ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN

ECONOMICS

Ciclo 35

Settore Concorsuale: 13/A1 – ECONOMIA POLITICA

Settore Scientifico Disciplinare: SECS-P/01 - ECONOMICS

Essays in Applied Microeconomics

Presentata da: ***Vito Stefano Bramante***

Coordinatore Dottorato | **Supervisore**

Andrea Mattozzi | **Emilio Calvano**

Esame finale anno 2024

Abstract

This thesis investigates the interplay between algorithms, market entry, and bundling practices within digital marketplaces, using Amazon as a central case study. It leverages innovation from machine-learning both as the object of study and as a tool to extract insights in the world of e-commerce. Combined, these papers shed light on how digital platforms and their algorithms impact market dynamics, competition, and regulatory concerns.

The first paper, ***Algorithms in the Wild: Experimental Evidence from an Online Marketplace***, explores the potential for algorithmic collusion in online marketplaces. By deploying an experimental setup using a custom-designed repricing algorithm, the study assesses the competitive effects of off-the-shelf commercial repricers. The findings demonstrate that the commercial repricers are often less sophisticated than expected, and that even simple algorithms can exhibit collusive behaviors under certain market conditions. This analysis highlights how algorithm-driven pricing strategies can distort competition, which is especially concerning in an era where algorithmic decision-making is pervasive across digital platforms.

The second paper, ***Competition on Hybrid Platforms: Evidence from Amazon US***, provides a detailed analysis of Amazon's dual role as both a marketplace and a competitor. It examines the effects of Amazon's entry into product markets where it previously allowed third-party sellers to dominate. The study shows that Amazon's entry tends to lower prices but leads to significant competitive pressure for existing sellers, driving some out of the market altogether. This analysis is crucial for understanding how platform monopolies use their market position not only to facilitate trade but also to control competitive outcomes, shaping market structure in their favor.

Finally, ***Bundling Services in Digital Markets: FTC vs Amazon Inc***. investigates Amazon's logistics practices, particularly the tying of its Fulfillment by Amazon (FBA) service with access to its marketplace's most valuable customer segment—Prime members. Through a comprehensive dataset, the paper shows that sellers who switch to FBA experience marginal improvements in visibility but must lower prices and, in many cases, accept lower revenues. These findings suggest that Amazon's bundling of services creates barriers for third-party sellers, raising important questions about antitrust policy and platform competition.

Algorithms in the Wild: Experimental Evidence from an Online Marketplace

Click here for most recent version.

Vito Stefano Bramante*

Emilio Calvano†

Giacomo Calzolari‡

Maximilian Schaefer§¶

December 14, 2023

Abstract

Can off-the-shelf repricing algorithms used in online marketplaces learn collusive strategies that harm consumers? To shed light on the sophistication of commercial repricing technology, we deploy our own repricing software on an online platform. We implement a EXP3 repricing algorithm and compare its performance against the artificial intelligence algorithm of a selected commercial repricer. We start by establishing a performance benchmark for myopic pricing strategies when faced with a mechanical repricing rule that undercuts rivals' prices. When competing against the mechanical rule, our EXP3 algorithm achieves a better performance than the commercial software. Additionally, our EXP3 algorithms out-competes the commercial repricing software in a direct competition. These results cast doubt on the sophistication of the selected commercial repricing software. Designing algorithms that allow for intertemporal trade-offs is a prerequisite for collusion to arise. In simulations, we show that forward-looking strategies can be learned at low costs. This provides the basis for a more in-depth investigation of forward-looking algorithms, and, hence, collusion in future iterations of this work.

JEL Codes: D21, D43, D83, L12, L13
Keywords: Algorithms, Collusion, Platform Markets

*University of Bologna, vitostefano.bramant2@unibo.it
†University of Rome - Tor Vergata, emilio.calvano@gmail.com
‡European University Institute, giacocalzolari@gmail.com
§Yale University & University of Bologna, maximilian.schaefer@yale.edu
¶This research is funded by research grants from the Center for Technology, Innovation and Competition at the University of Pennsylvania, the Digital Economics Research Network, the European University Institute Research Council Funding, and PRIN 2017.

1 Introduction

The use of repricing software has become ubiquitous in recent years: One example is the case of gasoline markets in which the widespread adoption of repricing software by large companies has attracted the attention of both academia (Assad et al., 2020) and the media.[1] However, the use of repricing software is not limited to large firms. The emergence of peer-to-peer marketplaces such as Amazon, Airbnb, Booking, and eBay has given rise to a rich landscape of software companies that offer affordable off-the-shelf repricing solutions for small sellers.[2]

In the context of online platforms, sheer speed is often advertised as the main benefit of using repricing software: Automatically monitoring and reacting to rival prices allows to relentlessly undercut slower sellers, who set prices manually. This, in turn, allows sellers to capture consumer demand because online platforms prominently feature the seller with the lowest price. While fierce price competition is likely to benefit customers through lower prices, the possibility that repricers might be powered by intelligent algorithms has also raised concerns that they might autonomously learn sophisticated strategies to charge higher prices, i.e. that they might autonomously learn to collude.[3]

The study of Calvano et al. (2020) has provided the proof-of-concept that even simple artificial intelligence algorithms have the capability of sustaining collusion by autonomously learning to punish rivals who deviate from the collusive agreement. This lends credibility to the concerns related to algorithmic collusion, especially since many repricing companies advertise the use of artificial intelligence algorithms.

However, empirical evidence assessing the real-world effects of repricing software remains scarce. While the finance literature has started studying the subject of algorithmic trading already one decade ago, it does not directly speak to collusion (Hendershott et al., 2011; Chaboud et al., 2014). Additionally, it is questionable how findings from sophisticated financial markets, where sellers and buyers both use software, extend to consumer mass markets, where typically only the seller-side uses software. Assad et al. (2020) is among the first studies finding evidence consistent with a chilling effect on competition in a consumer market.

One shortcoming of existing real-world studies is the imperfect ability of researchers to identify the algorithms in use, as companies are naturally reluctant to reveal their pricing technology. This raises questions about the sophistication of the adopted repricing technologies and, ultimately, leaves open the possibility that reported findings might not capture the effect of machine intelligence. For example, it has been pointed out that high prices might be the consequence of the

[1] See https://www.economist.com/finance-and-economics/2017/05/06/price-bots-can-collude-against-consumers and https://www.wsj.com/articles/why-do-gas-station-prices-constantly-change-blame-the-algorithm-1494262674 (last accessed: August 22, 2022).

[2] A Google search for "algorithmic repricer" on August 22, 2022 returned $52,800$ results.

[3] In fact, the possibility that algorithms might have the ability to collude has already drawn the attention of competition authorities around the globe. For example, the topic of algorithmic collusion has been discussed at the 7th session of the FTC Hearings on competition and consumer protection (November 2018). Furthermore, white papers of the OECD (2017) and the British CMA (2021) also discuss the subject.

algorithm's failure to optimize, providing an interesting contrast to the prevailing narrative (Cooper et al., 2015).

Our study sets out to address this shortcoming by leveraging the environment created by peer-to-peer marketplaces, which allows researchers to implement real-world repricing software in a controlled yet realistic environment. To this end, we create two seller accounts on a online marketplace, stock them with goods, and engage in sales activity. In addition, we create our own repricing software, which allows us to create arbitrarily sophisticated opponents against which we let commercial repricing software compete.

One major benefit of our setting is that it enables the controlled implementation of commercial repricing software and the extensive monitoring of the market environment, ruling out possible alternative explanations for increasing prices, such as changes in demand and supply conditions.[4]

In this article, we present the results from an experimental protocol in which we compare the performance of our own algorithm against the performance of a selected commercial repricing company that advertises the use of Artificial Intelligence algorithms.[5] Our own algorithm is a so-called EXP3 algorithm, which belongs to the class of *no-regret* reinforcement learning algorithms. No-regret algorithms guarantee a payoff close to the payoff of the ex-post optimal myopic strategy.

To compare the performance between our own algorithm and the commercial algorithm, we start by characterizing the optimal strategy of a seller when confronted with an opponent that always undercuts her prices. In the online marketplace selected for our experiment, the cheapest seller is prominently advertised to customers arriving on the product page. The strategy to marginally undercut rival prices to be the cheapest seller and to capture consumer demand constitutes a plausible competitive benchmark in the myopic-seller setting we consider. In fact, it is an option that many repricing companies offer as a default repricing rule.

After establishing the optimal response, we let both our own and the commercial repricing software compete against the opponent always undercutting rivals' prices. Our EXP3 algorithm is more successful in learning a strategy close to the optimal strategy than the commercial algorithm. Additionally, our algorithm outperforms the commercial repricing software in direct competition in the sense that it obtains the

[4]The benefits of this research approach, however, are to be weighed against the complications researchers face when dealing and interacting with actual markets. Approaches based on simulations, lab experiments and empirical analysis, developed so far, do not require to interact directly with actual market. We had instead to address a series of issues that are uncommon to academic research.

[5]As appears to be common in the industry, the selected company does not provide information on the type of algorithm used.

advertised seller position more often than the commercial repricing software. These results suggest that the level of sophistication of the selected commercial repricing software is not particularly evolved.

While our experimental protocol provides a measure to benchmark the performance of algorithmic repricers for myopic strategies, the case of algorithms able to implement intertemporal trade offs is not yet covered. The ability to incorporate the future consequences of current decisions is a prerequisite for proper collusion, which requires balancing the trade-offs between the short-term benefits of subverting the collusive agreement with the associated long-term costs of lower prices resulting from the subversive action.

In this context, the main challenge is to develop an algorithm that is capable of integrating intertemporal trade-offs while learning in a reasonable time to avoid incurring prohibitive learning costs. Algorithms that quickly learn strategies that allow for intertemporal trade-offs pose a significant challenge that researchers in laboratory environments can afford to ignore.[6] We present findings based on an algorithm that successfully learns the optimal intertemporal strategy against the mechanical algorithm relentlessly undercutting rivals' prices. This proof-of-concept from the laboratory setting still needs to be validated in a real world environment.

Looking ahead, the case of algorithms capable of implementing intertemporal trade-offs needs to be elaborated further. It remains to be determined whether it is possible to quickly train forward-looking algorithms that can successfully compete against a wider variety of strategies - not only the algorithm relentlessly undercutting rival prices. Our goal is to assess the feasibility of algorithmic collusion in real world market places using an algorithm that learns quickly and that accommodates a wide range of possible rival strategies.

In future work, we aim to deploy a revised experimental protocol, which will cover the myopic and forward-looking case, on a broader variety of repricing companies to provide a more comprehensive overview about the typical capabilities of off-the-shelf repricing solutions offered to small sellers in online marketplaces.

The remainder of the article is structured as follows: Section 2 introduces background information about the online marketplace and the commercial repricer we selected for our experiment. Section 3 introduces our current experimental protocol, the optimal myopic strategy, and the myopic EXP3 algorithm. Section 4 presents the results obtained from the experimental protocol. Section 5 deals with algorithms able to learn intertemporal trade-offs. Section 6 concludes.

[6] For example, Calvano et al. (2020) allow for several hundreds of thousands of periods of learning.

2 Background

The Online Marketplace

Our experiment is implemented in an online marketplace. We create two seller accounts and engage in real sales activity. We chose a cheap to supply product to implement our experiment. When customers search the product, they are displayed a picture of the product together with the price of the advertised seller. The online marketplace does not directly disclose how the advertised seller is selected. However, in the context of our experiment, the seller with the strictly cheapest price was almost certainly the advertised seller as is shown in Figure 1.

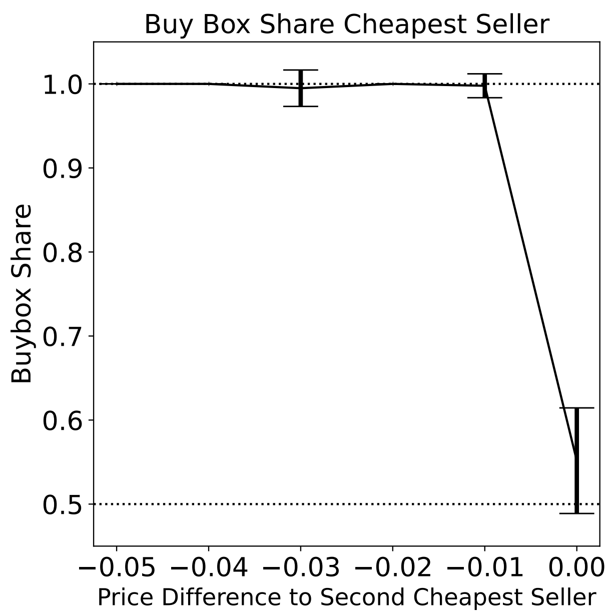

Figure 1: Share of Advertised Seller Position

Notes: The graph shows the share of the advertised seller position for the cheapest seller as a function of the price difference to the second cheapest seller. The two dotted horizontal lines show the 50 and 100 percent thresholds, respectively. The data were obtained from one of our experiments. The error bars show the 95% confidence intervals. We observe no variation for price differences of two and five cents.

Consumers can select the offer of the cheapest seller by clicking on a salient link which allows them to buy the product. The prices of non-advertised sellers can be accessed by clicking on a non-salient link, which leads to the list of non-advertised sellers. There is no reliable information on the aggregate share of consumers ordering products from non-advertised sellers. However, it is clear that being the advertised seller increases the likelihood of sales significantly.

The Commercial Repricing Software

The repricing company we selected for our experiment offers off-the-shelf repricing solutions for several online marketplaces. It offers a variety of repricing features, such as rule based repricers and AI-driven repricers. Typically, rule-based repricers allow to undercut specific rivals such as the advertised sellers, or a designated seller (which can be identified by her seller id).

The company does not disclose any information about the nature of the AI-driven repricing algorithms. However, we were able to select among various options such as AI-driven repricers designed to maximize sales or profits. The results of the present article are based on the AI-algorithm to maximize profits.

The repricer is connected to the seller account through a designated API. In practice, the seller can connect her account to the repricing solution by following a series of easy-to-implement steps. Once connected, the seller can design own mechanical repricing rules or select prespecified mechanical or algorithmic repricing rules. Different rules can be applied to different products.

One noteworthy feature is that *all* repricing rules require the specification of a minimum and maximum price. When the minimum price is reached, the seller has to specify which action the repricer should take, which could be to either stay at the minimum price, or to revert to the maximum price. Throughout the experiment, we instruct the algorithm to revert to the maximum price in case the minimum is reached. The significance of this choice will become clear later.

Own Repricing Solution

We hired a software developer to create the architecture to interact with the seller accounts using the same API as the one used by commercial repricing companies. The software allows us to implement algorithms of arbitrary complexity using near real-time data from the online marketplace. We use the Python programming language to program our own algorithms. The software solution also allows us to closely monitor the market: We obtain a full picture of all prices and the identity of the advertised seller at one minute intervals. To the best of our knowledge, this granularity of data is unprecedented in academic research examining online marketplaces.

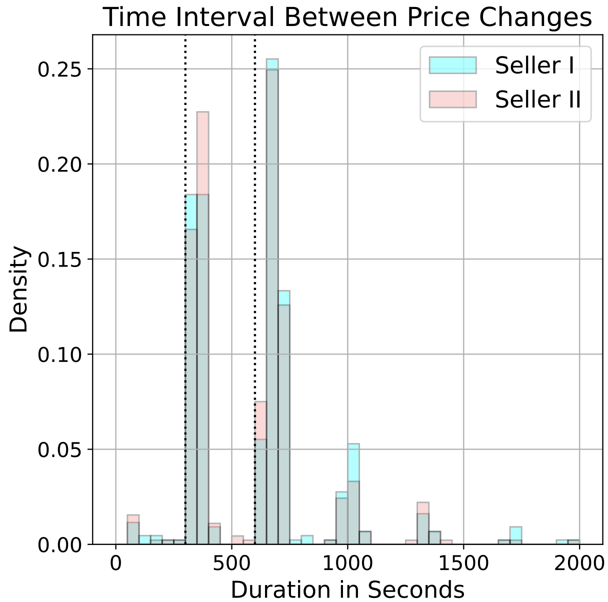

Figure 2: Distribution of Time-Intervals Between Price Events

Notes: The two dotted vertical lines mark the five and ten minutes time intervals, respectively. The data were obtained from one of our experiments.

Latency in Price-Setting

Two important aspects for competition are price-stickiness and whether rivals implement actions simultaneously or sequentially. Our data reveal that the typical time-interval between any two price changes is either five or ten minutes. Hence, once a new price is set, the seller is committed to this price for a given amount of time. Figure 2 shows the distribution of the time interval between price changes for both sellers. Our data also reveal that the prices of all sellers for a given product are typically updated simultaneously.

Price-stickiness is important because it implies that a seller who only reacts to price changes will need five minutes to implement a new price. This provides an advantage to first-movers. For example, player A who sets a new price p_t^a and who knows that rival B only passively reacts by undercutting the new price by an amount x such that $p_{t+1}^b = p_t^a - x$ can anticipate the rival move and instruct that $p_{t+1}^a = p_t^a - (x + \epsilon)$. By repeating this strategy $\forall t$, player a can ensure that she always has the cheapest offer and, thus, is the advertised seller. The minimum increment for price changes is one cent, i.e. we have that $\epsilon = 0.01$.

3 Experimental Design

In this Section, we describe our experimental approach to test the sophistication of the commercial repricing software. We start by introducing the overall design which relies on a benchmark strategy. Subsequently, we will discuss the myopic best response to the benchmark strategy and provide a detailed description of our own artificial intelligence algorithm.

Benchmark Strategy and Overall Design

The goal of the project is to assess the sophistication of commercial repricing software. Naturally, this assessment has to be done taking into consideration the environment within which the commercial repricing software will be deployed. One key feature of our online marketplace is that, by prominently displaying the cheapest seller no matter how small the price advantage, it implements an almost ideal version of Bertrand competition.

Assuming, for simplicity, that i and j denote the two cheapest sellers, we have

$$\Pi_i(p_i, p_j) = \begin{cases} p_i - c & \text{if } p_i < p_j \\ \frac{1}{2}(p_i - c) & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j \end{cases}, \tag{1}$$

where we assume that the position as the advertised seller is equally shared in the scenario of price parity (Figure 1 corroborates this assumption).[7] Note that Equation (1) assumes symmetric marginal costs; we will maintain this assumption throughout the following exposition.

The Bertrand environment provides a strong incentive to undercut rivals. Thus, it appears natural that a seller would instruct a mechanical repricer to undercut the cheapest rival price by one cent while simultaneously setting a price floor which captures the marginal costs. In other words, it appears natural that a seller would instruct a repricer to implement the Bertrand strategy. Thus, our benchmark strategy is defined as the Bertrand reaction function (p_{-i} denotes the set of rival prices):

$$BR_i(p_{-i}, c) = max\{min\{p_{-i}\} \ - \epsilon, c\}. \tag{2}$$

[7] We tested this assumption more rigorously by analyzing how the advertised seller position is split between sellers in an experiment in which both sellers set the same price for a long period of time: The advertised seller position is nearly perfectly split between both sellers, as is shown in Figure 9.

The general approach of our design is to establish a best response to the Bertrand strategy and to assess the performance of the commercial repricer and our own repricer relative to this best response. This approach provides a metric for the performance of repricing software, which is derived from a sensible myopic benchmark strategy.

Myopic Best Response to Benchmark Strategy

The myopic best response strategy to the Bertrand reaction function exploits the price-stickiness discussed in Section 2 and the ensuing advantage for first-movers. The strategy we will introduce can be thought of as the best response of a strategic but myopic seller, who wants to ensure herself a large share of the advertised seller position, even in the event when she is confronted with a seller implementing a Bertrand strategy.

We call the best response to the Bertrand strategy the *relentless cycling strategy*. The relentless cycling strategy consists of continuously lowering the price of the seller implementing the strategy by $\epsilon + 0.01$. By continuously lowering her own price, the seller exploits the price latency and the fact that the Bertrand strategy is only reactive: In period t, the Bertrand strategy will implement a price change such that $p_{t+1} = p_t - \epsilon$, while the relentless cycling strategy will implement a price change such that $p_{t+1} = p_t - (\epsilon + 0.01)$. This way, the player with the relentless cycling strategy will have a by one cent lower price and obtain the advertised seller position.

Eventually, the relentless cycling strategy will reach the minimum price, which will lead to a price-reset to the allowable maximum price. As a result, the seller will loose the advertised seller position. However, the Bertrand strategy will instruct to follow the price increase, which will re-initiate the downward dynamic previously described. The relentless cycling strategy ensures full control over the advertised seller position in all periods, except when the cycle is re-initiated.

The reset to the maximum price follows the template of commercial repricing companies that typical offer either the option to stay at the minimum price or to reset to the maximum price. Note that resetting the maximum price is necessary to exploit the price-stickiness on the downward trajectory, which is the key feature of the relentless cycling strategy.

While the relentless cycling strategy has a forward-looking component because it anticipates the next move of the Bertrand strategy, it is still to be understood as a myopic strategy in the sense that it does not compute and solve intertemporal trade-offs. As we will discuss in Section 5, it is typically not optimal to wait until reaching the minimum price before reverting to the maximum price. However, this requires to solve the trade-off between the benefits of a higher average price and the immediate loss associated with loosing the advertised seller position when resetting.

We conclude the discussion of the relentless cycling strategy by noting that the combination of Bertrand best response and relentless cycling does not constitute a Nash-Equilibrium. From the perspective of the Bertrand repricer, the best response to the relentless cycling strategy would be to stick to the minimal price once it is reached and not to follow price increases. Not following price increases is a setting a seller might choose. Note, however, that such a strategy requires anticipating future behavior of the relentless cycling strategy and foregoing an immediate best response (undercutting a higher price).

The EXP3 Algorithm

For our own repricer, we employ the EXP3 algorithm, which belongs to the class of no-regret reinforcement learning algorithms. No-regret algorithms minimize the expected loss of a sequence of actions relative to an ex-post optimal action, i.e. they minimize the "regret" from not knowing the ex-post optimal action ex-ante.

More precisely, denote by $u_t(a_j)$ the payoff from taking action $j \in 1, \cdots, K$ in period $t \in 1, \cdots, T$. The objective of the EXP3 algorithm is to minimize the regret function

$$\frac{1}{T} \sum_{t=1}^{T} \Big(u_t(a_j) - u_t(a^\star) \Big), \tag{3}$$

where $a^\star$ denotes the ex-post optimal *single* action over all rounds T. In each period t, the algorithm only observes the payoff of the action selected during that period. EXP3 is a popular online-learning algorithm because of its low informational requirements and ease of implementation.

Algorithm 1 describes the implementation rules for the EXP3 algorithm. Initially, equal weights w are assigned to each action. Over the periods, the weights are updated such that actions that yielded high rewards are chosen with higher probability. The parameter $\gamma \in [0, 1]$ is a tuning parameter governing the exploration-exploitation trade-off, which is common to all reinforcement learning algorithms. The higher the parameter γ, the more likely it is that the algorithm will choose an action at random, favoring exploration over exploitation. Throughout the analysis, we set $\gamma = 0.1$, which is a typical choice in the computer science literature.

In Equation (3), the payoffs for the same action are allowed to change over time. This implies that the weights of the algorithm might not change monotonically, which, in turn, implies that the algorithm might not converge towards playing one single action over time. However, it can be shown that the regret from Equation (3) is bounded by $\sqrt{Klog(K)E_t(u_t(a^\star))}$ (Auer et al., 2002).

Algorithm 1: EXP3 ALGORITHM

Initialization: Set $w_{t=0}(a_j) = 1 \quad \forall j$

Algorithm: For $t = 1 \cdots T$:

1. $\forall j$, set $p_t(a_j) = (1-\gamma)\frac{w_t(a_j)}{\sum_{j=1}^{K} w_t(a_j)} + \frac{\gamma}{K}$

2. Draw one action from the probability distribution p_t, denote the selected action by a'

3. Set $w_{t+1}(a') = w_t(a') \exp\left(\frac{\gamma u(a')}{K p_t(a')}\right)$

Notes: In each period t, only the weight $w(a')$ is updated. γ is a tuning parameter that governs the exploration-exploitation trade-off.

The bound derived by (Auer et al., 2002) suggests that the number of actions K has a negative effect on the performance of the EXP3 algorithm. To keep K small, we opt to use deviations from the current market price as our action set $\mathcal{A}$.

More precisely, we define $\mathcal{A} = \{-0.04, -0.02, 0.00, 0.02, 0.04\}$, and the price implemented by the EXP3 algorithm if action $a' \in \mathcal{A}$ is chosen is $p_{t+1} = min\{\mathbf{p_t}\} + a'$, where $\mathbf{p_t}$ denotes the price vector of all sellers in period t. Thus, the EXP3 algorithm increases or decreases the minimum price by an increment a'. Our reward function is given by $u(a') = (min\{\mathbf{p_t}\} + a') \times \mathbb{I}\{AS_{t+1} = True\}$, where the indicator function $I\{AS_{t+1} = True\}$ takes the value one if the seller controlled by the EXP3 algorithm has the advertised seller position in period $t+1$.

The action space we select ensures that the EXP3 algorithm will, by design, implement prices close to the competitive price. Under the specified reward function, the EXP3 algorithm seeks to maximize the price-weighted share of times it acquires the advertised seller position.

Finally, we need to specify an ad-hoc rule for the instance that the EXP3 algorithm selects a price weakly below the marginal costs. We instruct the algorithm to revert to the maximum price if the minimum price (i.e. the marginal costs) are reached. Note that, due to the definition of the action space, the EXP3 algorithm cannot itself decide to revert to the maximum price. The ad-hoc rule of reverting to the maximum price is also implemented for the relentless cycling strategy, ensuring a level playing field between both algorithms.

Additional Remarks

Our goal is to implement a realistic repricing software that achieves good performance in real markets, which are characterized by rivals who might change their strategy. In this respect, one concern is that the relentless cycling strategy and the EXP3 algorithm might be overspecialized to the scenario of a rival implementing a Bertrand strategy.

For instance, the relentless cycling strategy involves lowering the own price preemptively to prevent the reactive Bertrand-player to successfully undercut. Note that this strategy would be non-optimal if the Bertrand-player would decide to give up and implement a higher static price. In this case it would be optimal for the relentless cycling strategy to stop cycling and simply undercut the new higher price.

While the relentless cycling strategy that we implement does not account for this scenario, it would be easily feasible to modify the code to monitor the activity of the rival (for instance by counting the number of price changes in a given time-interval) and to instruct the algorithm to switch modes in case the rival becomes inactive. Because this case is trivial both economically and technically, we do not explicitly consider it in our experiments.

A similar reasoning holds with respect to the EXP3 algorithm: If the rival becomes inactive, the EXP3 can be instructed to undercut the rival. When the rival resumes activity, the EXP3 algorithm can simply be reactivated and continue using the weights it had learned prior to the period of inactivity.

4 Results

In this Section, we present and discuss the results of our experiments. We start with the scenario in which a Bertrand strategy competes against a relentless cycling strategy. The results of this first experiment will establish the performance benchmark for the two subsequent experiments in which the commercial repricing software and our EXP3 algorithm will compete against the Bertrand strategy. Finally, we will show the results of the experiment in which our EXP3 algorithm competes directly against the commercial repricing software.

All the experiments where conducted with a maximum allowable price of 2.65 and a minimum allowable price of 2.0. The minimum allowable price can be thought of as representing the marginal costs. The prices where chosen with the goal to ensure that the two algorithms competing against each other are below the minimum price of all other sellers active in the market. We opted for this approach in an attempt to make sure that the prices of other sellers do not interfere with our experiments.

Each experiment lasted for approximately one-and-a-half days.

Bertrand Strategy Against Relentless Cycling Strategy

Figure 3a shows the prices set by the Bertrand and the relentless cycling strategy over the course of the entire experiment. Note that the Bertrand strategy occasionally does reset the price cycle to a price just below 2.65. This is due to a professional seller who unexpectedly decreased her price which led to interferences with our experiment. Figure 3b shows one selected cycle of the experiment: It illustrates that the relentless cycling strategy manages to prevent the Bertrand strategy from successfully undercutting most of the times.

Figure 3c shows the share of the advertised seller position that the relentless cycling strategy obtains. The share is obtained over rolling windows of 250 minutes, which corresponds to roughly two full cycles. The share of the advertised seller position is our success criterion. When computing the share, we remove instances in which the Bertrand strategy did not reset to a price above 2.5 to avoid that our results are affected by the prices of the professional seller.

The share of the advertised seller position of the relentless cycling strategy oscillates around 86.5%. The red dotted horizontal line indicates the share of the advertised seller positions that the relentless cycling strategy should theoretically attain. In theory, the relentless cycling strategy should always charge the lowest price except when it initiates a price increase to reset the cycle. Taking into consideration that a cycle takes 23 price updates to complete, we would therefore expect that the theoretical share of the relentless recycling strategy is approximately equal to $1 - 1/23 = 95.5\%$.

The main reason for the discrepancy between the observed and the theoretical value is that the pricing API used to control prices exhibits unexpected behavior. For instance, the time between consecutive pricing events varies and price updates might occur asynchronously, which explains why the Bertrand strategy is occasionally able to successfully undercut the relentless cycling strategy (as is shown Figure 3b). This results in the Bertrand strategy being able to claim a larger share of the advertised seller position than expected from theoretical considerations.

The discrepancy illustrates how real world technical issues hinder a perfect implementation of the intended strategies. As a result, hypothetical benchmarks may never be reached. Therefore, instead of the theoretically optimal share, we will use the average share of the advertised seller position observed for the relentless cycling strategy as the benchmark for the other experiments.

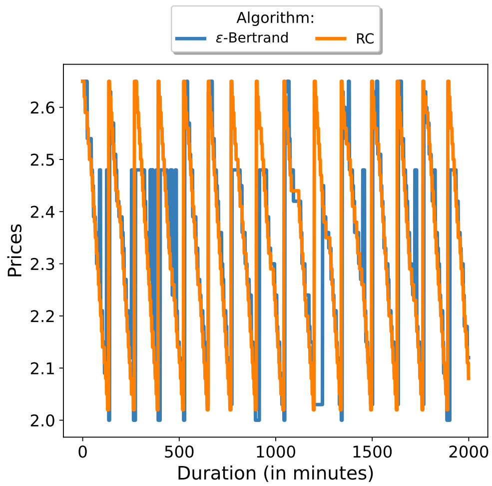

(a) Prices Paths Entire Experiment

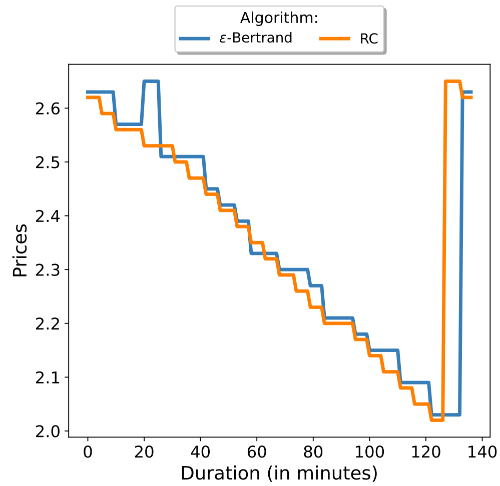

(b) Price Paths for Selected Cycle

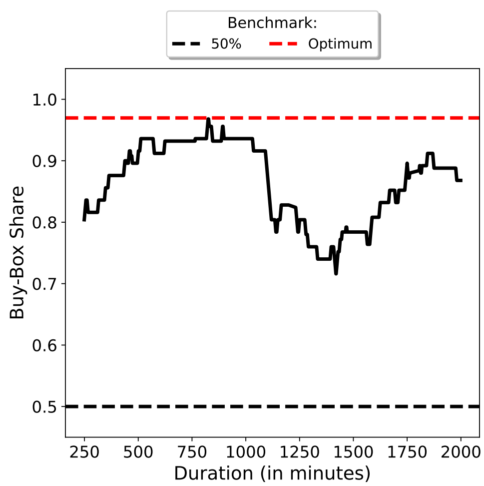

(c) Share of Advertised Seller Position for Relentless Cycling Strategy

Figure 3: Bertrand Best-Response vs. Relentless Cycling Strategy

Notes: In the upper-left panel, the Bertrand strategy occasionally resets the price cycle to a price of 2.48. This is due to another seller setting a price of 2.5, prompting the Bertrand best-response strategy to ignore the price of the relentless cycler.

Bertrand Strategy Against Commercial Repricing Software

Figure 4a shows the prices set by the commercial repricer and the Betrand strategy. Figure 4b shows the share of the advertised seller position obtained by the commercial repricing software. The share obtained by the commercial repricer is consistently below 50%. This suggests that the Bertrand strategy is sufficient to outperform the AI of the commercial repricing company.

The prices observed in Figure 4a show that the commercial repricer is occasionally holding constant, or increasing the prices. This also explains why the cycle length observed in the experiment between the commercial repricer and the Bertrand strategy is significantly longer (approximately 500 minutes) than in the previous setting.

Theoretically, one advantage of extending the cycle length is that it decreases the costs associated with resetting the cycle to the maximum price, as this is associated with loosing the advertised seller position. Note, however, that this strategy would only make sense if the rival would cooperate in extending the cycle length by not always undercutting. As is evident from our experiment, extending the cycle length against a Bertrand strategy is clearly sub-optimal.

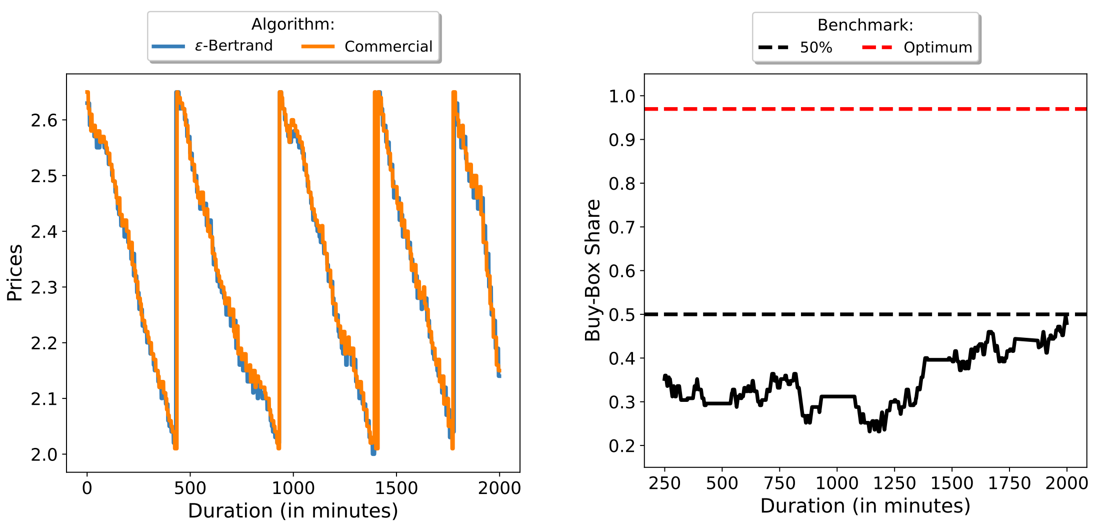

(a) Price Paths Entire Experiment (b) Share Advertised Seller of Commercial

Figure 4: Commercial Repricing Software vs. Bertrand Strategy

While the performance of the commercial repricing software appears sub-optimal from the perspective of a myopic seller, the same cannot be said about a forward-looking seller. For instance, the commercial repricing software might try to deplete the stock of the rival to obtain a monopoly position in the allowable price range that was specified by the seller (in our case from 2.65 to 2).

We note that there is no guarantee that the the rival will ever allow the stock to deplete, in which case the observed scenario is clearly sub-optimal. Additionally, if the strategy is indeed aimed at depleting the rival's stock, the observed price cycles implemented by the commercial software appear unnecessarily complex: The same could be achieved by allowing the Bertrand strategy to undercut a low static price.

Finally, it might be the case that the commercial repricer did not have sufficient time to learn and adapt to the Bertrand strategy. We cannot rule out this possibility. However, such concerns can be easily addressed by running longer experiments. Additionally, as we will show next, it is possible for an artificial intelligence algorithm to successfully adapt to the Bertrand Strategy within the time-span of one-and-a-half days.

EXP3 Against Bertrand Strategy and Commercial Repricing Software

Figure 5a shows the performance the EXP3 algorithm against the Bertrand strategy. The share of the advertised seller position captured by the EXP3 algorithm is consistently above 50% and increasing over time.

Figure 5a shows the evolution of the probabilities with which the EXP3 algorithm chooses a certain action. In light of the discussion in Section 3, the optimal action is to choose to decrease the lowest price by four cents. Clearly, the algorithm learns to predominantly play the optimal action. Expectedly, the probabilities to choose an action that would weakly increase the currently lowest price converge to zero.

Decreasing the price by only two cents does not guarantee the advertised seller position because it will lead to a situation in which the Bertrand and the EXP3 player charge the same price. Note, however, that players charging the same price will be randomly assigned the advertised seller position. Therefore, the action of decreasing the lowest price by only two cents will yield occasional positive payoffs. This is the reason why the probability associated with choosing a price reduction of two cents does not decay to zero.

Figure 6a shows the performance of the EXP3 algorithm against the commercial repricing software. According to our success measure, the EXP3 algorithm substantially outperforms the commercial repricing software: The share of the advertised seller position shows a clear upward trend and is consistently above 50%.

Figure 6b shows the evolution of the probabilities with which the EXP3 algorithm chooses a certain action when confronted with the commercial repricer. Compared to the scenario in which the EXP3 algorithm competes against a Bertrand Strategy,

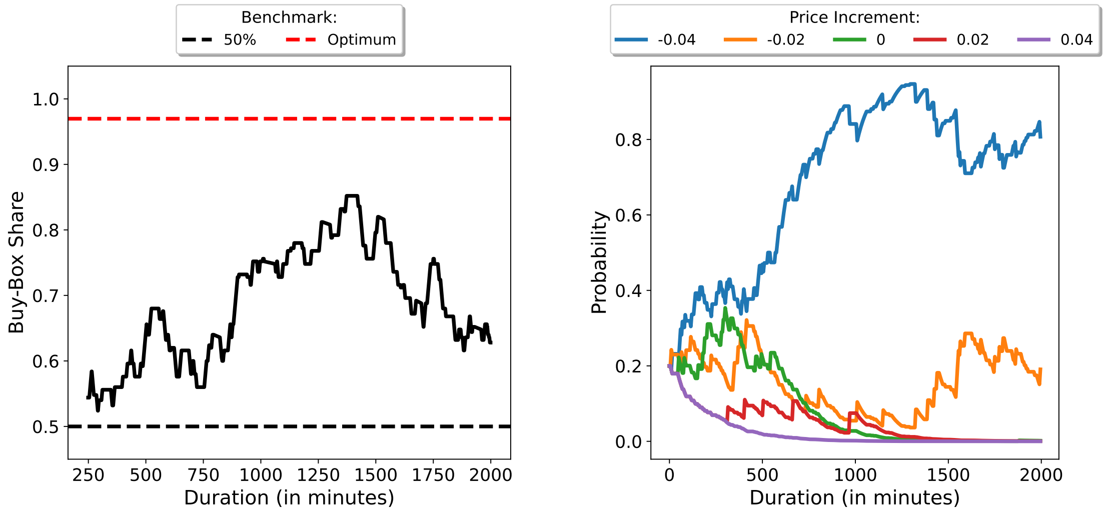

(a) Share Advertised Seller of EXP3

(b) Evolution of EXP3 Probabilities

Figure 5: EXP3 vs. Bertrand

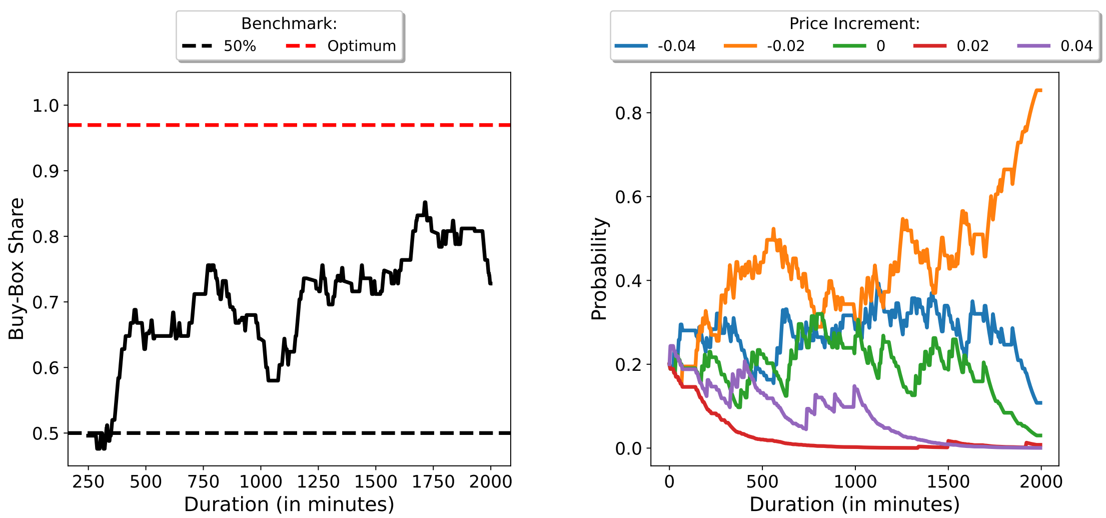

(a) Share Advertised Seller of EXP3

(b) Evolution of EXP3 Probabilities

Figure 6: EXP3 vs. Commercial

there does not appear to be one single action that the EXP3 algorithm predominantly chooses – although there appears to be a trend in favor of a price reduction of two cents towards the end of the experiment.

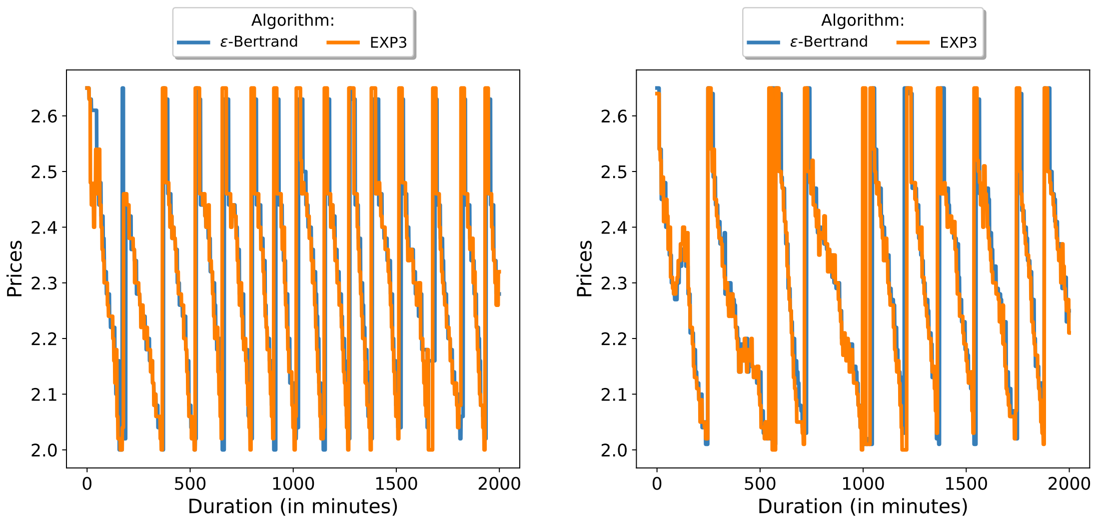

(a) Price Paths EXP3 vs Bertrand

(b) Price Paths EXP3 vs Commercial

Figure 7: Price Paths EXP3 Experiments

Over much of the experiment, the weights are consistent with a mixed strategy in which actions that lead to weak price reduction are chosen with approximately equal probabilities. Note that this "uniform" mixed strategy play does not appear to impact the performance of the EXP3 algorithm significantly compared to the later stages in which a single preferred action emerges. This appears consistent with the pricing pattern observed in Figure 4a for which we noted that the commercial repricer occasionally raises prices. In light of this, the EXP3 strategy of mixing across actions leading to a weak price decrease appears rational.

5 Forward-Looking Algorithms

So far, our research design has been limited to developing a protocol to test the sophistication of repricing technology assuming myopic players. Using one selected commercial repricing software and our own EXP3 algorithm, we demonstrated the applicability of our protocol. Our results indicated that the selected commercial repricing software is outperformed by simple repricing rules as well as our own software.

Our findings might be driven by a mismatch between the potentially forward-looking objective function of the commercial software and the myopic objective function our protocol has been developed for. While we argue in Section 4 that the more plausible explanation for our results is the lack of sophistication of the commercial repricer, a more rigorous approach is needed. Particularly in view of the fact that other commercial repricing technologies, which we want to test in future iterations of this work, might be more sophisticated.

The ability to learn strategies that make the myopically profitable action of undermining the collusive agreement unattractive is a key aspect of tacit collusion. Solving intertemporal trade-offs is therefore a prerequisite that algorithms must fulfill in order to be able to collude. In the remainder of this Section, we briefly discuss the main issue associated with forward-looking algorithms that learn in real time. We then review the class of Q-learning algorithms, and, finally, present results that highlight the possibility to efficiently learn intertemporal trade-offs.

Speed of Learning

Calvano et al. (2020) show that simple Q-learning algorithms are able to learn strategies that punish a deviation from the collusive agreement for a finite number of periods. One caveat associated with their findings is that the algorithms undergo extensive learning. The main results in Calvano et al. (2020) are based on algorithms trained over several millions of periods. While such extensive learning might be acceptable in markets characterized by a very high-frequency of price events (such as financial markets), it appears prohibitive in our setting.

A fast speed of learning appears crucial for repricing companies. Sellers experiencing extensive learning periods with sub-optimal behavior will likely discontinue the use of the repricing service. While offline learning, where the algorithm is trained on historical data before being deployed in the real environment, might be a partial remedy, such algorithms will typically lack the ability to optimally adjust to the idiosyncrasies of the market environment in which they will be deployed.

Digression: Q-learning Algorithms

Q-learning algorithms are designed to solve Markov decision problems. Given a state space S and an action space A, they find the optimal policy, i.e. a mapping $S \to A$ that maximizes the sum of discounted payoffs $\sum_{t=0}^{t=T} \delta^t u(a_t, s_t)$. Markov decision problems are characterized by the property that the probability to reach a certain state s' in period $t+1$ only depends on the state and action taken in period t, s and a, i.e. the (transition) probability to reach state s' given s and a is given by $prob(s'|s,a)$.

To apply the Q-learning algorithm, one starts by initializing a Q-matrix with dimensions $S \times A$. Each cell of the Q-matrix contains initial guesses for the Q-values, denoted by $Q(s,a)$. The Q-values describe the continuation value of choosing action $a \in A$ when in state $s \in S$. Once the Q-learning algorithm has been applied, the optimal policy can be read-off from the Q-matrix by finding $\operatorname{argmax}_{a \in A} Q(s,a) \forall s$, i.e. by determining for each state s which action maximizes the continuation value.

Algorithm 2: ϵ-GREEDY Q-LEARNING ALGORITHM

Initialization: Randomly set $Q(s,a), \forall(a,s)$, and initialize a state s

Algorithm: For $t = 1 \cdots T$, repeat:

1. With prob. ϵ select $a^\star \in A$ uniformly at random (exploration); with prob. $1-\epsilon$ select $a^\star = \operatorname{argmax}_{a \in A} Q_t(s,a)$ (greedy action)

2. Given $a^\star$ and s, select s' following $prob(s'|s,a^\star)$

3. Apply Equation (4) to update $Q_t(s,a^\star)$

One popular way to learn the optimal Q-matrix is the ϵ-greedy Q-learning algorithm: After initiating the Q-matrix and selecting an initial state s, learning proceeds in rounds t. In each round, the algorithm either selects one action $a \in A$ uniformly at random with probability ϵ, or, with probability $1-\epsilon$, it selects the greedy action, $\operatorname{argmax}_{a \in A} Q_t(s,a)$. Given the selected action, $a^\star$ and the current state, s, the next state is reached according to $prob(s'|s,a^\star)$. In each round, the Q-values are updated following:

$$Q_{t+1}(s,a^\star) = (1-\alpha)Q_t(s,a^\star) + \alpha\big[u(s,a^\star) + \delta max_{a \in A} Q_t(s',a)\big]. \tag{4}$$

α denotes the learning rate, which is selected by the researcher. It has been shown that, under fairly general conditions, ϵ-greedy Q-learning converges to the Q-matrix that solves the Markov decision problem (Watkins and Dayan, 1992). Algorithm 2 provides the pseudo-code to implement ϵ-greedy Q-learning.

Optimistic Q-learning

Optimistic Q-learning is a variant of Q-learning that follows Algorithm 2 with two important distinctions. Firstly, the Q-values are not randomly initiated but are initiated "optimistically", using the highest possible continuation value that is theoretically reachable in the environment. Secondly, optimistic Q-learning is performed without any exploration (i.e. $\epsilon = 0, \forall t$), i.e. chooses the greedy action at each step of the learning process. Optimistic Q-learning is popular as it is considered *sample efficient* and, hence, offers an avenue to circumvent issues related to extensive learning (Even-Dar and Mansour, 2001; Neustroev and de Weerdt, 2020).

Applying Optimistic Q-learning to Augment The Relentless Cycling Strategy

We now demonstrate, using a simulation approach, that an optimistic Q-learning algorithm can be used to find the optimal dynamic strategy against a myopic Bertrand player. Crucially, the learning process does only impose minimal costs on the seller. As discussed in Section 3, the relentless cycling strategy can only be considered myopically optimal. The reason is that the relentless cycling strategy only resets the price-cycle when it reaches the allowable minimum price.

By contrast, a forward-looking seller will choose the optimal reset-price. This optimal reset-price will solve the trade-off between the costs of resetting the price-cylce, i.e. loosing the advertised seller position, and the benefits of achieving a higher average price by resetting before reaching the minimum price. For example, if the consider the scenario of a price range between 2.65 and 2 and a Bertrand rival, it can be shown that the optimal reset price for a relentless cycling strategy is 2.35.[8]

For the following simulation, we let the state space S of our Q-learning algorithm contain all the 65 possible price points in interval $[2, 2.65]$. The algorithm is allowed to choose three action: (i) no price change, (ii) reset the price-cycle, and (iii) reduce the price by $\epsilon + 0.01$, where $\epsilon = 0.01$ is the amount by which the Bertrand strategy decides to undercut. Note that preventing the Bertrand strategy from successfully undercutting is embedded in the algorithm by design. Thus, whenever the algorithm decreases its price, it automatically exploits the main idea of the relentless cycling strategy.

For our simulation, we set $\delta = 0.999$ and $\alpha = 0.1$. To implement the optimistic Q-learner, we initiate all the Q-values at the continuation value that results from resetting the price at the optimum value, i.e. 2.35.

Figure 8 shows one representative run of the simulation. Figure 8a shows the reset price that the algorithm is currently implementing. Figure 8b shows the ratio between the discounted profit that would result from keeping the current reset price and the discounted profit that would be obtained if the algorithm would use the

[8] We use numerical optimization techniques to find the optimal reset price.

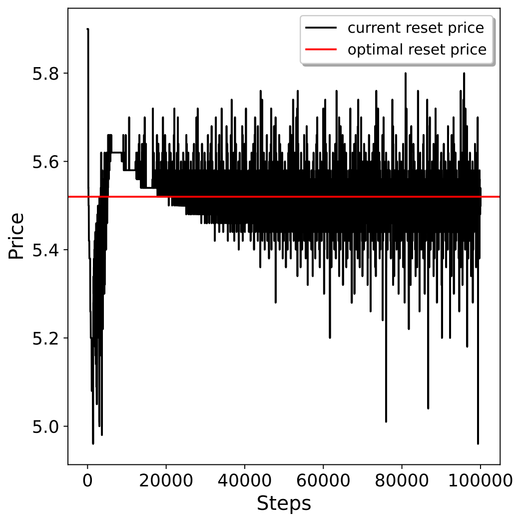

(a) Evolution Reset Price

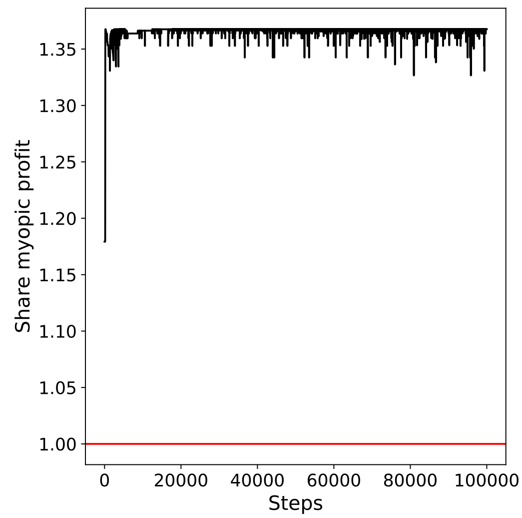

(b) Profit Compared to Myopic Strategy

Figure 8: Simulation Results Optimistic Q-learning

minimum price as a reset price. Note that using the minimum price as a reset price is the myopic optimal strategy. Therefore, the ratio indicates how much the seller would loose or gain from permanently implementing sub-optimal reset prices instead of simply playing the myopic strategy.

While Figure 8a reveals that the algorithm needs 5000 periods to settle on the optimal reset price (which corresponds to roughly 17 days in reality), Figure 8a shows that the typical learning process comes at almost no costs when taking the myopic optimal strategy as the benchmark. In fact, even in the learning stages, when the optimal reset price is not yet achieved, the optimistic Q-learning algorithm achieves a higher profit than the myopic benchmark. Thus, optimistic Q-learning, when combined with an ad-hoc implementation of the optimal myopic strategy, allows to learn dynamically optimal strategies against Bertrand strategies at no costs.

6 Conclusion

We develop a protocol to test the sophistication of commercial repricing software assuming myopic sellers. To this end, we create two seller accounts on an online platform and develop our own repricing software, which allows us to implement arbitrarily complex algorithms.

After establishing a performance benchmark based on implementing the best response against a seller using a Bertrand strategy, we let the commercial and our own software compete against this same Bertrand strategy. Our software is based on EXP3, which is a workhorse reinforcement learning algorithm for online learning

We find that our software comes closer to the performance benchmark than the commercial software, which suggests that a simple EXP3 algorithm is superior to the commercial software, at least in the myopic setting. This finding is further confirmed by an experiment in which we let the EXP3 and the commercial software compete directly: our own algorithm significantly outperforms the commercial software.

We discuss the limitations of our current findings, which could capture a failure of using the right benchmark. For instance, it could be that the commercial repricing software has a long-term objective that our protocol does not capture. We provide a proof-of-concept demonstrating the ability of algorithms to learn long-term startegies efficiently. This lays the groundwork for a protocol testing the sophistication of commercial repricing software assuming forward-looking sellers.

References

Assad, S., Clark, R., Ershov, D., and Xu, L. (2020). Algorithmic pricing and competition: Empirical evidence from the german retail gasoline market. *CESifo Working Paper*.

Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77.

Calvano, E., Calzolari, G., Denicolo, V., and Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–97.

Chaboud, A. P., Chiquoine, B., Hjalmarsson, E., and Vega, C. (2014). Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance*, 69(5):2045–2084.

CMA (2021). Algorithms: How they can reduce competition and harm consumers. *CMA Whitepaper*.

Cooper, W. L., Homem-de Mello, T., and Kleywegt, A. J. (2015). Learning and pricing with models that do not explicitly incorporate competition. *Operations research*, 63(1):86–103.

Even-Dar, E. and Mansour, Y. (2001). Convergence of optimistic and incremental q-learning. *Advances in neural information processing systems*, 14.

Hendershott, T., Jones, C. M., and Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? *The Journal of finance*, 66(1):1–33.

Neustroev, G. and de Weerdt, M. M. (2020). Generalized optimistic q-learning with provable efficiency. In *AAMAS*, pages 913–921.

OECD (2017). Algorithms and collusion: Competition policy in the digital age. *OECD Whitepaper*.

Watkins, C. J. and Dayan, P. (1992). Q-learning. *Machine learning*, 8(3):279–292.

7 List of Additional Figures

Figure 9: Share of Advertised Seller Position

Notes: The graph shows the share of the advertised seller position when both sellers set the same cheapest price. The experiment lasted for approximately two days. The error bars show the 95% confidence intervals.

Competition on Hybrid Platforms: Evidence from Amazon US

Vito Stefano Bramante*

Emilio Calvano†

Giacomo Calzolari‡

Yann Delaprez§

September 16, 2024

Abstract

This paper provides causal evidence on the competitive effects of Amazon's dual role as both a platform operator and retailer on its own marketplace. Leveraging data from over 9,000 staggered entries of Amazon as a first-party seller (FPS) across 18 product markets on Amazon US from 2018 to 2023, we analyze the impact of Amazon's entry on prices, market shares, and seller behavior. We find that Amazon's entry as an FPS results in a significant reduction in third-party sellers' (3PS) buy box shares and buy box prices, although these effects exhibit substantial heterogeneity across categories. Our findings suggest that Amazon's competitive advantage lies not only in pricing but also in leveraging its logistics arm, the Fulfilled by Amazon (FBA) program. Both FBA and Fulfilled by Merchant (FBM) sellers face diminished buy box shares following Amazon's entry, though FBA sellers appear to be price-constrained and unable to lower their prices further. Our study contributes to the growing literature on digital platform competition, offering insights into Amazon's ability to compete both as a retailer and a logistics provider.

JEL Codes:
Keywords:

*University of Bologna, vitostefano.bramant2@unibo.it
†LUISS University, emilio.calvano@gmail.com
‡European University Institute, giacocalzolari@gmail.com
§LUISS University, ydelaprez@luiss.it

1 Introduction

Amazon is the largest online marketplace in the world, with a 2023 revenue estimated at $574 billion, making it the third-largest public company globally by revenue.[1] Estimates suggest the platform serves over 300 million customers globally, with 165 million of them being Prime members in the U.S. alone.[2] The platform leverages their technological capabilities to match this customer base with roughly 2 million independent sellers, also referred to as third-party sellers (3PS, henceforth).[3]. Amazon sheer scope is rewarded by financial markets that valuate it at $ 1 trillion. To balance the perks, Amazon's pivotal presence in the digital landscape is also highlighted by the many investigations and interrogations, we counted twelve, that competition authorities have filed against the Seattle giant. Among their many concerns lies the central role that Amazon plays not only as a matchmaker but as a retailer, as a first-party seller (FPS, henceforth), that competes on its own platform against independent sellers that adhere to its own contractual terms and, in case of alleged non-compliance, are judged and punished according to its own policies.

The interests of competition authorities in the matter are to be found in the perils of this duality, often summarised by two different but complementary arguments. First, the platform might steer consumers toward those products sold by its retail arm; this mechanism can be enacted by manipulation of the search algorithm and by recommendation systems, both tuned with the purpose of increasing the likelihood that a consumer will land on a specific product-page. On Amazon, this concern amplifies due to the presence, in each product page, of the buy box, a spot dedicated to the best offer, chosen by an undisclosed algorithm, that contains the button ”Add to Cart” and ”Buy Now”; the specific concern here manifests itself via self-preferencing, the ability of the platform to raise the prominence of its own offers, despite the presence of other equally capable of attracting the same demand. Second, an information asymmetry argument is often suggested, whereby the platform can leverage data and observe choices operated by its partner/competitors that are not publicly shared among the sellers, enhancing its decision-making and leaving to third-party sellers the costs of experimentation.[4].

Alongside competition authorities Amazon's dual role has attracted the interest of the economics literature, with an effort that remains, for the most part, theoretical. The one empirical exception is Crawford et al. (2022), the closest work to ours, investigating Amazon's effect on competition in the Home & Kitchen category of the German marketplace, using Amazon's proprietary data. We contribute to this literature by shifting the focus on its ”.com” domain, Amazon US and expanding our attention to more than one category. Amazon North-America constitutes three times the revenue of Amazon international combined, and more than two thirds of its prime subscribers, their most loyal shoppers, are based in the US. We also pay particular attention to the interplay between Amazon's behavior as a FPS and its logistic services, represented

[1]https://www.forbes.com/sites/jasongoldberg/2024/02/12/amazon-dominates-2023-retail-landscape-with-record-q4-earnings/, last accessed on September 10, 2024

[2]https://www.businessofapps.com/data/amazon-statistics/, last accessed September 10, 2024

[3]than https://amzscout.net/blog/amazon-statistics/, last accessed September 10, 2024

[4]We add that the last argument is usually missing the intrinsic knowledge of the platform itself; for instance, the determinants of the search ranking and of the buy box, the algorithm determining the seller from which consumers ”Add to Cart” or ”Buy Bow”, are unknown to the public but might be known by the retail arm of Amazon Inc. enhancing their strategic decisions

by their Fulfilled by Amazon program, allowing third-party sellers to delegate most of the shipment process to the logistic branch of the platform in exchange of a fixed share of the transaction price. We are the first to do so.

In particular, we investigate the dual-role of Amazon on its own marketplace estimating the causal impact that Amazon entry as first-party seller has on competition in terms of market share, prices and market structure. We do so leveraging Amazon's staggered entry in more than eight thousands markets, and estimate a series of event study that credibly recover parallel-trend. Our data cover six years, from 2018 to 2023, and 18 categories of durable goods. We provide four set of results.

First, we find that Amazon is a "special" seller, able to provide unmatched competitive pressure in the marketplace. Upon its entry third-party sellers lose on average 20% of their buy box share, the buy box price decreases on average by 5% and Amazon sustains, on average, prices that are 6% lower than those winning the buy box before its entry[5]. The same metrics, compared to those of other large-third party merchants are nowhere to be matched. We measure the effects of Amazon's entry against the entry of the top 1% of large third-party merchants ranked according to two measure: the average number of ratings received[6], that we interpret as a proxy for transactions, and the average number of products listed on the platform, i.e. their storefront, throughout our period of study. To enhance readability we'll refer to the former as Big 3PS - CR and to the latter as Big 3PS - SF, throughout the text. Neither are capable of comparable results in the short-run and in long-run. Big 3PS - CR entry causes the decrease the buy box share of other 3PS by 19% in the first month and by 10% after eighteen months. Big 3PS - SF are only capable of causing the decrease of the share of 3PS by 2%. The buy box price of the products in which these large sellers enter are largely unaffected by the entry of Big 3PS in the short-run while possibly increasing in the long-run. Moreover their entry prices are substantially higher than Amazon's, suggesting that Amazon's strength lies in its ability to price-compete. This conclusion is largely mitigated by our second contribution.

We document a large heterogeneity in Amazon's competition across the eighteen categories we investigate. We find that Amazon's entry prices are lower than pre-entry buy box prices only in seven categories, spanning a range of -20% for groceries to -4% for sports & outdoors items. and its entry causes buy box prices to reduce only in six categories, with a range of -21% for groceries to -5.8% for arts, craft & sewing items. This finding, coupled with a consistently negative effect on the buy box share of 3PS in all categories, suggests that Amazon's entry strategy is not limited to price-competition. The algorithm that ranks offers and grants the buy box is unknown, we cannot therefore distinguish to what extent Amazon's outperforms third party-sellers in terms of quality of the offers and to what extent this result can be interpreted as evidence for self-preferencing.

Our third contribution is to complement the literature with evidences on how the Fulfilled by Amazon program (FBA) shapes third-party sellers ability to compete against the Seattle giant. In particular, FBA sellers, relying on Amazon's logistic services still suffer from Amazon entry, with a reduction in buy box share of 13% for FBA and of 16.5% for sellers fulfilled by merchants (FBM). In contrast, prices listed by FBA

[5]Throughout the text, we interpret buy box share as a close proxy to market share and buy box prices as a close proxy to transaction prices.

[6]We are referring to ratings received by the seller specifically and not the product. Sellers' rating are in a scale of one to ten. Products ratings are on a scale of one to five.

buy box winners are not affected by Amazon entry throughout the eighteen months. Amazon's entry does not cause a reduction in those third-party sellers winning the buy box. The reduction observed in buy box prices is, therefore, entirely explained by compositional change affecting the markets: more 3PS competing on prices, due to the increasing demand for the products they enter and Amazon's entry causing less competitive offers to lose buy box time. Amazon's entry though, does not causes prices to decrease, on the margin.[7] The compositional shift brought by Amazon joining the competitor's pool, on average, drives FBM sellers out of the buy box spots.

Finally, we build a predictive model of Amazon entry, that allows us, once again to highlight the prominent role of its logistic arm. Amazon's objective is largely driven by profit maximization concerns that share with third-party sellers a tendency to enter markets experience a fast demand growth. But its objectives depart from those of third-party sellers in two dimensions. Amazon internalizes the quality of the offers listed in each product by intervening as a first-party seller in markets where the buy box is absent for longer periods of time. These entries can be considered an intensive margin expansion: they guarantee higher revenues by direct sales, while preserving those coming from FBA fees. It also targets markets where FBA sellers hold the buy box for short period of times, no matter of the quality of the offers. These entries can be considered as an expansion on the extensive margin: Amazon increases its stream of revenues by trading off referral fees with direct sales.

Our work unfolds as follows. In section 2 we review the literature we contribute to; in section 3 we give an overview of the Amazon marketplace and of those aspect of its functioning relevant for our purpose. In section 4 we describe the data we leverage for our causal analyses which we outline in detail in section 5. The core of our work is in section 6: here we document all our results, ordered according to our outcomes of interest. We conclude by taking stock and discussing avenues of future research.

2 Literature

This paper contributes to two strands of the literature. First, we contribute to the literature seeking to understand how entry impacts competition. A large number of papers has studied the causal impact of entry on prices in markets such as retail (Busso and Galiani (2019)) or gasoline (Fischer et al. (2023)). For instance, Fischer et al. (2023) exploit 700 entries at different point in time over 5 years in the German gasoline market. They find evidences of heterogeneity in price responses within the price distribution with stronger effects at the bottom of the price distribution. Overall prices decrease after entry. We differ from this work along at least two dimensions. First, we provide evidence on the effect of entry on a large digital platform. In this setting, we show that entry lowers prices regardless of whether Amazon or another large seller enters. However, the impact of Amazon's entry is significantly larger. On average Amazon is able to set prices that are 6% lower than prices of 3PS that were winning the buy box before entry. We also record instances in which prices of large sellers increase after entry. Furthermore, we identify two novel margins through which entry can affect competition

[7]The question of how can Amazon sustain its out-of-the-ordinary price competition is outside the scope of this work. We point out though that 3PS relying on Amazon logistic services share part of the revenues with Amazon itself. To what extent it is able to cross-subsidize its low prices via the FBA program is an interesting avenue of future research.

that are relevant to competition policy. We find that entry by Amazon not only leads to lower prices but also decreases the average buy box share and is correlated with higher average FBA share which is a way to finance lower prices caused by entry. By contrast, after entry of large sellers the average buy box share decreases but substantially less. Our most conservative estimate correspond to a 2% of the average buybox share of rivals third party sellers. With these findings, we complement Crawford et al. (2022) which is the first paper studying the impact of Amazon entry on prices. The authors study the price effects of Amazon entry on prices and revenue in the Home and Kitchen market in Germany. They found that Amazon lowers prices of third party sellers by about 3% after 10 weeks. Our estimates are slightly higher as we find that buy box prices of third party sellers decrease by about 5%. We extend their findings by documenting price effects in 18 categories with important heterogeneity across categories. In fact, we find that only 6 categories experience a decrease in the price of the buy box. Moreover, we find that Amazon's entry prices are lower than pre-entry buybox prices in only 6 categories. Finally, their main result points toward an extensive margin expansion of the markets in which Amazon enters, countering the notion that Amazon revenues come from other third-party sellers. They argue that the increase in sales registered in markets where amazon becomes a first-party seller come from sales that would have not happened otherwise. By contrast, we find that Amazon causes a reduction in revenue of about 11.2% which more consistent with a "business stealing mechanism rather than a "market expansion" mechanism described in Crawford et al. (2022).

Second, with this paper we contribute also to the growing literature seeking to estimate the welfare consequences of Amazon. A key channel that has been studied is self-preferencing by Amazon in U.S. (Farronato et al. (2023)) or Europe (Waldfogel (2024)) through search rankings giving better ranks to Amazon private labels. To give an illustration, Farronato et al. (2023) find that Amazon private labels benefits from prominence 30% to 60% higher than the prominence granted to sponsored products. Another source of self-preferencing may be taking place through the buy box or Amazon FBA program. Raval (2022) uses Keepa data from 2020 to 2021 for several countries and categories and estimate a model of the buy box to document that Amazon algorithms favor its own products both trough the buybox and his FBA program. A limit acknowledged in his work is that entry is not considered. Here we take an opposite route, we do not model Amazon behavior but rely on theory-free approach to predict the determinant of Amazon entry. Also we exploit observed entry as a source of identifying variation. We find that Amazon entry raises the share of sellers relying on fulfilment by Amazon and lowers the share of sellers using their own fulfilment service. Moreover, after Amazon entry sellers relying on the FBA program experience a smaller drop in buy box share compared to sellers using their own fulfillment service.

3 The Amazon Marketplace

Amazon is one of the largest online marketplace in the world with, at the time of writing, a predicted revenues of 491 Billion $ for 2024 and, in the US, a market share of 60% in the online retail sector.[8] The marketplace acts as a storefront for a wide variety

[8]https://www.emarketer.com/insights/amazon-revenue/, Walmart is second at 94 Billions.

of products, and almost the entirety of the available catalogue can be sold by multiple sellers, competing for profits. For the purpose of our study we interchangeably use the terms product, market and ASIN, with the latter being the Amazon Standard Identification Number, the unique identifier of each item on the platform and in our data. We classify products according to Amazon's presence referring to *never treated* as those items that do not count Amazon among its sellers, *Treated* products for which we observe Amazon time of entry in our data, allowing us to isolate pre and post-entry characteristics, and as *always treated* those markets in which Amazon does operate as a first-party merchant but we are not able to observe its entry, due to the left-truncation of our panel. Table 2, that provides descriptives on the full panel, a subset of which we use for causal identification, shows how Amazon's presence in a market carries striking differences. Amazon operates as a FPS in 32% of the products in our sample which, it's worthwhile stressing, considers the upper end of the sales distribution, largely recognized to follow a power law (Brynjolfsson et al., 2011). In these markets it faces tougher competition selling against, on average, 8.29 competitors in always treated products, 6.85 in treated. The average number of sellers in never treated markets is, instead, sensibly lower with 1.57. Competition comes from third-party sellers, independent professional retailers that list their product on Amazon for a fixed fee of 39.99$/mo and a ad-valorem fee on closed transactions.

Whether Amazon's success as a FPS is driven by its attractiveness towards consumer or by its ability to self-select into attractive markets is an open question but two characteristics underlie its presence. On the one hand, the dynamism of the markets in which it's present. On average, in fact, only half of the products (2178) die each month when Amazon is present (AT) than those who end their life-cycle when Amazon is not a seller (4618). On the other hand, despite Amazon operating in 32% of the market, the (AT) column suggests that it faces an equal number of unique competitors than all operating in markets in which Amazon does not sell, the (NT) column.

The Amazon marketplace is also characterized by a prominence-mechanism called the buy box; an example is given in 1. The Buy Box is the section on a product page where customers can directly add an item to their cart; when doing so their cart is filled or their purchase is handled by a seller chosen by an algorithm. Sellers compete for the Buy Box, and Amazon uses various factors to determine which seller wins it. Its algorithm is not public but previous work suggest that prices play the most prominent role followed by quality of the shipment service, prime membership and seller quality in the form of rating (Raval, 2022; Lee and Musolff, 2021; Chen et al., 2016). When a FPS, Amazon also competes for the top spot. As table 2 suggests, Amazon wins the buy box 66% of the time in always treated markets against a combined buy box share of 23% from all the other third-party sellers and 14% in treated markets, although this number accounts also for periods before entry.[9] In never treated markets the buy box share is held 84% of the time by 3PS.

Out of all services that Amazon offers to its partners, its Fulfillment by Amazon (FBA) program is the most relevant for our purposes. In exchange for a fixed fee on each transaction, third-party sellers can store their products in Amazon's fulfillment centers. Amazon handles the storage, packing, shipping, and customer service on behalf of these sellers. The FBA fee is a function of many product characteristics such as weight, size. When third-party sellers manage their own inventory, packing, and

[9]The two do not sum to one because we also account for buy box idle time, that is, periods in which the buy box is absent and hence no seller can qualify for it

shipping their products are Fulfilled by Merchant (FBM). The quality of the shipment service and of the customer service is a crucial factor for the buy box status (Lee and Musolff, 2021). In fact, the Never Treated column shows how a buy box share of 84% is coupled with an 86% of offers part of the FBA program. In always treated and in treated markets FBA enrollments among offers reduces to 20% and 57%, respectively. Table 2, therefore, depicts FBA as a crucial tool to win the buy box when Amazon is not a competitor, but not when Amazon is. The reasons why this is the case are not immediate to assess. One possibility is that Amazon fierce price competition coupled with (ever) rising FBA fees renders the fulfillment program not attractive to the many seller competing directly with the Seattle giant. Importantly, FBA sellers are automatically eligible for the Seller Fulfilled Prime program; entering the SFP grants sellers their Prime badge and having the possibility to see their offers advertised and listed to prime members. The Prime program grants customers free delivery and shipping within 24 or 48 hours. The FBA program and its bundling with SFP has been the subject of several authorities investigations, including the FTC case September 2023.

The many outlined differences between markets in which Amazon operates as a first-party seller (AT) and those in which it doesn't (NT) suggest that unobservable market and seller characteristics are also likely to play a crucial role. For this reason a comparison between these two strata of the marketplace would likely provide a biased assessment of Amazon's effect on competition. To circumvent this limitation we exploit Amazon staggered entry in some markets, the treated ones in table 2; we think about these as transitioning from the (NT) part of the marketplace to the (AT) one. Next, we describe the sample used for our causal estimations.

4 Data and Descriptives

4.1 Data

4.1.1 Sampling and Representativeness

Amazon US lists billion of products on its marketplace. At the time of writing, Keepa, our data provider, tracks a number of items that is just shy of 1.4 Billions; yet for computational reasons we restrict our attention on a smaller subset of the universe of products. In particular, we sample over products, selecting those that have *ever* entered the category-wise monthly list of the top 5000 bestselling products from April 2021 to April 2024.[10] Our choice is guided by two considerations. Although not representative of the universe of products, were we able to estimate our event studies on the latter, it is reasonable to assume that the largest effects can be appreciated in those products granting large revenues, and hence sales volume, before Amazon entry. To this design-driven choice we add that bestselling products are monitored more frequently and hence with highest quality, allowing us more statistical precision. We cover 18 different categories, which, following Keepa's taxonomy, we will denote as *root categories*, the highest level of product categorization. We focus on durable goods but, due to Keepa's limited tracking capability in those categories, we exclude apparel,

[10] To give a comparison point Raval (2022) the top 5,000 products; 2,500 products ranked between 5,001 and 25,000; and 2,500 products ranked below 25,000.

Table 1: Example of buy box

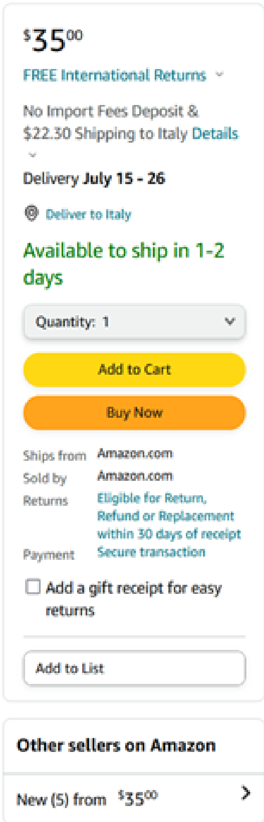

Table 2: Full sample - descriptive statistics

	(AT)	(NT)	(T)
N. of ASINS	172706.00	422809.00	25718.00
Age - (M)	49.56	21.81	51.67
Age at 1st Rec - (M)	25.75	7.76	23.96
Births/Month - (M)	610.01	1230.89	35.20
Births/Month - (Std)	228.44	770.31	15.58
Death/Mo - (M)	2178.75	4618.89	337.55
Deaths/Month - (Std)	7161.96	2948.44	1216.32
3PS BB Share - (M)	0.23	0.84	0.73
3PS BB Share - (Std)	0.28	0.21	0.30
Amz BB Share - (M)	0.66	0.00	0.14
Amz BB Share - (Std)	0.31	0.00	0.23
FBA Share - (M)	0.20	0.86	0.57
FBA Share - (Std)	0.26	0.29	0.35
N. of Sellers - (M)	8.29	1.57	6.85
N. of Sellers - (Std)	10.67	2.82	9.54
Sales Rank - (M)	82797.00	43978.00	69964.00
Sales Rank - (Std)	233197.00	137798.00	164239.00
Ratings - (M)	44.46	44.27	44.44
Ratings - (Std)	3.87	3.60	3.04
Reviews - (M)	4914.00	3280.00	4347.00
Reviews - (Std)	8636.00	7558.00	7630.00
New Offers - (M)	5.00	1.00	4.00
New Offers - (Std)	7.00	2.00	5.00
N. of ASINS in L2 - (M)	5770.00	6978.00	5230.00
N. of ASINS in L2 - (Std)	5321.00	6464.00	4572.00
N. of ASINS in L3 - (M)	1456.00	2102.00	1465.00
N. of ASINS in L3 - (Std)	1524.00	2707.00	1675.00
Number of Unique Sellers	337573.00	335224.00	203016.00
Seller N. Ratings - (M)	1507.22	1334.81	1749.14
Seller N. Ratings - (Std)	19391.19	14280.19	18877.73
Seller Rating - (M)	86.27	88.94	87.07
Seller Rating - (Std)	16.13	15.11	15.52
Seller Portfolio Size - (M)	1845.87	1350.83	1900.50
Seller Portfolio Size - (Std)	42008.99	41427.22	39208.12

Note: The figure on the left represents an example of a buy box; when clicking either of the "Add to Cart" or "Buy Now" button a customer is initiating a transaction with the seller the product is "Sold by", in this case Amazon.com; the buy box also reports the fulfillment status, reading "Ships from" Amazon.com for FBA sellers and "Ships from" accompanied by the name of the merchant for FBM. Clicking on the bottom arrow, on "New (5)" opens a tab containing all competing non-buy box winner offers. The table on the right provides descriptive statistics for the full sample.

music physical products and luxury items.[11] Still, we are able to cover not only the majority of the root categories but, importantly the most profitable ones, with the sole exception of the clothing category.[12] The raw data contain contain both market-level data, that is, information that are common to a product and to all sellers participating in a market at a specific point in time. It also contains seller-specific and offer-specific data, as captured at a given time. Keepa collects information using a time resolution that is a middle-ground between an infinite one, e.g. push-notification where the collection is triggered by an event, and a fixed one, where the collection is triggered at equidistant points in time. We build, from the raw data, a panel with observations at

[11] In particular we omit the following categories Clothing, Shoes & Jewelry,: CDs & Vinyl, Software, Books, Video Games, Magazine Subscriptions, Digital Music, Apps & Games, Movies & TV, Video Shorts, Vehicles. Moreover we omit non-fungible items from the Collectibles & Fine Art and Handmade Products categories. We also exclude Amazon's own products from the categories Audible Books & Originals, Kindle Store and Alexa Skills

[12] See. here for the most profitable categories: url (accessed 26 August 2024) `https://www.junglescout.com/resources/articles/amazon-product-categories/`.

a monthly frequency. Given the peculiarity of Keepa's collection, most continuous dependent variables and observable characteristics, are aggregated by mean of weighted averages, where the weights are given by the amount of time spent at a particular value. Table 2 gives an overview of the full panel. We observe at least a part of the history of 621233 products. Our empirical strategy, though, imposes us to get rid of all those products where Keepa's collection started after Amazon entry and for which we are not able to observe the periods around treatment time. We deem these products as always treated (AT) and we do not use these for our difference-in-differences.

The dataset used for the causal analyses of section 6 is a monthly panel running from January 2018 to December 2023. Being interested on the effect that Amazon has on third-party merchants, we drop from our sample Amazon-branded products. Our treatment definition, discussed in the previous section, requires us to exclude products in which Amazon has entered in the first six month of their life-cycle. We further restrict our sample to those products for which we observe their birth, and that therefore enters the panel during the six years time-span it covers. To ensure comparability, we also restrict the control group to product that have being listed on the platform for less than 78 months, grating this way a long-run comparison group to products that have experienced Amazon entry by age 60. Crucially, for computational reason, we cannot exploit the full control group, which we subsample drawing 10% of it at random. Table B.1, compares the subsample that we use as control for our causal analysis of section 6, to the full sample. All observable characteristics are comparable showing the ability of the 10% to be representative of the full sized control group. Finally, we restrict out attention on those products that experience Amazon's presence for the whole 18 months after entry. We do so because, as it will become clear in our analysis, the short run and long-run consequences of Amazon entry can differ substantially. We postpone an in-depth description of the data used for estimation to subsection 4.3. We move next, to define precisely the treatment under scrutiny.

4.2 Treatment definition

Our main source of identifying variation comes from Amazon's decision to enter a product market as first-party seller. Moreover, for each outcome of interest, we leverage entry of big 3PS to gauge both the sign and the magnitude of Amazon's effect. In this section we describe how we define entry for all sellers, how we define big 3PS and how our treatments spreads over time with respect to both the age of a product and calendar-time.

We infer merchants' dates of entry in a market, including Amazon's, starting from Keepa's disaggregated data and defining entry in a specific way that allows us to circumvent two limitations: firstly, an obvious one, the absence of explicit information about the date of entry and secondly, the left-truncation of the data due to the time in which the Keepa's scraper starts to track each product market. Therefore, for each product we select the first date in which we observe *any* price as marking *the beginning of the time-series*. For all sellers that have ever listed the product we compare *the first date of their first recorded price*, and interpret as entrants only those sellers for whom the latter is *at least six months older than the beginning of the time-series*. While the six months threshold is an arbitrary one, it allows us to handle the unavoidable left-truncation of the data at hand, preventing us from interpreting the date of the first price as the date of entry in all those cases where the *real* date of entry precedes the beginning of the time-series.[13] The chosen threshold, six months, mirrors the pre-treatment periods in our estimated equations. Our definition is not without limitation. In particular it implicitly interprets as entrants all sellers that might have entered in the past but that have been inactive for more than the six months before the date marking the beginning of the time-series. We notice though, that this is a concern only for treated products that are older than six months and that the percentage of markets in which Amazon is inactive for more than 180 consecutive days is less than 10% in all categories with the exception of Groceries and Health & Household. Moreover, our findings suggest that six-months are enough for competitors to adjust to entry; we speculate that the same holds true for exits.

We benchmark the effect of Amazon's entries against what we consider to be a natural first benchmark: the entry of the marketplace's largest merchants. The absence of transactions in our data prevents us from using revenues as the variable upon which we define "big". To circumvent this issue we leverage what we believe to be two highly correlated variables: the *overall number of ratings* received by a single seller and the portfolio of products listed by each seller on the platform, also called *storefront*. For both, we define as big third-party sellers (Big 3PS, henceforth) those merchants in the top percentile of the distribution of the seller-specific averages (the within variation). The definition gives rise to two distinct treatments, the entry of Big 3PS - CR and the entry of Big 3PS - SF, according to the count of rating and to the storefront, respectively. Entry for these sellers follows the same definition used for Amazon and as such is subject to the same limitation. Whenever we encounter products for which we do not observe the months around entry, we classify them as *Always Treated* and drop them from the event study. The entry of Big 3PS has two additional peculiarity worth outlining. First, it can happen that two or more Big 3PS enter the same market. We handle this contingency by selecting only the first, hence earliest, entry. Moreover, it can happen that Big3PS enter *after* Amazon is already part of the market. To enhance comparison we avoid these products and consider only those instances where Big 3PS enter Amazon-less markets.

4.3 Descriptive Statistics

We leverage our panel to provide causal evidence on the effect of Amazon entry on competition in the form of market share, prices and market structure. In tables 3, 4 and 5 we provide descriptive statistics on our main variables of interest and on the observable characteristics we use as control.

Table 3 splits the sample according to the products used as control, the 10% random sample drawn to overcome computational limitations, and the treated products, before and after Amazon entry. For our causal identification we rely on 9742 treated products. Our proxy for market share, the buy box share hosts unsurprisingly almost only third-party sellers in the control group, on average for 92% of the time, not reaching 100% because the buy box can be absent if no seller is listing the product with an offer of sufficient quality. For treated products, instead, 3PS occupy the buy box 86% of the time before Amazon entry and 66% of the time after, a 20% difference that will

[13] While the definition is the same for Amazon entry and Big 3PS entry, it's worth specifying that our uncertainty of the exact date of entry is dependent on Keepa's API scraping frequency. Amazon's prices are tracked at a higher frequency than 3PS. 3PS prices and offer related characteristics are scraped 26.45 times per month, on average.

Table 3: Event Study Descriptives - Amazon Entry

	Control - 10%			Amazon (Pre)			Amazon (Post)		
	count	mean	std	count	mean	std	count	mean	std
N. of Products	17468	nan	nan	9742	nan	nan	9742	nan	nan
Age	763615	31.17	20.44	58452	24.50	14.26	185098	37.00	15.18
BB Price	724483	51.36	259.30	58149	55.40	164.61	184630	53.82	163.41
Min Price	713532	49.10	444.17	57596	53.81	164.70	180755	56.07	180.63
Share of FBA 3PS	663914	0.97	0.12	45478	0.73	0.32	145716	0.71	0.32
Count of offers	641430	1.23	1.52	58256	4.08	6.73	182822	4.73	7.08
3PS first detected	767024	0.09	0.77	58452	1.00	3.66	185098	0.88	3.11
3PS last update	767024	0.07	0.67	58452	0.78	2.95	185098	0.88	2.92
Sales Rank	754205	45095.74	187333.58	58225	120254.72	403851.46	184911	52664.39	193700.82
Sales Rank Drop 1	719946	0.95	81.45	53894	0.68	40.64	179712	0.94	50.52
Sales Rank Drop 3	687484	2.63	111.57	51291	2.82	169.54	178952	3.19	266.59
Rating	727763	44.59	3.41	52967	43.89	4.74	181515	44.35	3.80
Seller N. Ratings	756059	6383.65	12212.56	57642	31906.40	85834.30	177138	30442.31	81300.75
Seller Storefront	556418	476.28	2777.08	50088	25300.79	503751.67	167602	19545.80	106003.60
Entry Age	0	nan	nan	58452	28.00	14.16	185098	28.00	14.16
3PS BB Share	766332	0.92	0.26	58452	0.86	0.31	179770	0.66	0.43
Price of Entrant	0	nan	nan	0	nan	nan	145968	45.19	130.92
bb$_p rice_3 ps$	724483	51.36	259.30	58149	55.40	164.61	155848	53.99	163.01

be mirrored by our causal analysis.

As for prices, our second main object of interest we notice the difference between the average prices of third-party seller and the price of the entrant, Amazon in this case. The former average 54$ while the latter almost 9$ less, a notable difference that suggests how Amazon's price competition can be fierce. Amazon's prices are also notably lower than pre-entry buy box prices at 55.40. Moreover, the little difference between buy box prices and average third-party sellers prices, hints toward a price competition that predates Amazon's presence. Comparing the sales rank before and after Amazon entry we notice how treated products are rising in sales, featuring an increasing demand. This is confirmed by the lower sales rank drop at one month, that displays a value of 0.64 for treated product pre-entry, suggesting that the average sales rank of these products has worsened of 64% with respect to the previous month, less than the control group and than the treated group after Amazon entry.

Finally, treated products host on average four offers both pre and after Amazon entry against an average of 1.23 for products in the control group. To further corroborate the dynamic nature of treated products we couple this statistics with the average number of offers first detected in a month and the average number of last updates in a month. We observe one new offer each month, and 0.88 last updates each month, suggesting not only a high stock of offers, but a high turnover.

We stress the element of competition before and after Amazon entry because our purpose is exactly to attempt, via a suitable identification strategy, to isolate the variation that can plausibly be attributed to Amazon entry alone. Descriptive statistics, so far, suggest competition that while correlated with Amazon entry does not preclude the possibility evidence of pro-competitive behavior on behalf of the Seattle giant might be entirely attributable to competition that could exist without it.

In table 4, we report descriptive statistics on the entry of a large third-party merchants belonging to the top 1% of all sellers, excluding Amazon, that have received the highest average number of ratings. We will refer to these as Big 3PS - CR. The rationale for our focus on these sellers is to use the count of ratings as a proxy for their revenues, as the two notoriously correlates. We exploit 4983 products that fall in this definition. Contrasting with Amazon summary statistics, Big 3PS - CR show a larger price gap

Table 4: Event Study Descriptives - Big 3PS CR Entry

	Control - 10%			Big 3PS CR (Pre)			Big 3PS CR (Post)		
	count	mean	std	count	mean	std	count	mean	std
N. of Products	17468	nan	nan	4983	nan	nan	4983	nan	nan
Age	763615	31.17	20.44	85141	20.01	12.95	202036	49.56	21.08
BB Price	724483	51.36	259.30	81595	110.81	357.74	200011	92.19	330.22
Min Price	713532	49.10	444.17	79813	109.19	368.56	198569	90.54	326.03
Share of FBA 3PS	663914	0.97	0.12	58048	0.69	0.31	152936	0.63	0.29
Count of offers	641430	1.23	1.52	82781	3.52	4.67	198566	5.81	6.87
3PS first detected	767024	0.09	0.77	85141	0.87	3.01	202036	1.04	3.42
3PS last update	767024	0.07	0.67	85141	0.53	2.16	202036	1.03	3.17
Sales Rank	754205	45095.74	187333.58	83111	154971.23	412372.45	201996	69971.29	225331.34
Sales Rank Drop 1	719946	0.95	81.45	73344	0.36	7.83	196415	2.80	542.65
Sales Rank Drop 3	687484	2.63	111.57	66127	0.94	22.40	195789	16.86	3315.37
Rating	727763	44.59	3.41	69174	43.63	5.54	194756	44.74	4.18
Seller N. Ratings	756059	6383.65	12212.56	83330	30508.75	83593.63	200448	61551.38	104265.15
Seller Storefront	556418	476.28	2777.08	69014	28655.68	183136.30	197512	16326.92	160252.80
Entry Age	0	nan	nan	85141	31.86	13.85	202036	27.28	14.02
3PS BB Share	766332	0.92	0.26	85141	0.87	0.29	198386	0.81	0.33
Price of Entrant	0	nan	nan	0	nan	nan	99063	120.94	380.77
(Avg) 3PS Price	758463	47.15	219.02	84556	113.95	350.72	197561	97.89	329.21

Table 5: Event Study Descriptives - Big 3PS SF Entry

	Control - 10%			Big 3PS SF (Pre)			Big 3PS SF (Post)		
	count	mean	std	count	mean	std	count	mean	std
N. of Products	17468	nan	nan	14933	nan	nan	14650	nan	nan
Age	76826	31.66	20.81	150582	19.75	12.75	302975	47.86	19.84
BB Price	73202	44.66	222.98	144637	101.80	349.17	297848	110.87	402.06
Min Price	71934	41.17	198.23	141987	98.36	349.68	294404	111.37	455.75
Share of FBA 3PS	67430	0.97	0.12	115310	0.84	0.27	242416	0.77	0.30
Count of offers	65111	1.19	1.31	146122	2.36	3.64	282123	3.87	5.67
3PS first detected	77289	0.09	0.73	150582	0.50	1.98	302975	0.62	2.75
3PS last update	77289	0.07	0.63	150582	0.31	1.42	302975	0.62	2.58
Sales Rank	76062	48247.21	178701.27	147085	85059.52	334001.01	302834	59204.90	252284.13
Sales Rank Drop 1	72763	1.10	50.13	130553	0.48	9.75	293672	1.84	440.72
Sales Rank Drop 3	69520	3.50	116.04	118716	1.46	33.39	292510	11.53	2712.63
Rating	73443	44.56	3.38	130522	44.03	4.32	290984	44.58	4.02
Seller N. Ratings	76233	6576.67	12134.01	146978	24804.53	82191.97	298708	35153.69	89776.10
Seller Storefront	56007	405.02	1833.09	106519	19517.90	941461.68	281680	27698.45	144995.23
Entry Age	0	nan	nan	150582	32.46	13.56	302975	27.59	13.88
3PS BB Share	77235	0.93	0.25	150582	0.90	0.27	302578	0.89	0.28
Price of Entrant	0	nan	nan	0	nan	nan	79300	191.31	514.38
(Avg) 3PS Price	76332	42.72	203.74	149752	106.52	349.48	301387	114.38	393.77

between their average price and the buy box price, quantified at 28$. Before and after entry, the change in buy box share of 3PS, in this case, all merchants but the entrant, is of 6%, from 0.81 before to 0.81 after. These sellers still compete in markets that host, on average 5.81 offers and, they enter in market where demand is growing at an even faster pace than those targeted by Amazon. The sales rank grows, on average only of 36% with respect to the 64% previously observed for markets that attract Amazon's competition. The demand growth is evident also from the sales rank, roughly halving after entry. These markets host a comparable share of sellers adhering to the FBA program 63% pre-entry and 69% post. We count 276 unique Big 3PS - CR.

Finally, we look at table 5, where we report summaries of observable characteristics related to the entry of merchants belonging to the top 1% of all sellers, but Amazon, when ranked by their average storefront, their portfolio. The differences with observables reported in table 3, magnify even further. We rely on these sellers to have a further point of comparison that departs from revenues and focuses on the scope of a third-party seller. Prices for these sellers are sensibly higher with respect to buy box prices and are, even higher in a relative sense with respect to those found in the other two tables. Even in the aggregate the difference in buy box share of other 3PS sellers barely moves from 0.9 to 0.89. We count 406 unique Big 3PS - SF.

Figure 1 plots the distributions of entry for our three treatments over time, as found in the sample we use for our analyses of section 6. The top panel depicts entries over product ages, with all three type of sellers, Amazon, Big 3PS - CR and Big 3PS - SF, entering early in a product life-cycle. The same figure also depicts how staggered our treatment is. Each bar in the plot represents the number of treated products by age, which we will compare to non-treated products of the same age. The bottom figure in the panel reports entries by calendar-time in the form of year-months. We notice that a large spike for Big 3PS - SF in the first quarter of the second year of the COVID-19 pandemic, suggesting that these entries might relate to brick-and-mortar store moving their operations online. We also notice two spikes for Amazon, corresponding to black Friday of 2021 and Prime day of 2022.

Figure 1: Entry of Amazon and Big 3PS over age and calendar-time.

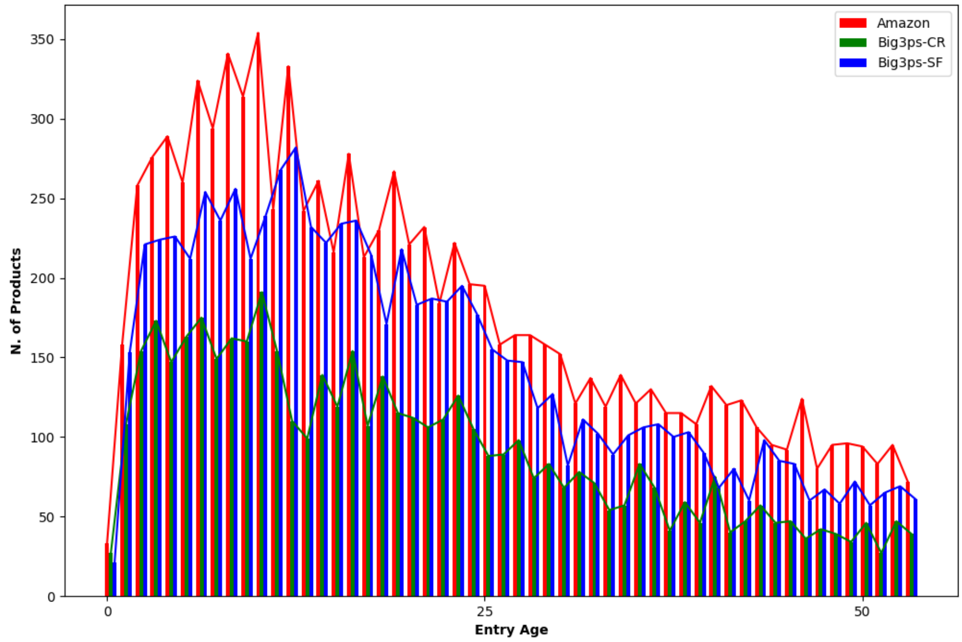

1.1 Event study of entry on buy box share of 3PS

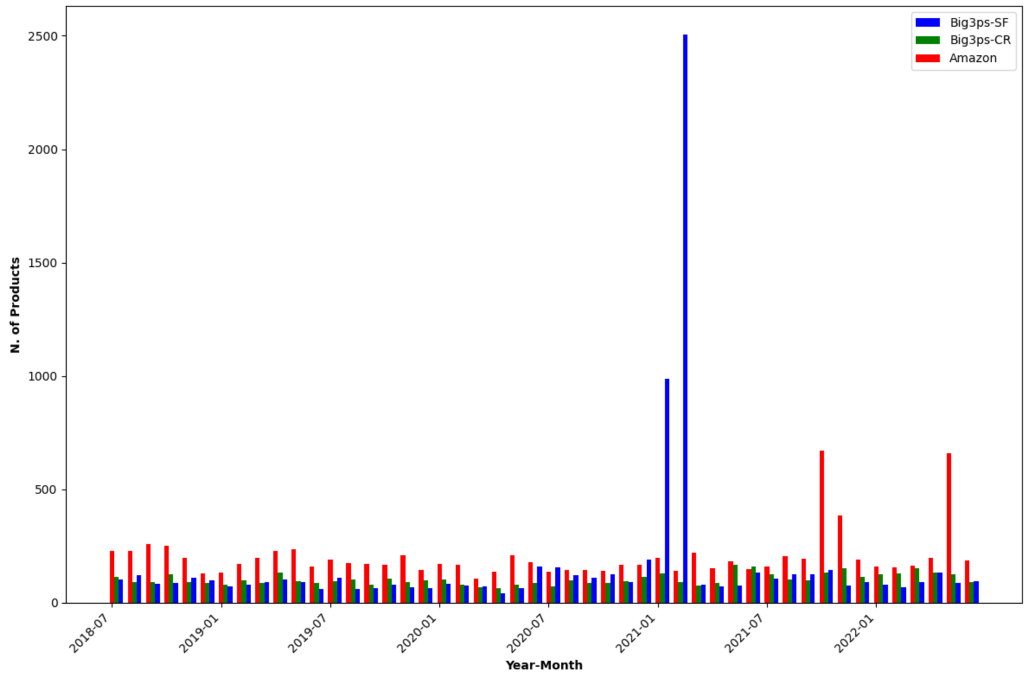

1.2 Static ATT by Root Category

Notes: Panel 1.1 Depicts the frequency of products in which Amazon, the top1% of large third-party merchants ranked by number of ratings received, Big 3PS (CR), and the top1% of large third-party merchants ranked by storefront, enter by age of the product at the month of entry. Figure 1.2 reports the frequency of products in which Amazon, the top1% of large third-party merchants ranked by number of ratings received, Big 3PS (CR), and the top1% of large third-party merchants ranked by storefront, enter by calendar-time, that is the year-month at time of entry.

5 Identification Strategy

We aim to investigate the causal impact of Amazon on competition. To do so, we exploit the staggered entry of Amazon in different product markets as a key source of identifying variation. To do so, we use a staggered difference-in-difference comparing several outcome variables related to competition on the Amazon platform, around entry age, in product markets (ASINs) in which Amazon entered, to a control group of product markets in which Amazon is absent. Our estimand is the Average Treatment Effect on the Treated, depicting the impact that Amazon entry in a market has, on average, on our target variable of interest. The baseline equation we estimate is the following:

$$y_{i,a} = \alpha_0 + \alpha_i + \alpha_a + \alpha_c + \alpha_t + \sum_{e=1}^{60} \sum_{l=-6,\neq -1}^{18} \delta_{e,l} \cdot \mathbb{1}\{E_i = e\} \cdot D^l_{i,a} + X_{i,a}\beta + u_{i,a}, \quad (1)$$

where i is *item* (product) and a is *age*. As target variable $y_{i,a}$ we use several outcomes that measure competition in terms of prices and market structure and for which we postpone the details for each corresponding subsection in section 6. α_0 is a constant, α_i is an item fixed effect, α_a is an age fixed effect, α_c is a level-2 category dummy and α_t is a year-month specific variable. $\mathbb{1}\{E_i = e\}$ is an indicator function of whether amazon has entered market i at age e establishing the cohorts[14] upon which we assume the heterogeneity is path dynamics. The indicator function is interacted with $D^l_{i,a}$ an indicator function of relative period l with respect to age of Amazon entry and $X_{i,a}$ is the entry corresponding to pair $\{i, a\}$ in a matrix containing time-varying controls such as ratings or the number of reviews. Our parameters of interests are $\delta_{e,l}$ capturing the impact of Amazon entry at age e in relative period l on $y_{i,a}$. For instance, $\delta_{10,-6}$ measures the effect of Amazon entry in product markets that are 10 months old, 6 months before entry actually occurred. Notice, that for each cohort we use a time window ranging from -6 to 18. This is a larger time span than comparable studies. To give a reference point, Crawford et al. (2022) study a time span ranging from 5 months before Amazon entry to 10 months after. We follow Crawford et al. (2022) by defining cohorts based on product age rather than calendar time; as they show the dynamic effect is not homogeneous across products at different stages of their life-cycles. By defining cohorts, or groups, by age we assume that, beside the cross-sectional unobservable heterogeneity accounted for by product fixed-effects, the evolution of the outcome and hence the treatment dynamics, are shared across products that saw Amazon joining as a first-party seller when they were of the same age. Moreover, by constructing the panel using age as the time dimension, we compare treated products with non-treated products at the same stage of their life-cycle.[15]

Equation 1 is estimated using the method proposed by Callaway and Sant'Anna (2021) (CS, henceforth), which is our preferred method. In appendix A we compare three estimators, a standard Two-Way Fixed-Effect, the Interactive Weighted estimator

[14] We use "cohort" and "group" as synonyms. Both terms refer to the variable that defines the observation with a shared and homogenous dynamic effect when sharing the same treatment time.

[15] For instance, define cohort $e = 10$, made by all products that saw Amazon's entry as competitor when they were 10 months of age. When computing the treatment effect $\delta_{e,l}$, the reference period will be $l = -1$ and hence age 9. The control group for this cohort will also have age 9 as reference period, allowing us to control for unobservable characteristics shared by products of the same age, rather than by products living in the market at the same point in time

proposed by Sun and Abraham (2021) and our preferred method. Both estimators proposed by Callaway and Sant'Anna (2021) and Sun and Abraham (2021) have been proposed to address the implicit homogeneity in treatment effects across all units and time periods that TWFE estimators assume and that can lead to incorrect inference when treatment effects are actually heterogeneous. We follow the innovations proposed by Callaway and Sant'Anna (2021), part of a more concerted effort to improve the credibility of the difference-in-difference method and surveyed in Roth et al. (2023). The parameter δ_k in equation 2 is commonly interpreted as the ATT that Amazon entry has on target variable $y_{i,a}$, k months after joining the market. This interpretation, though, is tied to the implicit aforementioned homogeneity assumption. To resolve it and accommodate heterogeneity in the cohort-specific evolution of the treatment effect, Sun and Abraham (2021) interact each cohort with relative treatment time, multiplying the number of ATT estimated from K to KxE. A similar but different solution is used by Callaway and Sant'Anna (2021); they recover an ATT that is specific for group-time combination where, in their nomenclature group is a synonym for cohort of treatment while time refers not to relative time, but to each value of the time-dimension in the panel, in our case, ages. Their estimator, therefore increase the number of ATTs from K to ExA. Both methods, therefore, suggest different way to aggregate each estimated ATT, to summarize the treatment effect, which we delegate to appendix A.

All the results on dynamic effect in the main text are presented as event-study plots. Relying on Callaway and Sant'Anna (2021) as the estimation method, the plot shows an aggregation that averages each group-time specific ATT into a relative-age-specific treatment effects. It expresses the average effect that the treatment, Amazon entry, has on the outcome variable for products that have experienced it for k periods. We refer the reader to equation (3.4) in section 3.1 of Callaway and Sant'Anna (2021). All estimates are carried via the `did` package in R.

For each outcome of interest we also highlight differences in the effect that Amazon entry has on each single root category separately. To do so we rely on a static specification, where the average treatment effect on the treated is aggregated from the group-time combinations on which the CS estimator is based on. These estimates depict the weighted average of all post-entry treatment effect. We refer to equation (3.7) in Callaway and Sant'Anna (2021).

Importantly, by choosing Callaway and Sant'Anna (2021) as our preferred estimator, we counter recent literature that uses the dynamic TWFE as their main estimator and relies on Sun and Abraham (2021) to validate their estimates and explore possible biases due to negative weighting (Crawford et al., 2022; Fischer et al., 2023). Our choice is driven by the fact that the $\delta_{e,l}$, the cohort-specific treatment effect, when estimated via an unconditional specification that relies only on unit and time fixed-effects and no controls, should result in the same value using CS and SA, we find this not to be the case. In order to decide which one of the two to rely on we notice that both method address the potential biases due to heterogeneity in the dynamic of the treatment effect by estimating each single $\delta_{e,l}$ in a way that should be equivalent to estimating separately each two-by-two difference-in-difference.[16] Motivated by the differences in the estimates provided by the two methodologies, and upon experimentation, we find this to be true only for Callaway and Sant'Anna (2021) because unlike SA, in the uncondi-

[16]For instance, on the subsample containing only observations in the control at reference period $l = -1$ and at target period l, and observations in the treated group at reference period $l = -1$ and at period l treated at the same period e

tional case, it relies on a comparison in mean, rather than an OLS estimation for each two-by-two, that should, under these conditions, once again be equivalent. This finding is not surprising given our setting. Our panel is subject to compositional changes due to products entering and leaving the market creating an unbalance that neither of the two contributions addresses directly, but that plays a role in our estimates.[17] We mitigate the composition of the panel by, in the control group, retaining for our sample only products that last at least for two consecutive years (24 months) out of the six we include. In the treatment group, motivated by our interest in short vs long-term effects we restrict our sample only to products where the entrant is active for all the 18 months period that follows its entry. Choosing Callaway and Sant'Anna (2021) as our main method also supports our need to control for seasonality and category-specific trends that require us to introduce controls in equation 1. We, therefore, rely on a *conditional* parallel trend and only Callaway and Sant'Anna (2021) show how their estimator is consistent and unbiased under this specific assumption, while Sun and Abraham (2021) estimator has to rely on an *unconditional* parallel trend assumption.

6 Effect of Amazon Entry

6.1 Buybox

6.1.1 Buybox Share

We start our analysis from the Buy Box. As introduced in Section 4, the Buy Box allows Amazon to direct users' attention to a box containing both the "Add to Cart" and "Buy Now" buttons. According to the FTC complaints against Amazon Inc., filed on September 2, 2023[18], 98% of transactions happen via the Buy Box. We therefore consider the time a seller spends in the top spot as a natural proxy for market share and ask to what extent, upon its entry, Amazon is able to capture it.

We do so by estimating specification 1, using as a dependent variable the sum of the fraction of time spent in the Buy Box by each third-party merchant m participating in market i. Denote by $M_{i,a}$ the set of all third-party sellers participating in product market i when the latter is of age a. Moreover, denote by $bb_{i,m,a}$ the total number of seconds that merchant $m \in M_{i,a}$ spends in the Buy Box of product i in a given month, and by $TotSec_{i,a}$ the total number of seconds in month a for product i. Our outcome variable is constructed as $y_{i,a} = bb^{3PS}_{i,a} = \frac{\sum_{m \in M_i} bb_{i,m,a}}{TotSec_{i,a}}$. Importantly, from the set $M_{i,a}$, we exclude the time spent in the Buy Box by Amazon and those intervals of time where no seller qualified for the Buy Box—that is, where there is no Buy Box available for product i. The result is the share of time spent each month in the Buy Box by any of the third-party merchants populating the market.

Figure 2.1, therefore, depicts in an event-study plot the effect that Amazon's entry has, on average, on the market share of third-party sellers, net of the time in which the Buy Box is absent. In the event study plot we overlay results from the same specification estimated on three different treatments: the entry of Amazon, in red, the entry

[17]We are currently working on a simulated environment with a data generating process tailored to show and magnify the role that panel composition plays on the two methods. A dedicated section will be added to this work in a future iteration

[18]https://www.ftc.gov/legal-library/browse/cases-proceedings/1910129-1910130-amazoncom-inc-amazon-ecommerce, last accessed August 30, 2024

of large third-party merchants according to their count of ratings received, in green, and the entry of large third-party merchants according to the number of product sold on the platform, in blue. Beginning by Amazon entry we appreciate how, the Buy Box share of third-party sellers, decreases on average 15% by the second month and continues to decline throughout the period, reaching 23% after 18 months. Upon visual inspection the parallel trend holds as pre-entry coefficients overlap the zero-effect line. As for the entry of a Big 3PS, the Buy Box share of 3PS also decreases, on average, but unlike what we observe following Amazon's entry, the effect is not as persistent. Specifically, the blue line describes the dynamic effect of Big 3PS entry according to their portfolio. Upon entry, they acquire at most 4% of the market share by their second month of tenure, but their share quickly reverts to around 2%. We emphasize that this effect cannot be attributed to their exit, as we select only those markets where Big 3PS remains for the full 18 months, mirroring our choice for Amazon's entry. A different pattern, however, is observed for Big 3PS based on the number of ratings received (green line), which shows that these sellers are capable of acquiring up to 18.5% of the market share, closer to Amazon's, by their second month after entry. However, this share halves over the subsequent 18 months.

Figure 2: Effects of Amazon and Large Sellers' Entry on Buy Box Share of Third-Party Sellers

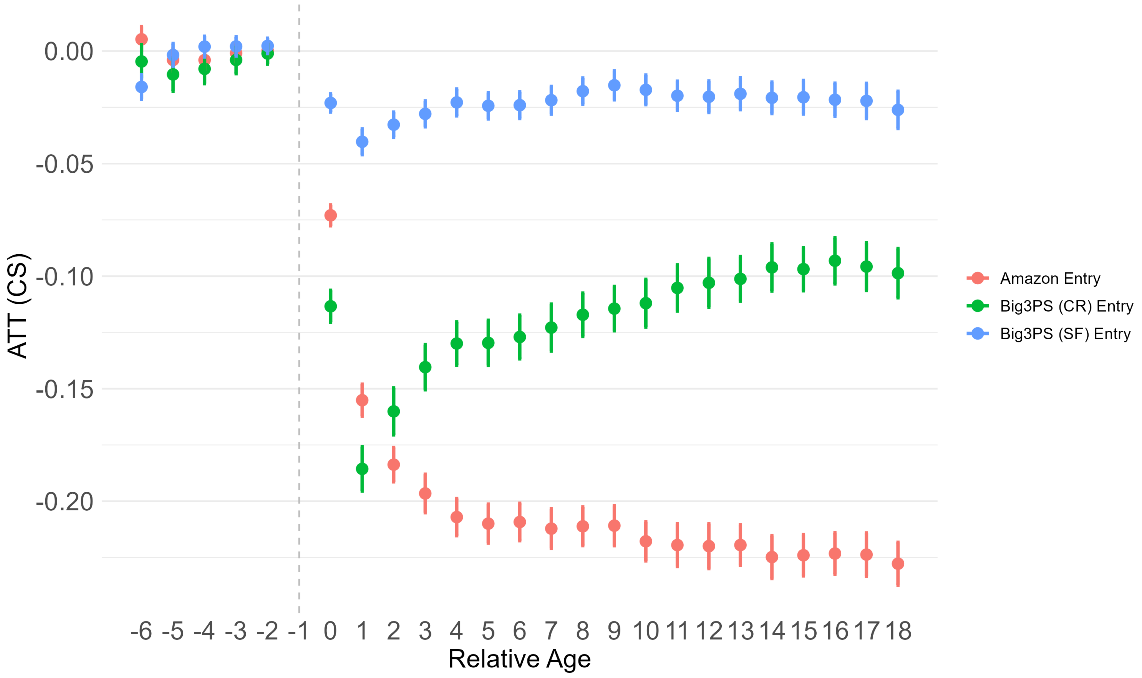

2.1 Event study of entry on buy box share of 3PS

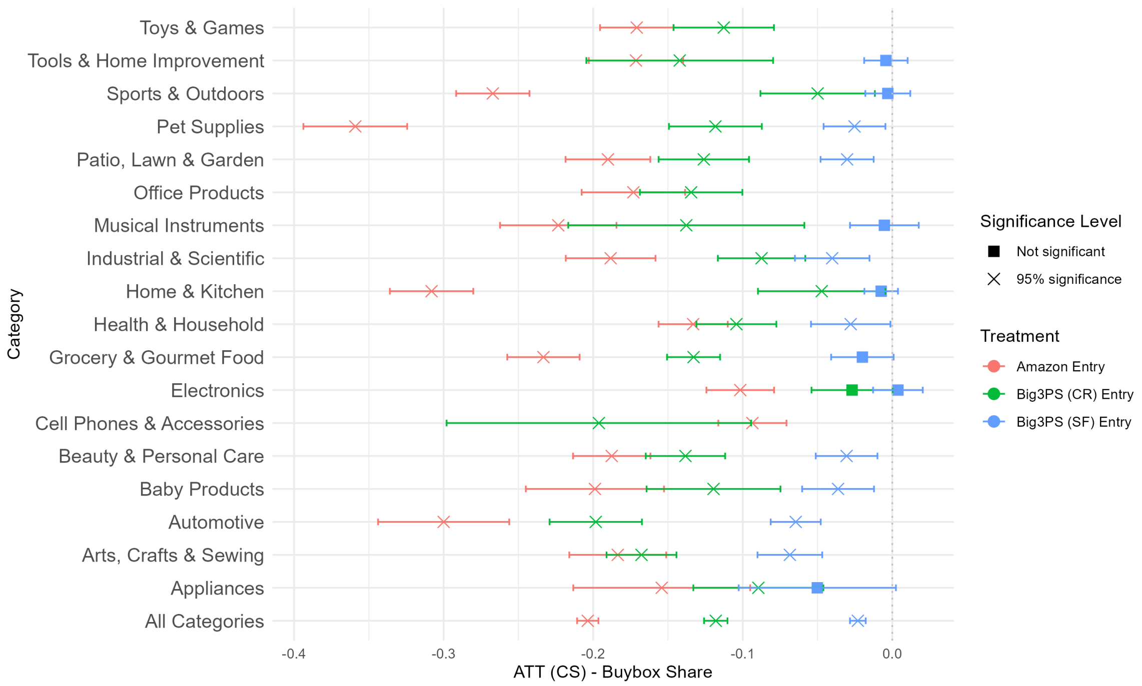

2.2 Static ATT by Root Category

Notes: Panel 2.1 Compares the effect of entry by large third-party merchants to to Amazon's. The horizontal axis represents the age relative to entry, and the vertical axis shows the Average Treatment Effect on the Treated of entry on the buy box share of all third-party sellers. Figure 2.2 reports the static ATT, aggregated as explained in section 5, by root category. All estimations are carried out using the doubly-robust estimator of Callaway and Sant'Anna (2021). Standard errors are clustered at the product level.

We also decompose the effect by root category. In particular, we leverage equation (3.7) in Callaway and Sant'Anna (2021) to obtain a single parameter reporting a weighted average of all group-time average treatment effects, with weights proportional to group size. This can be interpreted as the Average Treatment Effect on the Treated (ATT) of a static version of equation 1. Figure 2.2 shows the static ATT for each category separately, including "All Categories," where the ATT is simply the weighted average of the coefficients that in the event study plot of Figure 2.1 are aggregated by relative time to entry. In all categories, Amazon's red dots show a statistically significant decrease in the Buy Box share of 3PS, but with magnitudes varying from 36% in Pet Supplies to 9% in Cell Phones & Accessories.

It is also interesting to benchmark Amazon's effects against those of Big 3PS. In particular, Amazon's advantage is evident in Sports & Outdoors, Pet Supplies, and Home & Kitchen, where it outperforms Big 3PS (CR) by more than 20%. These differences will be a recurring theme throughout our analyses and highlight how extrapolating from one category can be misleading.

In the aggregate, and in the short run, Amazon's ability to acquire market share is comparable to the top 1% of 3PS as measured by their number of ratings. What differentiates Amazon is, therefore, its ability to retain its market share over the long run. Amazon's ability to acquire, on average, a fifth of market share and to retain it, can be the result of multiple competing mechanisms both competitive and anti-competitive. Among the latters, we outline self-preferencing, an anti-competitive behavior that, given offer characteristics, would raise the prominence of Amazon's offers even though 3PS' would guarantee an equal amount of revenues and hence fees, tilting the playing-field in favour of Amazon's retail branch. The same findings though can also be explained by competition, on either prices or quality of the offer.

In Figure 3.1 we expand on the latter, reporting, in an event study plot, the effect of Amazon entry on the buy box share of 3PS, split by fulfillment service. We estimate equation 1 twice, with outcome $y_{i,a} = bb_{i,a}^{FBA}$, the fraction of time in which a third-party FBA sellers has won the buy box of product i of age a and $y_{i,a} = bb_{i,a}^{FBM}$, , the fraction of time in which a third-party FBM sellers has won the buy box of product i of age a. Counter to expectation, in the average market, a seller relying on Amazon logistic service and a seller relying on independent logistic services, face a comparable loss of share. By the fourth month after treatment, the former lose 14% of the market, while the latter 17%.

Figure 3: Effects of Amazon Entry on Buy Box Share of 3PS by logistic service

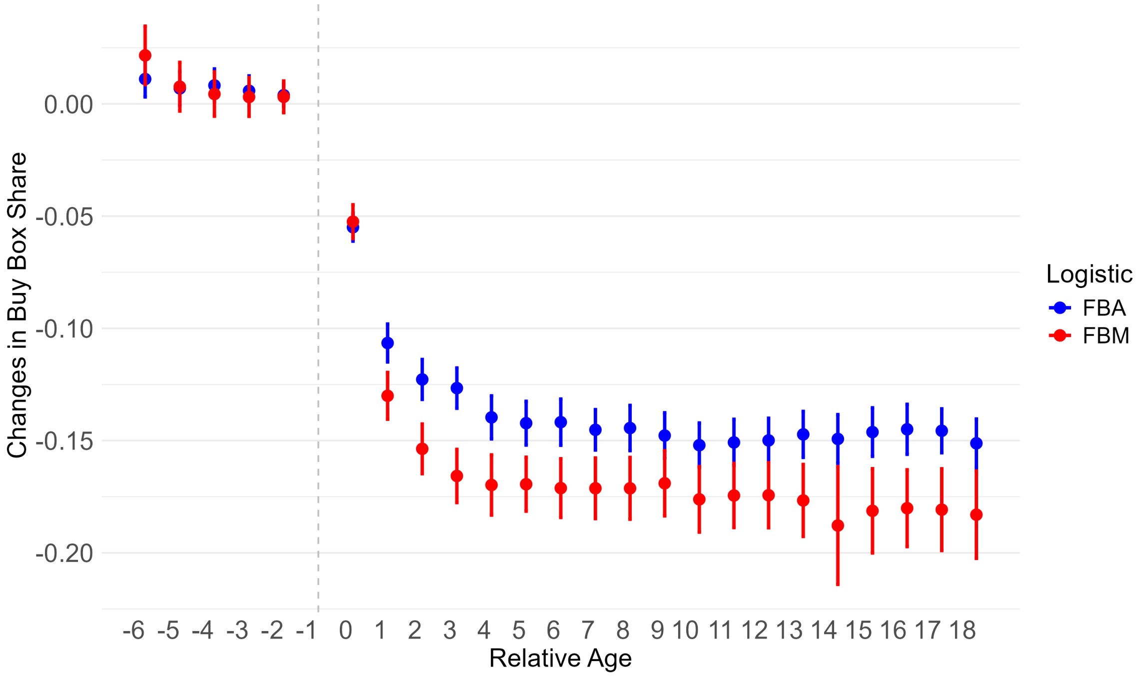

3.1 Event study on Buy Box Share of FBA and FBM Sellers

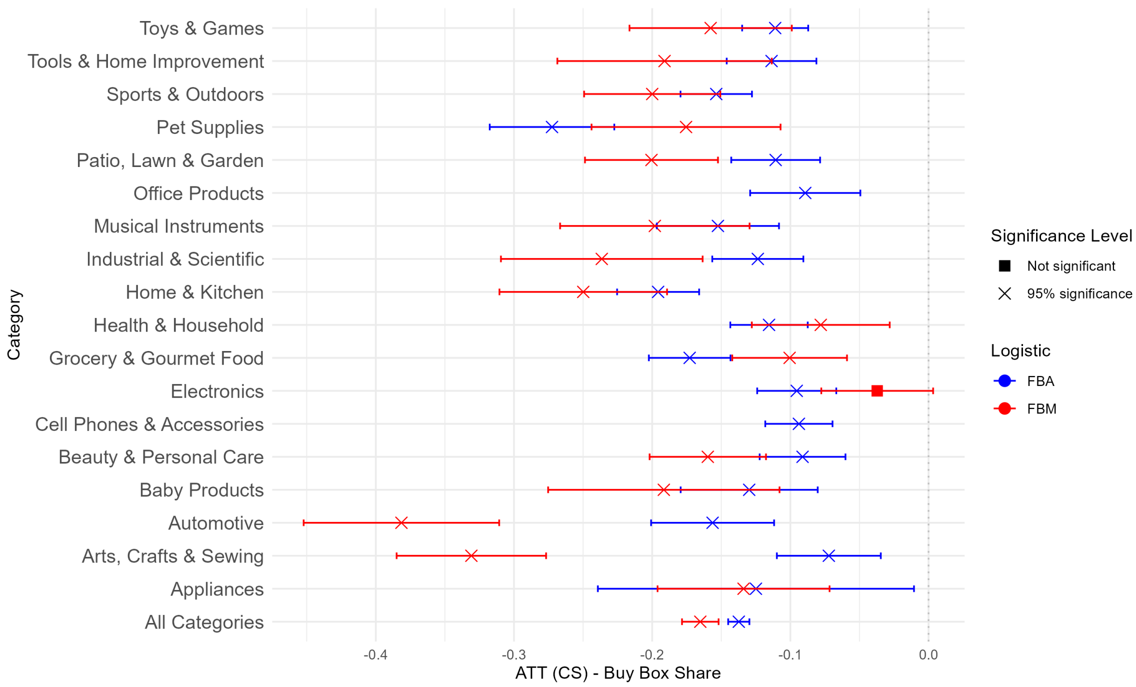

3.2 Static ATT by Root Category

Notes: Figure 3.1 shows the event study plot of specification 1 on the buy box share of FBA and FBM sellers, separately. Figure 3.2 depicts the static ATT aggregated using the method discussed in section 5. Standard errors are clustered at the product level.

6.1.2 Buybox Price

A crucial dimension over which sellers compete on the Amazon marketplace are prices. We take advantage of our rich dataset to explore the effect that Amazon entry has on buy box prices, that is, on the average price that, at any given moment, any one of the sellers occupying the featured spot, lists. Due to the crucial role that the buy box plays, Keepa, our data provider, monitors and scrapes it at a high frequency recording both the seller and the price. For our purpose we estimate equation 1 using as outcome $y_{i,a} = \log(p^{bb}_{i,a} \cdot \text{CPI}_t)$, where the price $p^{bb}_{i,a}$ is really an average of all prices that the featured offer has had in month a for product i, weighted by the time, in seconds, the product has been listed at a particular price [19]. As per convention, we account for inflation by deflating prices. We interpret buy box prices as a proxy for transaction prices, assuming it as the price that any consumer converting a visit to the product page to an actual purchase would pay, had the visit happened at any moment during month a. At this stage, we are interested in the ATT of Amazon's entry on transaction prices, to understand the extent to which it benefits consumers.

Similarly to buy box shares, in figure 4.1 of panel 4, we present event study plots that overlay Amazon entry and the entry of the top 1% of third-party merchants based on their number of times they have been rated and based on their portfolio of products sold in the Amazon digital marketplace, in green and blue, respectively. The dynamic effect of Amazon entry is negative and statistically significant throughout the 18 months after treatment, reaching its highest, in magnitude, at -6.6% by the third month. From the fourth to the eighteenth month, instead, the magnitude of the effect decreases suggesting a dynamic pattern that differs in the short-run and in the long-run. Products exposed to Amazon's competition for eighteen months, in fact, see their buy box prices lower, on average of -2.5%, if compared to pre-entry. Upon visual inspection the parallel trend assumption is satisfied; all standard error bands in the six-months preceding entry overlap with the zero line, suggesting the difference in buy box prices between treated and control group are not significantly different from zero.

Once again, we benchmark the effect of Amazon's entry with the effect that the first appearance of a top 1% third-party seller have on buy box prices. Figure 4.2 depicts in blue the effect of Big 3PS according to their number of ratings and in green the on of Big 3PS according to their storefront, with Amazon's effect still in red. A difference between short-run and long-run effect is present for Big 3PS entries as well, but the effect is negative and statistically significant at the 95% only when a seller belonging to the top 1% most rated merchants joins the market, reaching its highest at -2.2% by its second month; by month 8 the effect is not different from zero, and reverts from negative to positive by month 11, until the end of the event study window of eighteen months. Merchants in the top 1% according to their portfolio, instead, have a negligible negative impact even in the short run, of less than 1% in the first month after entry and not significantly different from zero until the tenth month, when their effect also reverts from negative to positive.

We notice that the dynamic effects depicted in figure 4.1 hide large heterogeneities

[19] For instance, if a product has been listed at a price of 10$ for 10 days and a price of 13$ for 20 days, the price we will use is (10x10+13x20)/30=12 and not the simple average (10+13)/2 = 11.5$. In our code, 10 days and 20 days are really the duration between two timestamps measured in seconds, and the denominator is the number of seconds in the month, which, of course depends on the month

across categories. In figure 4.2 we plot the static ATT, resulting from the Callaway and Sant'Anna (2021) estimator. The negative effect on prices is driven only by 5 categories out of 18: Pet Supplies, Health & Household, Grocery & Gourmet Food, Beauty & Personal Care and Art, Crafts & Sewing. For the remaining categories the static ATTs are either not significantly different from zero or even positive in Musical Instruments, Electronics and Cell Phones & Accessories.

Amazon, therefore distinguishes itself once again from the marketplace top competitors. Its entry causes transaction prices to decrease in the short run and in the long run, albeit at different magnitudes, with a more pronounced advantage for consumers in the short-run. Importantly, the effect is not persistent and it appears to be on a mean-reversal path by month 6. In order to understand whether the change in buy box prices is entirely due to Amazon's prices or whether Amazon's triggers a downward spiral in prices, we study entrants' prices and buy box prices of third-party sellers.

Figure 4: Effects of Amazon and Large Sellers' Entry on Buy Box Prices

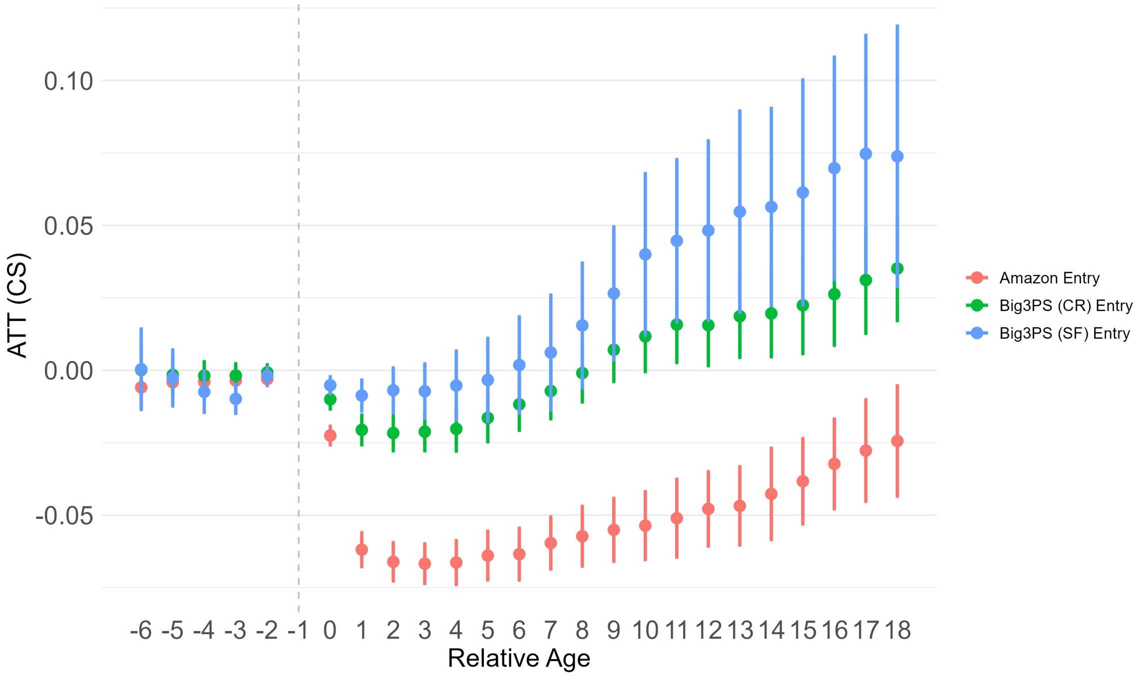

4.1 Event study of entry on buy box prices

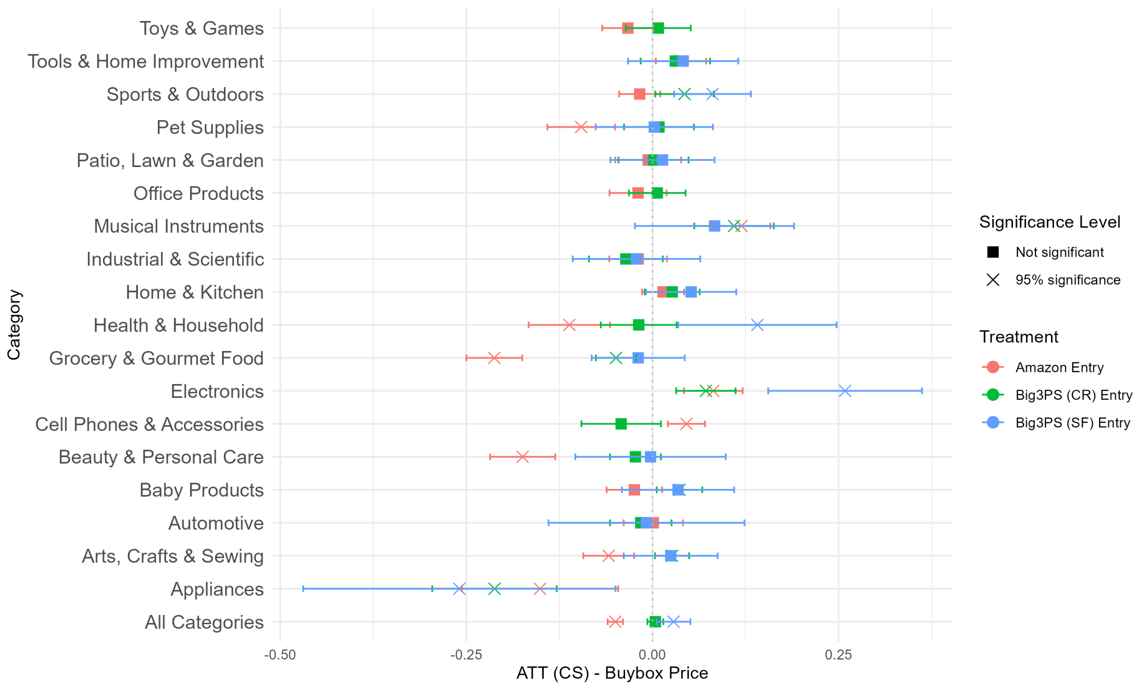

4.2 Static ATT by Root Category

Notes: Panel 4.1 compares the effect of entry by large third-party merchants to Amazon's entry using Callaway and Sant'Anna (2021). On the horizontal axis we display the age relative to entry, and the vertical axis shows the dynamic ATT on Buy Box prices as defined in the text. 4.2 depicts the static ATT, aggregated as discussed in section 5 of the doubly robust Callaway and Sant'Anna (2021) estimator, separately for each root category. Standard errors are clustered at the product level.

6.2 Prices

6.2.1 Entrant Prices and Buy box Prices of 3PS

We dissect buy box prices by studying separately entrants and third-party sellers that wins the buy box. Our goal is to understand the extent to which Amazon entry contributes to competition. We begin with entrant prices. We estimate equation 1, with outcome $y_{i,a}$ containing Amazon's prices for all post-entry periods in treated markets, while pre-entry periods and prices in the control group are filled with buy box prices. To formalize, define as $Pre_{i,a}$ a product-specific dummy variable taking value of one if the entrant $k = \{\text{Amz}, \text{3B-CR}, \text{3B-SF}\}$ is not a competitor in market i at age a and zero otherwise. The dependent variable used in 1 is $y_{i,a} = \log(p^{Ent}_{i,a} \cdot CPI_t) = \log(p^{bb}_{i,a} Pre_{i,a} + p^{k}_{i,a}(1 - Pre_{i,a}) \cdot CPI_t)$, the log of the deflated price of the entrant k in product i at age a.

In Figure 5.1, the event study plot depicts how Amazon's prices, by the fourth month of tenure in the market, are 6.5% lower than buy box prices before entry. Moreover, the decrease is persistent throughout our period of observation, until the eighteenth month. When compared to other large merchants, the difference become even starker. In the same figure, both type of large merchants do not mirror Amazon's price competition but, on the contrary, list prices that are well above the ones listed in the buy box before entry. Big 3PS (CR), on average, enter markets by listing product with prices of 5% higher and raises them over time reaching an 18% increase by month eighteen. Big 3PS (SF) instead, report prices that are, on average, 40% higher, suggesting once again that portfolio size is associated with sellers interested only in having an online Amazon store, rather than to actively compete.

When looking separately at the static ATT of each root category, in figure 5.2, it becomes apparent that, upon entry, Amazon does not compete on prices everywhere. Only for eight out of eighteen categories we find a a statistically significant negative coefficient, ranging from -4% for Sports & Outdoors to -28% for Appliances. For Musical Instruments, Cell Phone & Accessories and Electronics, Amazon prices are higher than buy box pre-entry, showing 17%, 14.1% and 12.7%, respectively. For the remaining seven categories coefficients are not different from zero at the 95% level of statistical significance, suggesting that in these categories, on average, Amazon does not compete on prices. We notice that in all the categories mentioned, Amazon causes 3PS time in the buy box to decrease as depicted in figure 2.2. Our findings, therefore, suggest that Amazon's entry strategy does not necessarily relies on lower prices, but it's tailored to the market it targets. We also point out that the algorithm that determines the featured seller is unknown; Amazon does not disclose the parameters that determine the way in which offers are ranked. We point to Raval (2022) contribution, suggesting that an Amazon premium exists only in categories that are not featured in this work, allowing us to conclude that the seven categories in which prices are either not lower or even higher than buy box prices before entry, Amazon must compete by increasing the quality of the delivery service, despite the presence of FBA sellers.

We now look at how Amazon entry affects the buy box prices of third-party sellers. To do so, we use as outcome the average price that 3PS list when winning the buy box, weighted by the time spend as featured sellers. Formally, $y_{i,a} = \log(p^{bb,3PS}_{i,a} \cdot \text{CPI}_t)$ where $p^{bb,3PS}_{i,a}$ is the average price of any third-party seller winning the buy box of product i of age a, weighted by the time, in seconds, each price spends as the buy box price and divided by the total seconds spent in the buy box of product i of age a by any 3PS.

Figure 6 depicts, in an event study plot, the effect of Amazon entry on 3PS buy box prices, in black, on buy box prices of FBA 3PS in blue and of buy box prices of FBM in red; we will discuss the last two in two paragraphs. Despite the misleading scale of the plot, the effect largely overlaps with the pattern found for all buy box prices in figure 4.1. It is negative in the short-term, with a -6% decrease already in the first month after treatment, but it reverts over the long run; by the month six the effect is - 5.3%, by month twelve -2% and by month eighteen is not significantly different from zero. The effect on what we interpret to be a proxy for transaction prices appears transitory, benefiting consumers only in the short run.

We ask to what extent a decrease in buy box prices of 3PS is due to them engaging in price competition with Amazon, or whether it is due to a compositional change in the market structure. Amazon's ability to penetrate, on average, 20% of buy box time translates to less time spent in the buy box by 3PS. We would not appreciate a reduction in prices if, to Amazon entry, would correspond a proportional shrinkage of the buy box share of the remaining third-party seller; this is the reason why we opt for weighted averages. A price decrease can be due only to either, 3PS engaging in price competition to curb Amazon's buy box share, or to Amazon's 20% of buy box share shifting in a disproportional fashion from 3PS with less competitive price offers than the Seattle giant, reducing the weighted average price. To answer this question we split third-party seller average prices by logistic service, to appreciate how they change, within quality of the service.

Figure 6 overlays the event study plots resulting from estimating equation 1 with the weighted averages of third-party sellers prices, split by fulfillment service, on the left-hand side. A stark difference between the two types of fulfillment emerges. Buy box prices of FBA 3PS are unaffected by Amazon entry as all coefficients are not significantly different from zero. Buy box prices of FBM third-party sellers, instead, increase steadily and with large magnitudes over the eighteen months after treatment. The error bands also show how that the latter estimate rely on a small sample; we will soon corroborate so, when looking at market structure. Importantly, in neither cases prices decrease. The cause of the decrease in buy box prices appreciated in figure 4, therefore, is to be found in the compositional change that Amazon's share of buy box time triggers in the remainder of the market, rather than in 3PS engaging in price competition. We, therefore transition an analysis of market structure next. We close this subsection, though, by highlighting the lack of elasticity to Amazon's presence of FBA sellers. The flat response suggests that FBA sellers might be price constrained, that is, unable to engage in price competition. For them, reducing prices might simply mean negative profits. While worthy of further investigation, we refrain from expanding on this point.

Figure 5: Comparison of Amazon and Large Sellers' Prices with Buy Box Prices

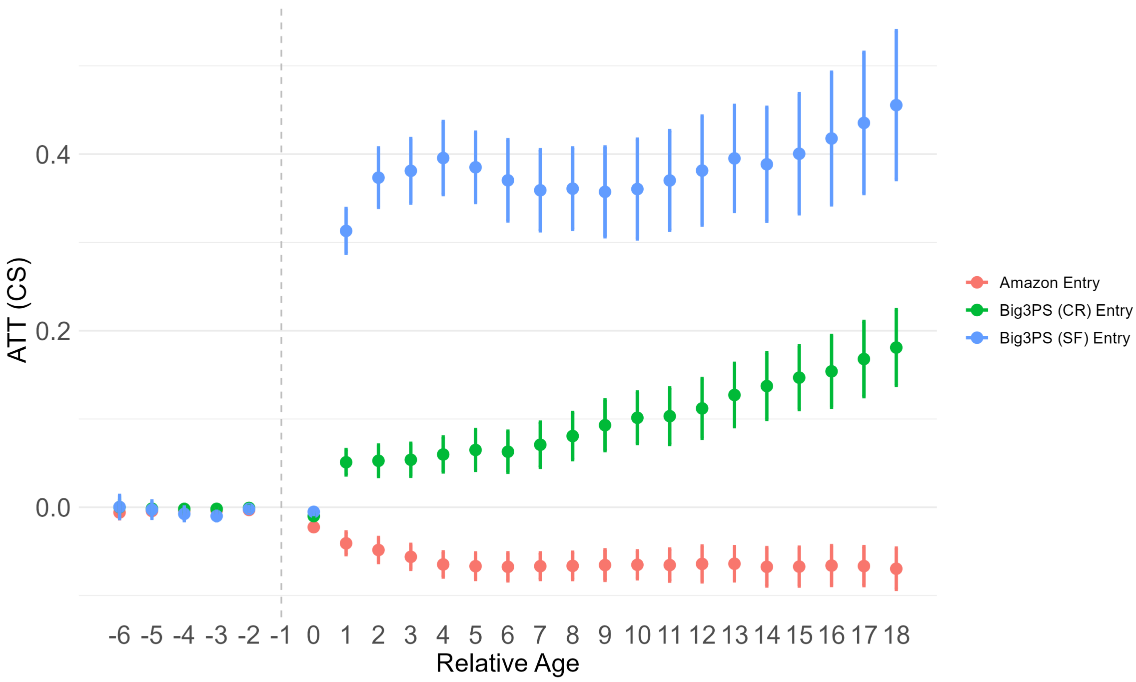

5.1 Amazon Entry by estimation method

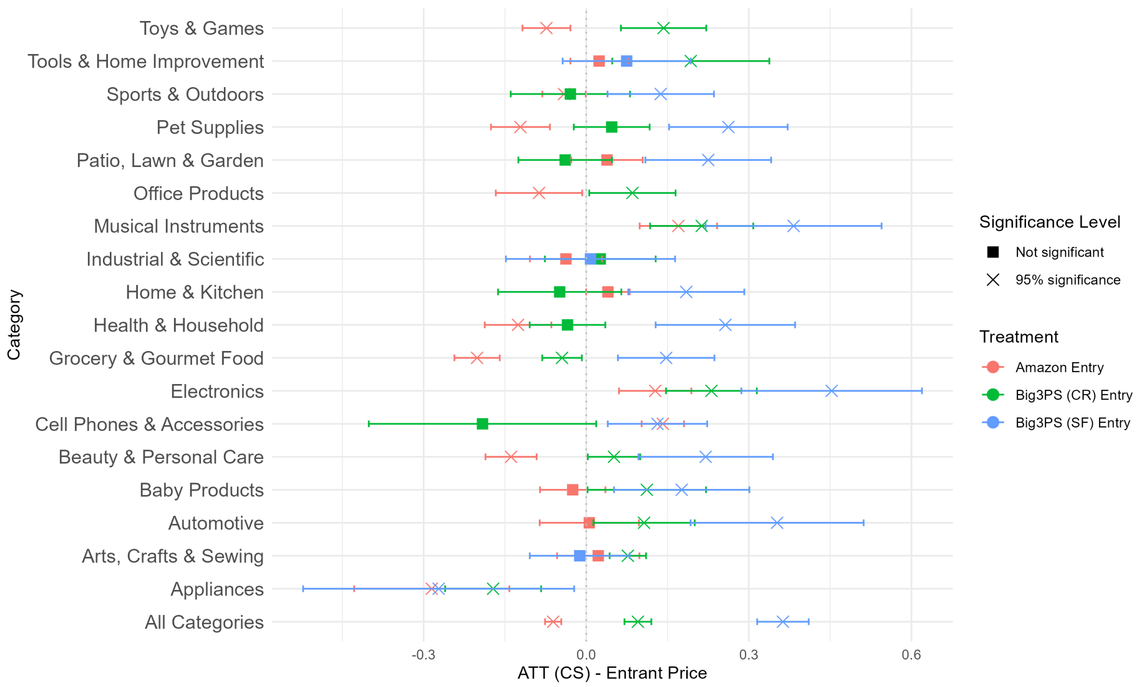

5.2 Large Sellers vs Amazon Entry

Notes: Panel 5.1 compares the prices of the entrant to the buy box prices pre-entry, by estimating equation 1 using Callaway and Sant'Anna (2021). 4.2 depicts the static ATT, aggregated as discussed in section 5 of the doubly robust Callaway and Sant'Anna (2021) estimator, separately for each root category. Standard errors are clustered at the product level.

Figure 6: Effects of Amazon Entry on Buy Box Prices of Third-Party Sellers

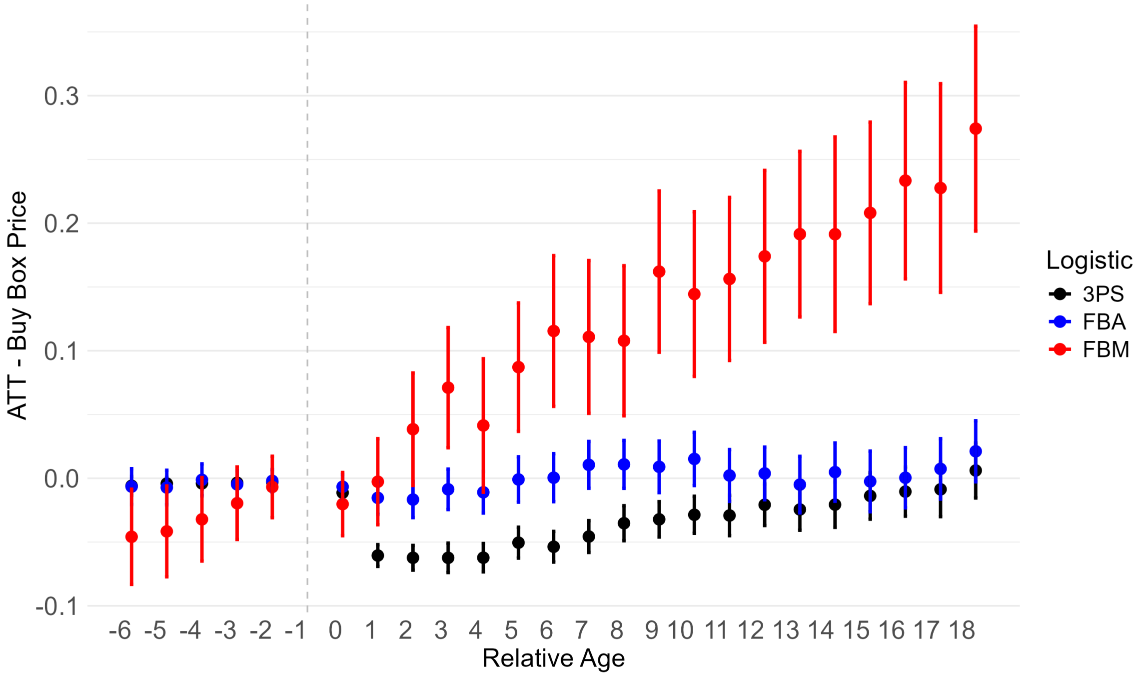

6.1 Event study plot of Amazon Entry on 3PS buy box prices

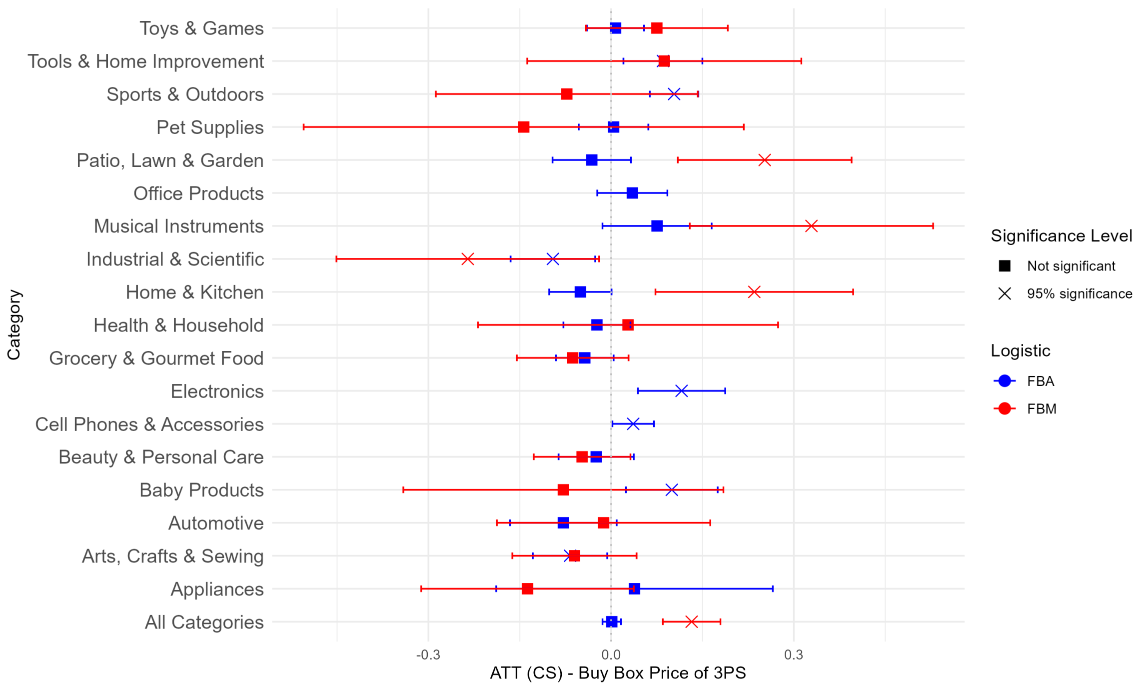

6.2 Static ATT by category

Notes: Panel 6.1 shows the event study plot of equation 1 estimated on buy box prices of third-party merchants using Callaway and Sant'Anna (2021) for all 3PS, and by logistic sector, separately. 6.2 depicts the static ATT, only of FBA and FBM sellers, aggregated as discussed in section 5 of the doubly robust Callaway and Sant'Anna (2021) estimator, separately for each root category. Standard errors are clustered at the product level.

6.3 Market Structure

Upon entry Amazon can also contribute to change the incentives that 3PS have to join a product market or to continue to be part of it. Acquiring a market share, in the form of time spent in the buy box, of about 20% on average, inevitably implies that other 3PS will lose share, but it says very little regarding who gives up those shares. In this section we evaluate how Amazon's entry contributes to changes in the composition of the market structure. We evaluate both descriptively and causally the evolution of market participants putting particular emphasis on the role that the Fulfilled by Amazon (FBA) program plays in granting third-party sellers competitive power. We do so once again through the lenses of the buy box, used as a proxy for market share. A characteristic that distinguish online marketplaces from brick and mortar stores, in fact, are the low cost of entry for all sellers wishing to expand horizontally to new products. While we observe 361253 unique seller-ASIN combinations, only 178478 (49.4%) have spent a non negative amount of time in the buy box, during the life-span of the products in our data. Only half of the sellers participating in the market as a whole, are capable of obtaining, at any point, positive market share.

Evaluating a product's market structure on the Amazon marketplace, therefore, requires to carefully distinguish who are the sellers actually at play. Many 3PS can participate to a market merely to promote their products on the platform, allowing them to have an online store while paying very little fixed costs if compared to what building an e-commerce website would entail. This is especially true for sellers not participating to the Fulfilled by Amazon program (FBA) and commonly referred to as Fulfilled by Merchant (FBM) sellers. Like their FBA peers, these merchants pay a fixed monthly fee of 39.99$ and an ad-valorem closing fee for each item sold, function of the transaction price and of the category the product belongs to, but they are not subject to FBA fees, rendering their presence on Amazon less costly, albeit arguably less effective. In panel 7, figure 1 depicts the number of offers detected for the first time by Keepa in treated products by relative treatment time. FBM offers, in red, are roughly stable through time before and after Amazon Entry; FBM entries appear unaffected by Amazon's presence as a retailer. Less FBA offers, instead, are first detected as relative treatment time passes; they bunch around Amazon entry, suggesting positive expected profits due to an increase in demand, but decrease as the market gets saturated. Figure 7.2, instead, depicts the buy box share of 3PS in treated markets, by relative treatment time. Overall, as previously discussed, 3PS concede shares to Amazon, but FBM sellers do so disproportionately more. In white, within each bar, we report the within-period relative share for FBA and FBM sellers; the latter lose almost 10% of buy box share from in the four months between reference period, -1, to the third full month in which Amazon has joined the market. The constant inflow of FBM offers, despite the low share dedicated to FBM offers underlies the importance of distinguishing among these two type of sellers and hint toward the interplay that the logistic branch of the Seattle giant has with the success of a third-party seller on the marketplace. We start, therefore, by the share of FBA and FBM sellers in the buy box.

Figure 8.1 depicts the event study plot of eq. 1 estimated on the buy box share normalized post-entry by the sum of buy box share held by third-party sellers only, excluding Amazon. Post-entry we are simply re-scaling the buy box share as it was entirely held by third-party sellers. The figure depicts describe the change in the composition of the buy box share due to Amazon entry. The effect on FBA share is absent;

no coefficient is significantly different from 0 throughout the eighteen months after entry, while the effect on FBM, while weak in terms of pre-trend, is still suggestive of a negative effect even in the composition of the buy box share that, after Amazon entry, remains in the hands of 3PS. While we recognize the presence of a negative differential pre-trend in the share held by FBM sellers in the treated and control group we notice that the magnitudes are large enough to assert that, recovering the parallel would not compromise the significance of the negative effect. While less confident on the proper magnitude, we are still confident in the compositional shift that Amazon causes in the buy box, which negatively impacts FBM sellers, possibly listing less convenient offers, while not affecting the share of FBA sellers.

Finally, figure 8.2 depicts descriptively, for treated products, the buy box share won by incumbents sellers and by recent entrants. The former are defined as sellers competing in a product market up to the month before the entry of Amazon as a FPS; the latter defined as entrants in the previous three months, e.g. months 7,8 and 9, for relative age 10. In both cases, unsurprisingly, the share decreases post-entry. Both type of sellers though show nuances. Incumbents, despite Amazon entry are still able to retain roughly 75% of the market after eighteen months. Entrants on the other hand are quickly able to accumulate, collectively, 10% of buy box share by the month of Amazon entry; post-entry they stop, on average to accumulate buy box share while losing roughly a third of it in the long-run.

Figure 7: Difference in the definition of market structure between all participants and buy box winner

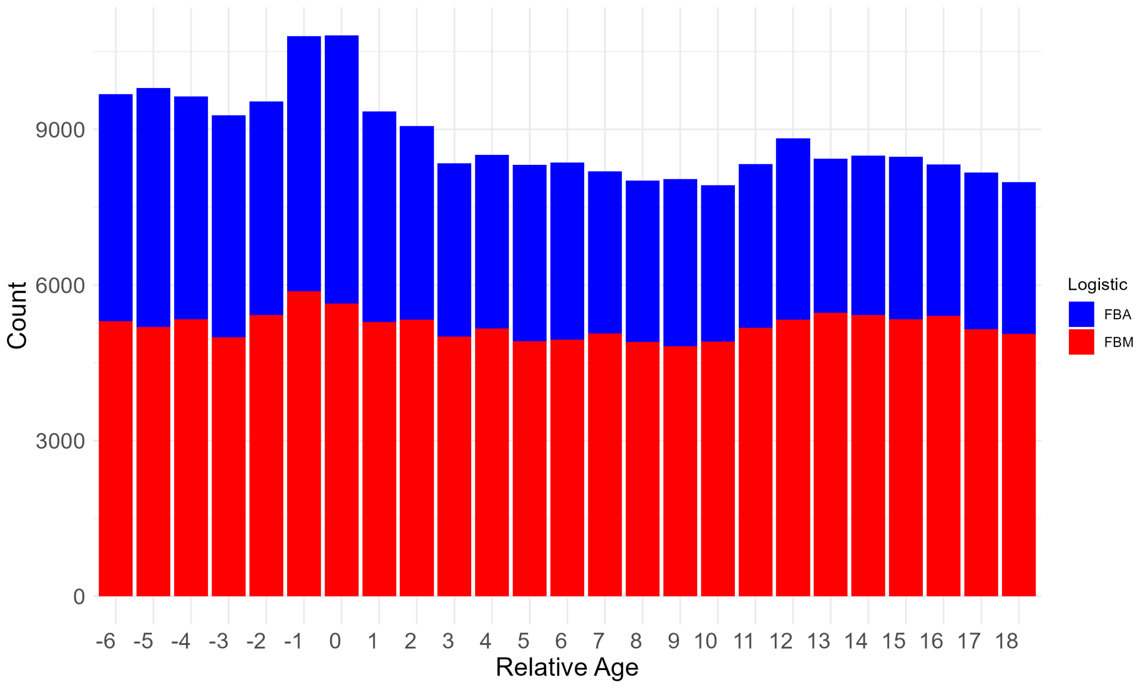

7.1 Number of FBA and FBM Offers at First Detection

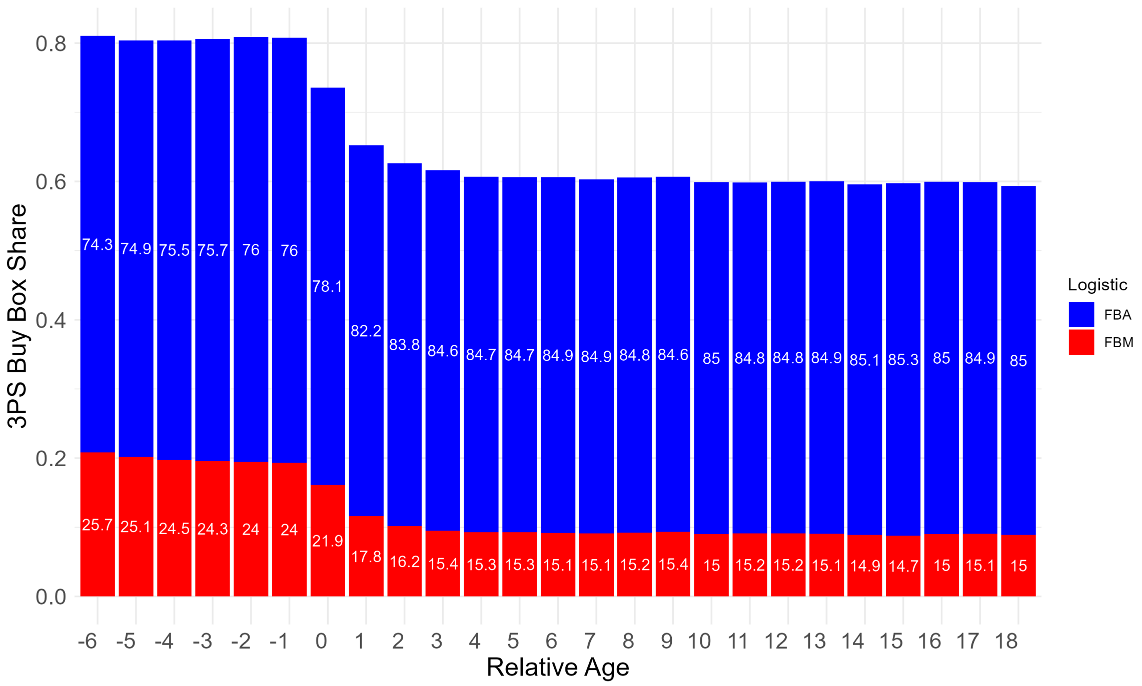

7.2 Buy box share by FBA and FBM

Notes: The panel describes the market structure according to the inflow of new-born offers and to changes in the composition of buy box winner. Figure 7.1 depicts the number of offers detected for the first time in treated products by relative treatment time, split by FBA and FBM. The number of new FBM offers is stable while the number of new FBA offers decreases, bunching around treatment time. Figure 7.2 depicts the buy box share that FBA and FBM third-party sellers win by relative treatment time. In each bar we display the within-period share, to highlight how, after treatment, FBM offers are relatively less featured in the buy box.

Figure 8: Composition of Buy box Share and Incumbents vs Entrant

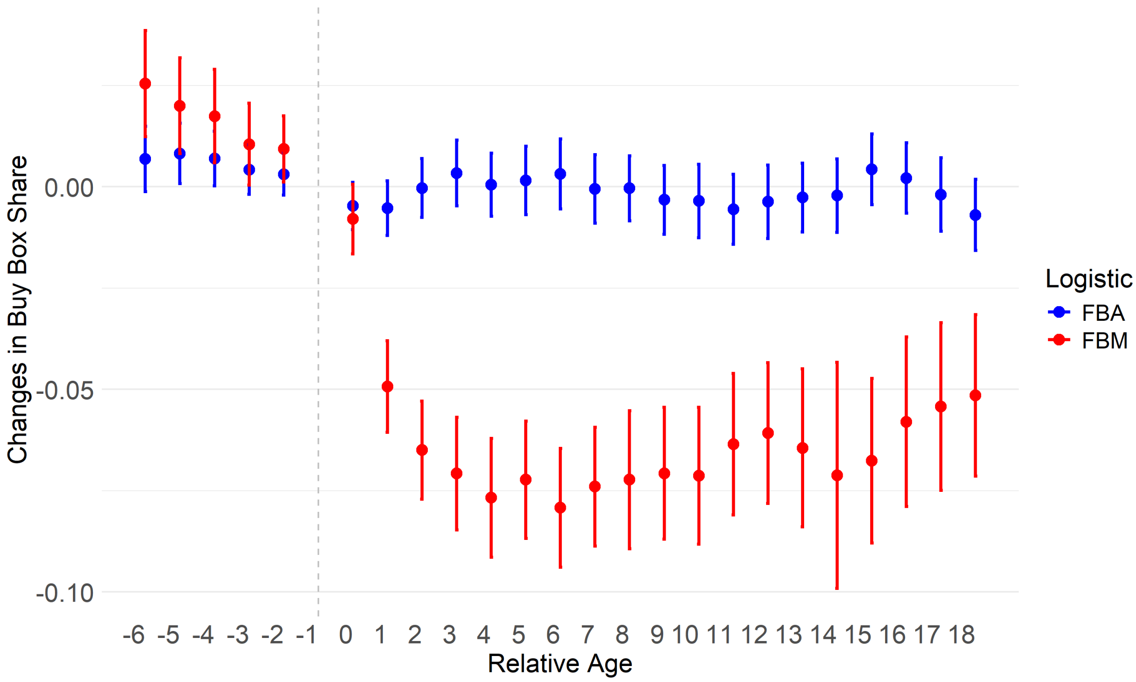

8.1 Event study plot of Amazon Entry on 3PS buy box share composition by logistic

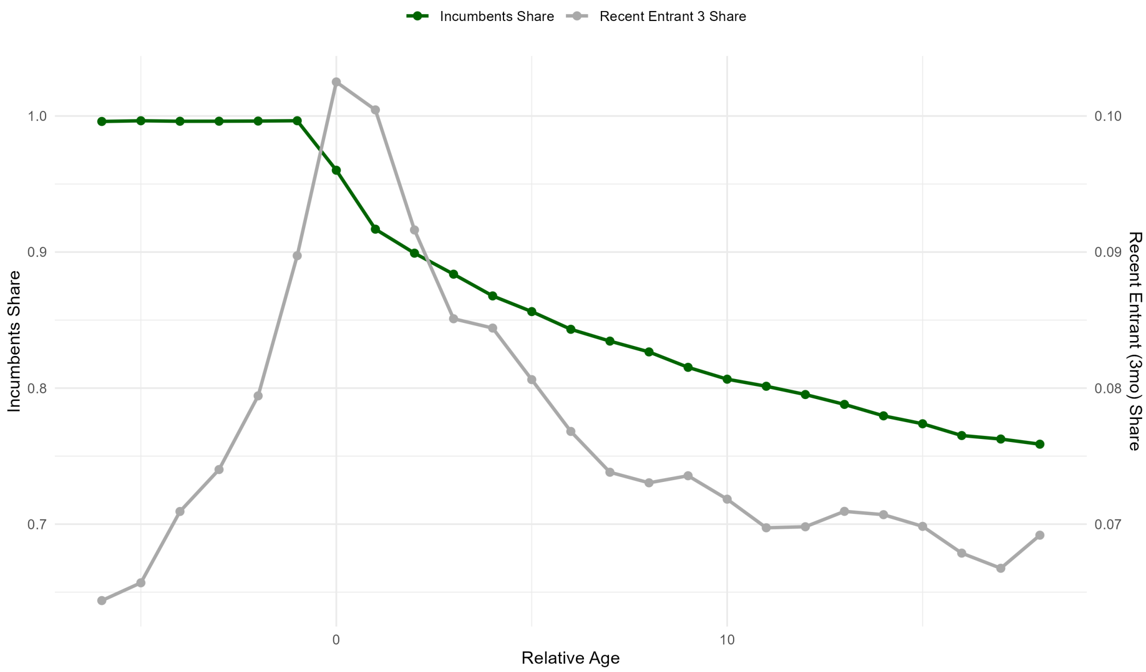

8.2 Incumbents and entrants buy box share in treated products

Notes: Figure 8.1 depicts the event study plot of the dynamic treatment effect of Amazon entry on the buy box share of 3PS normalized to 100 after entry, describing the change in the share of buy box gained or lost by sellers adhering to different logistic services. Figure 8.2 depicts, descriptively, the buy box share in treated products by incumbents sellers, sellers preceding amazon entry, and the buy box share of seller entered in the previous 3 months, for each relative age.

6.4 Revenue

Finally, we look at the effect of Amazon entry on revenues of 3PS. In the absence of transaction data we can only rely on an estimate[20]. In particular we interpret buy box prices as transaction prices and exploit the power-law relation between sales and sales rank multiplying prices with the inverse of the sales rank. We mediate the share of sales for each seller via buy box share, obtaining our proxy of revenue. Formally, we estimate equation 1 using as dependent variable $y_{i,a} = log(1 + RevenueProxy)$ where $RevenueProxy = (p^{3PS}_{i,a} \cdot CPI_t) \cdot bb^{3PS}_{i,a} \cdot log(1 + SalesRank_{i,a})^{-1}$. This approximation does not come without its limitations. In particular, relying on the buy box share to weight the fraction of sales that we attribute to 3PS, we are implicitly assuming a uniform distribution of purchases within the month. Customer arrivals are often non-uniform due to factors such as promotions, holidays, weekends, special sales events, and other demand shocks that cause spikes or dips in sales. Moreover, customers tend to purchase more during certain days of the week and during specific times of the month. Assuming uniform arrivals ignores these behavioral patterns. For all these reasons we consider these estimates tentative and leave refinements to the future. To stress the limitations just outlined we refer to our dependent variable as "Revenue Proxy".

Figure 9.1 depicts the dynamic effect of Amazon entry on 3PS proxy revenues. Upon entry Amazon causes a reduction of 11.2% that remains stable throughout the eighteen months that follow. Third-party sellers, therefore, give up one tenth of their revenues after treatment, with an effect that persists, on average to similar magnitudes in the long run; 3PS that have competed against amazon for eighteen months, see their revenues reduced, on average, by 13% with respect to pre-entry. While tentative, this result challenges previous estimate describing as neutral on third-party sellers revenue. When differentiating third-party sellers on the fulfillment service they rely on, we find the same stable and persistent dynamic pattern and similar magnitudes, with FBA sellers experiencing slightly less losses as captured by the static ATT of -8.7% against the one of FBM sellers of -11.3%.

The same static ATT are reported in figure 9.2 alongside those related to the eighteen category object of our study. We highlight how the hetereogeneity in the effect is evident both across categories and across fulfillment services. Amazon entry causes FBM sellers to lose more revenues than their FBA peers in nine categories out of eighteen, while the static effect is comparable in the remaining nine. The decrease observed in the aggregate appears to be largely driven by twelve categories, while for FBA sellers, in half of the categories the effect is not significantly different from zero.

[20]we are currently working on the possibility to recover sales estimates. Keepa contains low-frequency data on stock variation. Interpreting decreases at low distance in time, say 24 hours, we can obtain the sales of each product in the very short term and leverage the power-law that relates sales and sales rank. This method is similar to the one used in Gutierrez Gallardo (2021) and, as a robustness, in Lee and Musolff (2021). Modelling sales through this method would address limitations related to the assumption of uniform arrival of customers within the month.

Figure 9: Revenues for 3PS

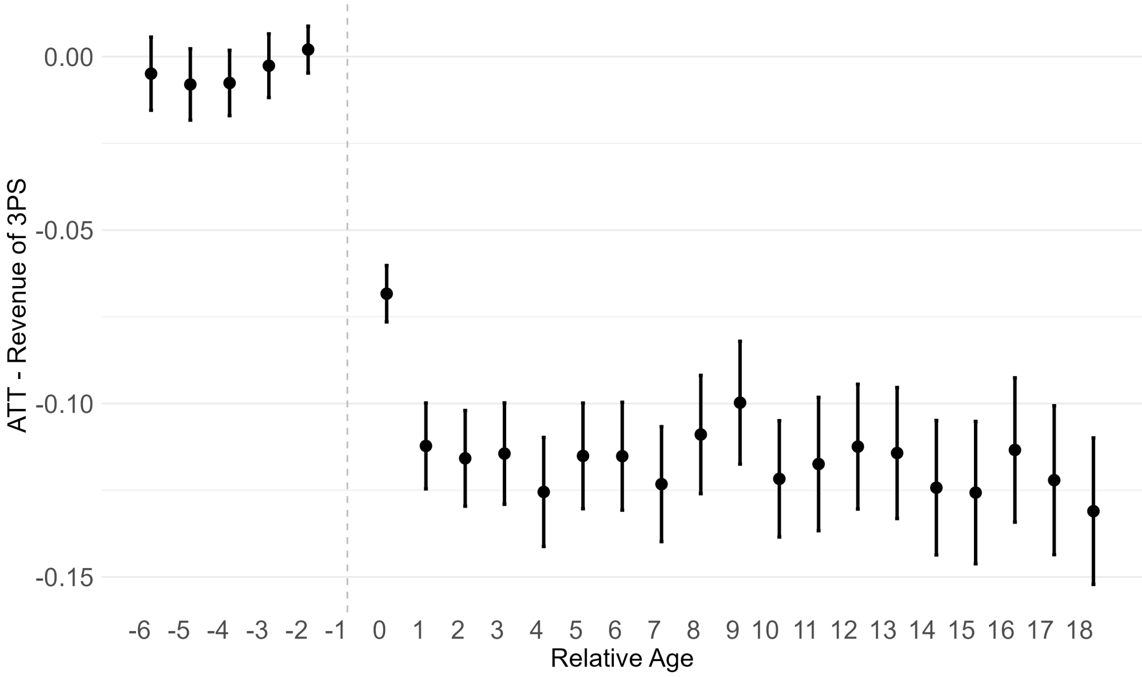

9.1 Revenue by 3PS

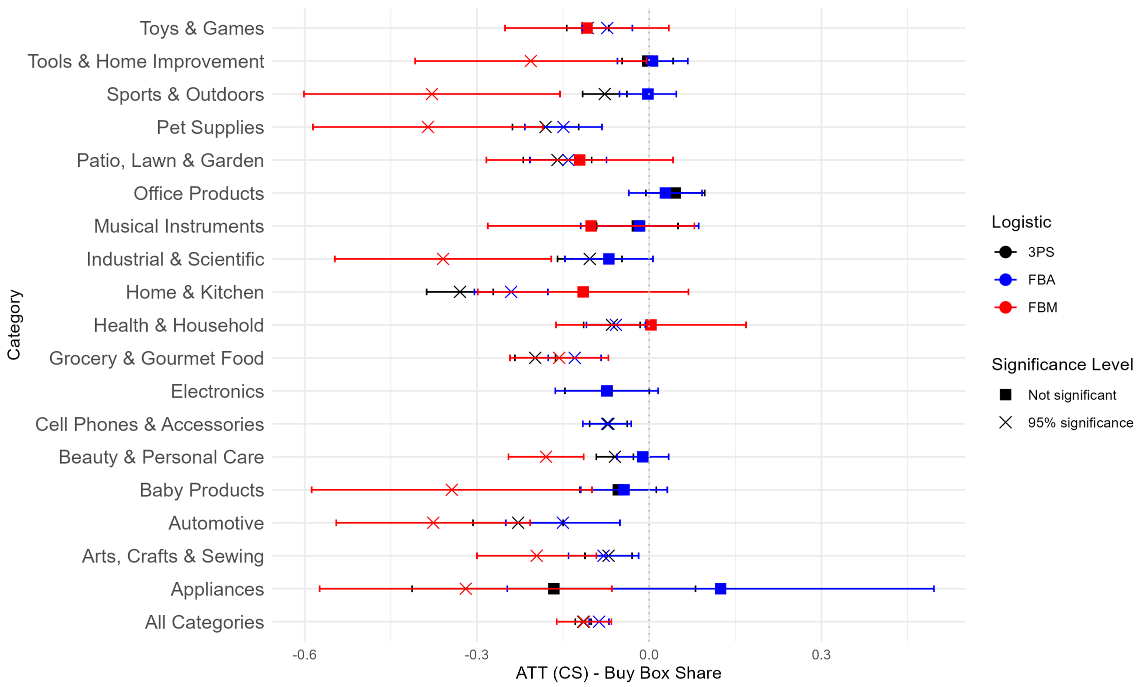

9.2 Revenue by FBA and FBM

Notes: Panel 9.1 shows the event study plot of equation 1 estimated on revenue proxy of third-party merchants using Callaway and Sant'Anna (2021). Figure 6.2 depicts the static ATT, on the same outcome, for 3PS, FBA and FBM sellers, aggregated as discussed in section 5 of the doubly robust Callaway and Sant'Anna (2021) estimator, separately for each root category. Standard errors are clustered at the product level.

7 Predicting Amazon Entry

What are the markets that Amazon choses to enter remains a crucial question to explain the interplay between its logistic arm and the platform. Amazon's quarterly results suggest that revenues that *"include commissions, related fulfillment and shipping fees, and other third-party seller services"* constitute about a quarter of yearly revenues.[21] Direct competition with their source of revenues reveals a paradox that challenge the rationality behind the platform's decision. In this section we build a model of entry that predicts Amazon's behavior. We do so leveraging machine-learning. We rely on a gradient boosting algorithm to understand the market characteristics associated with entry.

We choose a `XGBoost` classifier, a tree-based gradient boosting algorithm known for its efficiency and high predictive accuracy in classification tasks. The model's feature set includes a combination of product-level, seller-level, and competition-related variables. We favour this algorithm over an OLS regression for its ability to capture non-linearities and interactions that might be, and will be, important to understand our behavior of interest. While a common misconception is that these algorithms operate as a black-box, we employ SHAP, SHapley Additive exPlanations, to interpret the XGBoost model and identify the most important features driving predictions. SHAP provides a unified approach to explain the output of the machine learning model by calculating the contribution of each feature to the final prediction.

The data is divided into training and testing sets using an 80/20 split. Given the class imbalance inherent in the dataset, where non-entry cases outnumber entry cases, we employed a sampling strategy stratified by entry age to ensure proportional representation of each class across the training and test sets. For each entry age we "freeze time" collapsing market and seller characteristics to averages at the month before entry, what in the event studies above would be the reference period. For each entry age, we sample from the control group a random subset of market sharing the same age of the treated. For these markets we also collapse averages to the six months previous. This procedure allows us to avoid, data-leakage, a common term in data science, that simply prevents out-of-sample selection to leverage data that are the result of the target variable.

Hyperparameters are optimized using randomized search cross-validation. We sample hyperparameters from a defined distribution testing a total of 50 sets across three cross-validation folds, yielding a total of 150 model evaluations. The results displayed in figure 10.1 refer to the best model, the one with higher acuracy. The Receiver Operator Characteristic (ROC) curve represents the performance of the XGBoost classifier in predicting whether Amazon enters a market as a first-party seller. The curve plots the True Positive Rate (TPR) against the False Positive Rate (FPR) across different threshold settings of the model. The orange line, contrasted against a random classifier as represented by the diagonal line shows the ability of the algorithm to discriminate true positive correctly. The area under the curve (AUC) is 0.93, indicating good model performance; a perfect classifier would have an AUC of 1.0, while a random classifier would have an AUC of 0.5.

[21] https://www.marketplacepulse.com/stats/amazon-third-party-seller-services-sales, last accessed September 10, 2024

Figure 10: Predictive model of Amazon Entry

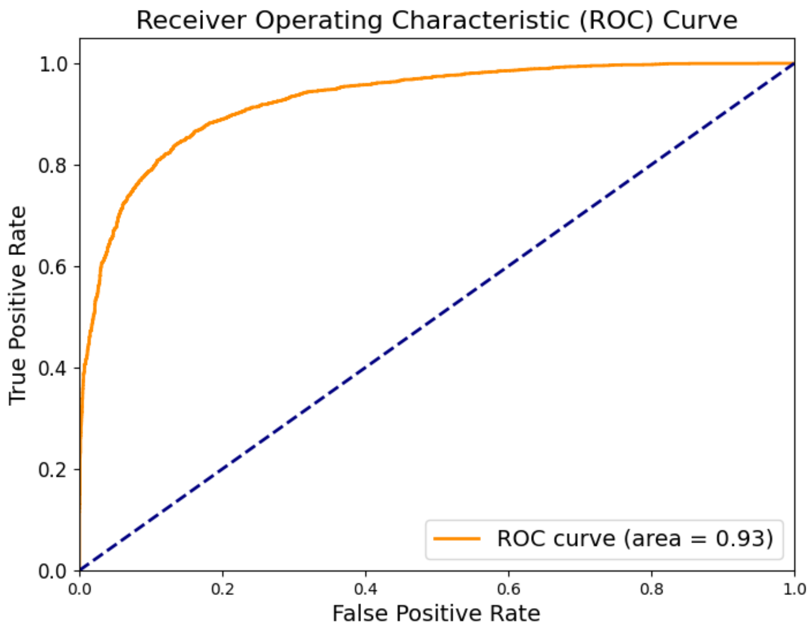

10.1 ROC Curve of the Xgboost model

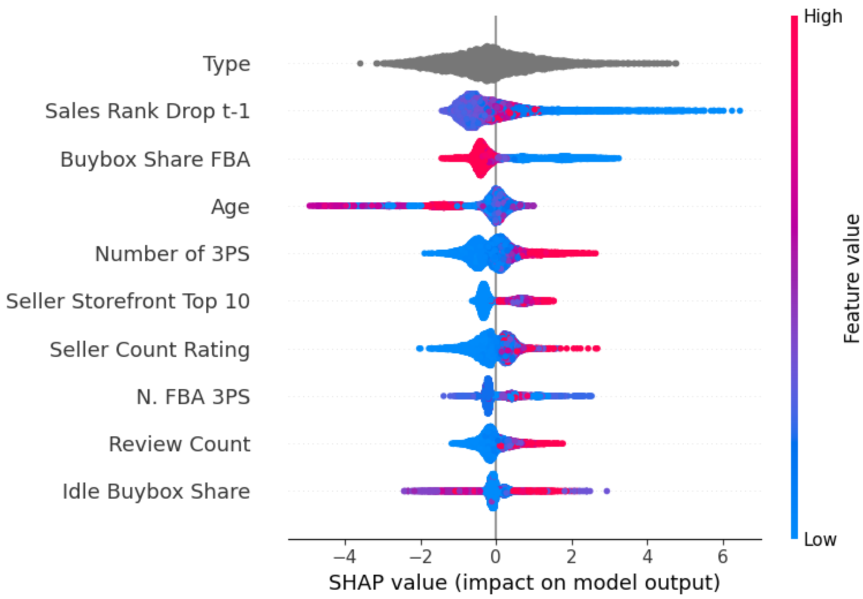

10.2 Beeswarm plot of the top 10 features

Notes: Panel 10.1 Depicts the ROC curve of the best xgboost model resulting from a 3-fold Cross-Validation. 4.2 Depicts a beeswarm plot of the top 10 most relevant predictors of Amazon Entry.

We focus out attention of the figure 10.2 where a beeswarm plot depicts the top ten covariates that contribute the most to discriminate true Amazon entry and enhance the accuracy of the model.[22] In the plot, features are ordered by their overall contribution

[22]In this section we sometimes refer to covariates as features

to the prediction, where the Type of a product shows the highest impact. ”Type” captures categorical distinctions among the products and allow us to account for demand patterns that naturally influence the decision to either enter or open a new product. The broad distribution of SHAP values for this feature suggests that its effect varies widely across observations, represented by each dot in the plot, impacting both the likelihood of entry (positive SHAP values) and non-entry (negative SHAP values).

Importantly two observables stand out: *Sales Rank Drop t-1* and *Buybox Share FBA*. The former emerges as a critical feature with a SHAP value displaying a wide spread, indicating that fluctuations in the sales rank one period prior strongly affect the model's output. The x-axis, depicting SHAP values, orders observations according to their contribution to positive classification (Amazon entry) or to negative classification. The color gradient of each dot, reveals how the magnitude of the feature value affects the prediction. In this case, low values of sales rank drop (in blue), that represent a fast increase in sales, consistently push the model towards predicting Amazon's entry. Importantly, low values of sales rank drop are spread throughout the positive values of SHAP, suggesting that Amazon enters more in products with fast-increasing sales.

The buy box share won by FBA sellers before entry is also a relevant feature. It shows high SHAP values associated with lower buybox shares. This feature directly correlates to Amazon's logistics services. Low share of buy box wins for FBA sellers push the prediction towards a positive classification, that is, Amazon entering the market, whereas high share of buy box won by FBA sellers are associated with a low probability of Amazon joining. In particular, we contrast the buy box share held by FBA sellers before entry with the idle buy box share, the last entry in the beeswarm plot of figure 10.2, that contains the share of time, in a month, in which the buy box was absent from the market, signifying either an out-of-stock product or a lack of sufficntly high-quality offers. Figure **??** depicts an interaction plot: markets are scattered throughout the plane ranked on the x-axis by the buy box share of FBA sellers and on the y-axis by their SHAP value, their contribution to a positive classification. The color gradient of each dot represents, for each observation/market the value of its idle buy box share. The plot suggests how Amazon tends to enter more either in products where the revenue stream is low, either because the market share of FBA sellers is low, or because their quality is not high enough. In both cases the platform intervenes by entering a market as first-party seller. It appears to do so also in markets in which the idle buy box share is low, but the presence of FBA sellers in the buy box is scant.

Figure 11: Interaction of Buy box Share of FBA sellers with Idle Buy box share

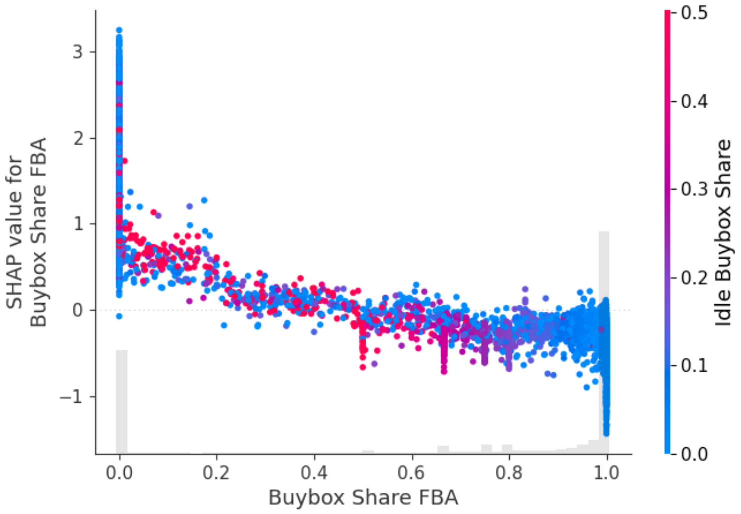

Notes: THe figure depicts a scatter plot of the SHAP values for the buy box share of FBA sellers and the Idle Buy box share.

Amazon's entry therefore, depicts a mixture of objectives: on the one hand it increases the quality and its revenue stream sourced by FBA fees, making sure that high-quality offers are available across all-markets; but on the other hand is targeting markets in which the revenue stream is low, not because of poor quality but because of a low presence of FBA sellers.

8 Conclusions

Our analysis reveals that Amazon's entry as a first-party seller exerts a unique and sustained competitive pressure on third-party sellers, significantly reshaping market dynamics. First, we show that Amazon's entry leads to a reduction in the buy box share of 3PS by an average of 20%, with some markets experiencing a more profound impact. Interestingly, while large third-party sellers can temporarily capture market share, they are unable to sustain it at the same level as Amazon over the long term. Second, our results indicate that Amazon's price competition is most pronounced in the short run, where buy box prices decrease by an average of 5% following Amazon's entry. However, these price reductions are concentrated in certain categories, with no significant effects in others, suggesting that Amazon's competitive strategy is not uniform across markets.

We document that Amazon's logistics service, plays a crucial role in shaping competition. Although FBA sellers continue to lose buy box share post-entry, their prices are unaffected by Amazon's entry, suggesting a differential impact on sellers depending on their fulfillment method. The extent to which FBA sellers are price-constrained is a crucial point of further discussion. To what extent Amazon benefits consumer by lowering prices should be weighted against the margins enjoyed by the Seattle giant through a quantitative dominance of the presence of its logistic arm. While certainly providing better quality, the exclusionary access of FBA sellers to Prime customers might hinder the possibility of third-party logistic services to develop a competitive landscape. driving down costs and hence, prices. to what extent this would benefit consumers more than Amazon price-competition is a crucial point that we leave to future research.

Overall, our findings contribute to the understanding of Amazon's dual role on its marketplace, with implications for competition policy. While Amazon's entry can benefit consumers through lower prices, the platform's self-preferencing and market power may raise concerns about long-term market competition and third-party seller viability. Future research could explore the role of Amazon's logistics services in cross-subsidizing its retail operations and the broader implications of such cross-subsidization for competition and consumer welfare.

References

Bergé, L. (2018). Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm. *CREA Discussion Papers*, (13).

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.

Busso, M. and Galiani, S. (2019). The causal effect of competition on prices and quality: Evidence from a field experiment. *American Economic Journal: Applied Economics*, 11(1):33–56.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230. Themed Issue: Treatment Effect 1.

Chen, L., Mislove, A., and Wilson, C. (2016). An empirical analysis of algorithmic pricing on amazon marketplace. In *Proceedings of the 25th international conference on World Wide Web*, pages 1339–1349.

Crawford, G. S., Courthoud, M., Seibel, R., and Zuzek, S. (2022). Amazon entry on amazon marketplace. *CEPR Discussion Papers*, (17531).

Farronato, C., Fradkin, A., and MacKay, A. (2023). Self-preferencing at amazon: evidence from search rankings. Technical report, National Bureau of Economic Research.

Fischer, K., Martin, S., and Schmidt-Dengler, P. (2023). The heterogeneous effects of entry on prices.

Gutierrez Gallardo, G. (2021). The welfare consequences of regulating amazon. *Available at SSRN 3965566*.

Lee, K. H. and Musolff, L. (2021). Entry into two-sided markets shaped by platform-guided search. *Job Market Paper, Princeton University*.

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

Roth, J., Sant'Anna, P. H., Bilinski, A., and Poe, J. (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199.

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

A Appendix A: Robustness

A.1 Main specifications with different methods

The literature on difference-in-differences is quickly expanding and in an attempt to provide results using up-to-date methods we have tested three main estimators: a standard Two-Way Fixed-Effect estimator, the Interactive Weighted estimator proposed by Sun and Abraham (2021) and the Doubly-Robust estimator of Callaway and Sant'Anna (2021). In this appendix we report the results of all three methods on the three specifications that motivate and kick-start our in-depth analysis in the main text: buy box share, buy box prices and entrant prices.

For each outcome, we report a standard dynamic Two-Way Fixed-Effect (TWFE, henceforth) estimation of equation 1 used in the main text:

$$y_{i,a} = \alpha_0 + \alpha_i + \alpha_a + \alpha_c + \alpha_t + \sum_{k=-K}^{K} \delta_k \cdot D_{i,a+k} + X_{i,a}\gamma + u_{i,a} \tag{2}$$

that omits the interaction of cohort-dummies with relative age, remaining otherwise unchanged. We also report the estimates of our main specification using the method proposed by Sun and Abraham (2021) (SA, henceforth).

All the main results are presented as event-study plots. The content of the plots varies depending on the estimation method used. When the TWFE method is employed, the coefficients reported correspond to the δ_k terms in equation 2, with error bands representing the 95% confidence intervals for each period-specific estimate, testing whether the coefficients are statistically different from zero. All estimates are carried via Bergé (2018) `fixest` package in R. When Sun and Abraham (2021) is used, instead, the coefficients shown in the plot are an aggregate of all the $\delta_{e,l}$ in eq. 1 that refer to the same relative treatment age l, weighted by the sample shares of each cohort e in the same relative age l. This aggregation form their *interaction weighted* (IW) estimator; we refer the reader to equation (28) in section 4.1 of Sun and Abraham (2021). IW estimators are also accompanied by period-wise error bands. All estimates are carried via Bergé (2018) `fixest` package in R. When the Callaway and Sant'Anna (2021) is the method of estimation, the plot also shows and aggregation that averages each group-time specific ATT into a relative-age-specific treatment effects. Their interpretation equals the one ot the other two methods, expressing the average effect on the target variable of Amazon's entry k periods after. We refer the reader to equation (3.4) in section 3.1 of Callaway and Sant'Anna (2021). All estimates are carried via the `did` package in R.

Figure 12: Impact of Amazon Entry on Buy Box Share of Third-Party Sellers by method of estimation

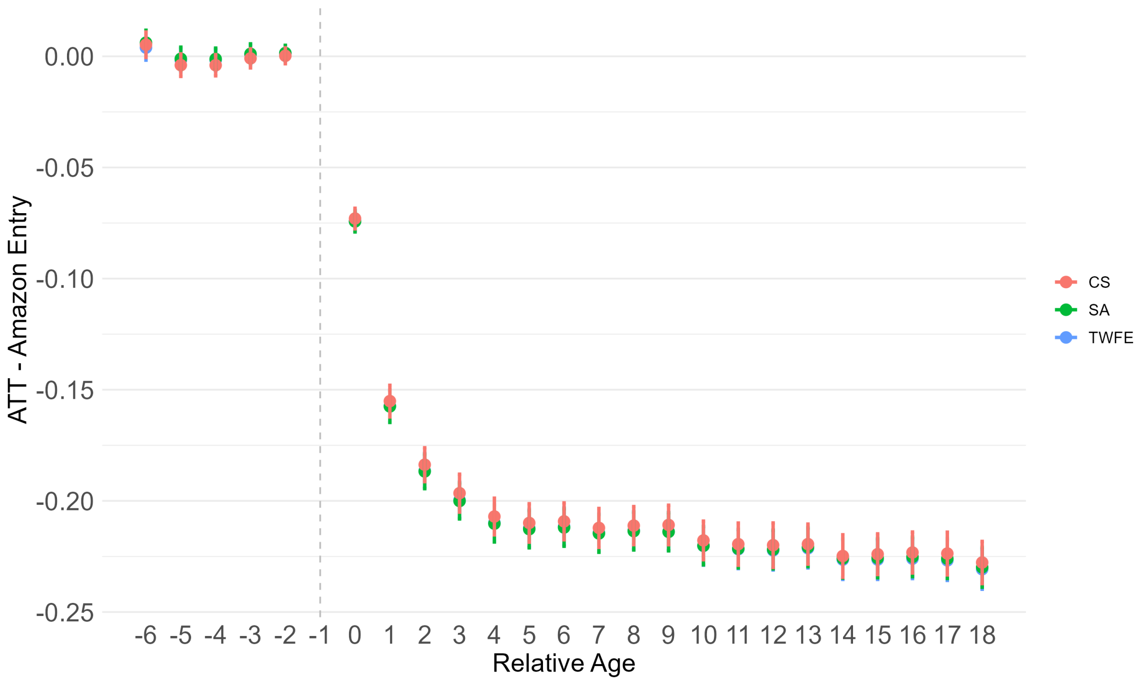

Notes: This figure presents the dynamic effects of entry by Amazon estimated using TWFE, Sun and Abraham (2021) and Callaway and Sant'Anna (2021). The latter is reported in 2 in the main text. In all methods standard errors are clustered at the product level.

Figure 13: Impact of Amazon and Large Sellers' Entry on Buy Box Price of Third-Party Sellers

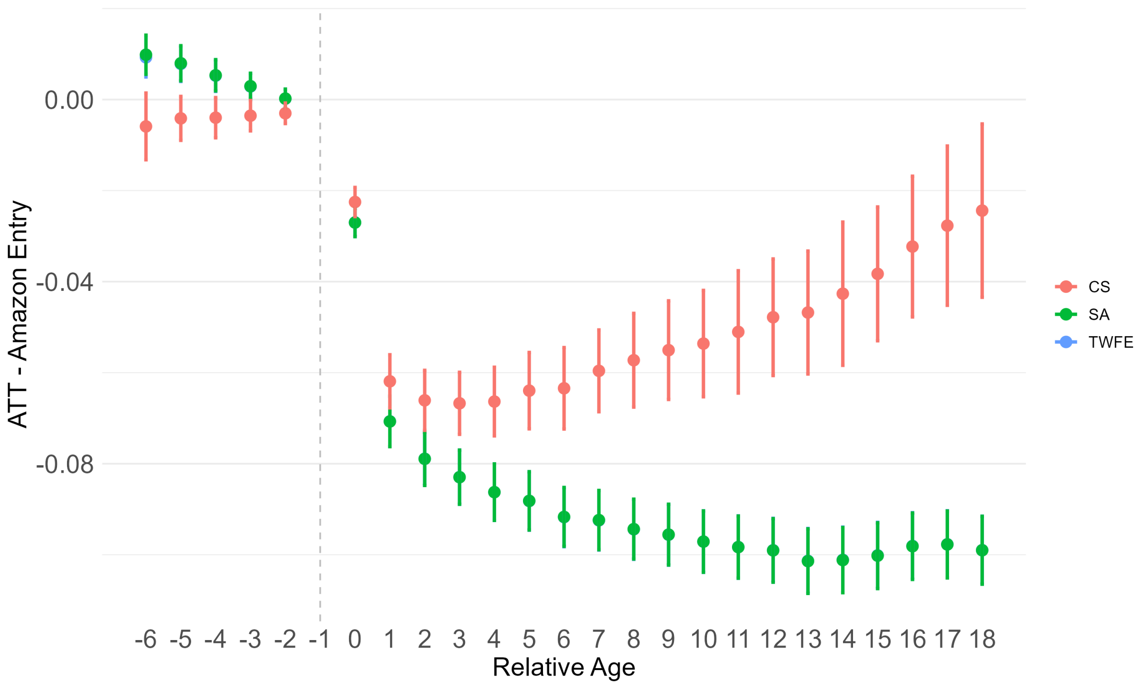

Notes: This figure presents the dynamic effects of entry by Amazon estimated using TWFE, Sun and Abraham (2021) and Callaway and Sant'Anna (2021). The latter is reported in 4 in the main text. In all methods standard errors are clustered at the product level.

Figure 14: Amazon Entry compared to buy box prices by method of estimation

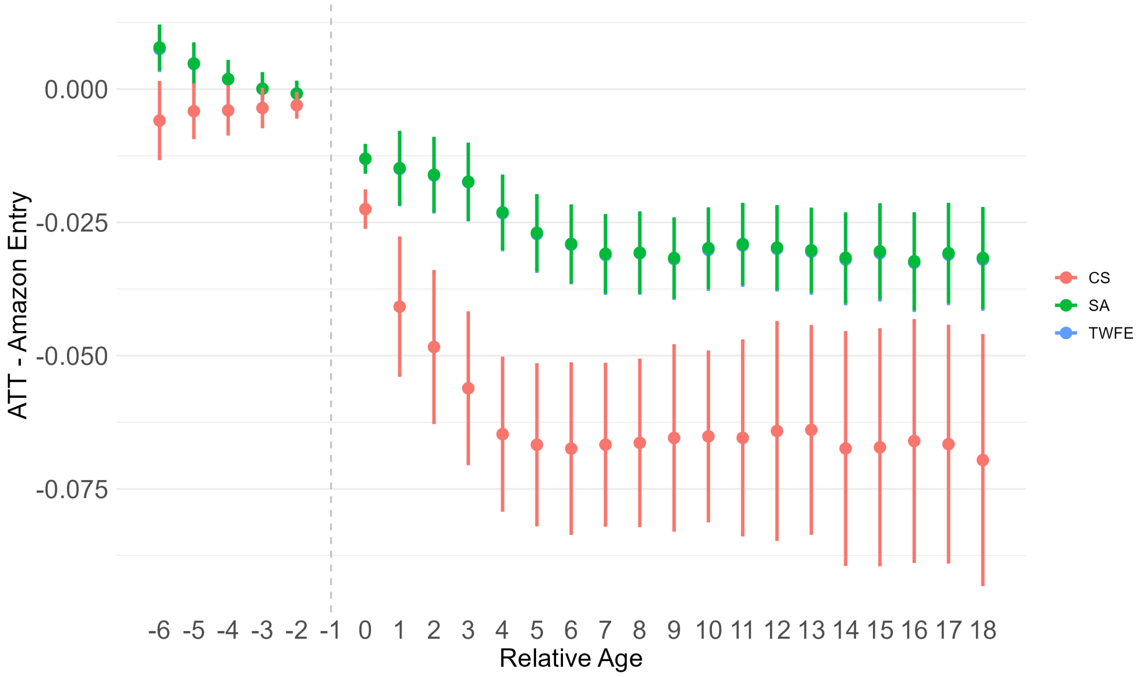

Notes: This figure presents the dynamic effects of entry by Amazon estimated using TWFE, Sun and Abraham (2021) and Callaway and Sant'Anna (2021). The latter is reported in 5 in the main text. In all methods standard errors are clustered at the product level.

A.2 Comparison of sample 10% control group and full control group

Table B.1: Event Study Descriptives - Control

	Control - 10%			Control - Full		
	count	mean	std	count	mean	std
N. of Products	37409	nan	nan	378845	nan	nan
Age	1014171	27.53	21.65	10318296	29.08	25.13
BB Price	947408	49.39	240.65	9639183	50.32	261.76
Min Price	922110	46.13	224.45	9388748	47.09	272.48
Share of FBA	1024672	0.86	0.32	10428315	0.86	0.32
Count of offers	821574	1.20	1.40	8380231	1.21	1.54
Share of 3PS Entry	1024672	0.01	0.06	10428315	0.01	0.06
Share of 3PS Exit	1024672	0.01	0.07	10428315	0.01	0.07
Sales Rank	993401	46015.60	187728.13	10113349	46367.14	198381.03
Sales Rank Drop 1	935248	1.14	113.43	9524478	1.37	435.79
Sales Rank Drop 3	866725	3.31	140.37	8830703	4.71	1456.98
Rating	937085	44.48	3.80	9535204	44.46	3.74
Seller N. Ratings	1005905	6143.20	12373.97	10245369	6100.34	12348.41
Seller Storefront	721146	477.10	3167.00	7353832	481.16	2994.23
Entry Age	0	nan	nan	0	nan	nan
3PS BB Share	1022285	0.91	0.28	10406220	0.90	0.29
Price of Entrant	0	nan	nan	0	nan	nan
(Avg) 3PS Price	1008071	44.72	197.26	10252274	44.64	197.90

Bundling Services in Digital Markets - FTC vs Amazon Inc.

Vito Stefano Bramante*

September 16, 2024

Abstract

This paper examines the Federal Trade Commission's case against Amazon Inc. regarding alleged anti-competitive practices, particularly the tying of logistics services to marketplace access, which has been integral to Amazon's dominance in digital markets. Using data from Keepa.com, we investigate the impact of Amazon's Fulfilled by Amazon (FBA) program on third-party sellers operating within its marketplace. Our dataset includes detailed information on over 32,000 sellers and 300,000 products from January 2018 to December 2021. The results indicate that while switching to FBA increases buy box share marginally, it leads to a 5% decrease in prices and is associated to a reduction in seller revenues. These findings suggest that Amazon's logistics program is used by sellers to remain competitive, particularly in the face of rising competition from Amazon and other FBA sellers. However, the switch does not improve product ratings, challenging Amazon's claim that the program enhances service quality. This study provides empirical evidence on the economic impact of Amazon's logistics practices, contributing to the ongoing debate on competition in digital marketplaces.

JEL Codes:
Keywords:

*University of Bologna, vitostefano.bramant2@unibo.it

1 Introduction

In September 2023 the Federal Trade Commission of the United States, along with 17 U.S. states, sued Amazon for allegedly acquiring monopoly power through platform policies, deemed anti-competitive, enforced on third-party sellers operating within the boundaries of its own digital marketplace. The case builds around two pillars that, the FTC states, interplay in self-reinforcing attempt to maintain dominance over two complementary markets: the logistic market upstream, and the e-commerce market downstream. In the logistic market, Amazon is accused of tying enrollment to its logistic services with access to the platform's prime customer base.

In this work we explore the empirical side of the FTC claims with regards to Amazon's logistic services. We leverage data from Keepa.com, an external data provider that monitors the US Amazon marketplace and provides product-specific and seller-specific data. We recover seller-specific data on prices, buy box share and observable characteristics of 32794 sellers operating in roughly 300'000 products.

We find that Amazon's logistic arm has gained prominence during the COVID-19 pandemic. During this period the fulfilled by Amazon program, that allows third-party sellers to delegate to Amazon's fulfillment centers the shipment process of purchased products, has increased twice in size. The success of the program, though is to attribute to the extensive margin: new sellers favoured the FBA program during the pandemic. In particular the FBA program renders adhering sellers automatically eligible to be Prime sellers, accessing a customer base estimated to reach 165 millions users. Sellers already populating the platform, on the other hand, appear to selectively chose FBA only for some of their products. Until the second quarter of 2019, also merchants in charge of their shipments, FBM sellers, were eligible for Prime status. When in 2019, Amazon silently closed this possibility, effectively tying Prime membership to FBA, the number of new FBM Prime offers reduced sensibly, especially during the pandemic.

To establish the average value of an FBA offer, therefore, we resort to a large sample of 25 million observations of seller-product pairs. We focus our attention to sellers that have, at least in one case added to an already existing FBM offer, a FBA one, for the same product. Sellers, in fact, are allowed to list the same product under multiple conditions, including the shipment type. We find that switching to fba, on average, is associated to an increase in the buy box share that is of negligible size when accounting for market characteristics; on the other hand is associated with a decrease in price of about 5% even when accounting for both market and seller characteristics. This result suggests that the FBA program might be selectively used by 3PS FBM sellers to shield against competition which might, on average, come either from Amazon or from other FBA offers, requiring FBM sellers to compete on prices anyway. To reinforce this finding, we find that switching is associated with a reduction in revenues, estimated by a rough approximation using the inverse of sales rank. It remains an open question to what extent sellers predominantly relying on the their own logistic services use the FBA program to explore new niches or to defend their threatened market share. We establish also that, counter to Amazon's claim, switching to FBA offers is not associated with an increase in rating for switching sellers.

We deem this a worthy exercise in at least two regards. Firstly, the FTC case is build around a heavy use of whistle-blowing. Evidence on the empirical and quantitative content of the FTC case against Amazon Inc. are absent as of today. Second, empirical evidence on the functioning of the largest marketplace in the world, in particular in relationship to its US market, are still in its infancy. We believe our rich dataset can provide stylized facts that can contribute to an informed discussion both in the economic circles and for competition authorities.

2 Literature

A key component of Amazon's business model is its vertical integration of e-commerce and logistics services a topic explored extensively in the economic literature (Hart and Tirole, 1990; Lafontaine and Slade, 2007). In the case of Amazon, the bundling of logistics services like FBA with access to its Prime membership may increase efficiency while foreclosing independent logistics providers or third-party sellers who cannot afford to participate in FBA. Moreover by tying access to Prime customers to its own logistics services, Amazon effectively raises the costs for sellers who want to remain competitive but do not wish to use FBA (Salop and Scheffman, 1983). This strategic behavior can reduce competition, both in logistics and in the marketplace, as sellers may find themselves locked into Amazon's ecosystem without alternative options.

Moreover, antitrust regulation dedicated to digital platforms is a growing area of focus in both academic and policy circles. Recent works, such as Khan (2017) and the on Digital Platforms (2019), argue that traditional antitrust frameworks may be inadequate for dealing with the challenges posed by digital monopolies like Amazon. Amazon's bundling of FBA with Prime membership eligibility brings these concerns into life; as digital platforms continue to consolidate market power, it become relevant whether their practices grant them market power by relaxing previous market friction or by introducing new practices that reduce consumer choice or stifle competition. The legal and economic literature on platform monopolies and bundling practices (Carlton and Waldman, 2002; Nalebuff, 2004) suggests that such strategies, when employed by dominant firms, can potentially violate antitrust principles. Our research contributes to this ongoing debate by providing empirical evidence of how Amazon's logistics practices affect competition in its marketplace.

3 Fulfillment Channels and the Seller Fulfilled Program

The Fulfilled by Amazon (FBA) program allows third-party sellers to delegate to the platform the entire shipment process. It requires the seller to ship the stock to an Amazon warehouse indicated by seller central, the dashboard dedicated to each seller, and from the arrival to the fulfillment centers the seller has no more duties. Once a product is sold, Amazon handles the entire fulfillment process, including picking, packing, shipping, customer service, and returns. This program provides sellers with the advantage of Amazon's extensive logistics network and gives their products automatic eligibility for Prime, as evidenced by a Prime badge beside the offer. FBA sellers pay several types of fees, depending on the services Amazon provides and the characteristics of the products they sell. One of the primary costs is the fulfillment fee charged per unit and determined by the size and weight of the item. For smaller, standard-size items, these fees typically range between three to five dollars per unit, while for larger or oversize items, the fees can be significantly higher.[1]

The "fulfilled by merchant" expression, instead, indicates that the seller will handle the shipment process independently, without relying on Amazon, but for mediating the transaction. These sellers do not share their revenues with the platform with the exception of the referral fees, common to all sellers, and computed as a percentage of the product's transaction price. The percentage varies by category but generally falls between 6% and 15%.

The Seller Fulfilled Prime (SFP) program, on the other hand, allows sellers to fulfill their own orders while still offering Prime shipping directly from their own warehouses or third-party logistics providers. To be eligible, a seller has to undergo a trial period and has to adhere to stringent delivery requirements such as, for instance shipping orders within a two-day window and maintaining high customer service standards, as monitored by Amazon metrics. In contrast to FBA, SFP sellers maintain more control over their inventory and fulfillment process but still benefit from offering Prime to their customers. The Prime program guarantees customers free delivery and shipments within 24 or 48 hours. It is particularly demanding for seller handling shipment on their own. That's why a market for independent logistic services that could guarantees Prime standards to be met, was developing along side the Prime program. These independent seller, or third-party logistic (3PL) were, of course, in direct competition with Amazon's own logistic arm.

Starting from the second quarter of 2019 all FBM third-party sellers submitting a request to become SFP, would receive the following message *"Seller Fulfilled Prime is not accepting new registrations at this time. Click the Join the Waitlist button if you would like to be notified when enrollment reopens."*.[2]. Amazon, according to the FTC case, suspended enrollment to the program, letting the waitlist grow. Figure 1 plots the number of FBM offers with a Prime badge, hence listed by SFP members at the time of first detection, that is, when first scraped by our data provider, a close proxy to entry. New FBM Prime offers, in our data, experience an inflow of at least 40'000 per month until the end of the second quarter of 2020, with spikes, at the beginning of 2019 reaching 100'000 new entries per month. The inflow of FBM Prime offers therefore, decreases abruptly at least in the bestselling products in the middle of the COVID-19 lockdown.

[1] In addition to fulfillment fees, FBA sellers are also responsible for storage fees, charged based on the amount of space dedicated to the stock, measured in cubic feet and the period of the year, with holiday season requiring higher rates.

[2] https://sellercentral.amazon.com/seller-forums/discussions/t/2767d78a4b59b781ee7519cc79000370, last accessed September 10, 2024

Figure 1: Number of FBM Prime offers by year-month

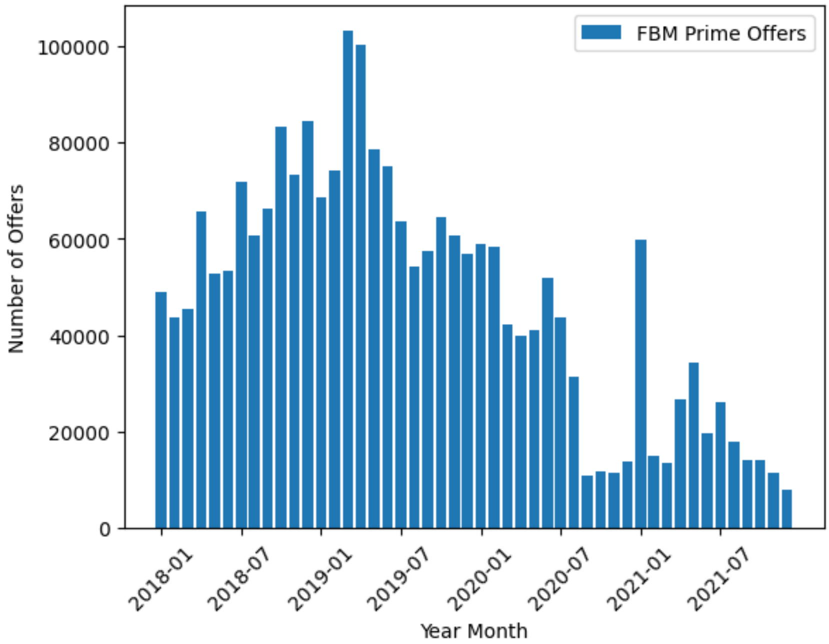

Notes: The figure plots the number of Prime offers fulfilled by merchants independently, or by their own chosen independent logistic service.

As for the stock of Prime offers, figure 2 reports the number of products with a prime badge, by month, split by logistic service of the listed offers; it remains constant for FBM Prime, suggesting that the growth is caused by the same limited number of SFP sellers expanding horizontally, on the intensive margin. On the other hand, Prime offers fulfilled by Amazon increase the rate of growth sensibly, especially during the COVID-19 pandemic, doubling in size in 2020 and 2021. The latter suggests an expansion on the extensive margin, new Prime sellers part of the FBA program.

Figure 2: Number of FBM Prime offers by year-month

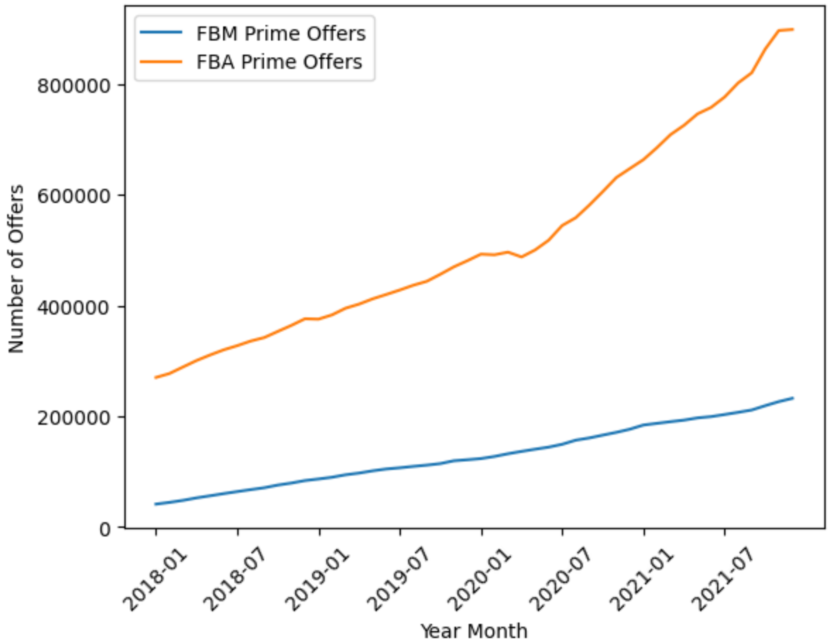

Notes: The figure plots the number of Prime offers fulfilled by merchants independently, or by their own chosen independent logistic service.

That the market is experiencing an expansion on the extensive margin is confirmed by a simple count of the number of sellers in our data over the period of study. Figure 3 plots the number of sellers present in out data by months, depicting an abrupt positive change that corresponds to the beginning of the COVID-19 lockdown, potentially driven by brick-and-mortar stores moving their operations online. In the same plot the left y-axis shows the increasingly prominent role of the FBA program over the years, by depicting the average share of products fulfilled by Amazon vis-à-vis those fulfilled by merchants. At the beginning of our period of study the sellers would, on average, fulfill 75% of their portfolio on their own, while delegating 35% of it to Amazon. At the end of 2021 the same figures change to 62% and 45%, respectively.[3]

[3] The two figures do not sum to one because a seller can list the same product both via FBM and via FBM.

Figure 3: Number of Sellers and Average Share of FBA vs FBM offers by Seller

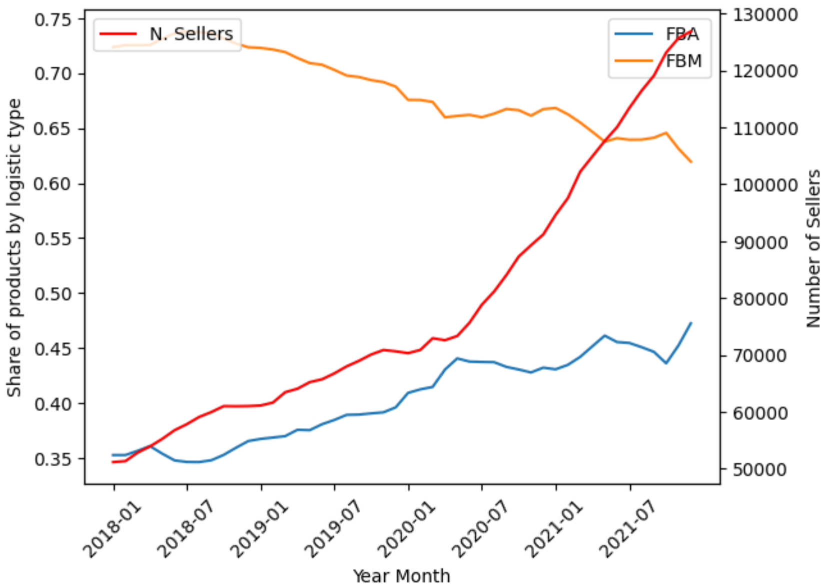

Notes: The figure plots the average share of offers listed as FBA and as FBM by seller on the left y-axis. It also plots the number of sellers on the right y-axis. Both are plotted by year-month.

To investigate the role that the FBA program plays for the individual sellers we focus our attention of those merchants that, at any point during the timespan of investigation, add to a product they already list and fulfilled independently a FBA offer, for the same product. For simplicity we refer to these sellers as *switchers*. Is it sufficient one switch to enter the dataset we use. In particular we denote a product a switched if the date in which the offer has appeared for the first time is greater or equal to the date of the fbm offer. These sellers and the product they list among the top 5000 bestselling products, are the subject of analysis in the remainder of this work. We describe their characteristics and the data we use for our event study in the next section.

4 Data

We leverage a monthly panel of 32794 switchers spanning period January 2018 - December 2021, containing seller-product pairs divided in 81776 products in which the seller added a FBA offer to an already present FBM one and 209924 in which the seller has not opted to do so. We restrict our attention to the 6 months preceding the switch and to the 12 months following it. We do not bin, we drop all the periods that precedes and follow. We observe, in each month the buy box share of the seller in the product market, the price listed and several observable characteristics related to the market and to the seller. For the market we observe the number of offers listed, the number of sellers, the sales rank of the product, the rating in a scale of 1 to 5, the number of product-specific ratings received by customers and whether Amazon competes in the product as a first-party seller. For sellers we observe the rating specific to the seller in a scale of 1 to 10 and the number of ratings received to its shop. We also observe and use as control, the level-2 category of the product, that is, the category it is classified under, within each root category, representing the highest level of products classification on the platform.

For the purpose of evaluating the effect of switching from FBM to FBA on the quality of the service we collapse the panel at the seller level and take the share of FBA offers over the total number of offer listed by a seller as our main variable of interest, while averaging their individual observable characteristics and the observable characteristics of the market it operates in, within our sample. We cannot fully observe all the products listed by each seller. The dataset in fact, is sampled over 18 categories, selecting for each of them all products that have been ranked above 5000 at any point in time, for the period 2021-2023.

Table 4 describes the dataset. We notice that all observable characteristics are largely comparable. Notably, all our dependent variables present a t-statistics above any thresholds of statistical significance on both sides of the distribution. We notice, though, that some differences, while not statistically relevant, can be appreciated. In particular the buy box share is higher for those cases in which the seller adds a FBA offer, occupying it 56% of the time against 37%. Moreover, pick and pack fees are on average lower for switchers, 6.4$ against 8.1$ and Amazon's presence as a competitor is lower in the first column. These differences might suggest a strategic nature behind the switch to FBA, to shield against competition from third-party sellers but suggest that the investment in adding a FBA offer might be less profitable when Amazon is a direct competitor. Moreover, the investment to switch suffer from the presence of FBA fees that might render it unprofitable. We also notice that the average of seller characteristics is different between the two columns. Switchers have accumulated less ratings than non-switchers, suggesting a shorter tenure on the platform, but tend to have a slightly higher rating than stayers.

Variables	Switched to FBA mean (s.e.)	Stayed FBM mean (s.e.)	Diff.	t-stat
BB Share	0.56 (0.45)	0.37 (0.44)	0.192	689
Full Price	46.7 (96.8)	53.6 (146.4)	-6.89	-107
Sales Rank	51668.6 (191957.6)	52887.8 (195521.8)	-1219.3	-10.2
Rating	44.3 (3.9)	44.1 (4.1)	0.281	114
Count Review	2222.3 (5933.6)	2242.4 (5851.8)	-20.1	-5.39
N FBA Offers	11.1 (23)	9.3 (19.1)	1.79	127
N FBM Offers	14.2 (19.7)	16.7 (17.9)	-2.46	-203
N Prime Offers	11.8 (23.4)	10.2 (19.4)	1.62	113
N Sellers	23.6 (34.8)	25 (28.8)	-1.37	-64.7
Pick and Pack Fee	6.4 (9.2)	8.1 (19.2)	-1.75	-263
Referral Fee Percent	13.6 (19.4)	13.6 (24)	-0.0387	-3.13
Amazon Competes	0.63 (0.48)	0.8 (0.4)	-0.168	-571
Seller Count Rating	39035.4 (106435.4)	47863.3 (118168.2)	-8827.9	-128
Seller Rating	93.5 (6.6)	91 (9.1)	2.52	565
Observations	2,893,564	25,343,198		
Individuals	81,776	209,924		

5 Identification Strategy

We recover the average treatment effect on the treated (ATT) of switching from a FBM offer to a FBA offer on the share of buy box time, prices and revenues. We do so by mean of a difference-in-difference strategy that we couple with an event study to address potential parallel-trend concerns. We use a Two-Way Fixed-Effect model, estimating the equation:[4]

$$Y_{ist} = \sum_{k=-6,k\neq -1}^{+12} \beta_k \cdot \text{Switch}_{i,s,t}^{k} + \gamma X_{i,t} + \delta M_{s,t} + \alpha_{i,s} + \theta_t + \mu_c + \varepsilon_{i,s,t} \tag{1}$$

where i indicates the product, s the seller, and t the time expressed in year-months. Our primary focus is on the dynamic treatment effect capturing the effect of the switch to FBA in the months leading up to and following the switch. The variable $\text{Switch}_{i,s,t}^{k}$ is an indicator for whether, in month k, the seller s has switched the fulfillment service to FBA for product i. The coefficients β_k measure the effect of the switch relative to the reference period, which we set to the month before the switch $k = -1$ and serves as the baseline period. We chose the months before the switch for reasons that relate to the practical aspects of switching to FBA. A seller opting to delegate its sales to Amazon has to physically ship the stock to a warehouse, after arranging the shipment either on its own or via Amazon chosen couriers. The terms $\gamma X_{i,t}$ and $\delta M_{s,t}$ represent control variables that account for product-specific and seller-specific characteristics, respectively. These controls help isolate the impact of the switch by accounting for other factors that might influence the outcome variable Y_{ist}. The fixed effects $\alpha_{i,s}$ control for unobserved, time-invariant characteristics specific to each product-seller pair, while θ_t denotes time fixed effects that account for temporal shocks or seasonality affecting all sellers or products in the same period. Finally, μ_c controls for category-level fixed effects, allowing us to account for differences across product categories. We couple the event study with a static version of equation 1, capturing the average effect over the post-treatment periods.

We add to these analyses a simple TWFE model of the rating of each seller in a month as a linear function of its own characteristics and the average characteristics of the markets it operates in. We do so to explore the association between the number of products in which each seller has added an FBA offer to an already present FBM one, and the perceived quality of the shipment service. As discussed introduction, a point of contention on behalf of Amazon is the inability of FBM Prime sellers to meet appropriate quality standards and that has caused the suspension of all enrollments, despite eligibility requirements. To investigate this point further we rely on the following specification

$$\text{Rating}_{s,t} = \text{Switched}_{i,s,t} + \gamma X_{s,t} + \delta M_{s,t} + \alpha_s + \theta_t \varepsilon_{s,t} \tag{2}$$

where we simply regress the rating of the seller s at year-month t on the number of product i in which she has switched, controlling for observables. We discuss the results next.

[4] all estimations are carried out via Bergé (2018) `fixest` package in R

6 Results

We begin by reporting the results for the static version of equation 1. Table 6 reports four dependent variables each with two columns showing results for the unconditional and the conditional specification, differing only in the product-specific and seller-specific controls. All standard-errors are cluster by the cross-ectional dimension of the panel, seller-product pairs.

The buy box share, our first outcome of interest, expresses the time spent in the buy box, the spot of the product page dedicated to the best offer, as a fraction of time in each month. In its case against Amazon, the FTC reports how roughly 98% of transactions are closed by the seller occupying the buy box making it a prominent outcome of investigation. Sellers adding the FBA offers gain buy box share by about 5%. This result is curbed once controls are introduced. The second column of table 6, in fact, suggests that, while still significant at 99%, switchers are able to grant only 1.2% more buy box after switching.

Prices, on the other hand, decrease of about 5% in both the unconditional and the conditional version of 1. This result is worth stressing in light of the FTC case, suggesting benefits from switching to the FBA program for consumers, on average. Moreover, as highlighted, FBA fees has seen a constant rise during our periods of interest, reaching in 2023 a share of one every two dollars of revenues going to the platform. A negative effect on prices might suggest a use of the FBA program that is used strategically by 3PS to either defend or increase their market share in selected markets. In this case the mechanism behind the result might be tied to competitive dynamics more than lower costs of logistic service, spilling over to consumers.

We also explore revenues estimated by multiplying the buy box share of each seller in each market by the inverse of its sales rank, recovering an estimate of sales. We then multiply the latter by the price and take the log. Switching to FBA is associated with a decrease in revenues of about 4%, suggesting once again, that FBM sellers might be ”forced” to enroll a particular product in the FBA program, possibly for increasing competition. Despite controlling for the numbr of FBA offers present in the market and the presence of Amazon as a direct competitor, resorting to FBA might not be enough to preserve market share, requiring the seller to decrease its price further.

Dependent Variables:	BB Share		log_price		Revenue Proxy		RP (Net of FBA Fees)	
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
fbm_to_fba_switched	0.0509***	0.0119***	-0.0562***	-0.0519***	-0.0430***	-0.0399***	-0.0506***	-0.0447***
	(0.0009)	(0.0009)	(0.0008)	(0.0009)	(0.0064)	(0.0064)	(0.0063)	(0.0063)
N FBA Offers		0.0028***		-0.0002***		0.0044***		0.0038***
		(2.05×10^{-5})		(1.52×10^{-5})		(0.0001)		(0.0001)
Sales Rank		-1.39×10^{-8}***		1.35×10^{-8}***		-8.1×10^{-7}***		-7.78×10^{-7}***
		(1.57×10^{-9})		(1.07×10^{-9})		(1.77×10^{-8})		(1.72×10^{-8})
Rating		-0.0006***		-0.0003***		-0.0004		-0.0006
		(7.1×10^{-5})		(5.1×10^{-5})		(0.0005)		(0.0005)
Seller Count Rating		-7.71×10^{-8}***		1.01×10^{-7}***		2.98×10^{-7}***		3.21×10^{-7}***
		(1.2×10^{-8})		(7.58×10^{-9})		(6.57×10^{-8})		(6.41×10^{-8})
Seller Rating		1.12×10^{-5}		-0.0004***		0.0011***		0.0010***
		(3.15×10^{-5})		(2.54×10^{-5})		(0.0002)		(0.0002)
Amazon Competes		-0.1895***		-0.0185***		-0.3151***		-0.2804***
		(0.0017)		(0.0009)		(0.0102)		(0.0100)
Fixed-effects								
asin-sellerId	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
L2 Category FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics								
Observations	26,766,971	24,457,904	26,344,105	24,124,376	26,270,299	24,124,376	26,173,440	24,058,093
R^2	0.78039	0.79253	0.98115	0.98077	0.64734	0.64448	0.64605	0.64496
Within R^2	0.00041	0.01287	0.00115	0.00141	8.34×10^{-6}	0.00441	1.22×10^{-5}	0.00400

Clustered (asin-sellerId) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Figure 4 depicts instead the dynamic effects behind the estimated of column (1) and (3), for buy box share and prices in the unconditional version of eq. 1, respectively. In both cases visual inspection suggests a possible violation of parallel trend but a clear change in the evolution of the outcome of interest upon switching. Given the magnitude of the estimates of the dynamic effect at different period post-switch we interpret the results as strong evidence for a causal relation, with a the true parameters possibly differing in magnitudes but not in signs. In both cases the event study plot a sudden change that is sustained over the long-run of the 12 months following the offer. Panel 4.1 focuses on the buy box share, with an increase in the first month after switch of 6.3% that slowly decreases during the course of the following months, to reach 4.8% by the end of the periods. As for prices, depicted in panel 4.2 show an immediate decrease in the short-run reaching the lowest point after six months. Both dynamic plots confirm that the static ATT capture the evolution of the outcome well.

We end by showing the results of equation 2. We stress that these results depict associations rather than causal estimates. We find no evidence for any association between the number of products in which a seller has switched and the ratings received. The coefficient of interest, switched, displays a negative sign but it's not significantly different from zero.

Figure 4: Event Study Plots of the effects of adding a FBA offer to Buy Box Share and Prices

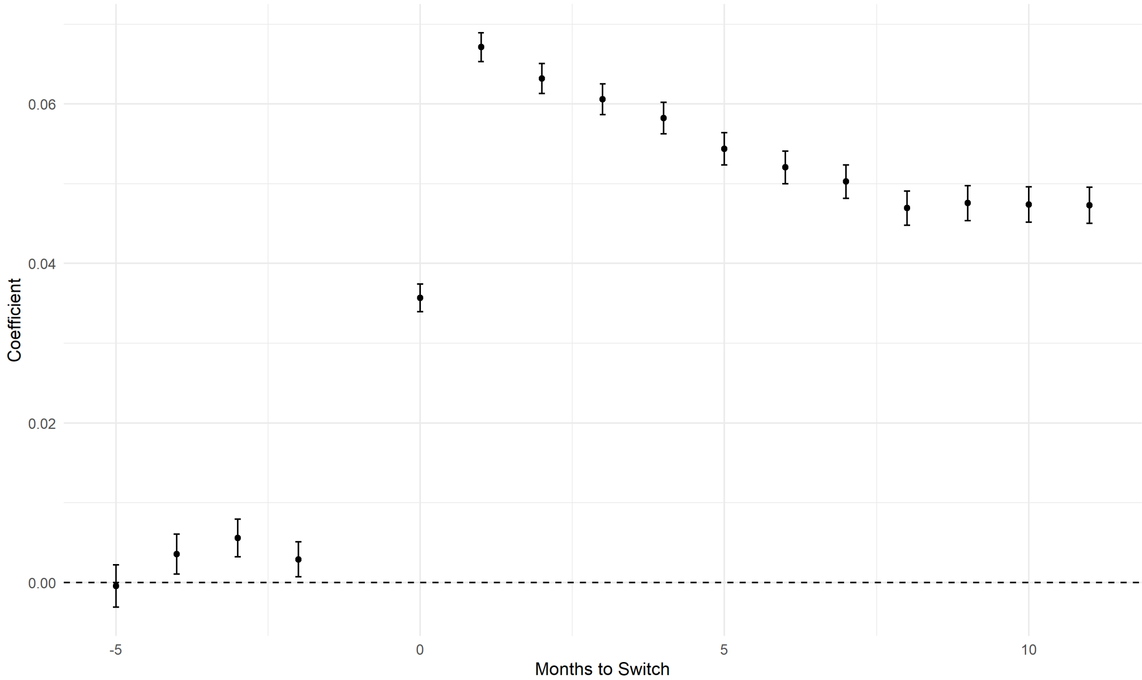

4.1 Buy Box Share

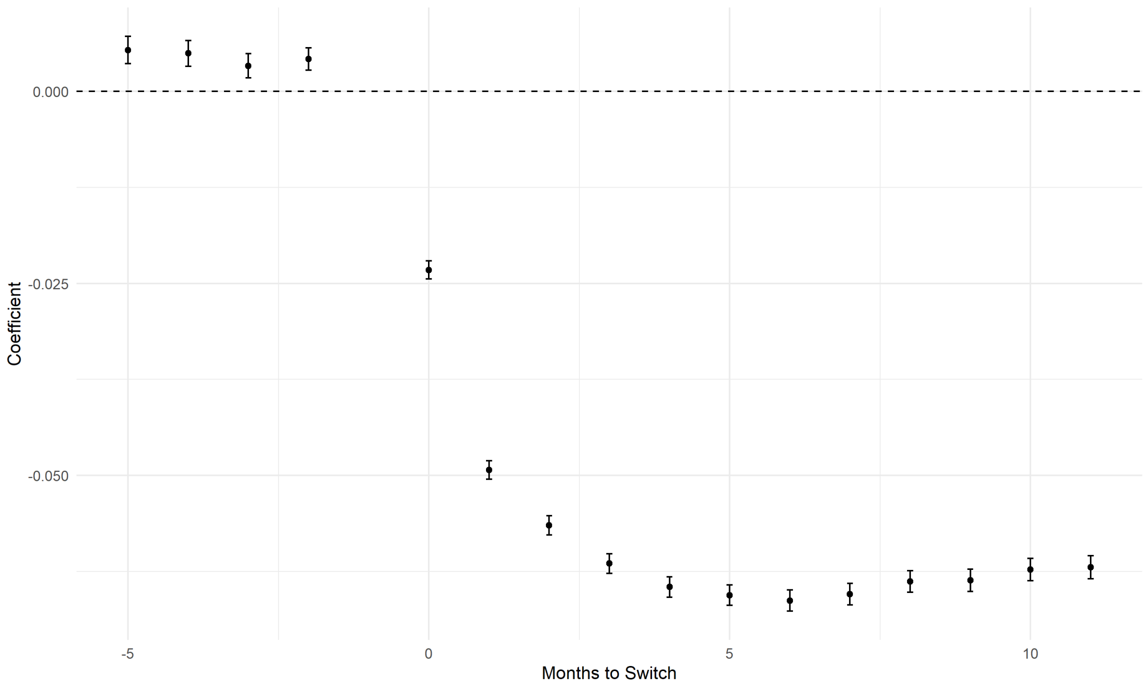

4.2 Prices

Notes: The panel contains event study plot showing the dynamic effect of adding a FBA offer to an already existent FBM offer for the same product, estimated via equation 1. The error bars show confidence interval at the 95% level of statistical significance, period-wise. Standard errors are clustered at the seller-product level. Figure 4.1 depicts the dynamic effect on the share of time spent in the buy box. Figure 4.2 depicts the dynamic effect on the log of prices, expressing elasticities.

Dependent Variable:	Seller Rating
Model:	(1)
Variables	
switched	-0.0143
	(0.0121)
Sales Rank	$-4.32 \times 10^{-7*}$
	(2.21×10^{-7})
Rating	-0.0007
	(0.0094)
Count Review	$-1.91 \times 10^{-5***}$
	(7.16×10^{-6})
N FBA Offers	-0.0225
	(0.0247)
N FBM Offers	-0.0175
	(0.0130)
N Prime Offers	0.0122
	(0.0196)
N Sellers	0.0168
	(0.0133)
Pick and Pack Fee	-0.0077
	(0.0102)
Amazon Competes	-0.0096
	(0.1076)
Log Price	-0.1377*
	(0.0721)
Fixed-effects	
sellerId	Yes
Year-Month FE	Yes
Fit statistics	
Observations	366,340
R^2	0.74027
Within R^2	0.00055

Clustered (sellerId) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

7 Conclusions

This paper sheds light on the relationship between Amazon's logistics programs and its marketplace dominance, focusing on the effects of the FBA program on third-party sellers. Our analysis of over 25 million seller-product pairs reveals several key insights. Sellers that switch to FBA experience an increase in buy box share, but the magnitude is economically insignificant when controlling for market characteristics. Switching to FBA is associated with a 5% reduction in prices, indicating a potential competitive strategy employed by sellers to defend market share, despite the rising costs of Amazon's fulfillment services. Contrary to Amazon's claims, we find no significant improvement in ratings for sellers who switch to FBA, suggesting that FBA does not necessarily enhance service quality as perceived by consumers. Finally, FBA adoption is correlated with a reduction in revenues, which might indicate that sellers feel compelled to adopt the program in response to increasing competition from Amazon or other FBA sellers. Our findings offer important implications for antitrust discussions, particularly in the context of the FTC's case against Amazon. The evidence suggests that Amazon's bundling of logistics services may lead to competitive distortions that harm third-party sellers, even though consumers may benefit from lower prices.

In future research, we plan to explore how the heterogeneity among sellers shapes the effects of switching to Amazon's FBA program. Sellers on Amazon's marketplace vary significantly in terms of size, experience, and the type of products they sell. Larger sellers, for instance, may already have their own logistical infrastructure, making the shift to FBA less attractive. Small sellers, especially those who are new to the platform, might leverage FBA to access Amazon's Prime customer base and acquire market share. Another direction for future research is a detailed examination of how policy changes, such as Amazon's decision to limit Prime eligibility to FBA participants in 2019, have affected the marketplace. We plan to employ a difference-in-difference methodology to compare seller outcomes before and after this significant policy shift. Specifically, we aim to analyze how the restriction of Seller Fulfilled Prime (SFP) impacted sellers who were previously able to fulfill their orders independently while still benefiting from Prime membership. By focusing on this policy change, we can isolate its effects on different seller groups, such as those who may have been forced into using FBA to retain Prime access.

References

Bergé, L. (2018). Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm. *CREA Discussion Papers*, (13).

Carlton, D. W. and Waldman, M. (2002). The strategic use of tying to preserve and create market power in evolving industries. *The RAND Journal of Economics*, pages 194–220.

Hart, O. and Tirole, J. (1990). Vertical integration and market foreclosure. *Brookings papers on economic activity. Microeconomics*, pages 205–286.

Khan, L. M. (2017). Amazon's antitrust paradox. *Yale Law Journal*, 126:710–805.

Lafontaine, F. and Slade, M. E. (2007). Vertical integration and firm boundaries: The evidence. *Journal of Economic Literature*, 45(3):629–685.

Nalebuff, B. (2004). Bundling as an entry barrier. *The Quarterly Journal of Economics*, 119(3):159–187.

on Digital Platforms, S. C. (2019). *Report of the Stigler committee on digital platforms*. University of Chicago Booth School of Business.

Salop, S. C. and Scheffman, D. T. (1983). Raising rivals' costs. *The American Economic Review*, 73(2):267–271.