



Data-Driven Limits: Optimizing Risk Management in E-commerce

Vinay Kumar Yaragani

vkyaragani@gmail.com

ABSTRACT

In the dynamic realm of e-commerce, managing risk is paramount to sustaining trust and operational integrity. This paper, presents a novel approach to balancing risk mitigation and user experience by deriving and implementing data-driven limits for buyers and sellers. We explore the spectrum of risk management actions, from reactive measures taken post-fraud to proactive mechanisms designed to prevent fraudulent activities. Our focus is on intermediary controls that operate in the gray area of uncertain behavior, where users are permitted to engage freely but within carefully calculated exposure limits. By analyzing extensive datasets, we develop predictive models to set dynamic limits on the number of items that users can buy or sell, thus curbing potential fraud without imposing undue friction on legitimate users. This strategy targets fraudsters who typically attempt rapid, high-volume transactions to exploit the platform before detection. Our findings indicate that implementing these data-driven limits effectively minimizes fraud exposure while fostering a growth-conducive environment for genuine users. This research highlights the critical role of data analytics in creating balanced risk management frameworks that protect e-commerce platforms and support sustainable user growth.

Key words: E-commerce Risk Management, Data-Driven Limits, Fraud Prevention, Predictive Analytics, User Behavior Analysis, Machine Learning

INTRODUCTION

The exponential growth of e-commerce has transformed the retail landscape, offering unparalleled convenience and accessibility to consumers worldwide. However, this rapid expansion has also brought forth a myriad of risks, including fraud, cybersecurity threats, and operational disruptions. Effective risk management is crucial for maintaining the integrity and trustworthiness of e-commerce platforms. Traditional risk management approaches, which often rely on reactive measures to address confirmed instances of fraud, are increasingly insufficient in the face of sophisticated fraudulent schemes. As such, there is a pressing need for more proactive and nuanced strategies to mitigate these risks.

Proactive risk management strategies aim to prevent fraudulent activities before they occur, employing mechanisms such as user verification and transaction monitoring to identify suspicious behavior. While these measures are essential, they can also introduce friction for legitimate users, potentially hindering their experience and satisfaction. This creates a delicate balancing act for e-commerce platforms: how to minimize fraud without imposing unnecessary barriers for genuine buyers and sellers. To address this challenge, it is imperative to explore intermediate risk management levers that can control exposure without overly restricting user activities.

The concept of data-driven limits in e-commerce is analogous to the credit limits set on credit cards, where both serve as mechanisms to manage risk and prevent potential misuse. Just as credit card companies impose credit limits to control the amount a cardholder can spend, thereby minimizing the risk of default and fraud, e-commerce platforms can implement dynamic limits on the number of items users can buy or sell. These limits are tailored based on historical data and predictive analytics, ensuring that while fraudsters are prevented from causing significant damage, legitimate users can continue their activities with minimal friction. This parallel underscores the importance of balancing risk management with user experience, allowing for growth and trust in both financial and digital marketplaces.

A. Objective

This paper introduces a data-driven approach to establishing dynamic limits on user activities, such as the number of items that buyers and sellers can transact on the platform. By analyzing historical data and employing predictive analytics, we can identify patterns that indicate potential fraud risks. These insights enable us to set personalized, adaptive limits that mitigate the impact of fraudulent activities while allowing legitimate users to operate with minimal disruptions. This intermediary strategy serves as a bridge between reactive and proactive measures, offering a balanced solution to the complex problem of e-commerce fraud.

Our research demonstrates that implementing data-driven limits significantly enhances the effectiveness of risk management in e-commerce. By curbing the ability of fraudsters to rapidly exploit the platform, we can reduce financial losses and protect the platform's integrity. Simultaneously, legitimate users benefit from a seamless experience, as they are not subjected to unnecessary verification steps or restrictions. This approach not only strengthens the overall security of e-commerce transactions but also promotes a healthy, controlled growth environment for all users.

LITERATURE REVIEW

The field of risk management in e-commerce has been extensively studied, with a significant focus on both reactive and proactive measures. Traditional reactive approaches, such as manual reviews and post-fraud investigations, have been foundational in identifying and mitigating risks after fraudulent activities have been detected. These methods, however, often suffer from delays and high operational costs, as highlighted by research conducted by Ngai, Hu, Wong, Chen, and Sun (2011). Proactive strategies, including real-time transaction monitoring and machine learning-based fraud detection systems, have emerged as vital tools in preventing fraud before it occurs. For instance, Bolton and Hand (2002) demonstrated the efficacy of statistical techniques in detecting outliers and anomalies in transactional data, paving the way for more sophisticated fraud detection mechanisms.

Recent advancements in data analytics and machine learning have further revolutionized risk management in e-commerce. The application of predictive modeling and big data analytics allows for more accurate identification of potential fraudsters and the prediction of fraudulent activities. There has been notable work emphasizing the role of machine learning algorithms in classifying transactions and predicting fraud risk. Moreover, the integration of these techniques into e-commerce platforms has shown promising results in reducing false positives and enhancing the overall accuracy of fraud detection systems, as discussed in the research by Bahnsen, Stojanovic, Aouada, and Ottersten (2016).

Despite these advancements, the challenge of balancing fraud prevention with user experience remains a critical issue. Proactive measures, while effective in detecting fraud, can inadvertently introduce friction for legitimate users, leading to negative user experiences and potential loss of business. Studies by Lai, Li, and Hsieh (2012) have highlighted the trade-offs between security measures and customer satisfaction, emphasizing the need for strategies that minimize user inconvenience. This necessitates the exploration of intermediary risk management levers that can provide a balanced approach, reducing fraud exposure while maintaining a smooth user experience.

The concept of setting dynamic limits on user activities, inspired by credit limit frameworks in the financial industry, offers a promising solution to this dilemma. Research in the financial sector, such as the works of Thomas, Crook, and Edelman (2017), has shown that credit limits are effective in managing credit risk and preventing overextension of credit. By drawing parallels to e-commerce, we can leverage similar data-driven methodologies to establish transaction limits for buyers and sellers. This approach not only controls exposure to potential fraud but also supports the growth and engagement of legitimate users, as researchers explored adaptive credit limits based on user behavior and transaction history.

In conclusion, the literature underscores the importance of a balanced risk management strategy in e-commerce. While reactive and proactive measures form the backbone of fraud prevention, the integration of intermediary controls, such as data-driven limits on user activities, offers a nuanced solution to the complex challenge of e-commerce fraud. By harnessing the power of predictive analytics and machine learning, we can create a robust risk management framework that safeguards the platform from fraud while fostering a positive user experience.

METHODOLOGY

A. Problem Statement

Unconstrained behavior on e-commerce platforms significantly increases vulnerability to breakout fraud, where fraudsters rapidly exploit the system to maximize their gains before detection. Without adequate controls, such as transaction limits, these platforms are exposed to high-volume fraudulent activities that can result in substantial financial losses and damage to reputation. Traditional reactive measures often prove insufficient, as they address fraud only after it has occurred, while stringent proactive measures can impede legitimate users, adversely affecting their experience. Therefore, there is a critical need for a balanced approach that leverages

data-driven limits to control exposure, thereby minimizing potential damage from fraudsters while ensuring smooth operations for genuine users.




Figure 1: Examples of breakout fraud

Breakout fraud in e-commerce typically manifests in three scenarios: new accounts created with malicious intent, established accounts with previously good history turning fraudulent, and account takeovers. In the first case, fraudsters create new accounts to quickly buy or sell a large number of items, aiming to maximize their gains before detection. This high-velocity activity often goes unnoticed initially due to the absence of a transaction history. In the second scenario, fraudsters exploit accounts with a good track record, suddenly ramping up activity to deceive buyers who trust their established reputation. Lastly, account takeovers involve malicious actors gaining control of legitimate accounts and engaging in a flurry of fraudulent transactions. All these behaviors are characterized by an explosion in activity and high-volume transactions, which are red flags for breakout fraud.

Implementing data-driven limits is essential to prevent breakout fraud. These limits control the number of transactions a user can perform within a certain period, thereby reducing the potential damage fraudsters can inflict. For new accounts, setting lower initial limits can help monitor their behavior without significantly impacting genuine users. For established accounts, adaptive limits based on historical activity can detect unusual spikes indicative of fraud. In cases of account takeovers, sudden changes in transaction patterns can trigger alerts and enforce limits to mitigate risks. By employing such mechanisms, e-commerce platforms can effectively control exposure to fraud, ensuring both the security of the platform and a positive experience for legitimate users.

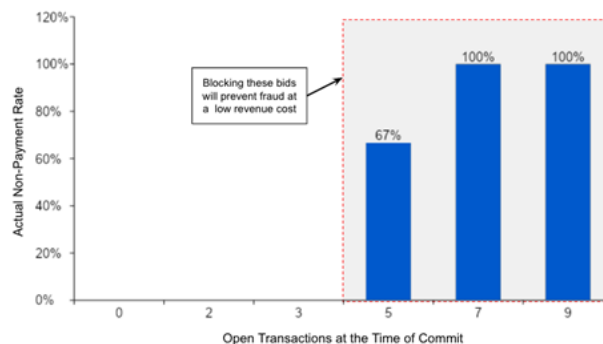


Figure 2: Example: propensity of unpaid buyer behavior with open bids

In our case study, we focus on the application of transaction limits for buyers participating in auctions, aiming to curb excessive bidding behavior and assess its impact on payment behavior. These limits are derived using a multifaceted approach that incorporates demographic information, user experience metrics, and past transaction history. By analyzing these factors, we can tailor limits that strike a balance between preventing over-engagement in auctions—potentially disruptive to the auction process—and maintaining a satisfactory user experience.

The implementation of these limits serves several purposes. Firstly, it aims to mitigate the risk of bidding manipulation or speculative behavior, which can distort auction outcomes and negatively impact seller trust. Secondly, by aligning limits with user demographics and past behavior, we aim to reduce instances of non-payment or delayed payments, thereby enhancing transaction reliability and seller satisfaction. This case study seeks to evaluate how such targeted limits influence buyer behavior and payment reliability, providing insights into effective risk management strategies in auction-based e-commerce environments.

In conducting an exploratory analysis focused on identifying variables correlating with non-payments from buyers, the study delves into various factors that may influence payment reliability. This analysis typically involves examining a range of demographic attributes, transactional history, and behavioral patterns of buyers.

By scrutinizing these variables, we uncover patterns or indicators that precede instances of non-payment. This initial exploration serves as a foundational step in understanding the underlying factors contributing to payment defaults in e-commerce transactions. It enables the identification of key predictors that may guide the development of more refined predictive models or risk management strategies aimed at reducing non-payment risks and enhancing transactional integrity in online marketplaces.




Figure 3: Correlation of Non payment rate to payment method and domestic/international demographics




Figure 4: Correlation of Non payment rate to Buyer Experience

While this novel approach of comparing future purchasing behavior between buyers who faced defects and those who did not offers valuable insights, it has inherent limitations. One significant issue is the inability to account for the various types of defects, which can range from minor issues, such as packaging problems, to major defects, such as faulty or incorrect products. Each type of defect can have a different impact on customer satisfaction and future purchasing behavior, necessitating a more granular analysis to capture these nuances accurately.

Leveraging comprehensive data on buyer history, demographics, user experience, and the number of open transactions, we developed various machine learning models to predict the propensity of non-payment. This predictive modeling approach encompassed several algorithms, including novel business rules tailored to specific risk factors, logistic regression for its robustness in binary classification, and ensemble methods like random forest to capture complex interactions between features. Additionally, decision tree algorithms were employed to provide interpretable insights into the decision-making process. Each model was rigorously tested and validated to identify the most effective predictors and achieve the highest accuracy in forecasting non-payment rates. This multifaceted approach aimed to develop a reliable risk assessment tool that can preemptively identify high-risk buyers and mitigate potential losses for e-commerce platforms.




Figure 5: ML Process to predict propensity of non payment

After achieving a satisfactory accuracy in predicting the propensity for non-payment, the next phase involved simulating the effects of implementing transaction blocks at various propensity thresholds. In this simulation, the open transaction count at each block level served as the transaction limit assigned to each user. The goal was to evaluate the performance of these blocks by balancing the trade-off between lost future revenue and the reduction in non-payment incidents. To ensure a fair comparison, we quantified the impact of non-payments in monetary terms, translating the number of non-payments prevented into dollar values by considering the churn and associated costs driven by non-payment events.

The simulation aimed to optimize the breakeven points where the reduction in gross merchandise volume (GMV) due to imposed limits was justified by the savings from avoiding non-payments. By analyzing different block thresholds, we could identify the optimal levels where the cost of lost potential revenue was minimized and outweighed by the benefits of reduced non-payment incidents. This approach provided a data-driven framework to implement transaction limits effectively, ensuring a balanced strategy that safeguards revenue while enhancing payment reliability on the e-commerce platform.

The simulation also facilitated the development of dynamic limits that adapt based on different risk buckets and user features, enabling a more tailored and responsive risk management approach. By categorizing users into various risk levels and analyzing their behavior over time, we established a system where users begin with an initial limit that can adjust according to their performance. As users demonstrate good behavior, such as timely payments and low transaction disputes, their limits increase, encouraging positive engagement and growth. Conversely, if users exhibit risky behavior or default on payments, their limits decrease, reducing their potential impact on the platform. This dynamic limit model ensures that users are incentivized to maintain good standing while providing a robust mechanism to mitigate risks progressively, ultimately safeguarding the platform from fraudulent activities and ensuring a balanced user experience.

Domestic Buyers w/ Payment Method						
		Lifetime Non-Payments				
		0	1-3	4-6	7-9	10+
Lifetime Paid Transactions	0	12	3	2	2	0
	1-3	16	10	4	2	1
	4-6	20	10	8	3	1
	7-9	20	12	8	3	2
	10-15	23	16	13	4	4
	16-19	27	18	14	8	6

Figure 6: Illustration of Limits for one segment

Upon validation, we found that less than 1% of buyers actually reached the imposed limits and would experience the blocks. To ensure a smooth user experience for legitimate buyers, we provided these users with the option to make immediate payments to lift the blocks. This approach allowed well-intentioned buyers to continue their transactions without significant disruption, maintaining a positive experience on the platform. At the same time, the system effectively controlled and curtailed the activities of buyers with potentially harmful intentions, preventing them from causing extensive damage. This dual strategy ensured that the majority of users encountered minimal friction, while still robustly protecting the platform from high-risk behaviors.

B. Results

An A/B testing experiment was conducted to evaluate the effectiveness of enforcing transaction blocks at predefined limits for the test group, while the control group experienced no such restrictions. The results demonstrated that the test group significantly reduced the non-payment rate, showcasing the efficacy of the imposed limits. Importantly, this reduction in non-payments came with minimal impact on future revenue, indicating that legitimate transactions were largely unaffected. This outcome highlights that implementing dynamic limits is a viable and efficient mechanism for preventing breakout fraud, safeguarding the platform, and maintaining overall revenue stability.

A/B test impact in US market	
Metric	Impact
Non Payments Dropped	10.4%
Subsequent Cancellations Dropped	7.5%
Paid Revenue Breakage	0.21%

Figure 7: Results from A/B test in one market

C. Future Scope

The successful application of data-driven limits in e-commerce to prevent breakout fraud presents an opportunity to extend this methodology to various other industries and domains. Financial services, for example, can benefit from similar strategies to manage credit risk by setting dynamic credit limits based on customer behavior and transaction history. Insurance companies could use predictive analytics to adjust policy terms and limits dynamically, reducing the risk of fraudulent claims. In telecommunications, service providers could implement usage limits that adapt based on customer payment patterns, preventing potential fraud and ensuring service continuity. The healthcare sector could leverage this approach to monitor and control fraudulent claims and optimize resource allocation based on patient history and behavior.

Further strengthening this methodology can be achieved through the advanced capabilities of Machine Learning (ML) and Artificial Intelligence (AI). Techniques such as deep learning can provide more nuanced insights into user behavior and risk patterns, enabling more accurate prediction models. Reinforcement learning could optimize limit-setting policies by continuously learning from new data and adjusting in real-time. Additionally, the integration of natural language processing (NLP) can enhance the analysis of unstructured data, such as customer reviews and support interactions, providing a more comprehensive risk assessment framework. These advanced ML and AI techniques will not only improve the precision of risk management strategies but also ensure their adaptability to evolving fraud tactics.

The framework developed for dynamic limit-setting in e-commerce can be adapted to address various other problems with appropriate modifications. In cybersecurity, similar models could be employed to dynamically adjust access controls based on user behavior and threat intelligence. In logistics and supply chain management, predictive analytics could optimize inventory limits and order quantities, minimizing risks associated with stockouts or overstocking. By customizing the input features and tuning the algorithms to the specific requirements of different industries, this robust framework can enhance risk management, improve operational efficiency, and drive strategic decision-making across diverse sectors. The flexibility and scalability of this approach underscore its potential as a versatile tool for mitigating risks and optimizing performance in a wide array of applications.

CONCLUSION

This paper demonstrates the efficacy of implementing data-driven transaction limits to optimize risk management in e-commerce platforms. By leveraging comprehensive buyer data and advanced machine learning models, we developed a dynamic system that adjusts limits based on user behavior, demographics, and transaction history. The A/B testing results underscored the effectiveness of these measures in significantly reducing non-payment rates while minimally impacting future revenue. This approach not only curtails breakout fraud but also ensures a seamless experience for legitimate users. The adaptable nature of this framework presents a promising opportunity for application across various industries, highlighting its potential as a versatile tool for enhancing risk management and operational efficiency.

REFERENCES

- [1]. E.W.T. Ngai, Yong Hu, Y.H. Wong, Yijun Chen, Xin Sun, The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature, *Decision Support Systems*, Volume 50, Issue 3, 2011, Pages 559-569, ISSN 0167-9236, <https://doi.org/10.1016/j.dss.2010.08.006>
- [2]. Bolton, Richard & Hand, David. (2002). Statistical Fraud Detection: A Review. *Statistical Science*. 17. 10.1214/ss/1042727940.
- [3]. Alejandro Correa Bahnsen, Djamila Aouada, Aleksandar Stojanovic, Björn Ottersten, Feature engineering strategies for credit card fraud detection, *Expert Systems with Applications*, Volume 51, 2016, Pages 134-142, ISSN 0957-4174, <https://doi.org/10.1016/j.eswa.2015.12.030>.
- [4]. Lai, F., Li, D., & Hsieh, C.-T. (2012). Fighting identity theft: The coping perspective. *Decision Support Systems*, 52(2), 353-363.
- [5]. Thomas, Crook, and Edelman (2017), *Credit Scoring and Its Applications (Mathematics in Industry)* 2nd Revised edition, 2017