

An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace

Le Chen
Northeastern University
Boston, MA USA
leonchen@ccs.neu.edu

Alan Mislove
Northeastern University
Boston, MA USA
amislove@ccs.neu.edu

Christo Wilson
Northeastern University
Boston, MA USA
cbw@ccs.neu.edu

ABSTRACT

The rise of e-commerce has unlocked practical applications for *algorithmic pricing* (also called dynamic pricing algorithms), where sellers set prices using computer algorithms. Travel websites and large, well known e-retailers have already adopted algorithmic pricing strategies, but the tools and techniques are now available to small-scale sellers as well.

While algorithmic pricing can make merchants more competitive, it also creates new challenges. Examples have emerged of cases where competing pieces of algorithmic pricing software interacted in unexpected ways and produced unpredictable prices [37], as well as cases where algorithms were intentionally designed to implement price fixing [5]. Unfortunately, the public currently lack comprehensive knowledge about the prevalence and behavior of algorithmic pricing algorithms in-the-wild.

In this study, we develop a methodology for detecting algorithmic pricing, and use it empirically to analyze their prevalence and behavior on Amazon Marketplace. We gather four months of data covering all merchants selling any of 1,641 best-seller products. Using this dataset, we are able to uncover the algorithmic pricing strategies adopted by over 500 sellers. We explore the characteristics of these sellers and characterize the impact of these strategies on the dynamics of the marketplace.

1. INTRODUCTION

For the last several years, growth in e-commerce has massively outpaced growth among traditional retailers. For example, while retail sales shrank 1.3% in the first quarter 2015 in the US, e-commerce grew 3.7% [21]. Although e-commerce only accounts for around 7.3% of the overall \$22 trillion in global retail spending projected for 2015, this percentage is projected to rise to 12.4% by 2019 [27]. Furthermore, these overall figures mask the disproportionate gains of e-commerce in specific sectors, such as apparel, media, and office supplies.

The rise of e-commerce has unlocked practical applications for *algorithmic pricing* (sometimes referred to as dynamic pricing algorithms or Revenue/Yield Management). Algorithmic pricing strategies are challenging to implement in traditional retail set-

tings due to lack of data (e.g., competitors' prices) and physical constraints (e.g., manually relabeling prices on products). In contrast, e-commerce is unconstrained by physical limitations, and collecting real-time data on customers and competitors is straightforward. Travel websites are known to use personalized pricing [25], while some e-retailers are known to automatically match competitors prices [40, 17].

While algorithmic pricing can make merchants more competitive and potentially increase revenue, it also creates new challenges. *First*, poorly implemented pricing algorithms can interact in unexpected ways and even produce unexpected results, especially in complex environments populated by other algorithms. For example, two competing dynamic pricing algorithms inadvertently raised the price of a used textbook to \$23M on Amazon [37]; reporters have noted that similar algorithmic pricing also exists in day-to-day commodities [9]. *Second*, dynamic pricing algorithms can implement collusive strategies that harm consumers. For example, the US Justice Department successfully prosecuted several individuals who implemented a price fixing scheme on Amazon using algorithms [5]. Unfortunately, regulators and the public currently lack comprehensive knowledge about the prevalence and behavior of algorithmic pricing algorithms in-the-wild.

In this study, our goal is to empirically analyze deployed algorithmic pricing strategies on Amazon Marketplace. Specifically, we want to understand what algorithmic pricing strategies are used by participants in the market, how prevalent these strategies are, and ultimately how they impact customer experience. We chose to focus on Amazon for three reasons: *first*, Amazon is the largest e-commerce destination in the US and Europe [16]. *Second*, Amazon is a true marketplace populated by third-party sellers, as well as Amazon itself. *Third*, Amazon's platform provides APIs that are specifically designed to facilitate algorithmic pricing [1].

To implement our study, we develop a novel methodology to collect data and uncover sellers that are likely using algorithmic pricing. We collect four months of data from 1,641 of the most popular products on Amazon. We gather information about the top-20 sellers of each product every 25 minutes, including the sellers' prices, ratings, and other attributes. We use this data to analyze changes in price over time, as well as compare the attributes of sellers. We focus on top selling products because they tend to have multiple sellers, and thus are likely to exhibit more competitive dynamics.

We begin by analyzing the algorithm underlying Amazon's *Buy Box*. This algorithm determines, for a given product being sold by many sellers, which of the sellers will be featured in the Buy Box on the product's landing page (i.e., which seller is the "default" seller). As shown in Figure 1, customers use the Buy Box to add products to their cart; sellers not selected for the Buy Box are relegated to a separate webpage. The precise features and weights used by the

Buy Box algorithm are unknown [13], yet the algorithm is of critical importance since 82% of sales on Amazon go through the Buy Box [38]. For our purposes, understanding the Buy Box algorithm is important because sellers may choose dynamic pricing strategies that maximize their chance of being selected by the algorithm.

Next, we examine the dynamic pricing strategies used by sellers in Amazon Marketplace. To identify pricing algorithms, we treat the *target price* of each product (e.g., the lowest advertised price or Amazon's price) as a time series, and use correlative analysis to identify specific sellers whose prices track the target price over time. Overall, we identify over 500 sellers who are very likely using algorithmic pricing.

Finally, we compare the characteristics of algorithmic and non-algorithmic sellers. We observe that algorithmic sellers appear to be more successful than non-algorithmic sellers: they offer fewer products, but receive significantly higher amounts of feedback (suggesting they have much higher sales volumes). Furthermore, algorithmic sellers "win" the Buy Box more frequently (even when they do not offer the lowest price for a given product), which may further contribute to their feedback scores. However, we also observe that the lowest price and the Buy Box for products with algorithmic sellers are significantly more volatile than for products without any algorithmic sellers. These rapidly fluctuating prices may lead to customer dissatisfaction [9].

In summary, this work makes the following contributions:

1. We present a comprehensive overview of dynamics on Amazon Marketplace, including the characteristics of sellers, and frequency of price changes.
2. Using Machine Learning (ML), we determine that, among all the variables we can observe, low prices are the most important feature used by the Buy Box algorithm to select sellers, but that customer feedback and ratings are also used.
3. We develop a technique to detect sellers likely using algorithmic pricing, and identify 543 such sellers.
4. We explore the properties of these sellers, showing they are strategic and successful; they have much higher levels of feedback than other sellers, and are more likely to be featured in the Buy Box.

To facilitate further study, we make our code and data available at

<http://personalization.ccs.neu.edu>

Outline. The remainder of this paper is organized as follows. § 2 covers background on Amazon and the Amazon Marketplace, and § 3 covers our data collection methodology. § 4 explores the algorithm that Amazon uses to select the Buy Box winner. § 5 presents our algorithm for detecting sellers using algorithmic pricing, and § 6 explores the characteristics and impact of these sellers. § 7 presents related work and § 8 concludes.

2. BACKGROUND

We begin by briefly introducing Amazon Marketplace. We focus on the features of the market that are salient to algorithmic pricing, including Third-Party (3P) sellers, the Buy Box, and finally the APIs offered by Amazon Marketplace Web Services.

2.1 Amazon Marketplace

Amazon, founded in 1994, is the largest e-commerce website in the US and Europe [27]. Although Amazon began as an online bookstore, it now sells over 20 categories of physical products (even fresh food in select cities [15]), as well as a wide range of digital goods (e.g., downloadable and streaming music, video, and




Figure 1: An example Buy Box on Amazon.

e-books). Overall, Amazon earned \$89B in revenue in 2014, and boasts 244M active customers [22].

Amazon inspires fierce loyalty among customers through their Prime membership program, which gives customers free 2-day shipping (or better) as well as unlimited access to digital streams for \$99/year. Amazon's success is further bolstered by their branded digital devices (Kindle e-readers, tablets, phones *etc.*), which push customers towards Amazon's shopping apps. Because of these customer retention efforts, 44% of online shoppers navigate directly to Amazon to make purchases, rather than using search engines or visiting competing online retailers [35].

3P Sellers and FBA. In addition to acting as a merchant, Amazon also functions as a marketplace for third parties. Amazon claims to have 2M Third-Party (3P) sellers worldwide who sold 2B items in 2014, representing 40% of all items sold via the website [3]. 3P sellers can opt to handle logistics (inventory, shipping, returns, *etc.*) themselves, or they can join the Fulfilled By Amazon (FBA) program, in which case Amazon handles all logistics.

The fee structure for 3P sellers is complicated, and involves five components [4, 6]:

1. **Seller Fee:** "Individual" sellers must pay \$0.99/item sold, or sellers may become "Pro Merchants" for \$39.99/month.
2. **Referral Fee:** Amazon assesses a referral fee on each product sold. The fees vary between 6-45% of the total sale price, depending on the product category. The vast majority of categories have a 15% referral fee. Amazon also enforces minimum referral fees of \$1-\$2/item.
3. **Closing Fee:** Amazon's closing fees vary based on product category, shipping method, and product weight. Media products (books, DVDs, *etc.*) have a flat fee of \$1.35/product. Other products have a \$0.45 + \$0.05/lb fee for standard shipping, or \$0.65 + \$0.10/lb for expedited shipping.
4. **Listing Fee:** High-volume sellers that list more than 2M Stock Keeping Units (SKUs, a seller-specified representation of an item) per month must pay \$0.0005 per active SKU.
5. **FBA Fee:** Sellers that use FBA must pay a \$1.04-\$10.34 packing fee per product depending on its size and type, plus variable per pound shipping fees ranging from \$0.39 for small media items, to \$124.58 for extremely heavy, irregularly shaped items.

As we discuss in § 5, these fees influence the dynamic pricing strategies used by 3P sellers.

2.2 The Buy Box

When customers purchase products from Amazon, they typically do so through the *Buy Box*. The Buy Box is shown on every product




Figure 2: Frequency of page updates.




Figure 3: Examples of price jitter (highlighted with arrows) in the Buy Box on a product page.

page on Amazon: it contains the price of the product, shipping information, the name of the seller, and a button to purchase the product. Figure 1 shows an example Buy Box.

However, many products on Amazon are sold by multiple sellers. In these cases, a proprietary Amazon algorithm determines which seller’s offer is displayed in the Buy Box. Formally, if product is being offered by n sellers with prices $P = \{p_1, \dots, p_n\}$, the Buy Box algorithm is a function $B(P) \rightarrow p_i$, with $p_i \in P$. As shown in Figure 1, offers from other sellers are relegated to a separate webpage (an example is shown in Figure 4).

Given the prominent placement of the Buy Box, it is not surprising that 82% of sales on Amazon go through it [38]. This has made the underlying algorithm the focus of much speculation by 3P sellers [13]. Although Amazon has released some information about the features used by the Buy Box algorithm (*e.g.*, prices, shipping options and speed) [7], it is unknown whether this feature list is complete, or what the weights of the features are.

Because “winning” the Buy Box is so critical for making sales on Amazon, sellers may use dynamic pricing strategies that give them an advantage with respect to being chosen by the algorithm. Thus, we use Machine Learning (ML) to examine the Buy Box algorithm in-depth in § 4.

2.3 Amazon Marketplace Web Service

Amazon offers an array of tools to help 3P sellers manage product inventory. The most sophisticated of these tools is the Amazon Marketplace Web Service (MWS), which is a set of APIs for programatically interfacing with the marketplace. MWS includes functions for listing products, managing inventory, and changing

prices.¹ MWS also has a subscription API, that allows sellers to receive near real-time price updates for specified products. Each update includes aggregated information about the lowest 20 prices offered for a product (or less, if there are fewer than 20 offers).

In addition to MWS, Amazon also has a web-based price matching tool for 3P sellers [8]. This tool allows a 3P seller to set a product’s price equal to the lowest competing offer. However, this tool only adjusts the product’s price *once*: if the lowest price changes again, the seller’s price is not automatically reduced as well.

Seller Platforms. The capabilities of MWS are clearly designed to facilitate dynamic pricing. Companies like Sellery, Feedvisor, Appeagle, RepricerIt, and RepricerExpress leverage MWS to offer subscription-based services for 3P sellers that combine inventory management with dynamic pricing capabilities. These services enable any merchant to easily become a 3P seller and leverage sophisticated dynamic pricing strategies. We discuss the types of strategies offered by these services in greater detail in § 5.

3. DATA COLLECTION

The goal of our study is to analyze the dynamic pricing strategies being used by sellers on Amazon. To achieve this goal, we require longitudinal data about sellers and their prices—ideally for a large number of products—in the marketplace. In this section, we describe our data collection process, including specific challenges that we needed to overcome to obtain useful, representative data.

3.1 Obtaining Sellers and Prices

We would ideally have liked to use the Amazon Marketplace Web Services (MWS) API to collect the seller and price information for products. Unfortunately, we found that the API did not meet our requirements for two reasons: the API does not return the identity of 3P sellers (just their chosen price), and the API is heavily rate-limited.

Instead, we used *web scraping* to obtain information on the active sellers and their prices. Specifically, for each product we examine, we crawled the *New Offers* page² (the page that is linked to in Figure 1 if one clicks on “2 new”, shown in Figure 4). This page lists all 3P sellers, their prices, their shipping costs, and their reviews (number of reviews and average score). Unfortunately, this information is paginated into 10 3P sellers per page; we describe below how we handle cases where there are more than 10 3P sellers.

In addition to scraping the *New Offers* pages for products, we also scraped the product pages themselves. We use the data from the product pages to analyze the Buy Box algorithm in § 4.

¹Amazon’s documentation stipulates that sellers may only update prices every 20 minutes [2].

²In this study, we only focus sellers who offer new items; used items are not covered, and we leave them to future work.




Figure 4: An example *New Offers* page on Amazon, listing all sellers for a given product.




Figure 5: Cumulative distribution of product prices.




Figure 6: Cumulative distribution of number of sellers per product.




Figure 7: Correlation between price and rank (1 is perfect correlation, -1 is anti-correlation).

Calculating Prices. It is important to note that the *New Offers* page lists both the *base price* and the *total price* (i.e., price including the lowest-cost shipping option) for each seller. Throughout the paper, when we refer to “price”, we are referring to the total price. We do this as Amazon uses total price when users explicitly sort products by price; users cannot search or sort by base price alone.

3.2 Determining Crawling Frequency

Because 3P sellers (and Amazon) can change their price at any time, we need to decide how frequently we will crawl each page. To do so, we create a high-resolution dataset that will help to illuminate the tradeoff between crawling resolution and frequency. Specifically, we randomly selected 5 products from the best-seller products³ and crawled their product page and the first 2 seller pages (covering up to 20 sellers) once per minute for 3 days.

We first examine how frequently sellers’ prices change and how frequently the Buy Box is updated. We plot the cumulative distribution of inter-update times for sellers, the Buy Box price, and the Buy Box seller in Figure 2. We observe that the updates are surprisingly dynamic: 40% of price changes occur within a minute of the previous price change, with a long tail of update times. To explore the origins of this high level of dynamicity, we plot a timeseries of the Buy Box price of an example product in Figure 3 (we observed similar behavior for other products, but do not include them due to space constraints).

We observe that the price appears to change five times in this timeseries, but that old prices sometimes briefly reappear after a price change. This result is likely due to Amazon’s distributed infrastructure; Amazon states that it can take up to 15 minutes for all systems to converge to the new price. Thus, the very rapid price “jitters” are likely caused by transient inconsistencies in Amazon’s infrastructure, rather than actual price changes by sellers.⁴

Using these results, we select a crawling frequency. As a tradeoff between number of products and crawling frequency, we choose to cover more products at longer intervals. As shown in Figure 2, most changes happen either on very short timescales (< 1 minute; likely Amazon inconsistencies) or very long timescales (> 30 minutes). We therefore choose a crawling frequency of every 25 minutes.

³<http://www.amazon.com/Best-Sellers/zgbs>, Best-seller products come from 23 departments from Amazon, such as Appliances, Beauty, Electronics, etc. Altogether there are 1,790 best-seller products (we exclude digital goods such as e-books, downloadable music, and gift cards).

⁴To further verify these results, we set up an Amazon Individual Seller account, listed several products, and changed their prices at specific times. We found that when prices are in an inconsistent state, a customer cannot add the item to their shopping cart (i.e., even though a customer may see an outdated price, the customer is not able to add the product to their cart at the old price).

3.3 Selecting Products

Next, we turn to selecting the products to study. Recall that we are aiming to study dynamic pricing; not all products are equally likely to have such sellers, so we focus on best-selling products since they are likely to have many competing sellers. We conduct two separate crawls that have different characteristics.

First Crawl (Crawl1). Our first crawl was conducted between September 15, 2014 and December 8, 2014. We select 837 best-selling products that had at least two sellers at the beginning of the crawl. For this crawl, we downloaded *all* seller pages, but did not download the product page (containing the Buy Box).

Second Crawl (Crawl2). We conduct a second crawl between August 11, 2015 and September 21, 2015. We select 1,000 best-selling products to study, and downloaded both the product page (containing the Buy Box) and the first two pages of 3P sellers (typically, but not always, containing the 20 sellers with the lowest prices). We choose to only download the first two pages of sellers, as we found the sellers who change their prices often (suggesting dynamic pricing algorithms) were within the first two pages 96% of the time. Thus, downloading only the first two pages massively reduces the amount of data we need to collect while still capturing most of the “interesting” behavior.

It is important to note that the first and second crawls cover different products, as the best-selling products change over time: there are 196 products in common between the two crawls. As shown in Figures 5 and 6, the overall characteristics of prices and sellers are very similar between the two crawls despite the time difference (details of these Figures are discussed in the next section).

3.4 Limitations

There are two noteworthy limitations to our dataset. First, our dataset is biased (by design) towards best-selling products. To briefly quantify this bias, we randomly sampled 2,158 products from a public listing of all Amazon products.⁵ We compare the product price and the number of sellers in Figures 5 and 6; as expected, we observe that our best-sellers show many more sellers than random products, as well as somewhat lower prices.

Second, we crawled data from Amazon using browsers that were not logged-in to Prime accounts. Although the exact number of Prime members is unknown, estimates place it at around 20–40% of all Amazon’s customers [23]. Thus, our dataset should accurately reflect what the majority of Amazon users see. However, Amazon may alter pages for Prime users, typically to highlight sellers and products that are eligible for expedited Prime shipping. Thus, some of our analysis and conclusions may not extend to Prime users.

⁵https://archive.org/details/asin_listing




Figure 8: Cumulative distribution of changes to the Buy Box price and winner, per product.




Figure 9: Probability of winning the Buy Box for sellers at different ranks.




Figure 10: Buy Box winner prediction accuracy for products with different numbers of sellers.

4. THE BUY BOX

We begin our analysis by exploring how Amazon’s systems evaluate sellers. *First*, we briefly examine Amazon’s seller ranking algorithm, and follow up by characterizing the dynamics and behavior of the Buy Box. In both cases, we observe that Amazon uses non-trivial strategies to evaluate sellers (*i.e.*, price is not the only factor that impacts ranking and selection for the Buy Box). *Second*, we conduct an in-depth investigation of the features and weights that drive the Buy Box algorithm. Understanding the Buy Box algorithm is crucial, since it may influence how sellers choose dynamic pricing strategies.

Note that in this section, we only use data from Crawl2, since it contains Buy Box winners and seller rankings.

4.1 Seller Ranking

As shown in Figure 4, Amazon explicitly *rank*s all sellers for each product on the *New Offers* page. However, the Buy Box winner is not necessarily the seller who is ranked the highest. Thus, we first examine the seller ranking algorithm as it offers clues as to how Amazon chooses to weigh various seller features.

We collect the rankings for all products in our dataset, and calculate Spearman’s Rank Correlation (ρ) between the ordered list of sellers returned by Amazon, and the list of sellers sorted by price, for each product in our dataset. If the lists perfectly correspond (*i.e.*, Amazon returns sellers sorted by price), then Spearman’s ρ will equal 1. Contrary to our expectations, Amazon does not always sort sellers by price. As shown in Figure 7, around 20% of products have correlation < 1 . This result gives us our first clue that Amazon’s systems take additional seller attributes into account (besides price) when making decisions.

4.2 Behavior of the Buy Box Algorithm

Next, we examine the empirical behavior of the Buy Box, starting with dynamics over time. Figure 8 plots the number of changes to the price and seller in the Buy Box for all products in the Crawl2 dataset over the six weeks of observation. We immediately see that most products see a number of changes: only 13% of products have static Buy Box prices over the entire period, while 50% of products have more than 14 changes. However, we see fewer changes to the Buy Box winner: the seller winning the Buy Box is constant for 31% of products. Thus, for many products, the Buy Box winner and price is highly dynamic; some products even experience hundreds of changes, or many more than one per day.

Next, we examine the relationship between seller rank (from the *New Offers* page) and the winner of the Buy Box. Figure 9 shows the fraction of sellers at different ranks that “win” the Buy Box, *i.e.*, are chosen by the algorithm. Rank zero means they are the first seller in the list. Surprisingly, only 60% of the top-ranked sellers win the Buy Box, and there is a long tail of sellers at higher ranks that win. Recall that we have already shown that Amazon does not

rank sellers solely on prices (see Figure 7). Taken together, these results show that Amazon’s systems take additional characteristics beyond price into account when evaluating sellers.

4.3 Algorithm Features and Weights

In the previous section, we demonstrated that the Buy Box algorithm uses features beyond just price to select the Buy Box winner. In this section, we use Machine Learning (ML) to try to infer some of the features and weights used by the Buy Box algorithm.

Model and Features. To facilitate our analysis, we model the Buy Box as a prediction problem. Specifically, for a product offered by n sellers, each of which is characterized by a feature vector, our goal is to predict which seller will be chosen to occupy the Buy Box. Given our dataset, we construct a feature vector for each seller containing the following seven features:

1. *Price Difference to the Lowest*: difference between the seller’s price and the current lowest price for the product.
2. *Price Ratio to the Lowest*: ratio between the seller’s price and the current lowest price of the product.
3. *Average Rating*: average customer rating⁶ of the seller.
4. *Positive Feedback*: positive feedback percentage for the seller.
5. *Feedback Count*: total feedback count for the seller.
6. *Is the Product FBA?*: true if the seller uses FBA.
7. *Is Amazon the Seller?*: true if the seller is Amazon.

According to Amazon’s documentation, as well as speculation from 3P sellers, other features are possibly used by the Buy Box algorithm [13, 7]. This includes sales volume, response time to customer inquiries, rate of returns and refunds, and shipping times. Unfortunately, we cannot measure these features, and thus cannot quantify their impact on the Buy Box algorithm. However, as we will show, even without these features we are able to achieve high prediction accuracy, suggesting that our data does capture the many important seller features.

Classifier Selection. We leverage a Random Forest (RF) classifier to predict whether a specific seller of a product wins the Buy Box. RF is an ensemble classifier that achieves low bias and low variance by aggregating the decisions from a large number of low-correlated decision trees with different feature combinations and bagging samples [36]. Furthermore, RF is an ideal classifier for our task because it outputs interpretable measures of feature importance (calculated as the average of the Gini index among all splits in the trees). In contrast, other classifiers, such as kernel-based SVM, are harder to interpret.

⁶Fully described in § 6.1, the rating and feedback of a seller come from customer surveys asking for a star rating (0–5 stars) and specific questions about the customer’s experience.

Feature	Weight
Price Difference to the Lowest	0.36
Price Ratio to the Lowest	0.33
Positive Feedback	0.10
Is Amazon the Seller?	0.10
Feedback Count	0.06
Average Rating	0.03
Is the Product FBA?	0.02

Table 1: Relative importance of different features in winning the Buy Box, as determined by our RF classifier.

Evaluation. Figure 10 shows the accuracy of our RF classifier at predicting the winner of the Buy Box (using 10-fold cross-validation) for all products in Crawl2, as a function of the number of sellers for a given product. Obviously, it is trivial to achieve 100% accuracy in the 1-seller case; however, we see that the classifier achieves 75–85% accuracy even in the most challenging cases with many sellers.

To put the accuracy results of our classifier in perspective, Figure 10 also depicts the accuracy of two naïve *baseline* classifiers. One baseline classifier always predicts that the seller with the lowest price will win the Buy Box (if there are multiple sellers offering the same lowest price, it chooses the lowest ranked one), while the other chooses the lowest ranked seller. Both baselines only achieve 50–60% accuracy, which reconfirms that price is not the sole feature used by the Buy Box algorithm, and also highlights the impressive predictive power of our RF classifier.⁷

Feature Weights. Finally, we examine the weights calculated for each feature by our RF classifier. Higher weights mean that the feature is more predictive of who will win the Buy Box. As shown in Table 1, the two price-based features are significantly more important than other features. However, the seller’s positive feedback and feedback count are also important metrics for winning the Buy Box. Interestingly, we observe that using FBA has low importance, which contradicts conventional wisdom about how the Buy Box algorithm functions [13]. However, it is possible that FBA is an important factor in cases where the customer is an Amazon Prime member, since FBA confers free shipping for Prime users.

Although we observe that “being Amazon” does confer some advantage, we caution that this does not necessarily mean that Amazon has tilted the Buy Box algorithm in their favor. Recall that Amazon’s Buy Box documentation states that the algorithm uses several features that we are unable to measure, such as sales volume [7]. It is possible that Amazon scores highly in these missing features, which manifests in our classifier as additional weight in the “Amazon is the Seller” feature.

Overall, we observe that Amazon’s algorithm for choosing the winner of the Buy Box is a combination of a number of undocumented features and weights. We are able to gain some visibility into this algorithm, with the results indicating that price, seller feedback, and feedback count are all important features. These results suggest that sellers who use algorithmic strategies to maintain low prices relative to their competitors are likely to gain a large advantage in the struggle to win the Buy Box.

5. DYNAMIC PRICING DETECTION

We now turn to detecting algorithmic pricing on Amazon Marketplace. We note that doing so is non-trivial, as external observers

⁷Our dataset includes at least 1K samples at each rank, thus the results are statistically significant.

Strategy	Threshold = 10		Threshold = 20	
	Sellers	Products	Sellers	Products
Lowest Price	726	544	426	408
Amazon Price	297	277	176	183
2nd Lowest Price	721	494	425	370
Total	918	678	543	513

Table 2: Number of sellers and products with detected algorithmic pricing, based on two different change thresholds. We use a change threshold of 20 unless otherwise stated.

such as ourselves are only able to measure the prices offered by sellers (and not their usage of the Amazon Marketplace API, *etc*). Moreover, we lack ground truth on which sellers are using algorithm pricing. Therefore, we build a detection algorithm that tries to locate sellers that *behave* like “bots”, *i.e.*, sellers where the prices they set and the timing of changes suggest algorithmic control.

5.1 Methodology

We hypothesize that sellers using algorithmic pricing are likely to base their prices at least partially on the prices of other sellers. This makes sense intuitively: for example, a seller who always wants to offer the lowest on a specific product must set their price relative to the competitor with the lowest price. Thus, we first define several *target prices* that the seller could match against. We motivate our selection of target prices by examining popular repricing software for Amazon Marketplace,⁸ and choose three target prices for each product: lowest price, Amazon’s price, and the second lowest price.⁹

For a given seller/product pair (s, r) , we construct a time series of the prices p_i offered by the seller at time t_i :

$$S_r = \{(t_0, p_0), (t_1, p_1), \dots, (t_m, p_m)\}$$

We also construct three target price time series, corresponding to the lowest price p_i^{low} , the 2nd lowest price p_i^{2nd} , and Amazon’s price p_i^{amzn} for r at each time t_i :

$$LOW_r = \{(t_0, p_0^{low}), (t_1, p_1^{low}), \dots, (t_m, p_m^{low})\}$$

$$2ND_r = \{(t_0, p_0^{2nd}), (t_1, p_1^{2nd}), \dots, (t_m, p_m^{2nd})\}$$

$$AMZN_r = \{(t_0, p_0^{amzn}), (t_1, p_1^{amzn}), \dots, (t_m, p_m^{amzn})\}$$

Note that when we construct the three target price time series, we exclude the prices offered by s . For example, if s always offers the lowest price for r , then LOW_r will actually contain the second-lowest price for r at each time t_i . This exclusion rule also prevents us from comparing Amazon’s prices against themselves. Finally, note that Amazon does not sell all products in our dataset, thus only a subset of seller/product pairs include $AMZN_r$.

Once we have constructed the time series corresponding to (s, r) , we calculate the similarity between S_r and LOW_r , $2ND_r$, and $AMZN_r$ (respectively) using Spearman’s Rank Correlation. When ρ is large, it means that the price changes contained in the pair of time series occur at the same moments, and that the magnitude of the price changes are relatively constant. We mark pairs with $\rho \geq 0.7$ (the empirical cutoff of a strong positive correlation) and $p\text{-value} \leq 0.05$ as *algorithmic pricing candidates*.

The final step in our methodology is to filter our candidates. Intuitively, if a seller exhibiting high correlation with the target price

⁸<https://sellerengine.com/sellery/>

⁹In fact, Sellery provides one more option: matching to the average price. However, this strategy is neither likely to be useful for winning the Buy Box, nor being competitive among the sellers.




Figure 11: Number of algorithmic sellers detected with different change thresholds.




Figure 12: CDF of absolute price differences between algorithmic sellers and target prices.




Figure 13: CDF of relative price differences between algorithmic sellers and target prices.




Figure 14: Example of 3P seller (in red) matching the lowest price of all other sellers.




Figure 15: A second example of 3P seller (in red) matching the lowest prices offer by two other sellers.




Figure 16: Example of Amazon (in red) setting a premium over the lowest price of all other sellers.




Figure 17: Example of Amazon (in red) matching to the lowest price over time.

also makes a large number of price changes, this provides more evidence that the seller is using algorithmic pricing. Conversely, if the number of price changes in the time series is small, then it is possible that the correlation is coincidental. Thus, we define the *change threshold* as the minimum number of price changes that must occur in a time series S_r for us to consider s as using algorithmic pricing.

Figure 11 shows the number of sellers that we consider to be doing algorithmic pricing when we apply different change thresholds. As expected, we observe that the number of sellers decreases rapidly as we increase the change threshold. Unless otherwise stated, in the remainder of the paper, we choose 20 as our change threshold since it represents a conservative threshold that is in the “knee” of the distribution.

5.2 Algorithmic Pricing Sellers

Now that we have described our methodology, we briefly examine the set of sellers that we find to be doing algorithmic pricing. Table 2 shows the number of algorithmic pricing sellers and the number of products they sell that we detect with change thresholds of 10 and 20. In this table, we merge the sellers and products from Crawl1 and Crawl2 and present the total unique numbers.

We immediately observe that many more sellers appear to be using the overall lowest price (and 2nd lowest price) as the target for their algorithmic pricing than Amazon’s price. However, it is important to note the different strategies are not necessarily mutually exclusive. For example, a seller matches to both Amazon and low-

est when Amazon is the lowest price, and a seller matching to the lowest price is likely to often match (*i.e.*, correlate strongly with) to the second-lowest price as well. The **Total** line shows the overall unique numbers of sellers and products we detect. In the case when the change threshold is 20, we see that 2.4% of all sellers in our dataset use algorithmic pricing. However, this is 38% of all sellers that have ≥ 20 changes for *at least* one product they sell.

To determine the gap between the prices offered by the suspected algorithmic sellers and the target prices, we plot two figures. Figure 12 examines the *absolute* difference between the algorithmic sellers’ prices and the corresponding target prices. We separate sellers matching to the lowest price and sellers matching Amazon’s price in this plot (we ignore the second lowest price in this plot as matching to the second lowest price is very similar to matching to the lowest). We observe that algorithmic sellers who match to the lowest price are very close to the lowest price: 70% of these sellers set their price within \$1 of the lowest price. However, only 40% of algorithmic sellers are within \$1 of Amazon’s price.

The fact that algorithmic sellers matching to Amazon tend to charge higher prices may be due to the required commission fees that Amazon charges. For example, if a 3P seller and Amazon share the same wholesale cost for a product, the 3P seller must charge a higher price to maintain the same profit margin. As described in § 2.1, Amazon’s commission fees are around 15% for most product categories. To see if we can observe algorithmic sellers that include these fees in their prices, we plot the *relative* difference between




Figure 18: Distribution of seller/product lifetimes for algorithmic and non-algorithmic sellers.




Figure 19: Number of products sold by algorithmic and non-algorithmic sellers.




Figure 20: Percentage of positive feedback for algorithmic and non-algorithmic sellers.

the algorithmic sellers’ prices and the target prices in Figure 13. As expected, we see very different behavior between algorithmic sellers matching Amazon’s price and the overall lowest price; we can observe a number of sellers who choose a new price that is 15–30% above Amazon’s price.

5.3 Price Matching Examples

We conclude this section by showing a few example products where we detected algorithmic pricing. First, Figure 14 shows an example where a 3P seller has a clear strategy to always match the lowest price across all other sellers. In the figure, we can see four other sellers that offer the lowest price over time, and the algorithmic seller (in red) always quickly matches their price.

Second, we observe several cases where the seller offering the lowest price is able to sell the product well above their reserve price. As shown in Figure 15, the algorithmic seller always matches the lowest price from the other two sellers. Although we the algorithmic seller is willing to sell the product for as low as \$12, the majority of the time they sell at prices up to 40% higher.

Third, we observe many cases where Amazon itself appears to be employing algorithmic pricing. Figure 16 shows a case where Amazon (in red) chooses their price to be a premium above the lowest price of all other sellers. In the figure, we observe that there are three other sellers that offer the lowest price at different points in time, but that Amazon is almost always slightly more expensive.

Fourth, we observe cases that Amazon adopts more complex pricing strategies than simply matching lowest prices. As shown in Fig 17, Amazon appears to have a ceiling at around \$9, above which they match the lowest price, but below which they sell the product at a small premium relative to the lowest price.

6. ANALYSIS

At this point, we have identified the sellers who are likely using algorithmic pricing. In this section, we compare and contrast the characteristics of algorithmic and non-algorithmic sellers. In particular, we are interested in answering the following questions: (1) How do the business practices of algorithmic sellers compare to non-algorithmic sellers? (2) What fraction of market dynamics are likely caused by algorithmic sellers? and (3) What is the impact of algorithmic sellers on the Buy Box?

6.1 Business Practices

To compare the general business practices between algorithmic sellers and non-algorithmic sellers, we examine the following four seller-level characteristics: lifespan of products, inventory size, feedback volume, and ranking in the seller page. Note that this list of characteristics is not comprehensive: as mentioned in § 4.3, we do not have access to several seller features such as return rate of products, shipping time, *etc.* Since Amazon itself plays a dual

role as a merchant and the host of the marketplace, we examine its role as a seller separately.

Product Lifespan. We begin by examining the *lifespan* of seller/product pairs in our dataset. The lifespan of a pair begins the first time we observe a seller offering that product, and ends the last time we observe that seller offering the product. Given our crawling methodology, the shortest possible lifespan is 25 minutes, while the longest are 3 and 1 months for Crawl1 and Crawl2, respectively.

Figure 18 shows the distribution of seller/product lifespans for both algorithmic and non-algorithmic sellers. We observe that algorithmic sellers are active in the marketplace for significantly *longer* periods of time than non-algorithmic sellers. For example, the median seller/product lifetime for an algorithmic seller is 30 days, while it is only 15 days for a non-algorithmic seller. As we show momentarily, our data suggests that algorithmic sellers have a high sales volume, so the long lifespans of their products further suggest that they have a large amount of inventory. Note that the vertical anomalies in Figure 18 around 1 month are artifacts caused by the different lengths of Crawl1 and Crawl2.

Inventory, Feedback, and Rank. Next, we compare the total number of products sold by algorithmic and non-algorithmic sellers. Since Crawl1 and Crawl2 focused on best selling products, the dataset may not contain all products sold by sellers. To obtain complete inventories, we conducted a separate crawl that exhaustively collected the entire inventory for 100 randomly selected algorithmic and non-algorithmic sellers, respectively. Note that the inventory for algorithmic sellers includes *all* products they sell, not just specific products where we detect algorithmic pricing.

Surprisingly, as shown in Figure 19, algorithmic sellers sell fewer unique products by a large margin. This suggests that algorithmic sellers tend to specialize in a relatively small number of products, perhaps focusing on items that they can obtain in bulk at low wholesale prices.

Next, we examine the feedback received by sellers from customers. On Amazon, customers may rate sellers on a 0–5 scale and also provide feedback about whether their experience with the seller was positive or negative. Amazon presents each seller’s average rating (0–5), percentage of feedback that is positive (0–100), and total amount of feedback on the *New Offers* pages. Note that Amazon does not display these stats for sellers with insufficient feedback (typically new sellers), and thus we ignore them in the following analysis (this only filter out 5% and 15% of algorithmic and non-algorithmic sellers in our dataset, respectively).

Figure 20 shows the cumulative distribution of positive feedback percentage for all sellers in our dataset. We observe that algorithmic sellers have slightly higher positive feedback than non-algorithmic sellers. However, almost all sellers have greater than 80% positive feedback; given this compressed value range, algorithmic sellers’ positive feedback advantage is more significant.




Figure 21: Amount of feedback received for algorithmic and non-algorithmic sellers.




Figure 22: Cumulative distribution of rank on the *New Offers* page for algo and non-algo sellers.




Figure 23: Cumulative distribution of Amazon's rank in the presence/absence of algorithmic sellers.




Figure 24: Number of price changes per seller/product pair.




Figure 25: Number of changes in the Buy Box for products with and without algorithmic sellers.




Figure 26: Probability of winning the Buy Box for algo and non-algo sellers at different ranks.

Figure 21 examines the amount of feedback received by algorithmic and non-algorithmic sellers. In this case we observe a stark contrast: algorithmic sellers acquire significantly greater amounts of feedback. There are two reasons why this could be happening: first, algorithmic sellers could have much higher sales volume than non-algorithmic sellers. Second, algorithmic sellers could be more aggressive about soliciting customer feedback.

Taken together, our results show that algorithmic sellers exhibit significantly different characteristics than non-algorithmic sellers: they sell fewer unique products, but in the products where they do compete, they participate in the marketplace for longer periods of time, and they acquire significantly larger amounts of positive feedback (suggesting they may have higher sales volumes).

The feedback and sales volume characteristics of algorithmic sellers put them at an advantage with respect to non-algorithmic sellers. As shown in Figure 22, when comparing all the products in their respective inventories, the rank of algorithmic sellers on the *New Offers* page tends to be significantly higher than that of non-algorithmic sellers. Note that neither CDF goes to 100% because there are cases where sellers do not appear in the top 20 list.

Amazon as a Seller. Given Amazon's dual role as merchant and host of the marketplace, it makes sense to examine their role as a seller separately. Table 3 shows the percentage of best-selling products that are sold by Amazon. "Overall" means that Amazon sold the product at some point during our crawls, while "Each Crawl" is the average percentage of products sold by Amazon during each snapshot. Although Amazon's inventory of products for-sale does appear to change over time, they still dominate the market, offering around 75% of all best-selling products over time.

	Algorithmic	Non-Algorithmic
Overall	72%	73%
Each Crawl	62%	63%

Table 3: Percentage of products with Amazon as one of the listed sellers.

Besides having an enormous inventory, Amazon is also known for their low prices. To demonstrate this, we plot Figure 23, which shows the rank of Amazon as a seller on the *New Offers* page for every product at each crawl. As we observed in Figure 7, the list is *roughly* ordered from least-to-highest price.

Overall, we observe that Amazon ranks highly for almost all products they sell. However, if we compare the products that have sellers doing algorithmic pricing to products without such sellers, we see a dramatic difference. In products with algorithmic sellers, we see that Amazon ranks in the top 5 sellers 88% of the time; in products without such sellers, we see this percentage increases to 96%. Thus, even Amazon tends to be ranked lower for products where algorithmic sellers are active.

6.2 Price and Buy Box Dynamics

Next, we move on to examining the dynamics of Amazon Marketplace. First, we look into price changes for each seller/product pair, followed by price and product changes in the Buy Box. Finally, we compare the frequency of algorithmic and non-algorithmic sellers winning the Buy Box.

Price Changes. Figure 24 shows the distribution of number of price changes per seller/product pair. We observe that for non-algorithmic sellers, the price *never* changes for 65% of the pairs. This suggests that either there is no competitive pressure to instigate changes, or the seller is unsophisticated.

However, if we examine the algorithmic sellers, we observe a very different distribution. Pairs with at least one price change are more common than not,¹⁰ and there is a long tail of products whose prices change 100 or even 1000 times. Unsurprisingly, algorithmic sellers are much more active in the marketplace.

¹⁰Readers may notice that the number of price changes for algorithmic sellers in Figure 24 goes below 20, which is the change threshold we set to detect those algorithmic sellers. However, this plot includes *all* the products sold by the algorithmic sellers, some of which have <20 changes during our crawl.

Buy Box. Next, we examine the impact of algorithmic sellers on the Buy Box. Figure 25 compares the number of seller (labeled as *Sid*) and price changes we observe in Buy Boxes for products that have algorithmic sellers, and products that do not. As expected, products with algorithmic sellers experience many more price and seller changes in the Buy Box: for example, 20% of products without algorithmic sellers have zero price changes, versus only 2% for products with algorithmic sellers. Unfortunately, this exposes customers to a great deal of volatility, which they may perceive to be confusing and undesirable [9].

Next, we examine whether algorithmic sellers are successful at winning the Buy Box. As shown in Figure 26, this is indeed the case: algorithmic sellers are more likely to win the Buy Box at all ranks except for the top one. Given the importance of winning the Buy Box, this result is quite interesting: as shown in Figures 12 and 13, algorithmic sellers tend to set their prices greater than or equal to the lowest price for a product. However, even though algorithmic sellers do not offer the lowest prices, they manage to win the Buy Box anyway due to their feedback and sales volume.

7. RELATED WORK

Theoretical Work on Price Competition. With easy access to the Internet and computation technologies, e-commerce sellers are able to adjust their prices automatically by setting algorithmic rules against other competitors in the market. These sellers are playing a *pricing game* in the marketplace. [11, 32, 14] model a pricing game played by the sellers and study the properties of its equilibria as a function of the dependencies among goods/services offered by the sellers. [19, 18, 10] extend the traditional price competition model proposed by Bertrand [14] to combinatorial settings. [12] models price competition in marketplaces where equilibria rarely exist.

Issues in Online Marketplace. Online marketplaces bring customers convenience, low prices, and a vast inventory of products. However, the anonymous nature of online marketplaces makes them vulnerable to manipulation and fraud conducted by unscrupulous parties. [26, 31, 33] study insincere sellers that generate opinion spam and artificial ratings to manipulate the reputation systems on online marketplace. [39] conduct an empirical analysis on the Seller Reputation Escalation (SRE) ecosystem that provides a shill-purchasing service for escalating business' reputations on the Taobao online marketplace.

Several empirical studies have also show that online marketplaces can cause privacy issues for consumers. Minkus *et al.* [30] reveal that attackers can correlate the highly sensitive user information from public profiles in eBay's feedback system with their social network profiles on Facebook. Similarly, [34] discovered personal information and detailed shopping habits leaking from online merchants to payment providers (*e.g.*, PayPal).

Auditing E-Commerce Algorithms. Automated algorithms are becoming increasingly ubiquitous on online marketplaces. However, the impact of these algorithms on users are often poorly understood, and not always positive. [20] studied Uber's surge pricing algorithm and revealed that an implementation bug was causing users to receive out-of-date pricing information. [25, 28, 29] uncovered instances of price discrimination and price steering on major e-commerce sites. Finally, Edelman *et al.* revealed the existence of systemic racial discrimination on AirBnB [24].

8. CONCLUDING DISCUSSION

E-commerce marketplaces have changed many aspects of how goods are bought and sold. Recently, these services have made algorithmic pricing (*i.e.*, using computer algorithms to automatically price goods) a realistic possibility for even small-scale sellers. However, the impact of algorithmic pricing on marketplaces and customers is not yet understood, especially in heterogeneous markets that include competing algorithmic and non-algorithmic sellers.

In this paper, we took the first steps towards detecting and quantifying sellers using algorithmic pricing on Amazon Marketplace. We collected large-scale data on products and sellers on Amazon Marketplace, and we make our code and data available to the research community.¹¹ We found that algorithmic sellers can be detecting using a target price time series, and we identify over 500 such sellers in our data set.

Our findings illustrate the power of algorithmic pricing in online marketplaces. Sellers we identified as using algorithmic pricing receive more feedback and win the Buy Box more frequently, likely suggesting higher sales volumes and thus more revenue than non-algorithmic sellers. Furthermore, we observe cases where algorithmic sellers change prices tens or even hundreds of times per day, which would be difficult for a human to maintain over time—especially one attempting to manage many products simultaneously—but is trivially automated. Clearly, the existence of cost-effective, user-friendly automation platforms like Sellery and Feedvisor is a win for sellers, especially smaller merchants who lack a dedicated programming staff.

However, there are also caveats introduced by algorithmic pricing. First, it is challenging for non-algorithmic sellers to compete with algorithmic sellers, which suggests an arms race that may terminate with all serious sellers adopting automation. The Buy Box algorithm exacerbates the disparity between algorithmic and non-algorithmic sellers, as it creates a largely winner-take-all marketplace where the Buy Box winner receives the vast majority of sales.

Second, increasing automation opens the door to intentional and unintentional market distortions. Although we do not observe any of these issues in our data, there are documented cases of algorithms pushing prices to unrealistic heights [37] and being used to implement price fixing [5]. We view our efforts to detect dynamic pricing as the first step towards long-term monitoring of algorithms in markets, with the ultimate goal of increasing transparency of these practices.

Finally, it is not clear what the impact of dynamic pricing is on customers. As shown in Figures 14–17, the presence of algorithmic sellers does not necessarily push item prices down to their reserves. Furthermore, as previously noted, algorithmic pricing can cause prices to fluctuate rapidly, which gives rise to the need for third-party price monitoring tools like CamelCamelCamel.¹² Arguably, this makes the shopping experience more complicated for customers, although more quantitative and qualitative work is necessary to truly understand how these factors impact customers.

Acknowledgements

We thank the anonymous reviewers for their helpful comments. This research was supported in part by NSF grants CNS-1054233, CNS-1319019, and CHS-1408345. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

¹¹ Available at <http://personalization.ccs.neu.edu>

¹² <http://camelcamelcamel.com>

9. REFERENCES

- [1] Amazon marketplace web service (amazon mws) documentation. Amazon.com, 2015. http://docs.developer.amazonservices.com/en_US/dev_guide/index.html.
- [2] Amazon marketplace web service (amazon mws) documentation – submitfeed. Amazon.com, 2015. http://docs.developer.amazonservices.com/en_US/feeds/Feeds_SubmitFeed.html.
- [3] Amazon sellers sold record-setting more than 2 billion items worldwide in 2014. Business Wire, January 2015. <http://www.businesswire.com/news/home/20150105005186/en/Amazon-Sellers-Sold-Record-Setting-2-Billion-Items>.
- [4] Fees and pricing. Amazon.com, 2015. <https://www.amazon.com/gp/help/customer/display.html?nodeId=1161240>.
- [5] Former e-commerce executive charged with price fixing in the antitrust division’s first online marketplace prosecution. U.S. Department of Justice, April 2015. <http://www.justice.gov/opa/pr/former-e-commerce-executive-charged-price-fixing-antitrust-divisions-first-online-marketplace>.
- [6] Fulfillment fees for orders on amazon.com. Amazon.com, 2015. <http://www.amazon.com/gp/help/customer/display.html?nodeId=201119410>.
- [7] Increase your chances of winning the buy box. Amazon.com, 2015. https://www.amazon.com/gp/help/customer/display.html?ref=hp_rel_topic?ie=UTF8&nodeId=201687830.
- [8] Match low price. Amazon.com, 2015. http://www.amazon.com/gp/help/customer/display.html?ref=hp_left_cn?ie=UTF8&nodeId=200832850.
- [9] ANGWIN, J., AND MATTIOLI, D. Coming soon: Toilet paper priced like airline tickets. The Wall Street Journal, 2012. <http://www.wsj.com/articles/SB10000872396390444914904577617333130724846>.
- [10] BABAIOFF, M., LUCIER, B., AND NISAN, N. Bertrand networks. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce* (2013), EC ’13.
- [11] BABAIOFF, M., NISAN, N., AND PAES LEME, R. Price competition in online combinatorial markets. In *Proceedings of the 23rd International Conference on World Wide Web* (2014), WWW ’14.
- [12] BABAIOFF, M., PAES LEME, R., AND SIVAN, B. Price competition, fluctuations and welfare guarantees. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation* (2015), EC ’15.
- [13] BENMAYER, T. The buy box cheat sheet. Feedvisor, May 2014. <http://feedvisor.com/the-buy-box-cheat-sheet/>.
- [14] BERTRAND, J. Book review of theorie mathematique de la richesse sociale and of recherches sur les principes mathematiques de la theorie des richesses. *J des Savants* 67, 2 (1883), 499–508.
- [15] BISHOP, T. Amazonfresh delays \$299/year prime fresh rollout, tweaks subscription model in california. Geek Wire, June 2015. <http://www.geekwire.com/2015/amazonfresh-delays-299year-prime-fresh-rollout-for-another-3-months/>.
- [16] BROHAN, M. Amazon builds up its european marketplace. Internet Retailer, April 2015. <https://www.internetretailer.com/2015/04/30/amazon-builds-its-european-marketplace>.
- [17] CASTILLO, M. Amazon hoax coupled with walmart’s price matching leads to ridiculously cheap ps4s. Adweek, November 2014. <http://www.adweek.com/news/technology/amazon-hoax-coupled-walmarts-price-matching-leads-ridiculously-cheap-ps4s-161527>.
- [18] CHAWLA, S., AND NIU, F. The price of anarchy in bertrand games. In *Proceedings of the 10th ACM Conference on Electronic Commerce* (2009), EC ’09.
- [19] CHAWLA, S., NIU, F., AND ROUGHGARDEN, T. Bertrand competition in networks.
- [20] CHEN, L., MISLOVE, A., AND WILSON, C. Peeking Beneath the Hood of Uber. In *Proceedings of the 15th ACM/USENIX Internet Measurement Conference (IMC’15)* (Tokyo, Japan, October 2015).
- [21] DENALE, R., LIU, X., AND WEIDENHAMER, D. Quarterly e-commerce retail sales, 2nd quarter 2015. U.S. Census Bureau News, August 2015. https://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf.
- [22] DURYEE, T. Amazon adds 30 million customers in the past year. Geek Wire, May 2014. <http://www.geekwire.com/2014/amazon-adds-30-million-customers-past-year/>.
- [23] DURYEE, T. Amazon may have up to 80 million high-spending prime members worldwide. GeekWire, September 2015. <http://www.geekwire.com/2015/amazon-may-have-up-to-80-million-high-spending-prime-members-worldwide/>.
- [24] EDELMAN, B. G., AND LUCA, M. Digital discrimination: The case of airbnb.com. Tech. rep., 2014.
- [25] HANNAK, A., SOELLER, G., LAZER, D., MISLOVE, A., AND WILSON, C. Measuring price discrimination and steering on e-commerce web sites.
- [26] LI, J., OTT, M., AND CARDIE, C. Identifying manipulated offerings on review portals. In *In Proceedings of the Conference on Empirical Methods in Natural Language Processing* (2015), EMNLP ’13.
- [27] LINDNER, M. Global e-commerce set to grow 25% in 2015. Internet Retailer, July 2015. <https://www.internetretailer.com/2015/07/29/global-e-commerce-set-grow-25-2015>.
- [28] MIKIANS, J., GYARMATI, L., ERRAMILI, V., AND LAOUTARIS, N. Detecting price and search discrimination on the internet. In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks* (2012), HotNets-XI.
- [29] MIKIANS, J., GYARMATI, L., ERRAMILI, V., AND LAOUTARIS, N. Crowd-assisted search for price discrimination in e-commerce: First results. In *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies* (2013), CoNEXT ’13.
- [30] MINKUS, T., AND ROSS, K. W. I know what you’re buying: Privacy breaches on ebay. In *Privacy Enhancing Technologies* (2014), PETS ’14.
- [31] MUKHERJEE, A., KUMAR, A., LIU, B., WANG, J., HSU, M., CASTELLANOS, M., AND GHOSH, R. Spotting opinion spammers using behavioral footprints. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2013), KDD ’13.
- [32] NASH, J. The Bargaining Problem. *Econometrica* 18, 2 (1950), 155–162.
- [33] OTT, M., CARDIE, C., AND HANCOCK, J. Estimating the prevalence of deception in online review communities. In *Proceedings of the 21st International Conference on World Wide Web* (2012), WWW ’12.
- [34] PREIBUSCH, S., PEETZ, T., GUNES, A., AND BERENDT, B. Purchase details leaked to paypal (short paper). In *Financial Cryptography* (2015).
- [35] RAY, J. D. Amazon is absolutely eviscerating other retailers online, new survey shows. Re/code, October 2015. <http://recode.net/2015/10/05/amazon-is-absolutely-eviscerating-other-retailers-online-new-survey-shows/>.
- [36] STATISTICS, L. B., AND BREIMAN, L. Random forests. In *Machine Learning* (2001).
- [37] SUTTER, J. D. Amazon seller lists book at \$23,698,655.93 – plus shipping. CNN, April 2011. <http://www.cnn.com/2011/TECH/web/04/25/amazon.price.algorithm/>.
- [38] TAFT, D. K. Amazon buy box: The internet’s \$80 billion sales button. eWeek, October 2014. <http://www.eweek.com/enterprise-apps/slideshows/amazon-buy-box-the-internets-80-billion-sales-button.html>.
- [39] XU, H., LIU, D., WANG, H., AND STAVROU, A. E-commerce reputation manipulation: The emergence of reputation-escalation-as-a-service. In *Proceedings of the 24th International Conference on World Wide Web* (2015), WWW ’15.
- [40] ZIOBRO, P. Target expands online price-match policy to include amazon, wal-mart. The Wall Street Journal, September 2015. <http://www.wsj.com/articles/target-expands-online-price-match-policy-to-include-amazon-wal-mart-1443636271>.