses to verify the robustness of our results, demonstrating that the observed increases are driven by buybox activation rather than alternative mechanisms.

Furthermore, our results indicate that as the number of sellers participating in a buybox increases, the buybox has a larger effect on conversion rates and orders. This finding highlights how competition within the buybox influences its overall effectiveness. We then explore how the implementation of the buybox reduces friction in the marketplace for customers. The increased conversion rate implies a more efficient customer journey. We find that the positive effects of the buybox are more pronounced on mobile channels, which typically have higher search frictions due to smaller screen sizes and more difficult navigation.[1] Similarly, we find that the buybox has larger effects on weekdays and during work hours, where customer search costs tend to be higher. These results are consistent with the buybox reducing search frictions.

Lastly, we explore the changes in marketplace outcomes following the buybox implementation. We first examine the buybox's impact on the average price paid for products and the average seller quality of the product sold. Buybox activation does not directly influence these metrics, suggesting that the buybox does not necessarily prompt sellers to post better offerings or enable the marketplace to curate cheaper, higher-quality options. We then investigate how competition levels in the buybox affect these outcomes, as the marketplace's ability to curate better offerings might depend on the number of sellers participating in a buybox. We find that as the buybox becomes more relevant for a product—i.e. when more sellers participate in a product's buybox—the average price paid for the product decreases, and the average seller quality of the product sold increases. These findings suggest that the buybox is an effective tool in inducing competition

[1] In this paper, search frictions refer to obstacles that make it difficult for buyers and sellers to find and engage with each other in the marketplace.

between sellers and results in better outcomes for customers. Moreover, the introduction of the buybox makes it easier for sellers to post a product in the marketplace. Consistent with this notion, we observe an increase in the number of sellers offering a product following the buybox implementation, suggesting reduced friction for sellers. Finally, we observe increased seller concentration within a product and within the subcategory following the buybox implementation, indicating that the buybox induces changes in the distribution of sales across sellers.

Our findings offer empirical evidence on the benefits of algorithmic curation in marketplaces, which positively impacts marketplace transactions and mitigates customer and seller frictions. The literature has previously studied tools and algorithms that help customers choose between different products, such as recommender systems and reviews. However, there is limited evidence about tools that help customers choose between different listings of a given product, which is an increasingly relevant problem due to the prominence of marketplaces where multiple sellers offer the same products.

We contribute to the burgeoning literature in operations management on algorithms in marketplaces. While most of this literature theoretically studies algorithms in marketplaces, we empirically examine the algorithm's implementation in the marketplace and outline its effects on different stakeholders. Furthermore, while empirical papers focusing on marketplaces primarily concentrate on *customer* behavior, the operations management literature in other areas extensively documents the importance of understanding and influencing *supplier* behavior for companies' success. We fill a gap in the literature by investigating how changes in the marketplace's curation algorithm impact seller behavior, customer behavior, and marketplace outcomes.

2. The Buybox Tool: Related Literature and Conceptual Framework

Our work examines the implementation of buybox as a tool for assortment curation in a leading marketplace. In this section, we establish the theoretical context for our study. We start by describing the buybox tool and then position our research question within the existing literature. Finally, we develop a conceptual framework to study the impact of the buybox tool in relation to other algorithmic tools typically offered by marketplaces.

2.1. General Overview of the Buybox Tool

A thriving marketplace attracts a large customer base and entices third-party sellers to list thousands of products. While customers benefit from a robust seller base and a diverse product assortment, an excessive number of listings for identical products can negatively impact the customer experience and create unnecessary friction.[2] Similarly, an overabundance of listings can diminish a seller's visibility, impeding their ability to identify relevant competitors and respond to market conditions. Marketplaces can use curation algorithms to filter out multiple offerings and highlight more attractive options. A buybox is a widely used

[2] In a marketplace, a listing is a detailed record that contains essential information about a product, including the title, description, images, and seller details. Sellers must create separate listings for each product they offer in the marketplace.

algorithmic curation tool that enables marketplaces to sort out listings of the same product and select the best value option to be prominently displayed and set as the default option.

To implement a buybox, marketplaces must create a standardized product catalog, allowing sellers to match their listings to catalog items. The marketplace's catalog listing for a product (i.e. the product page) includes a title, description, and standardized images. This product listing can be used for the product's buybox.

When the marketplace creates a buybox for a product, sellers offering that product can participate and compete to be listed as the default seller. The buybox algorithm scores the different options based on factors such as price, fulfillment options, and seller reputation and performance to determine the best value offering. The selected seller is set as the default seller, displayed prominently on the product's buybox page, and receives the default sale on the product page. While customers can still access other participating sellers who have not been chosen as the default by the buybox, the navigation path to select them is significantly more cumbersome.

Most marketplaces offer a search tool where customers begin their journey with keyword searches. This search tool (which is independent of the buybox tool) provides a popular entry point to product pages where customers can make their purchases.

Figure 1a shows the search result for a PlayStation 4 Slim 1 Terabyte in a typical marketplace before the buybox implementation. Each entry corresponds to specific sellers on the site who may offer the product at different prices or with varying delivery attributes. The search results include listings from different sellers, and some products that are displayed may not be relevant to the intended search.

Figure 1b presents a schematic view of the search result after implementing a buybox. In this updated search result, there is a listing with no specific seller mentioned (highlighted in yellow by us). If customers click on this buybox listing, they will be directed to the product page.

Figure 1c shows the product page with the details of the listing selected by the buybox algorithm.

These schematic views reflect displays commonly used by marketplaces and are similar to those employed by the focal marketplace examined in this paper. On the product page, customers can see the seller chosen as the default by the buybox algorithm, along with the corresponding price and fulfillment details. In the right column, customers can view the seller's name, the "Buy Now" and "Add to Cart" buttons, and other product details. When a customer clicks "Buy Now" or "Add to Cart," the sale goes to the seller selected by the buybox algorithm.

Other sellers participating in the buybox but not selected as the default are shown in less prominent sections of the product page. Customers must navigate to these less visible areas to find alternative sellers offering the product (see Figure 1c).

While the specifics of buybox algorithms vary between marketplaces and are often opaque to sellers, they share core features and serve similar purposes.

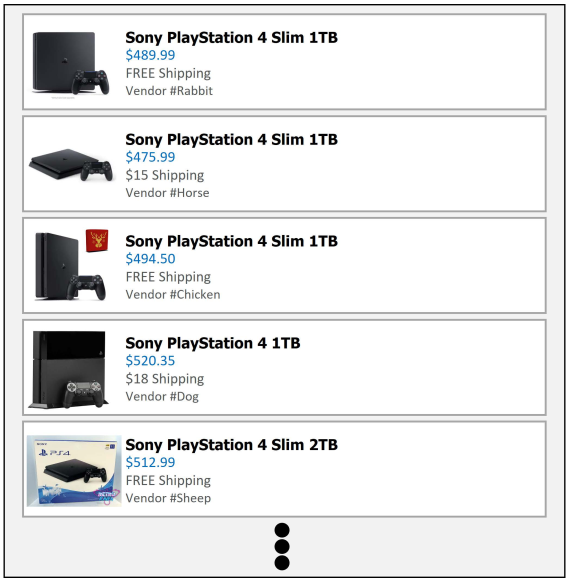

(a) Search Result before Buybox

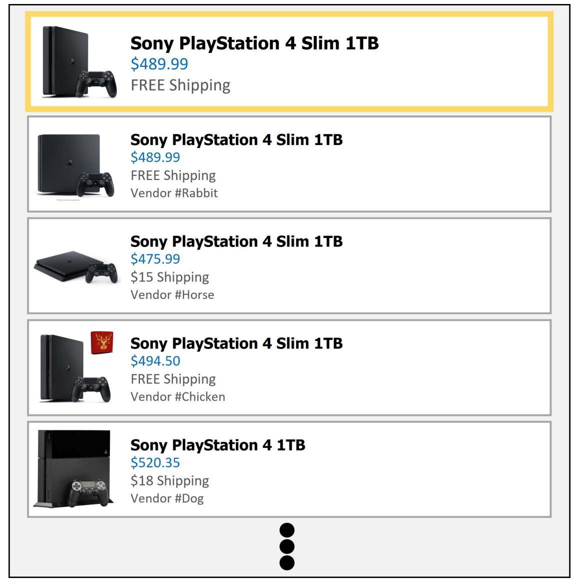

(b) Search Result after Buybox

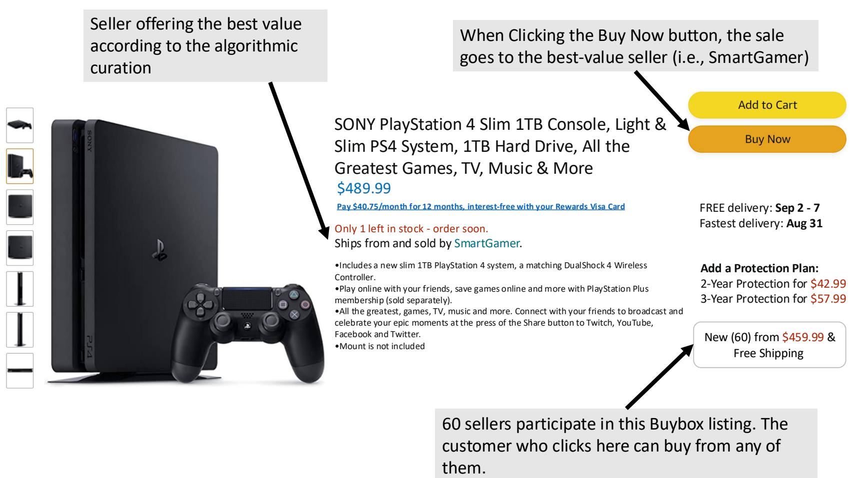

(c) Buybox Listing

Figure 1: Buybox Schematic Layout (Source: Amazon)

2.2. Related Literature

Our study of a buybox algorithm as an assortment curation tool contributes to four streams of literature: marketplace operations, search frictions in marketplaces, assortment planning and curation, and algorithmic tools in marketplaces.

2.2.1. Marketplace Operations The operations management community has recently focused significant attention on platforms and two-sided markets, with much of the literature examining platform policies and algorithms through a theoretical lens (Aouad and Saban 2023, Manshadi and Rodilitz 2022). Additionally, a growing body of empirical research has explored related issues in contexts such as dating, (Rios et al. 2023), and accommodation platforms (Li and Netessine 2020). These studies typically analyze environments where customers encounter assortments of differentiated products, such as dating profiles, (Rios et al. 2023), where an increase in listings offers a wider variety of choices. In contrast, our study examines a setting where a buybox algorithm sorts through different vendors' offerings of the same product, and additional listings do not necessarily expand product variety. Our work contributes to this literature by providing empirical evidence on the impact of adopting a buybox algorithm, which curates listings of the same product, on marketplace dynamics. Despite the algorithm's widespread use in practice, it has received limited attention in the operations management literature, with Gur et al. (2023) being a notable exception. Their study theoretically explored how platforms can use algorithms like the buybox, which modify sellers' visibility, to influence pricing in marketplaces.

The operations management literature has also emphasized the importance of considering customers' search costs when making changes to platforms. For example, Li and Netessine (2020) underscore that a high number of participants can hinder the market due to increased search costs. Bimpikis et al. (2020) demonstrate that consolidating auction end times reduces frictions for buyers, thereby increasing revenue in a liquidation auction marketplace. Algorithmic assortment curation using a buybox affects customer search frictions, reducing the search costs customers face when exploring the featured option selected by a buybox algorithm and increasing the costs of exploring other options for the same product. Search costs have been shown to vary across channels, with mobile phones having higher search costs than desktop (Ghose et al. 2013). Consistent with the role of the buybox in reducing search costs, we observe that the buybox has a greater impact on transactions on the mobile channel relative to the desktop channel.

Our research is also related to the growing empirical work in operations management literature that focuses on retail marketplaces (Cui et al. 2020, Zhang et al. 2024, Deshpande and Pendem 2023). Most of this literature examines how changes in the marketplace, such as delivery options (Cui et al. 2020) and pop-up stores (Zhang et al. 2019), affect customer behavior. While the significance of supplier behavior and relationships has been extensively explored in supply chain research, the empirical marketplace literature on the seller side remains limited, with a few notable exceptions such as Zhang et al. (2024). Zhang et al. (2024) provide valuable insights on how marquee seller entry affects marketplace dynamics. They show that the entry of high-quality marquee sellers can benefit existing high-quality sellers, whereas entry by low-quality marquee sellers benefits existing low-quality sellers. While their work highlights important quality-based competitive dynamics, our paper examines a distinct but complementary marketplace mechanism — algorithmic curation through a buybox tool. Though Zhang et al. (2024)'s insights on quality-based competition inform our understanding of seller behavior, we focus on how algorithmic curation tools can reduce marketplace friction and increase conversion. Together, these papers offer complementary perspectives on how marketplaces can strategically influence seller and buyer behavior through different tools and mechanisms.

Our work adopts a broad perspective and explores how algorithmic curation affects sellers, customers, and marketplace dynamics. We show that although the marketplace does not control the products sellers post to the platform, it can influence sellers' decisions through algorithmic curation tools. This work contributes to the broader stream of research that explores the role of incentives and algorithms in retail (Caro and Sáez de Tejada Cuenca 2023, Kesavan and Kushwaha 2020, Lee et al. 2022) and online platforms (Kabra et al. 2016, Allon et al. 2023).

2.2.2. Assortment Planning and Curation The operations management literature has extensively studied assortment planning and revenue management (Cachon and Kök 2007, Wang and Sahin 2018, Li et al. 2023). With the rise of e-commerce, the application of these techniques on online platforms has gained significant attention. Assortment planning algorithms are critical for marketplaces due to their extensive assortments, frequent price updates, and personalization opportunities (Lobel 2021). In this context, it is crucial for site operators to decide which products to offer and determine how to curate and display these assortments effectively on the site (Ferreira et al. 2022).

Most of the research in this area has focused on settings in which platforms can set prices themselves (Cohen et al. 2020). However, we consider an increasingly important setting where third-party sellers, rather than the marketplace operator, decide whether to offer a product on the marketplace and make their own inventory and pricing decisions. In such a setting, where the marketplace does not directly control the assortment, supply levels, and prices offered by sellers, the way products are presented to customers becomes critical. Marketplaces use assortment curation algorithms to present customers with the best options and influence sellers' offerings and prices. Therefore, it is important to understand the potential impact of these algorithms on customer and seller behavior.

While the assortment planning literature typically focuses on customer choices *between different products*, an important feature of our setting is that multiple sellers can offer the same product. Our work focuses on the buybox tool, which facilitates the choice *between different sellers of the same product*, adding a new dimension to the existing literature.

2.2.3. Related Algorithmic Tools Used in Marketplaces Our research is closely related to empirical work studying the use of algorithmic tools in marketplaces. A buybox algorithm has similarities to other tools commonly used in marketplaces, such as search ranking algorithms, sponsored search mechanisms, recommender systems, reputation systems, reviews, and platform endorsements. It also presents important differences from each of these tools, which we conceptualize in this section.

Most marketplaces provide search functionality that allows customers to enter a query and receive a ranked list of results. By ranking the retrieved results based on estimated relevance, the marketplace can significantly reduce search costs. Previous research has explored the role of the **search algorithm** in ranking these results. For example, De los Santos and Koulayev (2017) estimate a search model and propose a method to rank search results in a way that reduces customer search costs.

Like search ranking algorithms, a buybox algorithm scores different options and presents customers with the best possible choice. However, there are important differences between the two, offering complementary functionality. A search ranking algorithm primarily helps customers find products, while a buybox algorithm helps customers select a specific listing among the sellers offering a product. A buybox algorithm operates within products, comparing and scoring different offerings for the same product. Unlike search algorithms, which present an ordered list to customers, a buybox algorithm selects a single offering as the default option, which becomes the default choice if the customer decides to purchase without further search. In this way, a buybox algorithm simplifies the decision-making process for customers and acts as a nudge (Thaler and Sunstein 2021) that influences customer behavior without restricting their choices.

Sponsored search is another commonly used tool in marketplaces to direct customers to sellers. Sellers can pay the marketplace to be listed in prominent parts of the site. Previous research has studied sponsored search tools, yielding mixed results regarding their effectiveness. For instance, Joo et al. (2024) found that customers prefer organic listings over sponsored ones, but sellers benefit from using sponsored search to increase their visibility on the site. Similar to a buybox algorithm, sponsored search can steer customers toward specific sellers. However, a crucial difference is that with a buybox, there is no monetary payment from the seller to the platform operator. The implications for customers can also vary significantly. Moshary (2021) finds that customers exposed to sponsored searches purchase more expensive sponsored items. Notably, our study finds that increased competition within a buybox is associated with lower average prices paid by customers. While many marketplaces feature a combination of sponsored links and a buybox algorithm, our focal marketplace was not conducting sponsored searches during the period of analysis.

Another example of tools commonly used in marketplaces is **recommender systems**. Recommender systems help users find products that may be relevant to them, given their past behavior and the behavior of other customers. Previous work has explored the role of recommender systems. For example, Hosanagar et al. (2014) find that recommender systems increase the sales of the most popular items, reducing sales diversity. Other studies have investigated the effect of recommender systems on the intensity of competition between brands and the exposure of products to more customers (Li et al. 2018 and Li et al. 2020). Like recommender systems, a buybox algorithm tries to help customers find more relevant options, but a key distinction is that while recommender systems typically aid product discovery, a buybox algorithm operates at the product level and simplifies seller selection for a given product.

Reputation systems and reviews are typically provided by marketplace operators (Dellarocas 2003). They help reduce transaction risks by collecting information about products and sellers and creating a rating. Customers can use this information to inform their decisions in the marketplace. The buybox algorithm also uses this information to ensure a reputable seller is selected as the default seller.

Finally, another tool related to a buybox algorithm is the **platform endorsement** mechanism. The marketplace operator can create a badge that denotes an endorsement of a listing by the operator. If customers trust the operator, these badges can increase the visibility and attractiveness of the endorsed options. For example, Amazon uses the “Amazon Choice" badge to help customers discover high-quality products that other customers frequently choose for similar shopping needs. Bairathi et al. (2022) studied one of these platform endorsement tools using a field experiment on an online freelance platform and found that exposure to platform endorsement increases user searches and purchases both for endorsed and non-endorsed services, which they attribute to increased perception of seller quality. Like a buybox algorithm, the marketplace uses an algorithm (typically private) to grant the badge and select high-value options. Customers can use the endorsement information to guide them in their product or seller selection process. However, unlike a buybox, this endorsement does not select a single seller to be the default option for a product.

2.3. Conceptual Framework: Buybox vs Related Algorithmic Tools

The comparison with related marketplace tools helps us conceptualize the distinctive features of a buybox algorithm. This tool assists customers by organizing listings from different sellers of a product from a standardized catalog rather than requiring navigation through listings of different products. This algorithm automatically assigns a score to each seller’s listing participating in a product’s buybox using a proprietary rule that balances various attributes (including price). The marketplace then features the listing with the highest score as the default option for that product. This streamlines the customer experience and reduces search costs, as customers need not explore multiple listings for the same product unless they choose to. Sellers do not make side payments to the marketplace operator to be treated preferentially by the algorithm; instead, the marketplace assigns scores based on attributes it considers desirable for customers, encouraging sellers to compete more intensively in these areas.

There is limited evidence on the impact of such a mechanism on the adopting marketplace. To our knowledge, the only empirical study exploring a similar mechanism is Dinerstein et al. (2018). Their study focuses on a marketplace from around 2011 that relied heavily on auctions. This marketplace implemented an algorithm to rank seller’s listings by price, helping customers compare between the different options. However, the algorithm did not select a leading option as a default, nor did the marketplace provide a standardized product catalog or listings. Several features of contemporary buybox implementations were absent. Our research investigates a representative contemporary buybox implementation, contributing to the literature by analyzing its impact on marketplace stakeholders and options.

Based on the buybox algorithm’s distinctive features and related literature on search costs, we propose that the adoption of a buybox algorithm will positively impact the marketplace through two interconnected mechanisms: a friction-reducing mechanism and a competition-inducing mechanism.

We hypothesize a *reduction in frictions* because a buybox algorithm significantly lowers search costs for customers who accept the default option selected by the buybox tool, with other options available at higher search costs. What once required comparing various options is now simplified—not merely through reducing the direct effort of additional clicks, but by mitigating the cognitive load imposed by processing extra information. Because the buybox-selected offer has desirable attributes, customers are more likely to find satisfactory products, potentially improving conversion rates.

Drawing on the literature on limited attention (Kahneman 1973, DellaVigna 2009), customers are likely to engage in heuristic decision-making that minimizes cognitive effort. A buybox streamlines this process by prominently featuring a preferred seller, effectively curating the choice for consumers. This reduces the perceived need for further exploration, as the expected gains from identifying a marginally better option may not outweigh the mental costs of continued search. This framework can also help explain why consumers may not have consistently identified the best seller before the introduction of a buybox. Prior to its implementation, the attention costs associated with evaluating multiple sellers likely deterred exhaustive comparison.

Additionally, the introduction of a buybox algorithm reduces seller-side frictions by automating tasks such as creating product descriptions and other content. As stated earlier, for a buybox to be introduced for a product, the marketplace first needs to create a product page that includes product descriptions, images, and specifications. Once this product page is established, sellers can leverage it and simply associate their offerings by providing price and shipping details rather than manually creating the product content when listing the product. By streamlining the listing process, a buybox reduces friction for sellers, minimizing the time and effort required to participate in the marketplace.

We also hypothesize that a buybox will *intensify competition between sellers* offering the same product. This is due to both the reduction of friction and the algorithmic selection of a default buybox listing. Without a buybox, it is harder for customers to compare different listings for the same product, allowing some sellers with inferior offerings to make sales. With a buybox, more sellers are attracted to the site and compete to sell their products. The algorithm automates the comparison of options and selects the one with the most desirable attributes. Sellers with inferior offerings are less likely to be presented as the default option, reducing their sales probability. The algorithmic comparison fosters direct competition between sellers offering the same product, compelling them to provide better terms to increase their chances of being considered by customers.

Our paper provides evidence supporting these two complementary mechanisms. With data from both mobile and desktop channels, we assess the impact of the buybox on search costs. Search costs are particularly high on mobile channels, and the buybox's effect on reducing these costs is evident in our data. We also observe a reduction in transaction prices and an increase in seller quality with an increase in buybox competition, consistent with the competition-inducing mechanism. We posit that these interconnected mechanisms contribute to improved marketplace outcomes, and our findings support these predictions.

In summary, we contribute to the rapidly growing research area on assortment curation algorithms in online marketplaces by investigating the effects of a buybox algorithm on marketplace dynamics.

3. Empirical Setting and Data

To study how algorithmic curation affects marketplace dynamics, we obtained data from one of the largest marketplaces in the world that implemented the buybox during our analysis period. In this section, describe the implementation of the buybox in the focal marketplace, discuss our dataset, and provide summary statistics.

3.1. Focal Marketplace and Buybox Implementation.

Our research partner is the market leader in the studied country, hosting thousands of sellers and millions of customers across a wide range of product categories. Sellers must register with the marketplace to list their products. Upon registration, they create a listing for each product they wish to sell. A listing represents the most granular level of observation in the marketplace. A product can appear in multiple listings by different sellers. Each listing includes a title and product details such as the name, model, description, and images. It also contains seller-related information, such as the seller's name, reputation, and geographic location. The marketplace does not charge listing fees or limit the number of listings; instead, it collects a commission for each transaction.

In June 2019, the marketplace began implementing the buybox, gradually rolling it out across several product categories. Our dataset focuses on a major subcategory in electronics where the buybox was initially implemented for a small number of products and then expanded to others.

To implement the buybox, the marketplace creates a dedicated product page for the buybox of that product. Once the buybox for a product is established, the marketplace notifies sellers that the buybox is available. Sellers offering the product can choose to participate in the buybox and use the standardized product content. In our focal marketplace, as in many others, participation in a product's buybox is voluntary. There is no limit to the number of sellers that can participate in a particular buybox, and participating sellers can maintain an additional listing for the product outside the buybox. They can also opt out of a buybox they are currently participating in. Once at least one seller competes for the buybox, the buybox listing becomes active and visible to customers. Buybox listings tend to be prominently displayed in search results related to the product. Notably, during our observation period, the focal marketplace did not operate a sponsored search program and did not act as a seller in this product category.

As is common in many marketplaces, the specific scoring rule used in the buybox algorithm is not publicly disclosed or known to sellers. However, sellers are aware that the buybox algorithm considers factors such as seller reputation, delivery options, and the price charged for the product, among other elements. Physical proximity between the seller and buyers is an important factor, as delivery options influence the buybox algorithm. Consequently, the selection of the seller offering the best value can be location-dependent, meaning customers in different locations may see different sellers displayed as providing the best value at any given time.

Creating a product catalog and implementing a buybox algorithm across a category is challenging and labor-intensive. For instance, Sadinle et al. (2022) note that Amazon uses information provided by sellers to create product catalogs. However, this often leads to errors, necessitating extensive manual human auditing to ensure the accuracy of the product catalog information. The complex nature of buybox implementation prompted the marketplace to undertake the process sequentially within and across categories. The marketplace implemented the buybox gradually within categories, meaning not all products within a category had a buybox at any given time. In our analysis, we leverage the variation in the timing of buybox activation for different products to estimate the effects of algorithmic curation on the marketplace.

3.2. Data

Our data correspond to a subcategory within the electronics category that holds a significant share of the marketplace. The dataset includes detailed information on all listings, orders, visits, and the timing of the buybox rollout within this subcategory for the period spanning January 22, 2019, to February 25, 2020. The company began implementing the buybox in September 2019.

For each listing in this subcategory, we obtained data on the unique listing ID, the unique seller ID, and the corresponding product ID. Additionally, we noted whether the listing is participating in the buybox and the number of unique visits this listing receives each week.

For each order, we observe the order date and hour, the listing ID of the purchased item, the price of the item, the rebate status, and the magnitude of the rebate, if applicable.[3] We also observe the channel used for the order (desktop or mobile) and the fulfillment method.

We observe the date on which the marketplace created the buybox for each product and the buybox participation for each listing at a weekly level, enabling us to determine the buybox activation time for each product.

Since the buybox algorithm operates within a product, we conduct analyses at the product level by combining corresponding listings for each product. With buybox participation data recorded weekly, we aggregate data at the weekly level. The unit of observation for our analyses is a product-week pair.

We focus our analysis on the relevant listings most directly affected by this algorithmic curation implementation. While sellers can list new or used products in the marketplace, used products (accounting for 8.9% of the orders in this category) are all unique and are not eligible for the buybox. Consequently, we exclude listings pertaining to used products from our analyses. Some listings lacked information on the product sold due to missing data. Of all orders placed in this subcategory, 5.05% do not contain product information. We remove these listings from the analyses. The marketplace actively combats against fraudulent listings and promptly flags suspicious listings. We remove orders for listings flagged as fraudulent. Additionally, we

[3] The marketplace sometimes gives rebates to sellers in exchange for reducing the price of a product below a certain price level. These rebates are presented as discounts to customers. We observe which orders there was an active rebate for and the magnitude of the rebate.

excluded orders with prices significantly deviating from the norm, which could indicate errors or fraudulent activity. We calculate the 1st percentile and 99th percentile paid prices for each product. For each product that sold at least 100 units, we exclude orders with prices below the 1st percentile and above the 99th percentile. Our results remain consistent if we use alternative cutoffs, such as below the 0.1st percentile and above the 99.9th percentile prices. Finally, to focus on popular and relevant products, we restrict our sample to products that account for 90% of the orders for the subcategory.[4]

3.3. Summary Statistics

We aggregate our data at the product-week level. Our final dataset covers 57 weeks and includes 487 products and 24,976 observations (i.e., product-week pairs). Table 1 presents the summary statistics for key metrics.

	Definition	Variable	Mean	St.Dev	Median
Orders	Number of orders for a product	#	66.5	211.7	15
Conversion	Orders divided by Visits	%	0.83	1.03	0.58
Buybox	Buybox of a product is active	binary	0.16	0.36	0.00
NumSellers	Number of sellers selling a product	#	44	80.2	20
NumBuyboxSellers	Number of Buybox sellers selling a product	#	23.6	41.2	13
Avg.Price	Average price of products sold	Currency	1,020.5	854.1	792
Avg.Quality	Average seller quality level of products sold	#	6.5	1.4	6.9

These statistics correspond to 24,976 product-week observations except for NumBuyboxSellers. NumBuyboxSellers considers only the product-week combinations for which the Buybox is active. Note that if a product has no sales at a particular week, there is no corresponding Avg.Price and Avg.Quality values.

Table 1: Summary Statistics for the Main Variables

During the observation period, the marketplace created a buybox for 458 of the 487 products in our sample. From the sample of 458 products for which the marketplace created the buybox, 225 became active—namely, at least one seller participated in the buybox of that product—during our period of observation. Figure 2 shows the number of buyboxes created by the platform and the number of products with an active buybox for each week during the analysis period.

4. Evaluating the Impact of the Buybox

In this section, we describe our overall empirical approach and analyze the impact of buybox activation on marketplace activity, specifically conversion rates and orders.

4.1. Econometric Specification and Results

To estimate the impact of buybox activation on the outcomes of interest, we use its staggered implementation across different products. Specifically, we employ a generalized difference-in-differences (DiD) model with product and time fixed effects. As discussed in Section 3.2 and illustrated in Figure 2, the marketplace gradually introduced the buybox for various products within the focal subcategory, with the intention of

[4] Our results are robust to using alternative cutoff values.

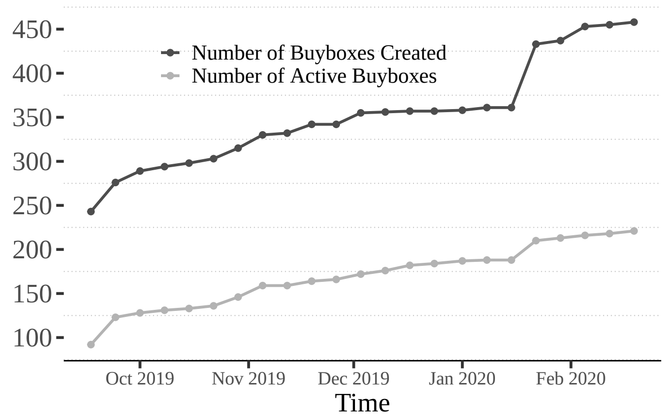

Note: This figure plots, for each week, the number of products for which the marketplace had implemented the buybox and the number of products for which the buybox was active.

Figure 2: Number of Created Buyboxes and Number of Active Buyboxes by Week

eventually implementing it for all products. The varying activation times of the buybox provide the necessary variation to estimate its effect on conversion rates and orders. We use the following specification:

$$Conversion_{p,t} = \alpha + \beta_1 Buybox_{p,t} + \gamma_p + \eta_t + \rho_p \times t + \theta OnRebate_{p,t} + \epsilon_{p,t} \tag{1}$$

Our preferred dependent variable is the conversion rate of product p during week t. To estimate the effect of the buybox on conversion rates, we include an indicator variable, $Buybox_{p,t}$, which equals 1 if the buybox for product p is active (i.e., at least one seller offers product p through the buybox) in week t, and 0 otherwise. Unlike traditional DiD applications, in our setting the treatment status $Buybox_{p,t}$ might change during the post-treatment period. For example, product p might have a buybox active in week $t-1$, but if no seller participates in week t, $Buybox_{p,t}$ would be 0 because the buybox would no longer be active in week t. These instances are rare, occurring in only 0.9% of cases.

We include product fixed effects (γ_p) to control for time-invariant, product-specific effects, and week fixed effects (η_t) to control for seasonal factors common to all products. Additionally, we account for product-specific, time-varying popularity by including product-specific linear time trends ($\rho_p \times t$).

We include the $OnRebate_{p,t}$ indicator variable, which denotes whether a rebate was applied to any order of product p during week t. Including $OnRebate_{p,t}$ controls for promotion-related factors that might influence product demand. We cluster standard errors at the product level to account for arbitrary correlations of observations for the same product. We estimate a similar specification for orders, using the logarithm of orders as the dependent variable for easier interpretation. The results remain robust when using levels.

We present the results of our analyses in Table 2. The coefficient of *Buybox*, our variable of interest, is positive and statistically significant for both dependent variables. These results indicate that the buybox activation is associated with increases in both conversion rates and orders. Specifically, the buybox activation

	(1) Conversion	(2) ln(Orders)
Buybox	0.194**	0.579***
	(0.062)	(0.114)
Observations	24,976	24,976
R^2	0.36	0.69

All columns include product FE, week FE, product-specific linear time trends, and OnRebate controls.
Robust standard errors are in parentheses, clustered at the product level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Effect of Buybox on Conversion Rates and Number of Orders

is associated with a 0.19 percentage point increase in conversion rates (approximately a 23% increase) and a 78% increase in the number of orders.

The large increase in orders observed with the activation of buybox cannot be immediately interpreted as a direct causal effect, as other factors may be involved. For example, if parts of the marketplace prioritize buybox listings—such as by displaying them higher in search results—this could lead to increased visits and orders through a different channel. In this case, the coefficient of interest would capture the combined effect of using the buybox mechanism and the contemporaneous actions taken by the marketplace, which are typical of a buybox implementation. In contrast, the conversion rate, which is our primary metric of interest, is less likely to be systematically influenced by contemporaneous actions because it normalizes sales per visit. This makes it a more reliable indicator of the tool's direct impact and allows for a clearer interpretation of the coefficients.

To further substantiate the directional impact of buybox on orders, we present two additional analyses. First, we demonstrate that an increased number of sellers participating in the buybox enhances its positive effect on orders. This finding suggests that the intensity of the buybox treatment is crucial, as greater participation is linked to a more significant impact on orders. Second, we highlight that the positive effect of buybox on orders is greater on the mobile channel compared to the desktop channel. The differences observed across channels, which align with a reduction in search frictions, offer a more direct measure of buybox's effects, assuming that any other contemporaneous actions impact both channels similarly. Additionally, we find a larger impact on orders during weekdays and work hours when customers are more time-constrained, indicating that the buybox effectively reduces search frictions.

4.2. Intensity of Treatment

As the number of sellers in a buybox increases, the treatment becomes more intense, potentially amplifying the buybox's impact. Sellers might have greater incentives to offer better terms to beat competitors and be selected as the default seller. Similarly, more sellers in a buybox increase the pool of competitive sellers, allowing the marketplace to choose a seller with better offerings. The better terms offered in a buybox can attract more customers and increase conversion rates. To examine the nonlinear relationship between the number of sellers in a buybox and the outcomes, we group each buybox into quartiles based on the $NumBuyboxSellers$ variable. We also create a separate category for instances in which the buybox has a single seller. In this case, the seller faces no competition, and the marketplace has only one listing to choose from for the buybox. We use the following regression to estimate the relationship between the number of sellers and the outcomes:

$$\begin{aligned} Conversion_{p,t} = \alpha + \beta_0 NumBuyboxSeller1_{p,t} + \beta_1 NumBuyboxSellerQ1_{p,t} + \\ \beta_2 NumBuyboxSellerQ2_{p,t} + \beta_3 NumBuyboxSellerQ3_{p,t} + \beta_4 NumBuyboxSellerQ4_{p,t} + \\ \gamma_p + \eta_t + \rho_p \times t + \theta OnRebate_{p,t} + \epsilon_{p,t} \end{aligned} \tag{2}$$

The dependent variable is the conversion rate of product p during week t. The independent variable $NumBuyboxSeller1_{p,t}$ is equal to one if there is only one seller in the buybox of product p on week t. $NumBuyboxSellerQ1_{p,t}$ takes value one if the number of sellers participating in the buybox of product p during week t is in the first quartile. We similarly defined $NumBuyboxSellerQ2_{p,t}$, $NumBuyboxSellerQ2_{p,t}$, and $NumBuyboxSellerQ4_{p,t}$.[5] We control for product fixed effects, week fixed effects, and product-specific linear time trends and include the $OnRebate$ variable.

As shown in Figure 3, the impact of the buybox on conversion rates and orders increases with the number of participating sellers. This suggests that the level of competition in the buybox influences its overall impact. Additionally, when only one seller participates, the buybox does not significantly impact orders and conversion rates. This result is reassuring: with a single seller, the algorithm lacks alternatives to curate from, resulting in a muted impact of the buybox.[6]

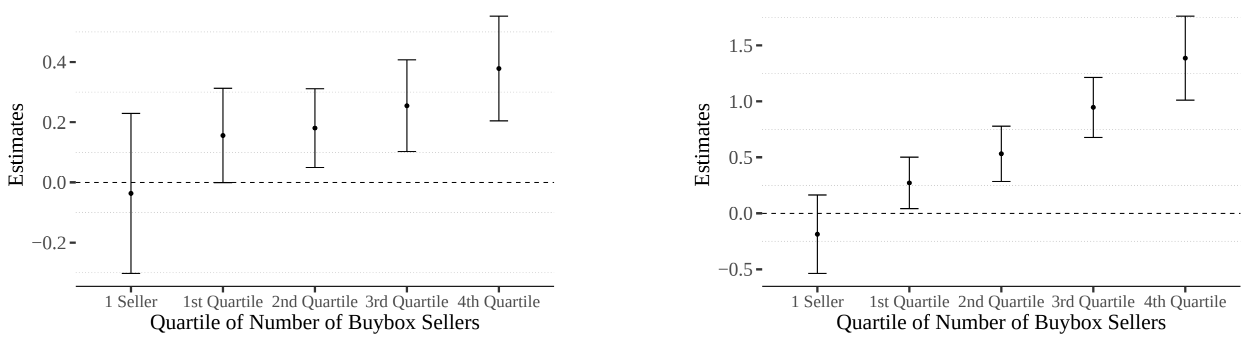

(a) Conversion **(b) ln(Orders)**

Figure 3: Nonlinear Effect of Number of Buybox Sellers

[5] The number of buybox sellers in the first, second, third, and fourth quartiles are in the (1,6], (6,13], (13,26], and (26, ∞) intervals, respectively. The base category corresponds to observations for which $Buybox_{p,t}$=0.

[6] As a separate robustness, we also included the number of sellers as a control to Equation 2, and our results are robust to the inclusion of this control.

4.3. Customer-Side Frictions

In this section, we explore whether the activation of the buybox reduces customer search friction and whether this reduction increases the number of orders.

Reducing search friction and enabling customers to find products smoothly is crucial for both offline and online retailers. In a marketplace, multiple sellers can offer identical products under different listings. Unlike other platforms (e.g., Airbnb or Tinder), a greater number of listings in product marketplaces does not necessarily indicate a broader assortment. This creates unnecessary search friction for customers, who must navigate various listings for the same product that offer little or no additional value while incurring substantial search costs.

Curation algorithms such as a buybox algorithm can alleviate search frictions by curating the best value alternatives and streamlining the purchasing process. The relevance and impact of such curation efforts may vary across retail settings. Prior research reveals that customers experience higher search costs on mobile channels compared to desktop channels, potentially due to smaller screen sizes or differences in internet speed (Ghose et al. 2013, Jin et al. 2019, Gallino et al. 2023, Jain and Tan 2022). Consequently, if a buybox algorithm helps reduce search frictions, its impact should be larger on mobile channels.

In our setting, mobile and desktop channels are designed identically in terms of interface and algorithms but differ in their level of search friction. To examine the buybox's effect on search friction, we use order-level data to observe the channel in which an order was placed. We compare the impact of the buybox on the number of orders separately for the mobile and desktop channels. We estimate a specification similar to Equation (1) with the addition of channel fixed effects and the interaction between the buybox and the mobile channel, where the unit of observation is at the product-week-channel level.

The results of this analysis are presented in Column (1) of Table 6 in Appendix A. Further details about the specification are also provided in the Appendix A. The coefficient for the interaction term "Buybox $\times$ Mobile" is positive and significant, indicating that the positive effect of the buybox on the number of orders is 9.5 percent higher on the mobile channel than on the desktop channel. This finding supports the hypothesis that the curation algorithm reduces search friction, and induces customers to place more orders.

In addition to exploring channel-level impacts, we conducted two additional analyses to examine how the effect of the buybox on orders varies by the time of day. First, we investigated whether the effects of the buybox on orders are more pronounced during weekdays than on weekends. Second, we assessed whether the effects are larger during work hours compared to non-work hours.

If a buybox effectively reduces customers search frictions, its impact on orders should be greater during periods with higher search friction. Customers are likely more time-constrained during work hours and weekdays than during non-work hours and weekends. Therefore, we expect the buybox to have a larger impact during work hours and weekdays. Our results confirm this hypothesis, showing that the buybox has larger effects on orders during work hours and weekdays. These findings provide further evidence that the buybox reduces search frictions. The details of these analyses are also presented in Appendix A.

Overall, the results on customer search frictions and the intensity of the buybox competition highlight that the buybox plays two primary and complementary roles: it reduces search frictions and induces competition among vendors, which improves the offering to customers.

4.4. Robustness and Sensitivity Analysis

In our setting, the buybox is created and activated in a staggered manner, which we leverage to estimate a generalized difference-in-differences model. The marketplace may strategically introduce the buybox across different products, and sellers may strategically decide whether to participate in the buybox of a product. In this section, we summarize several robustness analyses we conducted, with more details reported in Appendix B. Collectively, these analyses support the robustness of our findings.

First, we test for pre-treatment parallel trends using event study analyses (Autor 2003). These analyses show that the increases in conversion rates and orders following the buybox activation cannot be attributed to preexisting differential trends between treated and untreated products. Thus, the observed increases cannot be explained by the fact that products with an active buybox were already experiencing a surge in orders or conversion rates prior to treatment. The details of this analysis are provided in Section B.1.

Second, we assess whether there is any increase in orders or conversion rates when a product's buybox is created but not activated. Since created but not activated buyboxes are invisible to customers, their creation should not produce any observable changes in the marketplace. However, if the rollout timing for various products correlates with increased conversion rates and order numbers, we would expect to see increases in these metrics even when the buybox is created but not activated. Our findings show no such increase, which, along with the null pre-treatment trends result, suggests that the results should be attributed to buybox activation. The details of this analysis are provided in Section B.2.

Third, we repeat our analyses after removing products for which the buybox was never created during our period of analysis to ensure that the results are not influenced by these products. The details of this analysis are provided in Section B.3.

Fourth, we implement a matching algorithm to increase the similarity between treated and untreated products. By including covariates representing seller characteristics, order distribution, and prices before the first treatment in the marketplace in the matching algorithm, we aim to create a comparable set of treated and untreated products with a similar likelihood of getting an active buybox. By balancing these covariates, we enhance the similarity between treated and untreated products and control for factors that might influence the marketplace's or sellers' decisions to create or participate in a buybox. We estimate our main specification in the matched sample and find that the results are robust and similar to those in the full sample. The details of this analysis are provided in Section B.4.

Fifth, to account for sellers' strategic buybox participation decisions, we implement an instrumental variable (IV) approach. We use the intensity of sellers' buybox participation in other products as an instrument for the buybox participation in the focal product. Specifically, we calculate this instrument by identifying the other products offered by sellers of the focal product, counting how many products are eligible and competing for the buybox, and then calculating the competing-to-eligible products ratio (excluding the focal product) at a weekly level. This ratio is used as an instrument for the buybox variable, satisfying both the relevance condition, as it is correlated with the buybox participation, and plausibly the exclusion restriction, as it is unlikely to be related to other patterns directly affecting the focal product. Our IV results align with those from the main specification, suggesting that our findings are not driven by sellers' strategic buybox participation behavior. The details of this analysis are provided in Section B.5.

Sixth, we show that the buybox participation rates do not depend on sellers' order volumes. We first calculated the total number of orders fulfilled by each seller during the pre-treatment period to determine whether a particular seller is prominent in this subcategory. For each seller, we calculated the ratio of their buybox participation in their buybox-eligible products each week and averaged this ratio to create an overall participation ratio. The correlation between the seller's participation ratio and the number of orders fulfilled during the pre-treatment period is -0.017 (p=0.29). These results show that the buybox attracts sellers with varying order volumes at similar rates, indicating no selection into the buybox based on seller size measured by order volume.

Seventh, we conducted a seller-level analysis. We reran our analysis by disaggregating observations to the seller-product-week level rather than the product-week level used in our main specification. This approach allows us to include seller fixed effects and account for sellers' time-invariant characteristics that might influence both their buybox participation decisions and the number of orders they receive. The seller-product-week level analysis shows that participating in the buybox increases the number of orders a seller receives, even after controlling for seller fixed effects. The details of this analysis are provided in Section B.6.

Finally, we test the robustness of our findings using an alternative DiD estimator. Recent literature on staggered difference-in-differences points out that the two-way fixed effects estimator might be inconsistent under treatment effect heterogeneity (Callaway and Sant'Anna 2021). We implement the Callaway-Sant'Anna estimator and show that our results are robust. The details of this analysis are provided in Section B.7.

These analyses collectively reinforce our main findings and address potential concerns regarding the strategic creation of buyboxes by the marketplace and sellers' strategic participation decisions.

In the following section, we examine how the implementation of the buybox influences various marketplace dynamics. Specifically, we explore its impact on prices paid by customers, the quality of purchased listings, seller participation, and marketplace concentration.

5. Effects of the Buybox on Marketplace Outcomes

In this section, we first examine the impact of the buybox introduction on prices paid and the average quality of purchased listings. Next, we analyze the implications for seller participation. Finally, we explore the effects of the buybox on marketplace concentration.

5.1. Average Price Paid

To study whether the average price paid for a product changes after the buybox activation, we estimate Equation (1) where $Avg.Price_{p,t}$ is instead the dependent variable. $Avg.Price_{p,t}$ denotes the average price customers paid for product p in week t. Column (1) of Table 3 indicates that the buybox activation, by itself, does not significantly change the product's average paid price.

However, prices can be affected by the level of competition within the buybox. The level of competition in a product's buybox varies with the number of participating sellers. In a buybox with few sellers, there may be only mild incentives to lower prices to become the best-positioned seller. By contrast, as the number of sellers in a buybox increases, the incentive to reduce prices may intensify. Additionally, with more sellers participating, the curation algorithm has more options to choose from and can select a seller with a better offering. These factors might drive down average prices when more sellers participate in a buybox.

To study this, we explore how the number of sellers in the buybox impacts average prices paid with the following specification:

$$\begin{aligned} Avg.Price_{p,t} = \alpha + \beta NumBuyboxSellers_{p,t} + \beta_2 NumNonBuyboxSellers_{p,t} + \\ \gamma_p + \eta_t + \theta OnRebate_{p,t} + \epsilon_{p,t} \end{aligned} \tag{3}$$

where the dependent variable $Avg.Price_{p,t}$ is the average paid price, and the main independent variable $NumBuyboxSellers_{p,t}$ denotes the number of sellers in the buybox. $NumNonBuyboxSellers_{p,t}$ denotes the number of sellers offering product p in week t who do not participate in the product's buybox. We control for product fixed effects, week fixed effects, and include the $OnRebate$ variable. Unlike previous analyses, we do not control for product-specific linear time trends, as we aim to capture the effect of the change in the number of sellers in the buybox on average paid prices. The results of this analysis are provided in column (2) of Table 3.

We observe that the coefficient for $NumBuyboxSellers_{p,t}$ is negative and significant, consistent with the mechanism that the average paid price decreases as competition in the buybox increases. On average, 10 additional sellers participating in the buybox decrease the average paid price by 4.81 LC (local currency), which corresponds to a reduction of approximately 0.47% in the average paid price.

5.2. Average Quality

One of the metrics the buybox algorithm considers when selecting the best value listing is seller quality. Consequently, in addition to offering better prices for customers, the buybox activation may improve the

	(1) Avg.Price	(2) Avg.Price	(3) Avg.Quality	(4) Avg.Quality	(5) NumSellers
Buybox	12.098		0.110		13.462***
	(9.061)		(0.091)		(2.934)
NumBuyboxSellers		-0.481**		0.005***	
		(0.176)		(0.001)	
NumNonBuyboxSellers		0.071		0.001	
		(0.117)		(0.000)	
Observations	21859	21859	21693	21693	24976
R-sqr	0.97	0.96	0.44	0.34	0.89

We control for product FE, week FE, product-specific linear time trends, and OnRebate in columns (1), (3) and (5), and product FE, week FE, and OnRebate in columns (2) and (4).

Note: Robust standard errors are in parentheses, clustered at the product level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Effect of Buybox on Average Paid Prices and Average Quality Levels

average quality level of products sold in the marketplace by highlighting options from higher-quality sellers. Next, we investigate the impact of the buybox on average quality levels in the marketplace.

The marketplace classifies sellers into eight categorical, color-coded quality levels based on customer reviews and shipment performance. We converted this categorical scale to a discrete numerical scale, where the lowest quality level is 1, and the highest is 8. The average quality level of product p in week t ($Avg.Quality_{p,t}$) is calculated by averaging the seller quality scores for all sold items of product p in week t.

We first explore the impact of the buybox activation on the average quality of purchased products using Equation (1), where the dependent variable is instead the average quality level. Column (3) of Table 3 shows that the buybox activation does not significantly impact average quality levels. Next, following the same logic we used for prices, we explore whether competition levels in a buybox impact average seller quality. When many sellers participate in a buybox, the increased competition may motivate sellers to improve their performance, and the algorithm can select a higher-quality seller for the buybox. If these mechanisms are in play, the average quality of products sold would increase when more sellers join the buybox.

To analyze the effect of the number of sellers in the buybox on quality levels, we use the model described in Equation (3), where the dependent variable is instead the average seller quality level of the products sold ($Avg.Quality_{p,t}$). We control for product fixed effects and week fixed effects and include the $OnRebate$ variable. As seen in column (4) of Table 3, the coefficient for $NumBuyboxSellers$ is positive and significant. This result indicates that increases in seller participation in a buybox are associated with improvements in the average seller quality level for products sold. On average, 10 additional sellers in a buybox improve the average seller quality level of the products sold by 0.05 points, an increase of approximately 0.8%.

Overall, our results indicate that the buybox impacts market outcomes such as prices and quality levels primarily through increased seller participation in the buybox rather than its mere activation. This is evidenced by the significant impact of seller participation in the buybox on average paid prices and quality levels. As the number of participating sellers grows, competition within the buybox intensifies, allowing the algorithm to select from a wider pool of offerings. These findings suggest that a buybox algorithm's effectiveness in decreasing average paid prices and increasing the quality levels depends critically on increasing seller participation in a buybox. While the direct effect of the buybox activation on average paid prices and quality levels is statistically insignificant, this likely reflects the limited variation in the buybox variable as well as in the price and quality variables in our empirical setting.

5.3. Seller Participation

We now examine whether the buybox activation influences the number of sellers offering a product. Implementing the buybox involves standardizing the product catalog and creating product pages with detailed information. This reduces seller friction when listing products, as sellers no longer need to input all the data for the products they offer. Consequently, this may attract more sellers to offer products since the effort required to post them is reduced.

To test the buybox's impact on the number of sellers offering a product, we estimate the model described in Equation (1) where the dependent variable is instead the number of sellers for product p during week t. We calculate the number of sellers offering a product as follows: if a listing posted by seller s for product p receives at least one visit during week t, we consider that seller s is selling product p in week t. We control for product fixed effects, week fixed effects, and product-specific linear trends and include the $OnRebate$ variable. We also cluster the standard errors at the product level.

As shown in column (5) of Table 3, the buybox has a positive and significant effect on the number of sellers. Our results show that when the buybox becomes active for a product, approximately 13 additional sellers begin to offer this product. Thus, the buybox activation attracts sellers to the product, suggesting that the introduction of the buybox reduces marketplace friction on the seller side.

5.4. Impact on Concentration

We now turn our attention to a fundamental aspect of the marketplace: understanding whether the introduction of the buybox influences the distribution of sales among sellers. We investigate this effect from two perspectives: (1) how the sales distribution among sellers of a product changes when its buybox becomes active, and (2) the impact of the buybox on overall sales concentration within the marketplace.

We explore how orders are distributed across sellers by calculating Gini coefficients and Lorenz curves. In a Lorenz curve, the x-axis represents the cumulative percentage of sellers, while the y-axis denotes the cumulative percentage of orders. The Gini coefficient is derived from the Lorenz curve and equals the ratio of the area between the Lorenz curve and the 45-degree equality line to the area under the equality line (see Brynjolfsson et al. 2011 for a detailed discussion). For example, if each seller fulfills an equal number of orders, the Lorenz curve would coincide with the 45-degree equality line, and the Gini coefficient would be 0. As the distribution of orders becomes more concentrated among a few sellers, the Lorenz curve moves away from the equality line, and the Gini coefficient increases. These metrics are widely used in economics to calculate income inequality and have been applied to measure sales dispersion and diversity in retail (Brynjolfsson et al. 2011, Gallino et al. 2017).

In Section 4.1, we showed that when a product's buybox becomes active, the number of orders for that product increases. However, it is not immediately clear whether this increase is evenly distributed across all sellers of the product. To determine how the distribution of sales among sellers changes after a product's buybox becomes active, we examine changes in the product's Gini coefficient for orders.[7] We then estimated the model described in Equation (1), with the dependent variable being instead the Gini coefficient for the orders of product p during week t. As before, we control for product fixed effects, week fixed effects, and product-specific linear trends and include the $OnRebate$ variable. We also cluster the standard errors at the product level.

	(1) GiniOrders
Buybox	0.0212***
	(0.005)
Observations	21,438
R^2	0.70

We control for product FE, week FE, product-specific linear time trends, and OnRebate.
Note: Robust standard errors in parentheses, clustered at the week level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4: Effect of Buybox Introduction on Marketplace Concentration

We present the results of our analysis in Table 4. The coefficient for *Buybox* is positive and statistically significant, indicating a greater concentration of orders among sellers after the activation of a product's buybox. The mean value of the Gini coefficient for orders in our dataset is 0.852. Hence, the estimated effect corresponds to a 2.5% increase in concentration of orders among sellers after the activation of a product's buybox. Combined with our findings in Section 4.1, this result indicates that while the number of orders

[7] We calculated the Gini coefficient for each product and week where there is more than one seller, as the Gini coefficient is always equal to 0 when only one seller is present, making it uninformative in those cases.

for a product increases when its buybox becomes active, this increase is not equally distributed across all sellers.

Next, we analyze the overall impact of the buybox introduction on the marketplace, focusing on whether orders become more concentrated among fewer sellers after the buybox is introduced. For this analysis, we compare sales concentration among sellers before and after the activation of the first buybox in the marketplace.

We restricted the sample to the top-selling sellers who fulfilled 90% of all orders in this category during the pre-treatment period, resulting in 1,020 sellers in the sample.[8] We calculated the total number of orders each seller received per week across the products in the category and calculated a separate Lorenz curve for each week based on the percentage of total orders fulfilled by each seller.

The result of this analysis is presented in Figure 4. This plot indicates that the Lorenz curves for the weeks preceding the first buybox activation lie above the Lorenz curves for the weeks following the first buybox activation. This indicates that the marketplace becomes more concentrated after the buybox introduction, with fewer sellers accounting for the same percentage of sales. The Mann-Whitney test shows that the distributions of orders across sellers for the weeks before the first buybox activation are significantly different from those for the weeks after (p <0.001).

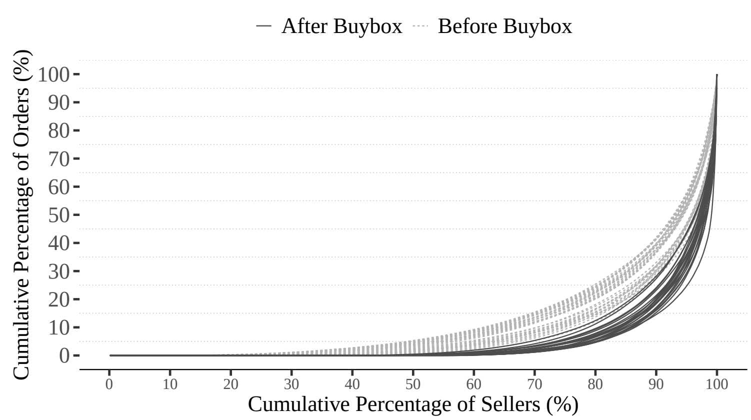

Figure 4: Market Concentration After Buybox Introduction

Finally, as an additional check on how the concentration of orders changed after the introduction of the buybox, we follow Brynjolfsson et al. (2011) and examine whether the relationship between sellers' order ranking (i.e., seller ranking) and the number of orders they receive each week changes after the first buybox becomes active in the marketplace. To this end, we estimate the following regression:

[8] Our results are robust to using an alternative 95% cutoff.

$$
\begin{aligned}
ln(SellerOrders_{s,t}) = \alpha + \beta_1 ln(SellerOrderRank_{s,t}) + \beta_2 AfterBuybox_t + \\
\beta_3 ln(SellerOrderRank_{s,t}) \times AfterBuybox_t + \epsilon_{s,t}
\end{aligned}
\tag{4}
$$

The $ln(SellerOrderRank_{s,t})$ represents the logarithm of the ordinal ranking of seller s among all sellers in the marketplace in week t based on the number of orders in the category. $ln(SellerOrders_{s,t})$ denotes the logarithm of orders seller s received in week t. $AfterBuybox_t$ is an indicator variable that takes value 1 if week t is in the period after the first buybox was activated in the marketplace. The results of these analyses are presented in Table 5.

	(1) ln(SellerOrders)
ln(SellerOrderRank)	-1.550*** (0.018)
AfterBuybox	0.738*** (0.166)
ln(SellerOrderRank)xAfterBuybox	-0.290*** (0.026)
Observations	50,614
R^2	0.84

Note: Robust standard errors are in parentheses, clustered at the week level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Effect of Buybox Introduction on Marketplace Concentration

The negative and significant coefficient of $AfterBuybox \times ln(SellerOrderRank)$ indicates that order concentration among sellers increases following the introduction of the buybox in the marketplace, with a smaller share of sellers receiving a higher percentage of orders. These results support and complement our previous findings from the study of Lorenz curves.

6. Conclusion

The prominence and significance of online marketplaces have surged remarkably in recent years. Traditional retailers like Urban Outfitters and Best Buy have evolved their online presence by establishing marketplaces (Waters 2021). Macy's is reportedly planning a similar move (Wassel 2021). This growing interest in online marketplaces has also driven the rise of startups specializing in marketplace services. For instance, companies such as Mirakl develop and integrate marketplaces into retailers' websites while offering management services. Given this context, it is reasonable to expect continued growth in algorithmic curation within marketplaces.

In this paper, we explore how multi-sided marketplace participants respond to the introduction of the buybox algorithmic curation tool. Our study demonstrates that algorithmic curation tools can benefit marketplace participants significantly. We provide novel empirical evidence on the impact of the buybox tool, highlighting that its activation substantially increases conversion rates and orders. Our findings indicate that the buybox reduces marketplace friction for customers and sellers. For customers, the positive effects of the buybox are more pronounced on mobile channels, consistent with a reduction in search frictions. On the seller side, we observe an increase in the number of sellers offering a product after the buybox implementation. We also document shifts in marketplace outcomes post-buybox: as competition within the buybox intensifies, average paid prices for products decrease, and average seller quality for sold products rises.

Despite these benefits, our findings highlight several potential risks and considerations for marketplace operators. For instance, new sellers may face a “cold start" problem, as they lack established track records of quality metrics needed to be selected as the default seller by the buybox algorithm. This could create entry barriers that potentially hinder marketplace growth over time. Additionally, as the buybox increases concentration at the product level, some sellers may struggle to maintain their market share while preserving sustainable margins since they need to offer competitive prices and maintain high service levels to be selected by the buybox algorithm. Since market thickness is crucial for marketplace health and growth, these concentration effects could affect the long-term sustainability of the marketplace.

The extent of these risks depends heavily on the buybox algorithm design. A strict selection approach that designates a single seller as the default option may amplify concentration issues, while more inclusive policies—such as probabilistic allocation that considers multiple high-quality, competitively priced sellers for the default position—could help mitigate these downsides.

Our findings have multiple implications for practitioners. The observed benefits suggest that a buybox is an effective tool for curating complex assortments, which explains its widespread adoption among leading marketplaces. As traditional retailers increasingly create marketplaces that feature third-party sellers, we anticipate that the use of buybox-like tools will expand beyond global marketplaces to include traditional retailers. Marketplace enablers like Mirakl provide retailers with tools they can leverage to implement a buybox mechanism.

Our research contributes to the literature on algorithms in platforms and retail by providing valuable insights into the impacts of buybox tools on marketplace outcomes. Much of the existing literature either takes a theoretical approach or overlooks the two-sided nature of platforms, where understanding stakeholder reactions to new algorithms is crucial for assessing effectiveness. This comprehensive view reveals both the direct benefits of reduced friction and the more complex competitive dynamics that emerge.

For marketplace operators, it is crucial to recognize that implementing a buybox algorithm is resource-intensive. It requires creating standardized product catalogs and product pages, which can be labor-intensive and costly. Our findings provide a measure of the potential benefits, helping practitioners decide whether to implement this tool. Moreover, our methodology can be used to evaluate the effects of a buybox across various marketplaces and categories and prioritize efforts accordingly.

Future research could expand on our study in several ways. Currently, we focus on a subcategory of electronics. Broadening this scope to include different product categories could provide a more nuanced understanding of the effects of buybox algorithms. We believe that the positive impacts of a buybox may be more pronounced in categories with a high density of sellers per product. By contrast, its effectiveness may diminish in categories dominated by a few major sellers. In categories like apparel, where individual products differ significantly and seller penetration is likely low, the benefits of a buybox may be less evident. The category in our study consists of high-cost items purchased infrequently. The overall impact of a buybox algorithm could be larger in categories with frequent repeat purchases, such as Consumer Packaged Goods. Our relatively short study period limits our ability to assess long-term dynamics. For example, future work examining longer time horizons could better illuminate intertemporal substitution effects. Extended timeframes would also allow researchers to assess how implementing a buybox algorithm shapes marketplace evolution through its effects on seller entry, exit, and customer behavior. Finally, more granular data on a buybox algorithm's selection process could further elucidate seller-side implications of these curation tools.

References

Allon, G., Cohen, M. C. and Sinchaisri, W. P. (2023), 'The impact of behavioral and economic drivers on gig economy workers', *Manufacturing & Service Operations Management* **25**(4), 1376–1393.

Aouad, A. and Saban, D. (2023), 'Online assortment optimization for two-sided matching platforms', *Management Science* **69**(4), 2069–2087.

Autor, D. H. (2003), 'Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing', *Journal of Labor Economics* **21**(1), 1–42.

Bairathi, M., Zhang, X. and Lambrecht, A. (2022), 'The value of platform endorsement', *Available at SSRN 4144605* .

Bimpikis, K., Elmaghraby, W. J., Moon, K. and Zhang, W. (2020), 'Managing market thickness in online business-to-business markets', *Management Science* **66**(12), 5783–5822.

Brynjolfsson, E., Hu, Y. and Simester, D. (2011), 'Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales', *Management Science* **57**(8), 1373–1386.

Brynjolfsson, E. and Smith, M. D. (2000), 'Frictionless commerce? a comparison of internet and conventional retailers', *Management science* **46**(4), 563–585.

Cachon, G. P. and Kök, A. G. (2007), 'Category management and coordination in retail assortment planning in the presence of basket shopping consumers', *Management Science* **53**(6), 934–951.

Callaway, B. and Sant'Anna, P. H. (2021), 'Difference-in-differences with multiple time periods', *Journal of Econometrics* **225**(2), 200–230.

Caro, F. and Sáez de Tejada Cuenca, A. (2023), 'Believing in analytics: Managers' adherence to price recommendations from a dss', *Manufacturing & Service Operations Management* **25**(2), 524–542.

Cohen, M. C., Lobel, I. and Paes Leme, R. (2020), 'Feature-based dynamic pricing', *Management Science* **66**(11), 4921–4943.

Cui, R., Li, M. and Li, Q. (2020), 'Value of high-quality logistics: Evidence from a clash between sf express and alibaba', *Management Science* **66**(9), 3879–3902.

De los Santos, B. and Koulayev, S. (2017), 'Optimizing click-through in online rankings with endogenous search refinement', *Marketing Science* **36**(4), 542–564.

Dellarocas, C. (2003), 'The digitization of word of mouth: Promise and challenges of online feedback mechanisms', *Management science* **49**(10), 1407–1424.

DellaVigna, S. (2009), 'Psychology and economics: Evidence from the field', *Journal of Economic literature* **47**(2), 315–372.

Deshpande, V. and Pendem, P. K. (2023), 'Logistics performance, ratings, and its impact on customer purchasing behavior and sales in e-commerce platforms', *Manufacturing & Service Operations Management* **25**(3), 827–845.

Dinerstein, M., Einav, L., Levin, J. and Sundaresan, N. (2018), 'Consumer price search and platform design in internet commerce', *American Economic Review* **108**(7), 1820–59.

Ferreira, K. J., Parthasarathy, S. and Sekar, S. (2022), 'Learning to rank an assortment of products', *Management Science* **68**(3), 1828–1848.

Gallino, S., Karacaoglu, N. and Moreno, A. (2023), 'Need for speed: The impact of in-process delays on customer behavior in online retail', *Operations Research* **71**(3), 876–894.

Gallino, S., Moreno, A. and Stamatopoulos, I. (2017), 'Channel integration, sales dispersion, and inventory management', *Management Science* **63**(9), 2813–2831.

Ghavami, F. (2024), 'Top 100 global online marketplaces total gmv year-over-year', *Digital Commerce* .
URL: *https://www.digitalcommerce360.com/top-online-marketplaces-data-stats/*

Ghose, A., Goldfarb, A. and Han, S. P. (2013), 'How is the mobile internet different? search costs and local activities', *Information Systems Research* **24**(3), 613–631.

Gur, Y., Macnamara, G., Morgenstern, I. and Saban, D. (2023), 'Information disclosure and promotion policy design for platforms', *Management Science* **69**(10), 5883–5903.

Hosanagar, K., Fleder, D., Lee, D. and Buja, A. (2014), 'Will the global village fracture into tribes? recommender systems and their effects on consumer fragmentation', *Management Science* **60**(4), 805–823.

Iyengar, S. S. and Lepper, M. R. (2000), 'When choice is demotivating: Can one desire too much of a good thing?', *Journal of Personality and Social Psychology* **79**(6), 995.

Jain, N. and Tan, T. F. (2022), 'M-commerce, sales concentration, and inventory management', *Manufacturing & Service Operations Management* **24**(4), 2256–2273.

Jin, F., Gao, F. and Li, J. (2019), 'Impact of free shipping threshold on pc and mobile channels: Evidence from an online retailer', *Available at SSRN 3426616* .

Joo, M., Shi, J. and Abhishek, V. (2024), 'Do sellers benefit from sponsored product listings? evidence from an online marketplace', *Marketing Science* .

Kabra, A., Belavina, E. and Girotra, K. (2016), Designing promotions to scale marketplaces, Technical report.

Kahneman, D. (1973), 'Attention and effort'.

Kaziukenas, J. (2021), 'Amazon tops six million third-party sellers'. (Accessed on 07/26/2022).
URL: *https://www.marketplacepulse.com/articles/amazon-reaches-six-million-third-party-sellers*

Kesavan, S. and Kushwaha, T. (2020), 'Field experiment on the profit implications of merchants' discretionary power to override data-driven decision-making tools', *Management Science* **66**(11), 5182–5190.

Lee, H. S., Kesavan, S. and Kuhnen, C. (2022), 'When do group incentives for retail store managers work?', *Production and Operations Management* **31**(8), 3077–3095.

Li, J. and Netessine, S. (2020), 'Higher market thickness reduces matching rate in online platforms: Evidence from a quasiexperiment', *Management Science* **66**(1), 271–289.

Li, L., Chen, J. and Raghunathan, S. (2018), 'Recommender system rethink: Implications for an electronic marketplace with competing manufacturers', *Information Systems Research* **29**(4), 1003–1023.

Li, L., Chen, J. and Raghunathan, S. (2020), 'Informative role of recommender systems in electronic marketplaces: A boon or a bane for competing sellers.', *MIS Quarterly* **44**(4).

Li, S., Lu, L. X., Lu, S. F. and Huang, S. (2023), 'Estimating the stockout-based demand spillover effect in a fashion retail setting', *Manufacturing & Service Operations Management* **25**(2), 468–488.

Lobel, I. (2021), 'Revenue management and the rise of the algorithmic economy', *Management Science* **67**(9), 5389–5398.

Manshadi, V. and Rodilitz, S. (2022), 'Online policies for efficient volunteer crowdsourcing', *Management Science* **68**(9), 6572–6590.

Moshary, S. (2021), 'Sponsored search in equilibrium: Evidence from two experiments', *Available at SSRN 3903602* .

Rios, I., Saban, D. and Zheng, F. (2023), 'Improving match rates in dating markets through assortment optimization', *Manufacturing & Service Operations Management* **25**(4), 1304–1323.

Sadinle, M., Bouyarmane, K., Galloway, G., Sam, S., Yuan, C. and Tutar, I. (2022), Semi-automated estimation of weighted rates for e-commerce catalog quality monitoring, *in* 'ECML PKDD 2022 Workshop on Learning to Quantify: Methods and Applications (LQ 2022)'.
URL: *https://www.amazon.science/publications/semi-automated-estimation-of-weighted-rates-for-e-commerce-catalog-quality-monitoring*

Thaler, R. H. and Sunstein, C. R. (2021), *Nudge: The final edition*, Yale University Press.

Wang, R. and Sahin, O. (2018), 'The impact of consumer search cost on assortment planning and pricing', *Management Science* **64**(8), 3649–3666.

Wassel, B. (2021), 'Macy's plans online marketplace launch for 2022, reconsiders 60 store closures'. (Accessed on 07/26/2022).

URL: *https://www.retailtouchpoints.com/topics/digital-commerce/macys-plans-online-marketplace-launch-for-2022-reconsiders-60-store-closures*

Waters, M. (2021), 'This is not pandemic boredom': Retailers are increasingly launching their own third-party marketplaces'. (Accessed on 07/26/2022).

URL: *https://www.modernretail.co/retailers/this-is-not-pandemic-boredom-retailers-are-increasingly-launching-their-own-third-party-marketplaces/*

Wooldridge, J. M. (2010), *Econometric analysis of cross section and panel data*, MIT press.

Wooldridge, J. M. (2015), 'Control function methods in applied econometrics', *Journal of Human Resources* **50**(2), 420–445.

Zhang, D. J., Dai, H., Dong, L., Wu, Q., Guo, L. and Liu, X. (2019), 'The value of pop-up stores on retailing platforms: Evidence from a field experiment with alibaba', *Management Science* **65**(11), 5142–5151.

Zhang, W., Elmaghraby, W. J. and Kabra, A. (2024), 'Marketplace expansion through marquee seller adoption: Externalities and quality implications', *Management Science* .

Online Appendix for Algorithmic Assortment Curation: An Empirical Study of Buybox in Online Marketplaces

Santiago Gallino
The Wharton School, University of Pennsylvania, Philadelphia, PA 19104
Email: sgallino@wharton.upenn.edu

Nil Karacaoglu
Fisher College of Business, The Ohio State University, Columbus, OH 43210
Email: karacaoglu.2@osu.edu

Antonio Moreno
Harvard Business School, Harvard University, Boston, MA 02163
Email: amoreno@hbs.edu

A. The impact of the Buybox on Customers' Search Frictions

In Section 4.3, we discuss how the introduction of the buybox may reduce search frictions. In this section, we explore the search friction reduction effect of the buybox by examining its impact on (i) mobile versus desktop channels, (ii) work hours versus non-work hours, and (iii) weekdays versus weekends.

To assess the impact of the buybox on orders placed in the mobile channel, we first calculate the number of orders received by product p, in channel c, in week t, denoted as $\text{OrdersChannel}_{p,c,t}$. We then estimate the following regression analysis where the unit of observation is at the product-channel-week level:

$$\begin{aligned}\ln(\text{OrdersChannel}_{p,c,t}) = \alpha + \beta_0 \text{Buybox}_{p,t} + \beta_1 \text{Buybox}_{p,t} \times \text{Mobile}_c + \\ \delta_p + \rho_p \times t + \phi \text{OnRebate}_{p,t} + \gamma_t + \omega_c + \epsilon_{p,c,t}\end{aligned} \tag{5}$$

We control for product fixed effects, week fixed effects, channel fixed effects, and product-specific linear trends, and include the OnRebate variable. We cluster standard errors at the product level. To explore how the effect of the buybox varies across the mobile and desktop channels, we include an interaction between the buybox variable and the Mobile variable. The variable Mobile_c takes value 1 if channel c is the mobile channel. The result of this analysis is presented in Column 1 of Table 6. The coefficient associated with the interaction term "Buybox × Mobile" is positive and significant, indicating that the effect of the buybox on the number of orders is larger in the mobile channel compared to the desktop channel. The increase in the number of orders after the buybox activation is 9.5 percent higher in the mobile channel compared to the desktop channel.

To assess the impact of the buybox on orders during work hours versus non-work hours, we first calculated the number of orders received by product p, at hour h, in week t, denoted as $\text{OrdersHourOfDay}_{p,h,t}$. We then estimated the following regression where the unit of observation is at the product-week-hour level:

$$\begin{aligned}\ln(\text{OrdersHourOfDay}_{p,h,t}) = \alpha + &\beta_0 \text{Buybox}_{p,t} + \beta_1 \text{Buybox}_{p,t} \times \text{Workhour}_h + \\ &\delta_p + \rho_p \times t + \phi \text{OnRebate}_{p,t} + \gamma_t + \omega_h + \epsilon_{p,h,t}\end{aligned} \tag{6}$$

Hours from 8 AM to 5 PM are classified as work hours, with all other periods classified as non-work hours. We control for product fixed effects, week fixed effects, hour-of-the-day fixed effects, and product-level linear trends, and include the OnRebate variable. Standard errors are clustered at the product level. To explore how the effect of the buybox varies across different times of the day, we include an interaction between the buybox variable and the work-hour variable, where Workhour_h takes value 1 if hour h is between 8 AM to 5 PM. The results of this analysis are presented in Column 2 of Table 6. The coefficient associated with the interaction term "Buybox × Workhour" is positive and significant, indicating that the effect of the buybox on the number of orders is greater during work hours than during non-work hours. The increase in the number of orders after buybox activation is 11.2 percent higher during work hours compared to non-work hours.

Next, we analyzed how the effect of the buybox varies between weekdays and weekends. To assess the impact of the buybox on orders during weekdays versus weekends, we first calculate the number of orders received by product p, on day-of-the-week d, in week t, denoted as $\text{OrdersDoftheWeek}_{p,d,t}$. We then estimate the following regression where the unit of observation is at the product-week-day-of-the-week level:

$$\begin{aligned}\ln(\text{OrdersDoftheWeek}_{p,d,t}) = \alpha + &\beta_0 \text{Buybox}_{p,t} + \beta_1 \text{Buybox}_{p,t} \times \text{Weekday}_d + \\ &\delta_p + \rho_p \times t + \phi \text{OnRebate}_{p,t} + \gamma_t + \omega_d + \epsilon_{p,d,t}\end{aligned} \tag{7}$$

We control for product fixed effects, week fixed effects, day-of-the-week fixed effects, and product-specific linear trends, and include the OnRebate variable. Standard errors are clustered at the product level. To explore how the effect of the buybox varies across weekdays versus weekends, we include an interaction between the buybox variable and the Weekday variable, where Weekday_d takes the value 1 if day d is a weekday. The results of this analysis are presented in Column 3 of Table 6. The coefficient for the interaction term "Buybox × Weekday" is positive and significant, indicating that the effect of the buybox on the number of orders is higher on weekdays than on weekends. The increase in the number of orders after buybox activation is 4.5 percent higher on weekdays compared to weekends.

B. Robustness Analyses

B.1. Parallel Trends

In this section, we probe the plausibility of the parallel trends assumption by examining whether the data exhibit pre-treatment trends. The econometric model described in Section 4.1 allows for a causal interpretation of the effect of the buybox activation on outcomes of interest under the assumption that, in the absence of buybox activation, orders for products with buybox activation and orders for products without buybox activation would have followed parallel trends. The parallel trends assumption is inherently untestable. To assess the plausibility of parallel trends during the post-treatment period, we test for the

	(1) ln(OrdersChannel)	(2) ln(OrdersHourofDay)	(3) ln(OrdersDofWeek)
Buybox	0.462***	0.190***	0.342***
	(0.099)	(0.046)	(0.074)
Buybox × Mobile	0.095**		
	(0.030)		
Buybox × WorkHours		0.112***	
		(0.020)	
Buybox × WeekDay			0.045***
			(0.011)
Observations	49,952	599,424	174,832
R^2	0.69	0.59	0.64

In Column (1), we control for product FE, week FE, channel FE, product-specific linear time-trend, and OnRebate. In Column (2), we control for product FE, week FE, hour-of-the-day FE, product-specific linear time trends, and OnRebate. In Column (3), we control for product FE, week FE, day-of-the-week FE, product-specific linear time trends, and OnRebate.

Note: Robust standard errors are in parentheses, clustered at the product level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6: Effect of Buybox on Customer-Side Frictions

existence of pre-trends. Then, we test whether the outcomes of interest sharply increase shortly after buybox activation. To this end, we estimate the following model, which includes variables capturing the time distance to and from the buybox activation:

$$
\begin{aligned}
Conversion_{p,t} = \alpha + &\sum_{k\in\{1,2,3,4,5,6,6+\}} \beta_{1,k} PreBuybox_{-k(p,t)} + \sum_{k\in\{0,1,2,3,4,5,6,6+\}} \beta_{2,k} PostBuybox_{k(p,t)} + \\
&\gamma_p + \eta_t + \rho_p \times t + \theta OnRebate_{p,t} + \epsilon_{p,t}
\end{aligned}
\tag{8}
$$

where the dependent variable is the conversion rate of product p on week t. The indicator variables $PreBuybox_{-k(p,t)}$ capture the time distance to buybox activation and take the value 1 if week t precedes the buybox activation for product p by k weeks. Similarly, the indicator variables $PostBuybox_{k,(p,t)}$ capture the distance from buybox activation and take the value 1 if week t follows the buybox activation for product p by k weeks.[9] $PreBuybox_{-6+,(p,t)}$ indicates that week t precedes the activation of the buybox for product p by more than six weeks. $PostBuybox_{6+,(p,t)}$ indicates that week t follows the activation of the buybox for product p by more than six weeks. Finally, similar to the main analysis, we control for product fixed effects (γ_p), week fixed effects (η_t), product-specific linear time trends ($\rho_p \times t$), and include the $OnRebate_{p,t}$ indicator variable that accounts for the rebate status of product p on week t. We repeated the same analysis for the logarithm of the number of orders.

We present the results in Figure 5 for conversion and orders. Consistent with the assumption of parallel trends, the coefficients associated with the PreBuybox variables during the six weeks preceding the activation

[9] For instance, if $PreBuybox_{-4,(p,t)}$ is equal to 1, the buybox of product p will be activated on week $t+4$. If $PostBuybox_{4,(p,t)}$ equals 1, the buybox of product p was activated on week $t-4$.

of the buybox are statistically insignificant. Moreover, we observe a sharp increase in the estimated effects during the week of the buybox activation for both orders and conversion rates. This indicates that the impact of buybox activation on these outcomes is immediate.

Conversion is arguably the more relevant of our two primary outcome variables, as it showcases the effectiveness of the buybox in reducing customer friction. But we observe similar patterns for orders. For orders, we also find that the coefficients associated with the PreBuybox variables during the six weeks preceding buybox activation are not statistically insignificant. The coefficient associated with the period preceding the buybox by more than six weeks appears significant. However, this is not concerning because this period covers distant periods relative to buybox activation.

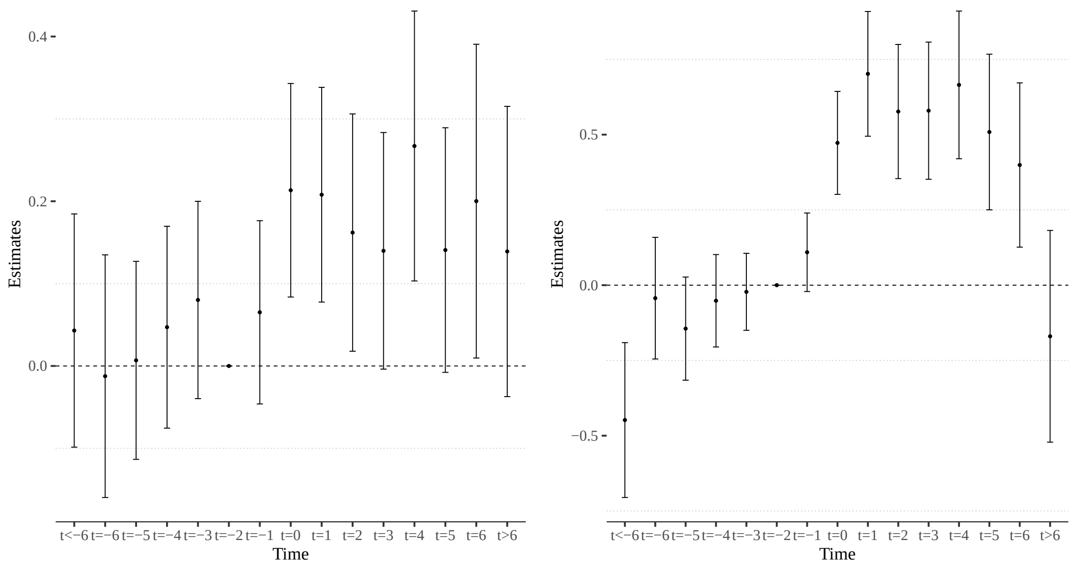

(a) Conversion **(b) ln(Orders)**

Figure 5: Effects Based on Distance to/from Buybox Activation

B.2. Rollout

As discussed in Section 3, product cataloging and buybox creation are labor-intensive processes. Due to this, the marketplace created buyboxes for different products in a staggered way. One can think the order of buybox creation might be endogenous, with the rollout timing for various products being positively correlated with increased conversion rates and orders. If the timing of buybox creation can explain the observed results, we would expect an increase in conversion rates and the orders for a product following the buybox creation, even before the buybox activation. To test for this hypothesis, we estimate the following regression:

$$\begin{aligned} Conversion_{p,t} = \alpha + \beta_1 Buybox_{p,t} + \beta_2 BuyboxCreation_{p,t} + \\ \gamma_p + \eta_t + \rho_p \times t + \theta OnRebate_{p,t} + \epsilon_{p,t} \end{aligned} \tag{9}$$

Similar to our main analysis, our dependent variable is the conversion rate of product p, during week t. To estimate whether buybox creation is associated with an increase in conversion rates for a product, we include an indicator variable, $BuyboxCreation_{p,t}$, which takes value 1 if the buybox of the product p is already created by week t, and 0 otherwise. We also include $Buybox_{p,t}$, which takes value 1 if the buybox of the product p is active at week t. We control for product fixed effects, (γ_p), week fixed effects, (η_t), product-specific linear trends ($\rho_p \times t$), and include the $OnRebate_{p,t}$ variable. Standard errors are clustered at the product level. We repeated the same analysis for the logarithm of the number of orders.

The results of this analysis are presented in Table 7. We observe that the coefficient associated with $BuyboxCreation$ is not statistically significant neither for conversion rates nor for orders. Reassuringly, in line with our main analysis, buybox activation (i.e., the $Buybox$ variable) has a positive and significant impact on both outcomes. Moreover, the coefficients associated with $Buybox$ and $BuyboxCreation$ are statistically different from each other both for conversion (F(1, 486) = 10.68, p= 0.0012) and the number of orders (F(1, 486) = 5.70, p=0.017). This result alleviates the concerns that our findings might be explained by the marketplace's strategic buybox rollout across different products.

	(1) Conversion	(2) ln(Orders)
Buybox	0.220***	0.515***
	(0.064)	(0.121)
BuyboxCreation	-0.051	0.128
	(0.040)	(0.076)
Observations	24,976	24,976
R^2	0.36	0.70

All columns include product FE, week FE, product-specific linear time trends, and the OnRebate control.
Robust standard errors are in parentheses, clustered at the product level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 7: Effect of Buybox Creation on Conversion Rates and Number of Orders

B.3. Excluding Products without the Buybox Option

Our sample includes 487 unique products. The buybox was created for 458 of these products during the study period. We conduct robustness analyses by excluding the products that never had a buybox created and repeating the analysis discussed in Section 4.1 on this sample. These analyses ensure that our results are not affected by these products. The results, presented in Table 8, confirm that the findings from this subset are consistent with our main results.

B.4. Matching

To ensure our analysis accurately estimates the effect of buybox activation, it is crucial that products with an active buybox (i.e., treated) and those without (i.e., untreated) are comparable. A matching approach increases the similarity between treated and untreated products and reduces bias, which can help obtain a more accurate estimate of the effect of the buybox.

	(1) Conversion	(2) ln(Orders)
Buybox	0.196^{**}	0.597^{***}
	(0.063)	(0.114)
Observations	23385	23385
R^2	0.32	0.69

All columns include product FE, week FE, product-specific linear time trends, and the OnRebate control.
Robust standard errors are in parentheses, clustered at the product level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: Effect of Buybox on Conversion Rates and Number of Orders When Product Lines without the Buybox option are excluded

To this end, we use one-to-one propensity score matching without replacement with a caliper of 0.2 to identify the closest untreated product to each treated product. We first estimate a logit regression to calculate the propensity scores as $e(X_p) = P(T_p = 1|X_p)$ where T_p takes the value 1 if the product p ever got an active buybox and 0 otherwise and X_p denotes the covariates of product p. For each product, we use the following five covariates for matching: average number of sellers, average paid price, average Gini coefficient for orders, age of the product on the platform, and average seller size[10] during the pre-treatment period. These covariates, which include details on the product's price, sellers, and age, help us capture product similarity from the customer's perspective and the likelihood of the product' buybox becoming active.

We conduct several checks to ensure that the matching algorithm improves the covariate balance between treated and untreated products. First, we show that matching significantly reduces the standardized mean differences between treated and untreated products (see Figure 6). Second, we use t-tests to assess the similarity of treated and untreated products in the matched sample. The results show that all p-values are greater than 0.05 (the minimum p-value is 0.45), indicating that treated and untreated units in the matched sample are not statistically different from each other.[11]

When we estimate the main specification discussed in Section 4.1 on the matched sample, we find that the results align with our main findings. As shown in the Table 9 below, we observe that conversion rates and orders increase when a product's buybox becomes active.

Finally, in addition to our main matching procedure, we explore several alternative matching specifications. These include adding covariates such as average quality levels, using different caliper sizes (0.1 and 0.25 times the standard deviation of the propensity scores), and implementing one-to-two and one-to-three matching instead of one-to-one. Our findings remain robust.

[10] Seller size is calculated based on order volume during the pre-treatment period.

[11] Results of the t-tests are as follows: t(273.91) = 0.76, p = 0.445; t(229.65) = 0.29, p = 0.770; t(272.09) = 0.13, p = 0.901; t(274.00) = 0.30, p = 0.766; t(249.32)= -0.46, p = 0.643 for average number of sellers, average paid price, average Gini coefficient for orders, age of the product on the platform, and average seller size, respectively.

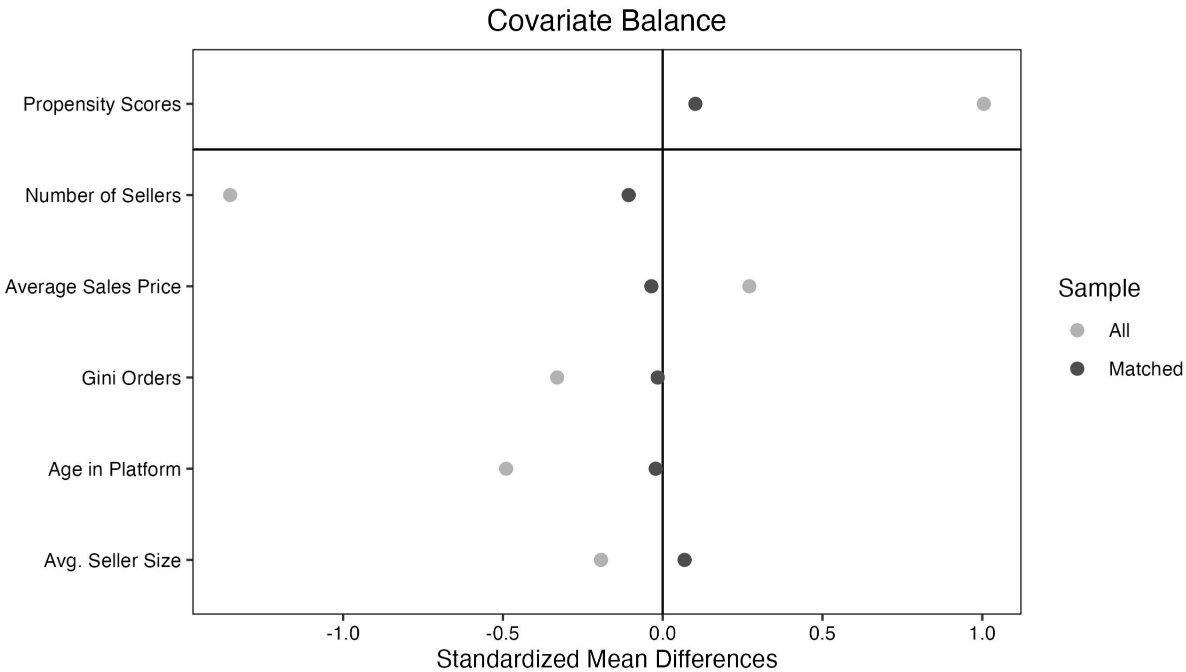

Figure 6: Standardized Mean Differences Before and After Matching

	(1) Conversion	(2) ln(Orders)
Buybox	0.265**	0.424**
	(0.083)	(0.141)
Observations	14,972	14,972
R^2	0.30	0.64

All columns include product FE, week FE, product-specific linear time trends, and the OnRebate variable.
Robust standard errors are in parentheses, clustered at the product level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 9: Effect of Buybox on Conversion Rates and Number of Orders Using the Matched Sample

B.5. Instrumental Variable

In Section 4.4, we showed that the buybox attracts sellers with varying order volumes at similar rates, suggesting that our results cannot be attributed solely to the involvement of top sellers or less successful sellers in the buybox. Moreover, we observe that many sellers who participate in a buybox typically do so for all their eligible products.[12] Next, we address the issue of sellers strategically selecting which products' buyboxes to participate in.

To address concerns about strategic buybox participation, we use an instrumental variable (IV) approach, leveraging sellers' buybox participation behavior for other products as an instrument for the buybox activation variable.

[12] 68.5% of the sellers who participate in the buybox do so for all their eligible products.

Specifically, we define the instrument as follows: for each week t and each product p, we identify the sellers offering the focal product p on week t. For these sellers, we examine the set of other products they offer on week t, excluding the focal product p. We then count the number of these products eligible for the buybox and the number of products for which the seller participates in the buybox on week t. Summing these numbers across all sellers of the focal product, we calculate the ratio of competing products to eligible products for product p on week t, excluding the focal product, referred to as $PercentageCompeting_{p,t}$.

For PercentageCompeting to be a valid IV, it must satisfy two conditions. First, the relevance condition: the PercentageCompeting variable should be correlated with the buybox activation of that product. We find that PercentageCompeting and the Buybox variable are significantly correlated ($r = 0.49,\ p < 0.001$), indicating that products with higher seller participation in other buyboxes are more likely to get an active buybox. Second, the exclusion restriction must also hold: the buybox participation rate of sellers for the focal product's other items should affect the focal product's conversion rates and orders only through its effect on the buybox activation of the focal product. This assumption is plausible because the buybox participation behavior of sellers for other products is unlikely to impact the conversion rates or orders of the focal product and is likely unrelated to other patterns affecting the focal product.

Buybox is a binary variable, which means that using the widely-used two-stage least squares method (2SLS) can lead to inconsistent estimates (Wooldridge 2010). Thus, following Wooldridge (2015), we implement the control function method (CF). CF is similar to 2SLS, but in contrast to 2SLS, it estimates a probit model in the first stage to account for the binary nature of the buybox variable. In the second stage, generalized probit residuals from the first stage are included as a control to account for endogeneity. We use this approach to estimate the effect of the buybox on conversion rates and orders. The details of the equations are as follows:

$$Buybox^*_{p,t} = \alpha_0 + \alpha_1 PercentageCompeting_{p,t} + \alpha X_{p,t} + \varepsilon_{p,t}, Buybox_{p,t} = 1\,[Buybox^*_{p,t} > 0] \tag{10}$$

$$Conversion_{p,t} = \beta_0 + \beta_1 Buybox_{p,t} + \beta_2 \widetilde{Buybox}_{p,t} + \gamma_p + \eta_t + \theta OnRebate_{p,t} + \epsilon_{p,t} \tag{11}$$

where $Buybox^*_{p,t}$ is a latent variable, $1\,[\cdot]$ is the indicator function, and $X_{p,t}$ includes product fixed effects, week fixed effects, and the OnRebate variable. $\widetilde{Buybox}_{p,t}$ is the generalized probit residual for product p on week t for Equation 10.[13]

The results are presented in Table 10. As shown in columns (1) and (2), the effect of the buybox on conversion rates and orders remains robust when using PercentageCompeting as an instrument for buybox.[14]

[13] We do not include product-specific linear trends due to convergence issues in the first-stage probit equation, but our results are qualitatively similar if we use the 2SLS method where we include product-specific linear trends.

[14] The Kleibergen-Paap F statistic is equal to 80, which is larger than the rule-of-thumb threshold of 10. This means that the instrument strongly predicts the buybox variable, and we do not have a weak instrument.

	Conversion	ln(Orders)
Buybox	0.236***	0.976***
	(0.020)	(0.034)
Observations	24976	24976

We control for product FE, week FE, and OnRebate.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10: Instrumental Variable Results

B.6. Product-Seller-Week Level Analysis

The buybox operates at the product level, and buybox participation data is recorded weekly. Therefore, we set the product-week pair as our unit of analysis throughout this paper. In this section, we assess the robustness of our findings at a more granular level of aggregation: the product-seller-week level. Conducting the analysis at this level allows us to control for seller fixed effects in addition to product and week fixed effects. Seller fixed effects allow us to control for sellers' inherent time-invariant differences on our outcomes of interest. While using the product-seller-week level as our unit of analysis enables us to control for seller differences, the estimates are noisier due to the sparser nature of orders at the product-seller level compared to the product level.

To assess the impact of participating in the buybox on the number of orders received by a seller s for a product p during week t, we calculated the number of orders received for each product-seller-week combination and estimated the following model:

$$\ln(\text{Orders}_{p,s,t}) = \alpha + \beta \text{Buybox}_{p,s,t} + \delta_p + \rho_p \times t + \phi \text{OnRebate}_{p,s,t} + \gamma_t + \omega_s + \varepsilon_{p,s,t} \tag{12}$$

$Buybox_{p,s,t}$ variable captures whether seller s participates in the buybox of product p on week t. We control for product fixed effects (δ_p), product-specific linear trends ($\rho_p \times t$), seller fixed effects (ω_s), and week fixed effects (γ_t). Additionally, we include the variable $OnRebate_{p,s,t}$, which indicates whether a rebate was applied to any order associated with product p of seller s during week t. We cluster standard errors at the seller level to account for arbitrary serial correlation in standard errors at the seller level. Our results are robust to clustering standard errors at the product level.

The results of this analysis are given in Table 11. Similar to our main findings, these analyses show that when a seller participates in the buybox of a product, this seller experiences a larger order volume.[15]

[15] We do not replicate seller-level results for conversion rates because the number of visits and the number of orders are sparse at the product-seller-week level. However, the results are noisier but robust.

	(1) ln(Orders)
Buybox	0.491***
	(0.022)
Observations	1,095,725
R^2	0.22

We control for product FE, seller FE, week FE, product-specific linear time trends, and OnRebate.
Note: Robust standard errors are in parentheses, clustered at the seller level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 11: Impact of Buybox Participation at the Seller-Product-Week Level

B.7. Alternative Estimators

Next, we assess the robustness of our findings to alternative estimators. In our main analysis, we leverage the staggered activation of the buybox to estimate a generalized difference-in-differences (DiD) model with product and time fixed effects (i.e., two-way fixed effects). Recent literature on staggered difference-in-differences points out that the two-way fixed effects estimator is consistent under treatment effect homogeneity. To explore possible heterogeneity in treatment effects across treated units and time, we implement the estimator proposed by Callaway and Sant'Anna (2021) as a robustness analysis. The estimator of Callaway and Sant'Anna (2021) shuts down the 2×2 DiD comparisons between already-treated and newly-treated units. It allows us to estimate consistent treatment effects even in the presence of heterogeneous treatment effects. The results of this robustness analysis are presented in Table 12. We observe that our main findings are robust to this alternative estimator.[16]

	(1) Conversion	(2) ln(Orders)
Buybox	0.429**	0.434*
	(0.144)	(0.183)
Observations	24,976	24,976

All columns include product FE, week FE, and OnRebate controls.
Robust standard errors are in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 12: Alternative Estimator for the Effect of Buybox

[16] In this analysis, we control for product fixed effects, week fixed effects, and include the OnRebate indicator variable. Moreover, the estimator requires the treatment status to be weakly increasing (i.e., once a unit is treated, it always stays treated). Hence, we assume that a product's buybox is always active once its buybox becomes active. This assumption is likely innocuous as only in 0.9% of cases did a product with an active buybox in the previous week exit the buybox.