Search Related Trust Incentives in Online Marketplaces

Ted Yuan and Brian Johnson

eBay Inc., 2065 Hamilton Avenue, San Jose, CA 95125, USA
{ted.yuan,bjohnson}@ebay.com

Abstract. Appropriate incentives and trust are vital for a clean, well-lit search experience in online marketplace like eBay. In this paper, we discuss the quality of the online information provided by sellers in an online marketplace (trust) and the levers available to incent sellers to maintain a clean, well-lit marketplace. We divide inaccurate information into two subclasses, unintentional misrepresentation due to seller error or marketplace error, and intentional fraud and misrepresentation. We confine our discussions to search related incentives and trust. The types of problems discussed are title keyword spam, miscategorization, structured metadata spam, price manipulation, and shill bidding/buying. We review and propose ways to detect spam and fraud using different models based on online transactions and user behavioral data. The incentives discussed are seller suggestions, listing removal, account suspension, and search recall and ranking actions.

Keywords. Trust in e-commerce search, Spam detection, Fraud detection, Bipartite cores, Belief propagation

1. Introduction

Online markets democratize e-commerce, allowing anyone to easily buy and sell merchandise with unprecedented selection and value. This democratization and opening up of e-commerce has been a tremendous boon to buyers and sellers. In many successful marketplaces it is the role of the market maker to provide incentives for good behavior and remove bad actors in order to maintain trust between market participants. In an online e-commerce marketplace such as eBay, most market participants are good buyers and sellers who trade various goods efficiently with little interaction. However, there are a minority of sellers who either inadvertently or intentionally misrepresent their inventory. It is these sellers that we focus on in this paper. They take advantage of an online trading platform to conduct unfair or illegal trades that result in bad user experiences and leave a negative image of the online marketplace. It is very important for an online marketplace to establish a safe trading platform and promote trust amongst its users.

The consequences of trust related bad user experiences are different, depending on the intention of people who are responsible. We focus on two types of bad sellers,

- **Fraudsters**. Fraud tends to target less experienced buyers. Fraud can generate not only financial loss to users who trade with them, but also loss of fees and insurance payout from the e-market.
- **Spammers**. Spammers mainly game a search and browse system to gain unfair display exposure over other competitive sellers who sell similar (or better) quality items.

Fraud often incurs direct financial losses. Spam degrades overall e-market search and browse performance.

Spam and fraud activity generated directly by an individual seller are termed "one layer" deep in this paper. Sellers who attempt to conceal their identity in order to remain well hidden or organize their suspicious activity across a network of accounts are termed "multi-layer". While continuing to minimize serious fraud, it is also important for online auction sites like eBay to fight spam in order to keep search and browse results clean so that buyers get the best possible experience.

Spam and fraud are related. Sellers with a lot of spam tend to have higher risk of committing fraud. For both spam and fraud, sellers use similar techniques to disguise themselves such as fake listings, transactions, and feedback.

We would like to remove sellers engaged in fraud from the site. For sellers who commit relatively minor policy violations, warnings and incentives such as search demotion are more appropriate.

2. Unintentional Misrepresentation

Not all bad activity is intentional. Sellers are not always aware of platform policies and may inadvertently violate policies. These sellers are easier to work with as they are not intentionally engaged in an adversarial game with the platform operator. We will not discuss unintentional misrepresentation in great detail. Many of the classification challenges are similar to intentional misrepresentation, but the problem is significantly less challenging as the sellers have good intentions and are not adversarial. We will provide only one example.

2.1. Miscategorizaton

Miscategorization is an area where many well intentioned sellers run afoul of best practices. Sellers may list inventory in the wrong categories due to lack of knowledge, lack of time, or even inaccurate platform recommendations. Many eBay sites (e.g. US, DE) have more than 10,000 categories. Sellers can become overwhelmed and sellers with large inventory stocks would sometimes like eBay to take care of categorization for them via our recommendation API. In addition, large sellers often list inventory on multiple e-commerce platforms and differences between platforms complicate their listing practices. They sometimes rely on the platform provider for category recommendations. In the case below (Fig. 1), car floor mats have likely been inadvertently listed in the Consumer Electronics category in the past, which now,

based on previous sales, incorrectly appears as a recommendation in the e-commerce platform selling flow. There is likely no advantage to a seller listing car floor mats in the Consumer Electronics category, although their items may sell even though they are disadvantaged in search. In cases like these, seller education and platform improvements are likely sufficient to solve this unintentional problem.

Fig. 1. An imperfect listing tool may suggest incorrect listing category

3. Intentional Misrepresentation

The examples below all violate eBay policy. Sellers engaging in these activities risk having their accounts limited or suspended, or their item removed from the site or disadvantaged in search. These are clear examples of where incentives play a very real role. If unscrupulous sellers think they can gain an advantage with little risk, they are often willing to engage in these activities. Without detection and enforcement, the e-commerce platform would be overrun with unscrupulous activities, and honest participants might feel compelled to game the system as well. Many of eBay's complaints are from sellers who feel their competitors are engaged in unfair practices that would degrade the site if allowed to continue, hurting all merchants.

An online marketplace can discover intentional misrepresentation through

1. User reports,
2. Rule based filters,
3. Algorithmic detection.

User reports are a great mechanism for detecting new types of fraud. For known types of fraud they are a great training signal for algorithmic detection, but without algorithmic enhancement they scale poorly. It is cost prohibitive to manually review every listing on large e-commerce platforms as they scale globally to many countries and hundreds of millions of items. Rule based filters are often somewhat successful, but they are also difficult to scale and maintain in an adversarial environment, where opponents quickly become aware of new rules and are actively engaged in finding new ways to skirt the rules. Algorithmic detection, particularly machine learning

techniques, are the state of the art. Large training sets and large sets of features scale nicely and can often provide confidence scores that are great features for search ranking and/or manual customer service review for egregious offenders.

The following sections provide examples of some of the types of intentional fraud found on e-commerce sites and some of the obvious steps available to combat this fraud. Descriptions are by necessity quite high-level as counter-fraud operations are by definition adversarial and some of the most motivated readers would be adversarial.

3.1. Miscategorizaton

Sellers might decide to list their iphone cases in the ipod product category instead of the appropriate accessory category. In the auction format these items may initially be difficult to distinguish from iphone auctions starting at $0.99. As transactions accumulate there are many features of the fraudulent inventory that will, alone or in combination, distinguish them from non-fraudulent inventory. In this example sale prices between accessories and products are often significantly different. These signals can be used to identify sellers of items that are outliers in the categories they are listed in.

3.2. Title Spam

Sellers may decide to engage in keyword spamming in order to increase the visibility of their inventory. Although this may increase visibility, in the long run it results in fewer bids, bins, views, sales when impression is normalized. In the case below (Fig. 2), the seller is not selling the Rosetta Stone language product. This item will probably not sell well for users searching for the Rosetta Stone products. It will have a low sale rate for the query Rosetta Stone.

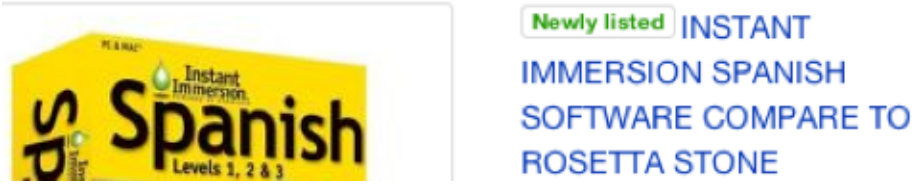

Fig. 2. A keyword spam example

3.3. Structured Metadata Spam

In this example (Fig. 3), the seller has listed brand Reebok in the title and brand Nike in the item structured metadata. Items are unlikely to have conflicting values for the Brand attribute. Sellers with high percentages of items with conflicting information can be identified.

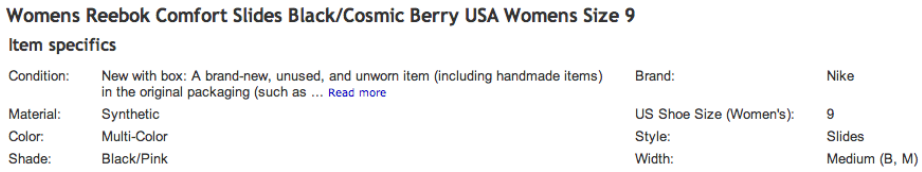

Fig. 3. A listing with conflict information, different brands in title and metadata

4. Spam – malicious search engine optimization

The intention of spam is to unfairly boost search and browse exposure of listings and feedback reputation scores of the spammer's accounts. Spam creators often have lower quality inventory and inferior or unclear inventory descriptions. A spammer's goal is to manipulate search and brows to gain unfair listing search display advantage over other competitive sellers.

Online marketplaces need to have ways to classify inventory by predicting the quality of individual listings. Listing quality may include measures like display quality, intrinsic value, click-through-rate prediction and sale probability.

Spam is often found in popular areas of an e-commerce site. On eBay, for example, we see more spam in fashion, accessories and popular electronics categories. In gaming the search engine of an e-commerce site, a spammer typically explores and gains some knowledge of the search and browse system, and uses that knowledge to manipulate their inventory. Platform providers can sometimes identify their own weaknesses by monitoring behavior offenders. Long ago eBay had a search ranking bug related to relevant token normalization. We quickly saw listing with egregious repetition of relevant term, e.g. "ipod 2 new in box, ipod ipod ipod ipod ipod ipod ipod". Problems like this are obvious and easy to fix, but monitoring know offenders is often highly effective in uncovering changes in more subtle seller practices.

Like other popular online search engines, relevance is an important factor to rank listings for a given user query. Relevance can be evaluated using similarity between a user query and listing textual attributes such as title, description, and structured data. These attributes are claimed by sellers and mandated by the online market's listing applications. In today's online marketplaces, it is not hard to find relevance gaming related spams in the forms of keyword stuffing, synonym stuffing, and intentional miscategorization that pollutes listing categories.

Ambiguous and false claims are another form of spam. A listing might be described as an "antique and vintage violin" bearing a European maker name but be sold in new condition and shipped from an Asian location.

In addition to gaming search relevance and false claims, below is an incomplete list of common spams with search engine manipulation,

1. Duplicates, or multiple listings of the same or similar items,
2. Inflated transactions and feedbacks from shill buyers,
3. Low item price but high shipping cost,
4. Unreasonably high initial prices but willing to accept best offer,
5. Items with graffiti images,

6. Raffles, off-site transaction advertisements, etc.

Spam degrades the quality of search results, reduces buyers' ability to find higher quality listings.

5. Spam and fraud detection with belief propagation

In dealing with specific types of spam or fraud, we need to develop unique classifiers that can separate spam as an outlier from the rest of listing population. The quality of such classifiers can be measured by comparing prediction precision and recall with actual labeled training data. For example, if a classifier is targeted to positively identify a certain type of spam, false positive rate and ROC curve [1] can measure the quality of the classifier. Because the impact and severity are different between fraud and spam, fraud detectors typically emphasize accuracy, whereas spam detectors tend to emphasize both coverage and accuracy.

For some online fraud and criminal cases with multi-layer user interactions, it may be better to use graph representations where users are represented as nodes, and interactions as edges connecting the nodes [2]. Groups of users can be identified and labeled (fraudster, accomplice, and normal user). To simulate the user graph over time, a time series iteration maybe run. At each iteration step, a belief propagation matrix is used to determine current state of a node based on the states of its surrounding nodes. Researchers have found that fraud sellers and accomplice buyers form recognizable "bipartite cores" in such user graphs. Members of fraud groups have many transactions with members of the accomplice group, but few transactions with members of their own group. There are limitations in the approach, some graph-based algorithms assume a stable user base such that the graph structure (nodes and edges) does not change much over time. The belief propagation matrix is assumed static as well.

Spam listings are more visible and frequent than fraud. Like fraud, spam can be created directly by a single seller, e.g. keyword stuffing and duplicates. These kinds of spam are direct, or only a "single layer" deep. "Two-layer" spam requires sellers and shill buyer participation, i.e., in gaming sale history and reputation multiple accounts are required.

Even though spam is more noticeable than fraud, spam is still a minority among the total listing inventory. We find that most spam can be traced back to small groups of sellers. This fact suggests aggregating listing signals at the seller level,

- For single layer spams, we may look into seller account history and target seller accounts with repeated spam policy violations,
- For two-layer spams, we could use algorithms similar to the graph-based belief propagation model. There may be short paths in finding bad spam sellers if we have a reliable bad buyer database. Bad buyer rings are designed to help spam sellers. Based on belief propagation, we should be able to find strong signals in the joint probabilities of interactions and transactions between spammers and bad buyers. Our preliminary investigation reveals that simple heuristic models, which are based on classification of mutual user transactions and feedbacks, can detect and reduce shills significantly. The models assume

belief propagation and look into recent suspicious, erroneous, faulty or incomplete transactions, as well as feedback messages between pairs of seller and buyer accounts.

While spam can be found and identified at listing item level, it may be more efficient to solve the spam issues at user level using a combination of seller incentive, listing demotion and account suspension. The promise is that cumulative spam and fraud concentration of a seller serves as good indicator of something underneath is going wrong, even some of the seller's listings are good, and some transactions and feedbacks of the seller are not exactly spam or fraud.

Generally, a spam predictor of a seller takes inputs from aggregation of listing level classifiers that are used to detect various types of spam and fraud in the seller's account, combined with other user account attributes to generate a quality or spam probability score for all active sellers. A normal seller would have a high quality score and low spam probability score. A spammer would have a higher spam score.

6. Shill activities

From Wikipedia, a shill is a person who publicly helps another person or organization without disclosing that he or she has a close relationship with that person or organization. Shill activities [3], which result in fake online transactions and feedback, are illegal, but well documented in auction marketplaces like eBay. eBay has explicit policies to prohibit shill activities on the site [4]. In many cases sellers and buyers involved in shill activities attempt to hide their actions so as to appear unrelated. It is often difficult to detect individual online shill activities in real time partially due to the ease of new buyer account creation. Shill not only incurs money loss from normal buyers but also leaves impression that a marketplace is unfair and unsafe.

On eBay, there are two primary forms of shill activities, one is shill bidding which inflates prices of online auctions; the other is shill buying which involves the buy-it-now fixed price listing format. New sellers may attempt to engage in shill buying in order to quickly accumulate a "successful" sale transaction history and feedback reputation. They hope to make their new seller account superficially look like a well-established account in the marketplace.

We have found that some bad sellers rely on networks of bad buyer accounts or account takeover in order to camouflage their activity. This is not surprising and is similar to using link farms to spam Google in the early days of search. Graph analysis can uncover these networks. Account activity can be used to identify account takeovers. The relationships between fraudulent sellers and buyers may not be immediately clear prior to trading activity. However, they can be discovered over time via historical transactions and user feedback data sources. One simple indicator of user relationships comes from account linkage data. Detection of more sophisticated networks often relies on post-hoc behavioral analysis. Sellers who sell to themselves do incur a transactional price penalty in that they must pay marketplace fees. It is possible to cancel a transaction to avoid fee payment, although high occurrence rates of cancelled transactions are a clear red flag.

Examples of user activities with shill buying include:

1. Implicit conspiracy between buyers and sellers, i.e., frequent purchases from a small group of buyers from the same seller,
2. Repeated positive feedback with many low dollar amount purchases,
3. Concealed bidding history detail to the public, i.e., excessive usage of private listings or other information hiding techniques (Fig. 4),
4. High rates of cancelled transaction (to avoid fees),
5. Organized and distributed services with buyer accounts providing paid services to clients who are eBay sellers.

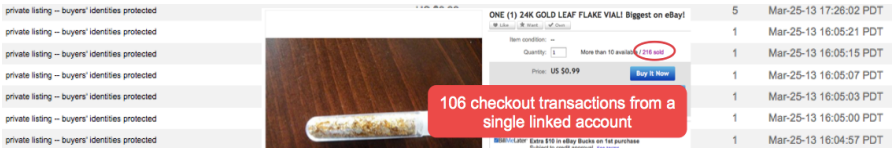

Fig. 4. Obscure private listing and transaction history

In dealing with shill activities that lead to fraud, we think the focus should be on both prevention and minimization of potential losses. In this regard, we will need

1. Techniques to detect fraud that rely on bad buyer rings,
2. Strategy to prevent and minimize future loss,
3. Removal of identified individual listing or offending user accounts.

To detect shill activities it is important to maintain an up-to-date and reliable bad buyer network. Once we have such bad buyer database, we can find sellers who interact with the bad buyers, and generate models and heuristics to evaluate seller accounts.

We do not believe bad buyer cliques (accomplices or fraud rings) surrounding sellers are static over a period of time. Some of the shill accounts may be left dormant for a long time and behave just like novice regular buyers. The fraud rings are in fact transient and dynamic with old shill accounts exposed and new accounts planted as future replacements. This may limit the effectiveness of detection techniques or algorithms that are not dynamic.

7. Effectiveness between Listing Policy Enforcement and Seller Search Incentive and Demotion

To fight against fraud and spam, an online marketplace typically implements many rules and policies [5] for online listings. Fraud and spam that can be associated with certain policy violations can be caught by automated policy enforcement applications, or reported by other users. The identified cases are constantly examined and verified by the online marketplace. If confirmed, the listings can be removed and the user accounts can be suspended.

The process for manually taking down offending items is hard to scale and maintain once an e-market grows beyond certain size, especially when spams are on the rise.

It is challenging to predict and react to many spam and fraud activities in real time. For example, individual shill buying accounts are often newly registered, or existing bad buyer accounts that are left dormant with minor defect but otherwise good account record. The fact that it is simple and requires no cost to create a new buyer account exacerbates the problem. In many cases shill activity is difficult to detect due to lack of significant evidence in user account history. Shill transactions are typically examined and identified after listings are sold.

As discussed above, we should aggregate listing spams and frauds and historical activities to seller level. We find the spam and fraud metrics of a user are highly consistent over time and correlate to current user behaviors, i.e., past sale performance of a seller can be a strong indicator of future transaction performance of the same account. We can compute seller quality based on the seller metrics.

With the help of seller quality, spam, and fraud probability models, we can reduce the overall display exposure of listings from suspicious sellers in search and browse applications. For instance, if a seller committed N (> threshold) spams in the recent week, we can reduce exposure of all his/her listings in the current search results and give more exposure to equal quality listings from other competitive sellers. We would continuously reduce the risk until the seller's spams count falls below the threshold.

Aggregated seller spam and fraud metrics can provide prior information to item-level spam and fraud detection classifiers. Seller statistics can also be used directly to demote listings from bad sellers according to probability score of spam and fraud prediction model. We believe the disincentive or demotion of listings from bad sellers by an online marketplace "helps" sellers improve their future listing quality and eventually reduce overall spam and fraud. Compared to harsh punishment like account suspension, demotion of search and browse exposure of spam on the e-commerce site forces seller to pay attention and gives them opportunities to improve.

8. Summary

In this article, we reviewed some known types of search related trust issues on eBay. We discussed the complexity in creating classifiers to detect particular spams and frauds like shill bidding and buying. We also compared two ways of improving trust and reputation to fight against spam and frauds, namely removing listings that are policy violations, and search and browse incentives and demotion based on user level spam prediction models. We believe both policy enforcement and seller incentives are needed to effectively reduce the overall spams and frauds in online auctions.

Here are some challenges and consideration of user level spam detection and demotion,

1. Prediction models to evaluate quality of listings and users on a vast and diverse inventory and user base,
2. Balance of accuracy and recall impact of spam predictors,

3. Data freshness and timely reaction, to compensate the fact that aggregation of bad user activities requires historical activity.

We believe that by developing online spam and fraud detectors and removing bad users, we can significantly improve the online marketplace user experience, while at the same time promoting healthy organic business growth.

References

1. Cost and benefit of classification models, receiver operating characteristic and ROC curve, http://en.wikipedia.org/wiki/Receiver_operating_characteristic
2. Duen Horng Chau, Shashank Pandit, and Christos Faloutsos: Detecting Fraudulent Personalities in Networks of Online Auctioneers. In: Proceedings of PKDD 2006, pages 103–114, Berlin, Germany, September 18-22, 2006.
3. Man fined over fake eBay auctions, http://www.bbc.co.uk/newsbeat/10508913
4. eBay shill bidding policy, http://pages.ebay.com/help/policies/seller-shill-bidding.html
5. eBay online policy, http://pages.ebay.com/help/policies/overview.html