American University Kyiv

A Capstone Project

THE IMPACT OF SEARCH QUALITY ON E-COMMERCE SALES
ВПЛИВ ЯКОСТІ ПОШУКУ НА ПРОДАЖІ В ЕЛЕКТРОННІЙ КОМЕРЦІЇ

by Mark Myropolskyi

Presented in Partial Fulfillment of the Requirements

for the Master's Degree

APPROVED BY:

Academic advisor: Professor Viktor Putrenko

2026

Table of Contents

Abstract

This capstone project examines the impact of internal search performance on customer behavior and e-commerce outcomes using real search data from EVA.UA. Internal search is treated not only as a navigation tool, but also as a mechanism that shapes product discovery, customer experience, and the allocation of economically valuable traffic. The study aims to understand whether query popularity, query characteristics, platform differences, or traffic concentration across demand segments drive search-related business performance.

The empirical analysis is based on a large real-world dataset that includes more than 42 million search events and more than 6 million unique queries, as extracted from the original raw data. After data cleaning, the study applies query-level and event-level analysis to evaluate search demand distribution, conversion patterns, add-to-cart behavior, platform differences, weighted and unweighted performance metrics, frequency deciles, and revenue concentration proxies. Numeric SKU-like queries and low-quality noise were filtered out to reflect genuine customer intent better.

The results show that search traffic is highly concentrated in a small subset of high-frequency queries. The top frequency decile accounts for a disproportionate share of total search traffic and revenue proxy. However, conversion differences across frequency segments are relatively small, which suggests that revenue concentration is driven more by traffic allocation than by substantially better query performance. The analysis also shows that query length has a non-linear relationship with conversion: medium-length queries tend to perform best. In contrast, very long queries yield unstable results due to a low sample size. Platform analysis reveals meaningful behavioral differences: web demonstrates the highest purchase conversion, iOS shows the highest add-to-cart rates, and Android underperforms on both conversion and downstream funnel efficiency.

The study concludes that internal search affects business performance not only through relevance, but also through visibility allocation and platform-specific friction. The findings suggest that the major business opportunity is not limited to further optimizing already dominant queries, but also to improving the exposure and handling of underutilized demand, especially on mobile platforms. These results contribute to understanding how internal search influences e-commerce performance and offer practical implications for search optimization, product discovery strategy, and customer experience management.

Executive Summary

This study examines how search behavior influences e-commerce performance, using large-scale real data from EVA.UA. The analysis focuses on the relationship between query frequency, user behavior, platform differences, and revenue distribution.

The results reveal a highly uneven distribution of value across search queries. The top decile of queries accounts for approximately 45% of total revenue, while the broader top segment (high-frequency queries) contributes up to 85%. This confirms a strong concentration of economic value within a small subset of queries.

However, the study's key finding challenges a common assumption in e-commerce. Despite large differences in traffic and revenue, conversion rates remain relatively stable across all query segments. This indicates that high-frequency queries do not generate more revenue because they perform better, but because they receive significantly more exposure.

In other words, revenue concentration in search is primarily driven by demand distribution and traffic allocation, rather than by conversion efficiency.

This insight highlights a structural limitation of traditional search systems, which tend to prioritize already popular queries. As a result, a substantial portion of demand — particularly in mid- and low-frequency queries — remains underutilized, despite demonstrating comparable conversion potential.

The analysis also identifies meaningful platform differences. The web platform consistently shows higher purchase conversion, while mobile platforms, particularly Android, underperform in both conversion and downstream funnel efficiency. This suggests that user experience and platform-specific friction play a significant role in shaping final business outcomes.

From a business perspective, the findings point to three key opportunities for improvement. First, mid-frequency queries represent the most immediate growth lever, combining meaningful traffic with strong conversion potential. Second, long-tail queries require more advanced search capabilities, such as semantic matching and intent recognition, to unlock their value. Third, improving mobile user experience, especially in the later stages of the funnel, can significantly increase overall conversion.

Overall, the study demonstrates that improving search performance is not only a matter of relevance optimization, but also of how visibility and traffic are distributed across queries. The results suggest that sustainable growth in e-commerce search requires a shift from frequency-based prioritization toward a more balanced, intent-driven approach.

CHAPTER 1. INTRODUCTION

1.1. Background and Problem Context

Search systems have become a central component of modern e-commerce platforms, serving as the primary interface between users and large-scale product catalogs. In many cases, search is responsible for a substantial share of total transactions, making it a critical driver of both user experience and business performance.

With the rapid growth in product assortment and user traffic, e-commerce platforms increasingly rely on algorithmic ranking systems to determine which products and queries receive visibility. These systems are often powered by machine learning models that optimize for relevance, click-through rate, or conversion.

However, a fundamental structural characteristic of user behavior complicates this optimization process: search demand is highly unevenly distributed. A small number of high-frequency queries account for a disproportionate share of total searches, while a long tail of less frequent queries represents a large but fragmented portion of demand.

This creates a systemic tension within search optimization:

- On one hand, focusing on high-frequency queries maximizes immediate impact due to scale
- On the other hand, long-tail queries may represent underutilized opportunities for value creation.

As a result, search systems do not simply retrieve information, but actively allocate attention and economic value across queries.

1.2. Problem Statement

Despite the central role of search in e-commerce, a key assumption remains largely untested in practice:

High-frequency queries are often assumed to be more valuable because they are more effective at converting users.

This assumption underpins many ranking and optimization strategies, where popular queries receive the majority of attention, engineering effort, and algorithmic prioritization.

However, this logic may be flawed.

If conversion efficiency is similar across query segments, prioritizing high-frequency queries may lead to systematic overexposure and underrepresentation of potentially valuable but less frequent queries.

This raises a critical question:

Does the current search system optimize for performance, or does it simply reinforce existing demand patterns?

1.3. Research Objectives

The primary objective of this study is to analyze how search query characteristics influence conversion performance and revenue distribution in an e-commerce environment.

The study aims to:

1. Quantify the distribution of search demand across frequency-based query segments.
2. Measure how revenue is distributed across these segments.
3. Evaluate whether high-frequency queries demonstrate superior conversion efficiency.
4. Analyze how platform differences (Web, iOS, Android) affect user behavior and conversion.
5. Identify structural inefficiencies in the allocation of traffic and economic value.

1.4. Research Questions

Based on these objectives, the study addresses the following research questions:

- RQ1: How is search demand distributed across query frequency segments?
- RQ2: How is revenue distributed across these segments?
- RQ3: Do high-frequency queries demonstrate higher conversion efficiency?
- RQ4: How do platform differences influence conversion outcomes?
- RQ5: To what extent does the search system reflect versus shape the distribution of economic value?

1.5. Research Contribution and Novelty

This study makes several contributions to both academic research and practical e-commerce strategy.

First, it provides empirical evidence that challenges the assumption that query popularity implies higher efficiency. By analyzing conversion metrics across frequency segments, the study demonstrates that performance differences are limited, while exposure differences are substantial.

Second, the study introduces a system-level perspective, showing that search engines act not only as retrieval mechanisms but as allocation systems that shape economic outcomes.
Third, it highlights structural bias toward high-frequency queries, which may lead to suboptimal utilization of available demand.

Finally, the study bridges the gap between theory and practice by linking concepts from:

- information retrieval,
- long-tail economics,
- and algorithmic bias

to real-world e-commerce performance data.

1.6. Practical Relevance

The findings of this research are directly applicable to modern e-commerce platforms. In particular, they provide guidance for:

- improving search ranking systems,
- designing AI-driven recommendation and retrieval models,
- Moreover, optimizing traffic allocation strategies.

The results suggest that focusing solely on high-frequency queries may limit growth potential and that redistributing visibility toward underutilized segments could increase total revenue without additional traffic acquisition.
This has important implications for companies investing in AI-powered search, personalization, and ranking systems.

1.7. Structure of the Study

The remainder of this paper is structured as follows:

- Chapter 1 – Introduction outlines the research context, the problem statement, the objectives, and the study's significance.

- Chapter 2 – reviews the relevant literature on search systems, e-commerce analytics, and demand distribution

- Chapter 3 – describes the methodology, dataset, and analytical framework

- Chapter 4 – presents the empirical results of the study

- Chapter 5 – Interpretation & Recommendations

- Chapter 6 – Conclusion

CHAPTER 2. LITERATURE REVIEW

2.1. Role of Search in E-commerce

Internal search has been widely recognized as a critical component of e-commerce systems, acting as a bridge between user intent and product discovery. Unlike traditional navigation, search allows users to directly express their needs, often reflecting a higher level of purchase intent (Jansen & Spink, 2006). Prior research suggests that users who engage with search are more likely to convert compared to those who rely solely on browsing, as search reduces cognitive load and shortens the decision-making process.

From a business perspective, search is not only a usability feature but also a revenue driver. Studies in information retrieval and e-commerce demonstrate that even small improvements in search relevance can lead to measurable increases in conversion rates and customer satisfaction (Croft et al., 2010). However, most of this literature focuses on relevance ranking and algorithmic performance, rather than on how search behavior translates into economic outcomes at scale.

In practice, search systems operate as allocation mechanisms: they determine which products and queries receive visibility. This means that search performance can influence not only individual user outcomes but also the distribution of traffic and revenue across the platform. Despite this, the relationship between search demand distribution and revenue concentration remains underexplored in academic literature.

2.2. Customer Intent and Query Characteristics

A central concept in search research is user intent. Queries are often categorized based on informational, navigational, or transactional intent (Broder, 2002). In e-commerce contexts, transactional queries—such as product names or specific brands—are typically associated with higher purchase likelihood.

Several studies have examined how query characteristics influence user behavior. Query length, for example, is often used as a proxy for specificity. Longer queries may indicate more precise intent, while shorter queries may reflect broader exploration (Baeza-Yates & Ribeiro-Neto, 2011). However, empirical findings are mixed. Some studies suggest that longer queries lead to higher conversion rates due to clearer intent, while others indicate that overly specific queries may reduce results quality if the system fails to match them effectively.

Another important dimension is query frequency. High-frequency queries (often referred to as “head” queries) typically represent common, well-understood user needs, while low-frequency queries (“long-tail”) represent more diverse, niche

demand. Traditional assumptions suggest that head queries are more valuable because they generate the majority of traffic. However, Anderson (2006) introduced the concept of the “long tail,” arguing that aggregated low-frequency demand can represent a significant share of total market value when properly served.
In e-commerce, this raises an important question: do head queries generate more revenue because they reflect higher intent, or simply because they receive more exposure? Existing literature does not provide a definitive answer, especially in large-scale real-world datasets.

2.3. Long Tail Theory and Demand Distribution

The long tail theory suggests that in digital markets, the cumulative demand for niche products can rival or exceed that of popular items (Anderson, 2006). This concept has been widely applied in areas such as media consumption, online retail, and recommendation systems.

In the context of search, the long tail manifests as a large number of unique, low-frequency queries. These queries are often more difficult for search systems to handle due to sparse data and limited historical signals. As a result, many systems prioritize head queries, where optimization yields immediate and measurable benefits.

However, recent research indicates that ignoring the long tail may lead to missed opportunities. Brynjolfsson et al. (2011) showed that online retailers benefit from broader product availability and better matching of niche demand. Similarly, studies in recommendation systems demonstrate that improving coverage of long-tail items can increase overall user satisfaction and engagement.

Despite these findings, there is limited empirical evidence on whether long-tail queries in e-commerce search are inherently less valuable or simply underrepresented due to system limitations. This distinction is critical for understanding whether search optimization should focus on head queries or on improving coverage of less frequent demand.

2.4. Search Quality and Conversion Performance

Search quality is traditionally evaluated using relevance metrics such as precision, recall, and ranking accuracy (Manning et al., 2008). While these metrics are important from a technical perspective, they do not always capture business impact.

In e-commerce, performance is often measured using behavioral metrics such as click-through rate (CTR), add-to-cart rate, and purchase conversion. These metrics reflect how effectively the search system supports the user journey from query to transaction. Research suggests that improvements in search relevance can increase these metrics, but the relationship is not always linear.

One challenge is that conversion is influenced by multiple factors beyond search, including pricing, product availability, trust, and user interface design. As a result, isolating the impact of search quality on conversion can be difficult. Some studies propose using funnel-based analysis to understand where users drop off, distinguishing between search-related issues and downstream friction (Huang et al., 2019).
This study adopts a similar approach by analyzing both add-to-cart and purchase behavior, as well as the transition between them. This allows for a more nuanced understanding of whether performance differences are driven by search quality or by later stages of the customer journey.

2.5. Platform Effects in E-commerce

Another important factor in e-commerce performance is platform context. User behavior differs significantly between web and mobile environments. Mobile users often exhibit higher engagement but lower conversion rates due to interface limitations, smaller screens, and more fragmented attention (Google, 2018).

Research indicates that mobile platforms tend to have higher add-to-cart rates but lower purchase completion rates, suggesting increased friction in the checkout process. In contrast, desktop users may demonstrate more deliberate behavior and higher final conversion rates.

These platform differences are important for interpreting search performance. A search system that performs well on one platform may not deliver the same results on another. Therefore, analyzing search outcomes without considering platform context may lead to incomplete or misleading conclusions.

2.6. Research Gap and Contribution

The academic literature provides important insights into both search systems and demand distribution in e-commerce environments. Studies in information retrieval emphasize ranking quality and relevance optimization (Manning et al., 2008; Baeza-Yates & Ribeiro-Neto, 2011), while research in e-commerce highlights the critical role of search in driving user engagement, conversions, and revenue (Jansen & Schuster, 2011). At the same time, the concept of the long tail suggests that aggregated demand from less frequent queries may represent a significant share of total market value (Anderson, 2006). More recent studies on algorithmic systems also identify the presence of exposure bias, where ranking mechanisms amplify already popular queries and products (Joachims et al., 2017).

Despite these contributions, there is a clear gap in empirical research that integrates these perspectives within a real-world business context. In particular, existing studies

do not fully address how search demand distribution, user behavior, and ranking mechanisms jointly influence economic outcomes in e-commerce systems.

Specifically, the following questions remain insufficiently explored:

- Whether high-frequency queries generate more revenue due to higher conversion efficiency or simply due to greater traffic allocation
- Whether long-tail queries are inherently less valuable or systematically underexposed by ranking mechanisms
- How platform differences (Web, iOS, Android) influence the relationship between search behavior and conversion outcomes
- How search quality interacts with downstream funnel stages, such as add-to-cart and purchase

Table 2. Summary of Key Literature on E-commerce Search and Conversion:

	Author	**Topic**	**Key Finding**	**Relevance to Study**
0	Jansen & Spink (2006)	Search behavior in e-commerce	Search users demonstrate higher purchase intent	Supports link between search and conversion
1	Broder (2002)	Query intent classification	Queries reflect informational, navigational, and transactional intent	Explains importance of query structure
2	Baeza-Yates & Ribeiro-Neto (2011)	Information retrieval fundamentals	Query structure affects retrieval performance	Provides theoretical base for query analysis
3	Croft et al. (2010)	Search engines and relevance	Small improvements in relevance can impact user outcomes	Supports importance of search quality
4	Manning et al. (2008)	Search evaluation metrics	Precision and recall define search effectiveness	Baseline for evaluating search performance

5	Anderson (2006)	Long tail theory	Long-tail demand can rival head demand in aggregate	Frames head vs tail concept
6	Brynjolfsson et al. (2011)	Long tail in online retail	Niche demand contributes significantly when properly matched	Supports hypothesis of underutilized tail
7	Agrawal et al. (2009)	Query frequency distribution	Search queries follow heavy-tailed distribution	Explains traffic concentration patterns
8	Huang et al. (2019)	User behavior in conversion funnel	Conversion is affected by multiple funnel stages	Justifies funnel-based analysis
9	Google (2018)	Mobile vs desktop behavior	Mobile users show higher engagement but lower conversion	Explains platform performance differences
10	Li et al. (2010)	User interaction with search	User interactions depend on query specificity	Supports query length analysis
11	Chapelle et al. (2012)	Click models and search evaluation	Clicks can be used as implicit relevance signals	Supports use of behavioral metrics
12	Ricci et al. (2015)	Recommender systems	Personalization improves discovery of niche items	Supports discovery beyond head queries
13	Shokouhi (2013)	Head vs tail query dynamics	Head queries dominate traffic but not necessarily intent	Directly relates to decile analysis findings

CHAPTER 3. METHODOLOGY

3.0. Methodology Approach

This study adopts a quantitative, data-driven approach to analyze the relationship between search behavior and e-commerce performance. The methodology is designed to bridge the gap between theoretical concepts in information retrieval and real-world business outcomes.
The research is based on the analysis of large-scale observational data, reflecting actual user interactions with an e-commerce search system. Rather than conducting controlled experiments, the study focuses on identifying patterns, relationships, and structural effects within naturally occurring data.

The methodological framework consists of three key stages:

1. Data Preparation and Cleaning
Raw search interaction data is processed to ensure analytical validity. This includes removing non-representative queries, filtering out noise (such as SKU-based searches), handling missing values, and aggregating data at the query level. The goal of this stage is to isolate genuine customer intent and create a reliable analytical dataset.

2. Segmentation and Structural Analysis
To analyze demand distribution, the study applies a frequency-based segmentation approach using deciles. This method allows for a more granular understanding of how search queries differ in terms of volume, behavior, and performance. Queries are grouped into low-, mid-, and high-frequency segments, enabling comparative analysis across the full demand spectrum.

3. Performance and Behavioral Analysis
The study evaluates search performance using behavioral metrics, including purchase rate, add-to-cart rate, and cart-to-purchase conversion. These metrics capture different stages of the customer journey and allow for identifying where value is created or lost.

Additionally, the analysis incorporates:

- Weighted and unweighted comparisons, to control for the dominance of high-frequency queries
- Query characteristic analysis (e.g., query length), as a proxy for user intent
- Platform-level analysis (Web, iOS, Android), to account for differences in user behavior and conversion dynamics
- Revenue proxy modeling, where purchase counts are used to approximate economic value in the absence of transaction-level revenue data

Overall, the methodology is designed to provide a structured, multidimensional understanding of search performance, focusing not only on relevance and conversion but also on how traffic distribution and system behavior influence business outcomes.

3.1. Data Source and Structure

This study is based on anonymized real-world data from EVA.UA, a large Ukrainian e-commerce platform. The dataset captures user interactions with the internal search system and subsequent behavioral events within the same session.
The original dataset consists of more than 42 million search events and over 6 million unique search queries. Due to the size of the data, it was processed in multiple files and analyzed using chunk-based processing in Google Colab.
Each observation in the dataset represents a search interaction and includes the following key variables:

- *query_text — the search query entered by the user*
- *session_id — unique identifier of the user session*
- *platform — platform type (Web, iOS, Android)*
- *had_purchase — binary indicator of whether a purchase occurred within the session*
- *total_add_to_cart_count — number of add-to-cart events within the session*

The dataset does not include product-level ranking positions, product prices, or exact revenue per transaction. Therefore, the analysis focuses on behavioral metrics (conversion and add-to-cart rates) and uses purchase counts as a proxy for economic value.

3.2. Data Cleaning and Preprocessing

To ensure the analysis reflects actual customer behavior, several data-cleaning steps were applied.
First, non-representative queries were removed. Specifically, numeric queries resembling SKU codes or internal product identifiers were filtered out. These queries are typically used by internal operators or power users and do not reflect typical customer intent.
Second, extremely low-frequency queries were excluded from the query-level analysis to reduce noise. Queries with very few observations tend to yield unstable conversion estimates and may distort aggregate results.

Third, missing and null query values were removed. Only sessions with valid search queries were retained for analysis.
Fourth, the dataset was aggregated at the query level to compute performance metrics such as total searches, total purchases, and add-to-cart counts per query. This allowed for consistent comparison across queries of different frequencies.
To validate robustness, both weighted and unweighted approaches were used. The weighted approach accounts for query frequency, while the unweighted approach treats all queries equally regardless of traffic volume.

Table 1. Example of Search Dataset Structure

query_text	***session_id***	***platform***	***had_purchase***	***total_add_to_cart_count***
shampoo for dry hair	12345	Web	1	2
face cream	67890	iOS	0	1
vitamin c serum	54321	Android	1	3

3.3. Analytical Framework

3.3.1 Analytical Dimensions:

1. Search demand distribution
2. Query performance (conversion and add-to-cart)
3. Platform differences
4. Revenue concentration (proxy-based)

3.3.2 Analytical Framework Explanation

The analytical framework used in this study is structured as a sequential and interconnected process that links search behavior to business outcomes.
The analysis begins with search demand distribution, which examines how user queries are distributed across frequency levels. This step identifies the structural characteristics of demand, including the presence of high-frequency (head) queries and low-frequency (long-tail) queries.

The second stage focuses on query performance, where behavioral metrics such as add-to-cart rate and purchase rate are analyzed. This allows the study to evaluate how effectively different queries translate user intent into measurable actions.
The third dimension introduces platform differences, comparing user behavior across Web, iOS, and Android environments.
This step helps isolate the impact of platform-specific factors, such as interface design and checkout friction, on conversion performance.
Finally, the analysis evaluates revenue concentration using a proxy based on purchase counts. This stage connects user behavior with economic outcomes and

allows the study to determine whether revenue distribution is driven by query performance or by traffic allocation.

Together, these four dimensions form an integrated framework that enables a comprehensive understanding of how search systems influence both user behavior and business results.

3.4. Key Metrics

To evaluate search performance and its impact on e-commerce outcomes, this study uses a set of behavioral metrics that reflect different stages of the customer journey, from initial search to final purchase.

- Search frequency *(searches = number of occurrences of a query) is used to measure demand and serves as the basis for query segmentation.*

- Purchase rate *(purchase rate = total purchases / total searches) indicates how effectively search leads to completed transactions and reflects overall conversion performance.*

- Add-to-cart rate *(add_to_cart_rate = sessions_with_add_to_cart / total_searches) captures user engagement by showing how often search results lead to product consideration, even if a purchase is not completed.*

- Cart-to-purchase conversion *(cart_to_purchase = purchases / sessions_with_add_to_cart) measures the efficiency of the final stage of the funnel and helps identify potential friction beyond search itself.*

Together, these metrics provide a structured view of user behavior, linking search demand, engagement, conversion, and overall business impact.

3.5. Analytical Applications and Visualization

To operationalize the analytical framework described above, this study applies a set of complementary analytical techniques. These methods are designed to examine search behavior from multiple perspectives, including demand distribution, user engagement, platform differences, and revenue allocation.

The following subsections present the key analytical components and their corresponding visualizations.

3.5.1 Query Segmentation and Decile Analysis

To analyze the distribution of search demand, queries were grouped into frequency-based segments.

Instead of a simple head–tail split, this study applies a decile-based segmentation, where queries are divided into 10 groups based on their frequency.

- Decile 9 — highest-frequency queries
- Decile 0 — lowest-frequency queries

This approach allows for a more granular understanding of how performance varies across the distribution.

Figure 1. Distribution of Search Queries by Frequency Decile

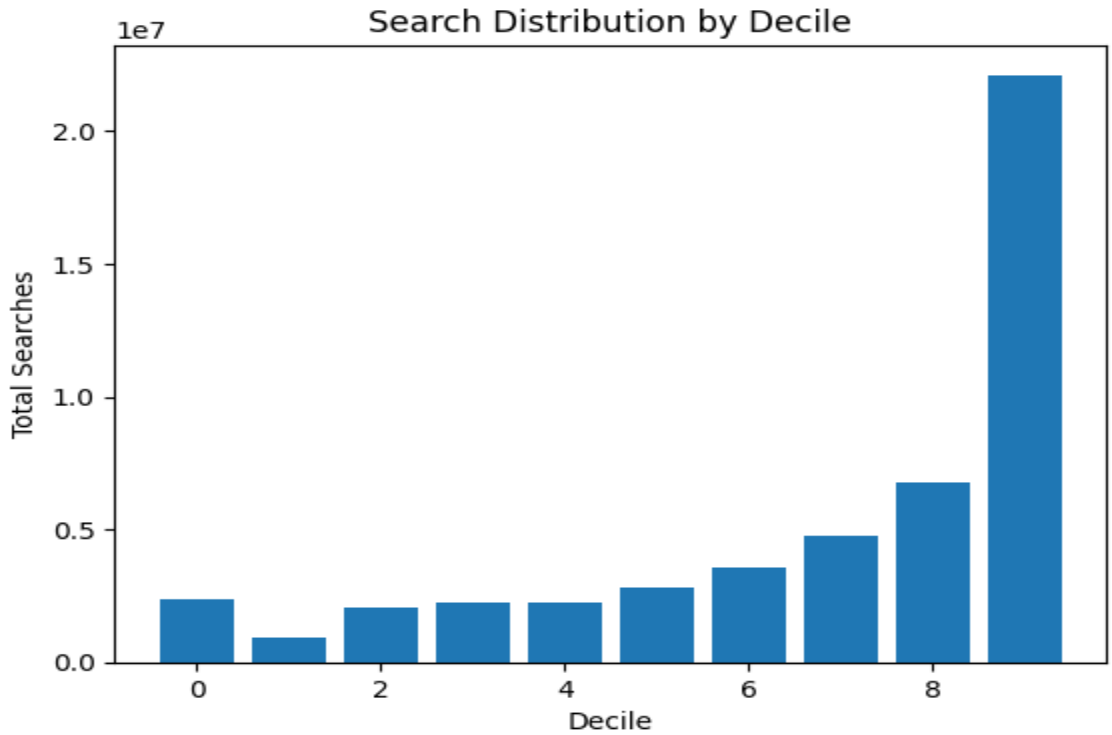

3.5.2. Weighted vs Unweighted Analysis

To avoid bias caused by highly frequent queries, the study compares:

- Unweighted metrics — each query has equal importance
- Weighted metrics — queries are weighted by their frequency

This distinction is critical because high-frequency queries dominate traffic and can distort aggregate results. By comparing both approaches, the study ensures that conclusions are not driven solely by head queries.

3.5.3. Query Length Analysis

Query length is used as a proxy for intent specificity. It is calculated as the number of words in the query.
The analysis examines how conversion and add-to-cart rates vary by query length.

Figure 2. Query Length vs Purchase Rate

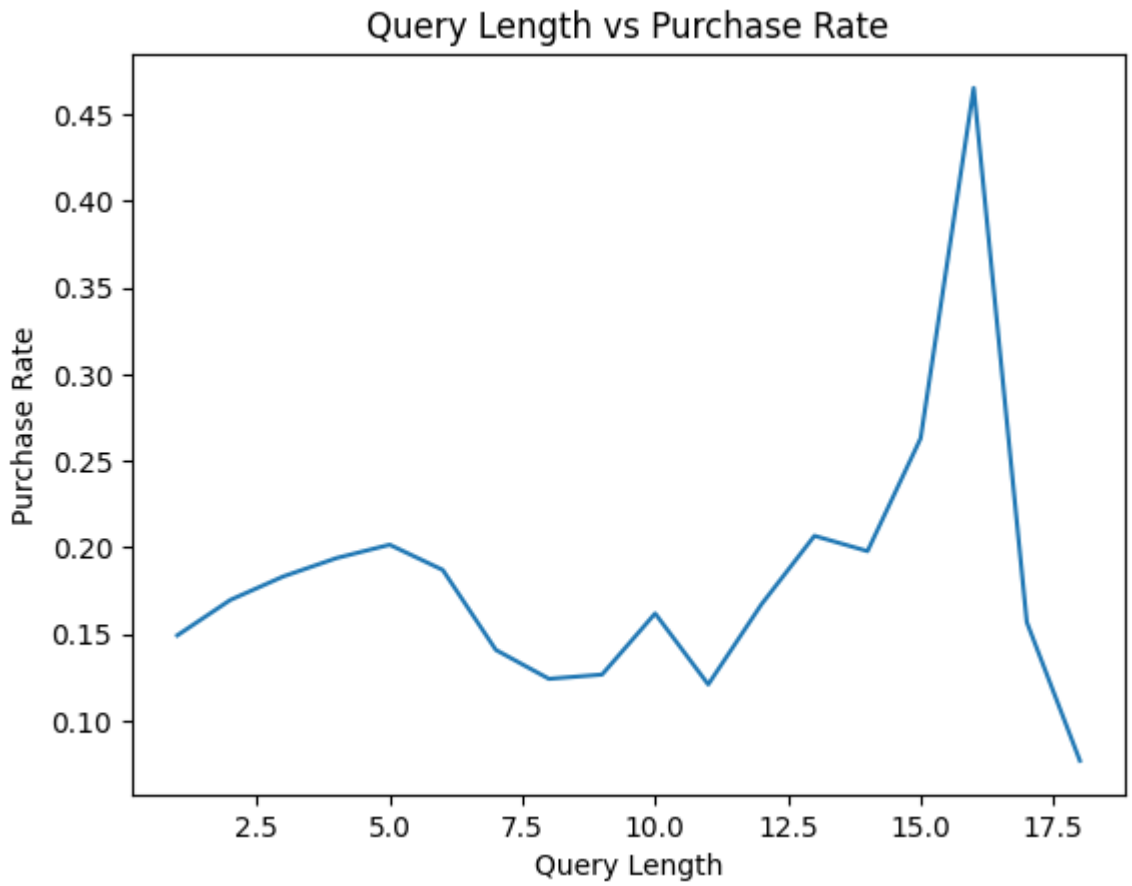

3.5.4. Platform-Level Analysis

To understand behavioral differences across platforms, the dataset is segmented into:

- Web
- iOS
- Android

For each platform, the following metrics are analyzed:

- search volume
- add-to-cart rate
- purchase rate
- cart-to-purchase conversion

Figure 3. Conversion Metrics by Platform

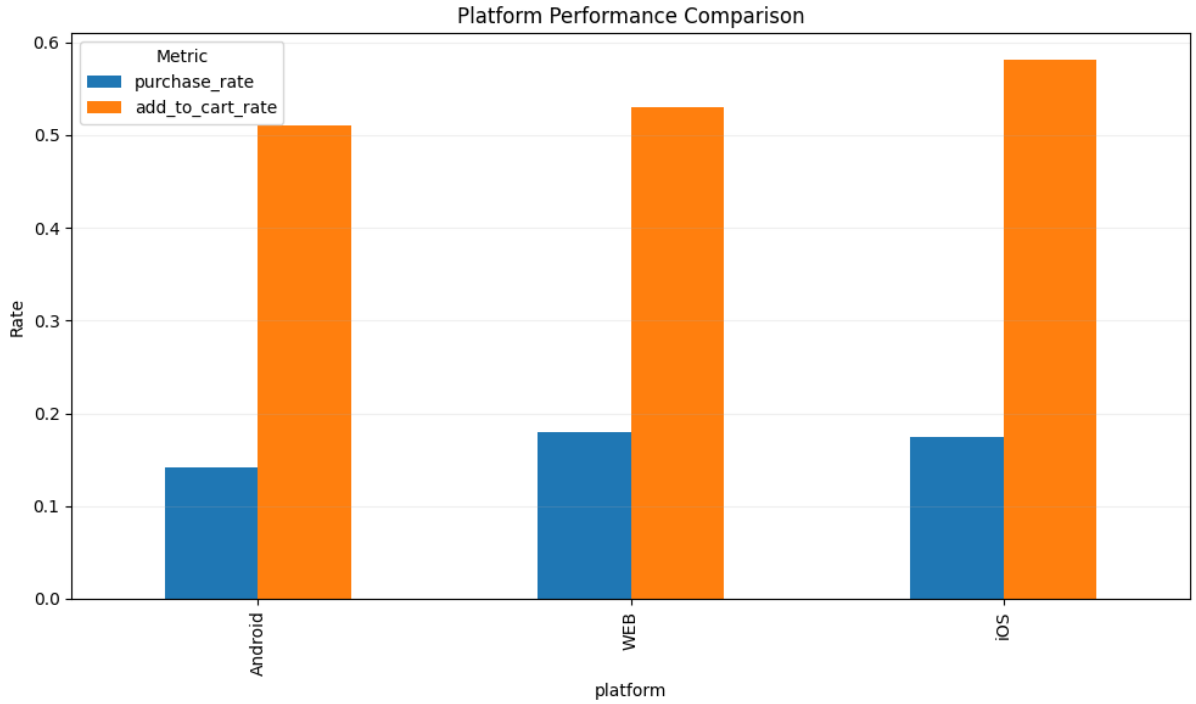

3.5.5. Revenue Concentration Analysis

To evaluate how economic value is distributed, a revenue proxy is calculated at the decile level.
Revenue share is defined as:
revenue_share = purchases_in_decile / total_purchases
This allows the study to assess whether revenue concentration is driven by better-performing queries or by traffic allocation.

Figure 4. Revenue Share by Query Frequency Decile

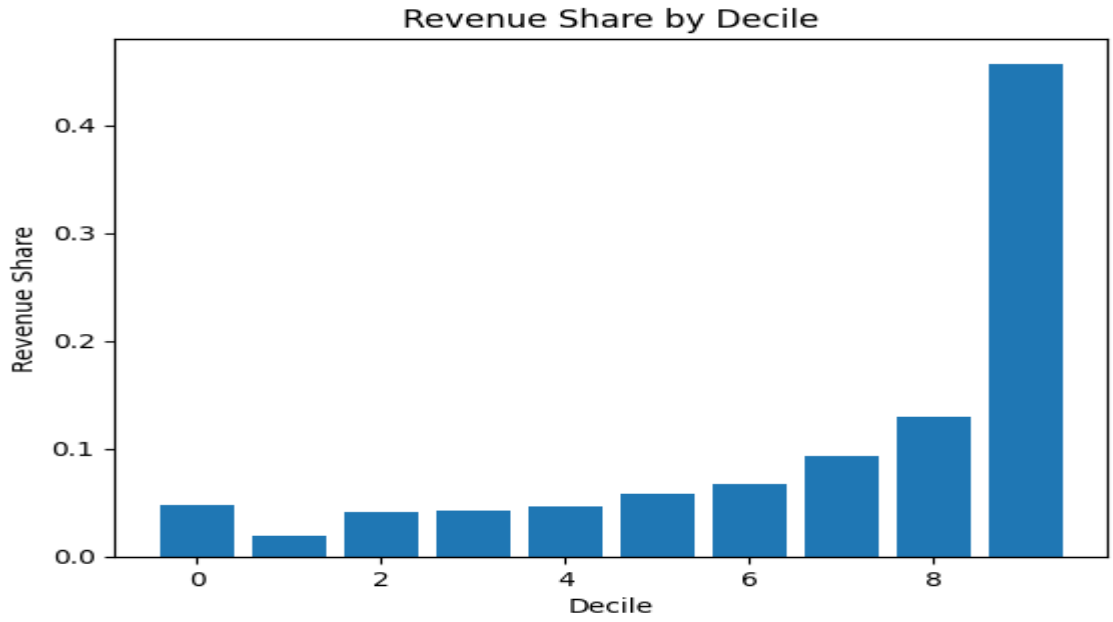

3.5.6. Limitations

Several limitations should be acknowledged.

First, the dataset does not include ranking position, product relevance scores, or detailed transaction values. This limits the ability to measure search quality at the algorithmic level directly.

Second, revenue is approximated using purchase counts, which may not fully capture differences in basket value.

Third, the analysis is based on a single platform (EVA.UA), which may limit generalizability.

Fourth, filtering and aggregation steps may influence absolute values, although relative patterns remain stable.

Despite these limitations, the dataset provides a large-scale, realistic representation of search behavior, enabling meaningful analysis of customer behavior and business outcomes.

CHAPTER 4. RESULTS

4.1. Interaction Between Query Frequency and Platform Performance

To examine whether platform differences are consistent across query segments, cart-to-purchase conversion was analyzed by frequency decile for each platform.

Figure 5. Cart-to-Purchase Conversion by Decile and Platform

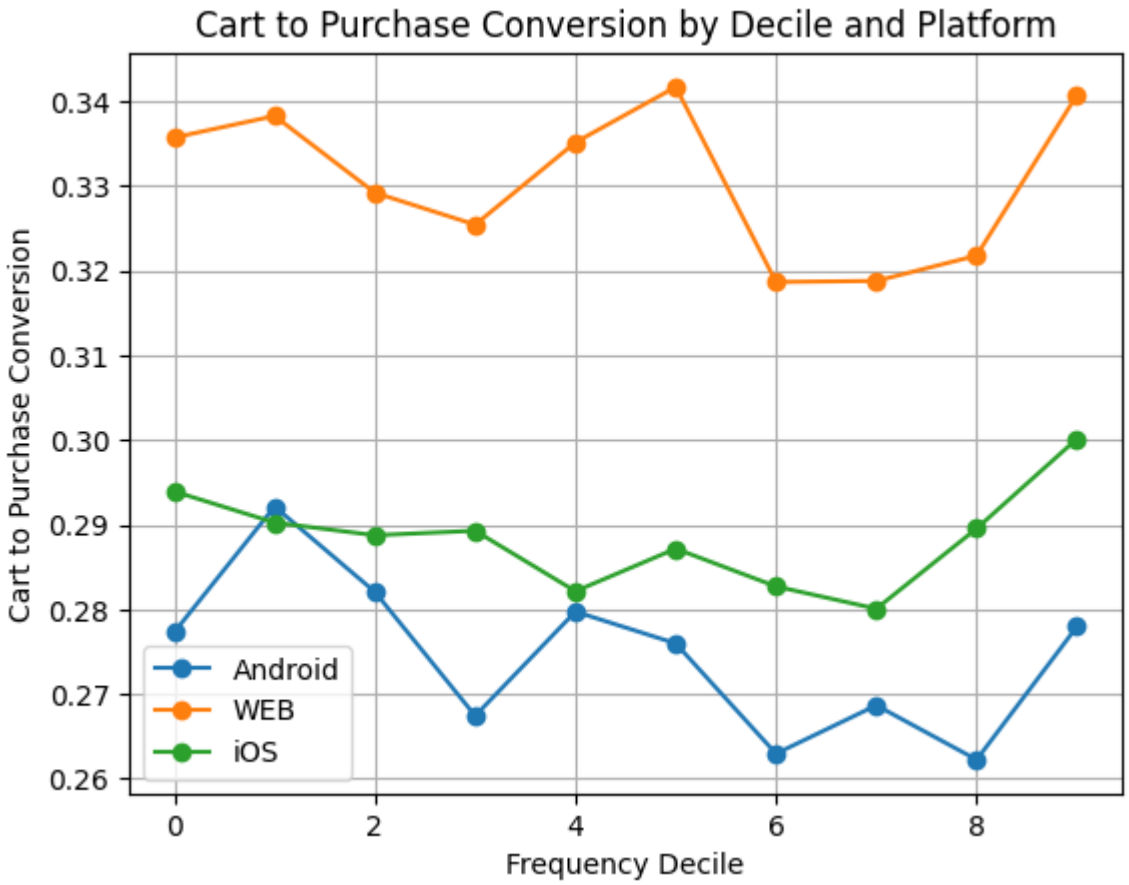

The results demonstrate that platform performance differences are both systematic and economically significant.

Across all deciles, the Web platform consistently outperforms both iOS and Android in cart-to-purchase conversion rates. The observed conversion rates for Web remain approximately 0.32–0.34, compared to 0.28–0.30 for iOS and 0.26–0.29 for Android.

This represents a gap of 5–7 percentage points between Web and mobile platforms.

Importantly, this difference is not uniform in its impact. The gap becomes particularly critical in higher-frequency deciles (Deciles 7–9), which account for the majority of total

search traffic. Given the scale of these segments, even small differences in conversion rates translate into substantial differences in absolute purchase volume.

From a system perspective, this suggests that:

- High-intent traffic does not compensate for platform-related inefficiencies
- Mobile platforms introduce friction at later stages of the funnel.
- The largest share of potential revenue is affected by suboptimal mobile conversion.

This indicates that platform inefficiencies are concentrated precisely where their economic impact is the highest.

4.2. Search Volume vs Conversion Efficiency

To test whether high-frequency queries are inherently more effective, the relationship between search volume and purchase rate was analyzed.

Figure 6. Relationship Between Search Volume and Purchase Rate

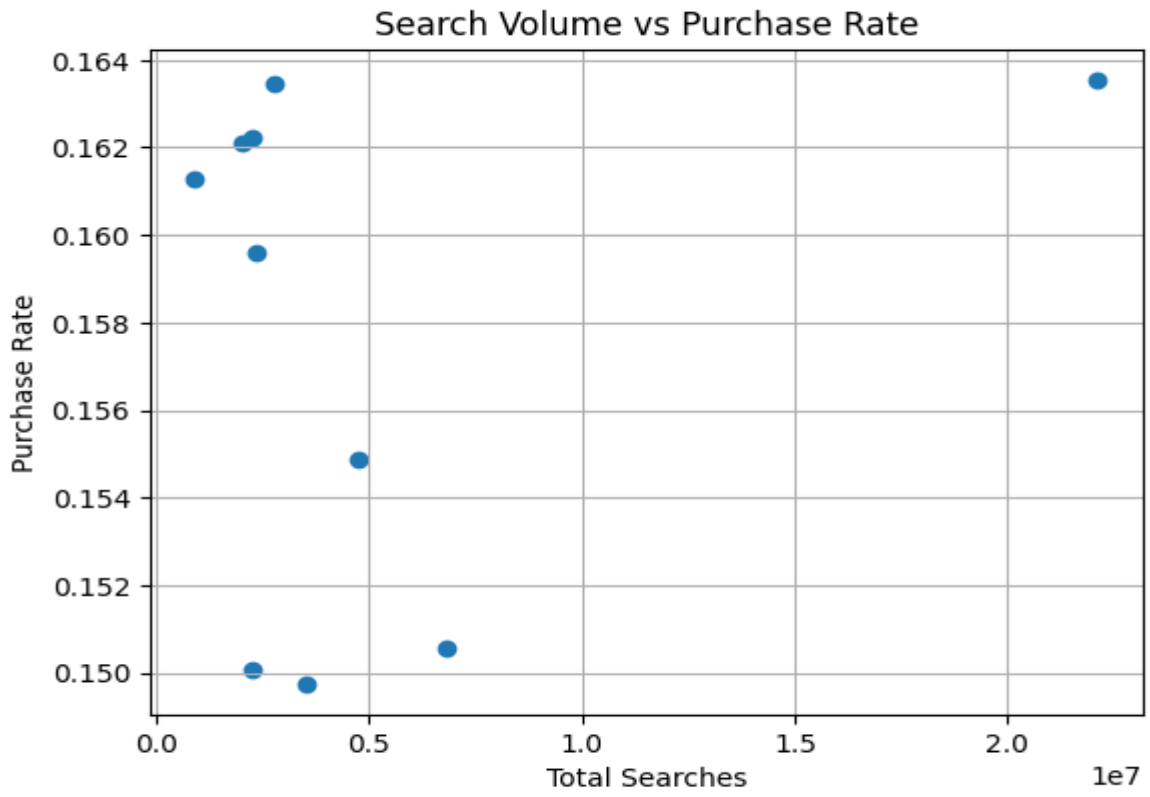

The results show no observable positive relationship between search volume and conversion efficiency.

Despite a more than 10-fold difference in total search volume between lower and higher deciles, purchase rates remain within a narrow range of approximately 0.149–0.164.

The scatter plot does not indicate any clear linear or monotonic trend. High-frequency queries do not consistently exhibit higher conversion rates than mid- or low-frequency queries.

This finding has critical implications:

- Query popularity is not a reliable indicator of conversion performance
- High-frequency queries do not demonstrate superior efficiency.
- Conversion behavior appears largely independent of query frequency.

The absence of correlation strongly suggests that revenue concentration is driven by traffic allocation rather than query effectiveness.

4.3. Stability of Funnel Conversion

To further analyze user behavior at the final stage of the funnel, cart-to-purchase conversion was examined across frequency deciles.

Figure 7. Cart-to-Purchase Conversion by Decile

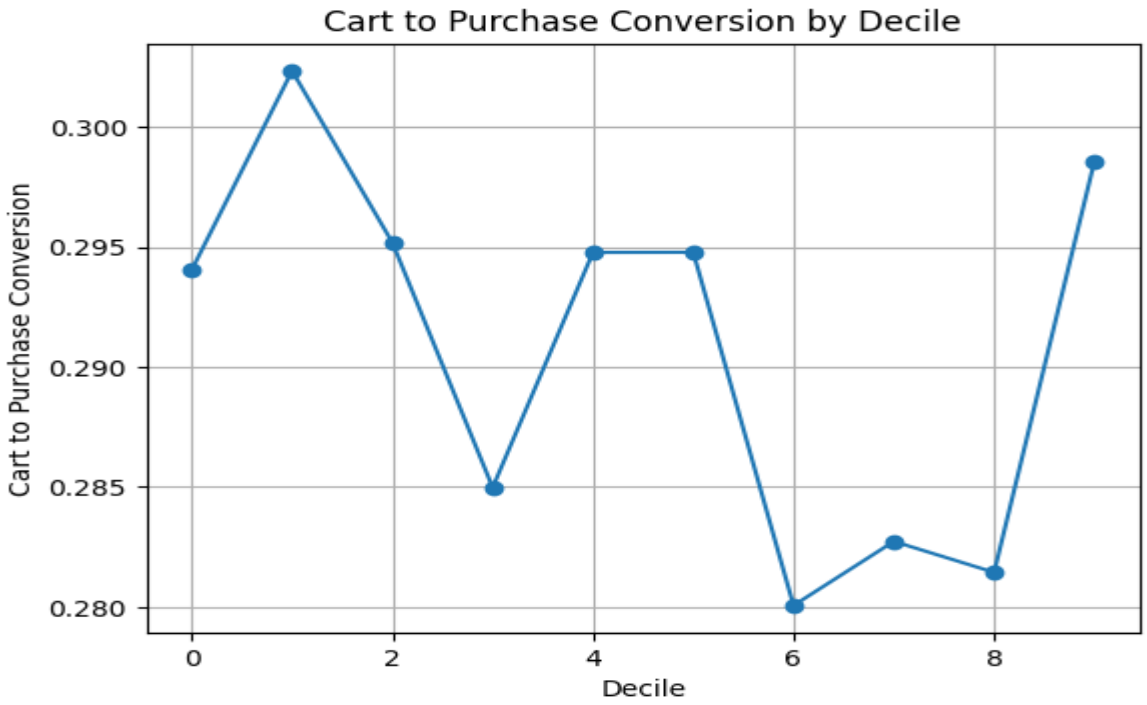

The results show that cart-to-purchase conversion remains remarkably stable across all deciles, fluctuating within a narrow range of approximately 0.28–0.30.

No consistent upward trend is observed in higher-frequency segments, despite their larger traffic volumes and presumed higher intent.

In fact, some mid-frequency deciles slightly outperform top deciles in conversion efficiency, although the differences remain small.

This challenges a widely held assumption in e-commerce analytics:

> High-frequency queries are often assumed to represent stronger user intent and, therefore, higher conversion rates.

However, the data suggests that once a user reaches the cart stage, the probability of completing a purchase is largely independent of query frequency.

This implies that:

- Checkout-related friction affects users uniformly across segments

- Conversion bottlenecks are systemic rather than query-specific
- Improvements in the checkout experience are likely to yield broad, cross-segment gains.

4.4. Revenue Efficiency Across Query Segments

To identify potential inefficiencies in value distribution, revenue efficiency was measured as the ratio of total purchases to total searches for each decile.

Figure 8. Revenue per Search by Frequency Decile

The results indicate that revenue efficiency is nearly constant across all deciles, remaining within a tight range of approximately 0.15–0.16.

This finding is particularly significant when compared to the distribution of search volume:

- The highest-frequency decile generates more than 10 times the number of searches compared to the lower deciles
- However, its efficiency is not meaningfully higher.

In other words:

> Large differences in traffic do not translate into proportional differences in performance.

This reveals a fundamental structural characteristic of the system:

- High-frequency queries dominate in absolute terms
- But not in relative efficiency

Mid-frequency deciles (approximately 4–7) demonstrate comparable efficiency levels, despite receiving substantially less traffic.

This suggests the presence of underutilized demand segments capable of generating similar value per interaction but underexposed within the system.

4.5. Interpretation

Taken together, the results provide a consistent and coherent explanation of value distribution in the e-commerce search system.
The analysis demonstrates that:

1. Search demand is highly concentrated in a small number of queries
2. Revenue concentration closely follows this demand distribution.
3. Conversion efficiency remains stable across query segments.
4. Platform differences significantly affect final conversion outcomes.
5. Mid-frequency queries exhibit competitive efficiency but limited exposure.

The most important implication is that:
Differences in performance do not drive the observed concentration of revenue; rather, they drive differences in exposure.
This suggests that the search system does not simply reflect user demand, but actively amplifies it through ranking and visibility mechanisms.

4.6. Implications

The findings of this study suggest that significant improvements in overall performance may be achieved without increasing total traffic.
Specifically:

- Redistributing exposure toward mid-frequency queries may unlock additional revenue

- Improving mobile conversion, particularly in high-frequency segments, represents a high-impact opportunity

- Optimization strategies should focus not only on top queries, but on the entire distribution of demand

CHAPTER 5. INTERPRETATION OF RESULT AND RECOMMENDATIONS

5.1. Managerial Implications

The findings of this study have direct implications for e-commerce product strategy, search optimization, and AI-driven ranking systems.

5.1.1. Rethinking Search Optimization Strategy

Most e-commerce platforms prioritize optimization of high-frequency queries, assuming that they represent the highest value opportunities.
However, the results suggest that this approach is suboptimal.
Since conversion efficiency is similar across query segments, focusing exclusively on top queries leads to diminishing returns.

Instead, companies should:

- expand coverage of mid-frequency queries,
- improve recall and relevance in underrepresented segments,
- and reduce dependency on a small subset of dominant queries.

5.1.2. Opportunity in Mid-Frequency Queries

Mid-frequency queries represent the most promising opportunity for growth.

These queries:

- demonstrate competitive conversion rates,
- reflect more specific user intent,
- and remain underexposed in the current system.

By reallocating visibility toward these segments, platforms can increase total revenue without increasing traffic acquisition costs.

5.1.3. Platform-Specific Optimization

The study highlights significant conversion gaps between platforms.
In particular:

- Web outperforms mobile in final conversion stages
- Mobile platforms show strong engagement but weaker completion rates

This suggests that optimization should not be uniform.
Instead:

- Mobile experience should focus on reducing checkout friction
- Payment flows should be simplified and accelerated
- UI/UX should be optimized for faster decision-making

Improving mobile conversion in high-frequency queries alone could generate substantial revenue gains.

5.1.4. Implications for Ranking Algorithms

The findings have direct implications for search ranking systems.
Traditional ranking approaches tend to reinforce popularity, leading to:

- overexposure of high-frequency queries,
- underrepresentation of long-tail demand,
- and suboptimal distribution of economic value.

To address this, ranking systems should incorporate additional signals beyond popularity.

5.2. AI-Driven Recommendations

Based on the findings, several advanced approaches can be proposed for improving search performance using AI.

5.2.1. Profit-Aware Ranking

Instead of ranking results purely by relevance or popularity, systems should incorporate expected economic value.
This can be achieved by optimizing ranking functions based on:

- predicted conversion rate
- expected margin
- probability of purchase

This approach allows the system to prioritize queries and products that generate higher value per interaction.

5.2.2. Exposure Rebalancing Algorithms

To reduce structural bias, platforms can introduce controlled exploration mechanisms.

For example:

- allocate a percentage of impressions to mid-frequency queries
- dynamically adjust ranking weights to improve coverage
- use multi-armed bandit or reinforcement learning approaches

This allows the system to discover and exploit underutilized demand.

5.2.3. Intent-Based Query Understanding

The study shows that query length and structure influence performance, but not linearly.
AI models can be used to better interpret user intent by:

- embedding queries into semantic vectors
- clustering similar intents
- mapping queries to product categories and user needs

This enables more accurate matching between user intent and product results.

5.2.4. Platform-Adaptive Ranking

Given the differences between platforms, ranking strategies should be adapted accordingly.

For example:

- prioritize fast-converting products on mobile
- simplify result sets for smaller screens
- adjust ranking weights based on device context

This allows the system to compensate for platform-specific limitations.

5.3. Limitations of the Study

This research has several limitations that should be acknowledged.
First, revenue was approximated using purchase counts, which does not account for price variability or profit margins.
Second, the analysis was based on aggregated data, which limits the ability to capture individual user behavior and session-level dynamics.
Third, the study focuses on observed correlations and does not establish strict causality.
Finally, external factors such as promotions, seasonality, and marketing campaigns were not explicitly controlled for.

5.4. Future Research Directions

Future research could extend this study in several ways:

- incorporating product-level pricing and margin data
- analyzing user sessions and behavioral sequences
- testing ranking interventions through A/B experiments
- applying causal inference methods to better isolate effects

Additionally, integrating real-time machine learning models into search systems represents a promising direction for further optimization.

Conclusion

The purpose of this study was to investigate how search query characteristics influence conversion performance and revenue distribution in an e-commerce environment.
The analysis focused on three key dimensions:

- query frequency (deciles),
- platform differences (Web, iOS, Android),
- and query structure (length).

The results provide a consistent and evidence-based answer to the research questions.
This study examined how search behavior influences e-commerce performance using real-world data from EVA.UA. The analysis shows that the primary driver of revenue in search is not conversion efficiency, but the distribution of demand across queries.
The results indicate that approximately 45% of revenue is generated by high-frequency queries. When aggregated at the segment level, the broader group of high-frequency queries (top deciles) contributes up to 85% of total revenue.

However, this dominance is not explained by significantly better conversion rates. In fact, conversion remains relatively stable across all query segments. This means that high-frequency queries generate more revenue primarily because they receive more traffic and visibility, not because they perform better.
This finding has a direct practical implication: current search systems are structurally biased toward already popular queries. As a result, a large portion of demand — particularly in mid- and low-frequency queries — remains underutilized.

From a business perspective, this creates a clear growth opportunity. Mid-frequency queries represent the most balanced segment, combining meaningful traffic with solid conversion potential. Improving ranking quality and product matching in this segment can lead to immediate and measurable revenue growth without increasing traffic acquisition costs.
Low-frequency queries, or the long tail, represent a different type of opportunity. These queries are highly diverse and often reflect more specific user intent. Their low contribution to revenue is not caused by weak demand, but by limitations in search systems, particularly in handling variation in language and intent. Unlocking this segment requires more advanced approaches, such as semantic search and query understanding.

The study also confirms that platform differences significantly affect performance. Lower conversion rates on mobile platforms suggest the presence of friction in the user journey, particularly in the later stages of the funnel. This indicates that improving mobile experience, especially in search interaction and checkout processes, can deliver additional performance gains.

Overall, the key conclusion is that search optimization should shift from a narrow focus on high-frequency queries toward a broader, segment-based approach. Growth will not come from further optimizing what already performs well, but from improving how the system handles underutilized demand.
In practical terms, companies should focus on three priorities: improving ranking quality in mid-frequency queries, investing in long-tail search capabilities, and reducing friction in mobile conversion flows. These actions provide a clear path to increasing revenue efficiency without relying solely on traffic growth.

Key Findings

First, the study confirms that search demand is highly concentrated. A small group of high-frequency queries generates the majority of search traffic. This concentration extends directly to revenue, with the top decile contributing nearly half of total purchases.
Second, despite this concentration, conversion efficiency remains stable across query segments. Purchase rates and cart-to-purchase conversion fluctuate within narrow ranges, showing no systematic advantage for high-frequency queries.
Third, the analysis reveals that platform differences significantly affect performance. Web users demonstrate consistently higher conversion rates, while mobile platforms, particularly Android, underperform at later stages of the funnel.
Fourth, mid-frequency queries demonstrate comparable revenue efficiency to top queries but receive significantly less traffic, indicating the presence of underutilized demand.

Core Insight

The central finding of this research is:
Revenue concentration in e-commerce search is primarily driven by traffic allocation rather than differences in conversion performance.
This means that high-frequency queries dominate not because they are inherently more effective, but because they receive disproportionately higher visibility within the search system.
As a result, the search engine does not merely reflect user demand, but actively shapes the distribution of economic value.5.1.3. Theoretical Contribution
This study contributes to the existing literature by providing empirical evidence that challenges the assumption that popular queries are more efficient.
Instead, it demonstrates that:

- query popularity is weakly related to conversion performance,
- conversion behavior is relatively uniform across segments,
- and system-level mechanisms (such as ranking and exposure) play a dominant role in value creation.

This aligns with and extends prior research on search ranking bias and long-tail economics by quantifying the gap between visibility and efficiency.

REFERENCES

Anderson, C. (2006). *The long tail: Why the future of business is selling less of more*. Hyperion.

Baeza-Yates, R., & Ribeiro-Neto, B. (2011). *Modern information retrieval: The concepts and technology behind search* (2nd ed.). Addison-Wesley.

Broder, A. (2002). A taxonomy of web search. *ACM SIGIR Forum, 36*(2), 3–10. https://doi.org/10.1145/792550.792552

Brynjolfsson, E., Hu, Y. J., & Smith, M. D. (2011). The longer tail: The changing shape of Amazon's sales distribution curve. *Management Science, 57*(8), 1373–1386. https://doi.org/10.1287/mnsc.1110.1371

Chapelle, O., Joachims, T., Radlinski, F., & Yue, Y. (2012). Large-scale validation and analysis of interleaved search evaluation. *ACM Transactions on Information Systems, 30*(1), 1–41. https://doi.org/10.1145/2094072.2094078

Croft, W. B., Metzler, D., & Strohman, T. (2010). *Search engines: Information retrieval in practice*. Addison-Wesley.

Google. (2018). How mobile has changed the way people shop. Retrieved from https://www.thinkwithgoogle.com

Huang, J., White, R. W., & Buscher, G. (2019). User behavior in web search. In B. Liu & H. Motoda (Eds.), *Encyclopedia of machine learning and data mining* (pp. 1345–1351). Springer.

Jansen, B. J., & Spink, A. (2006). How are we searching the World Wide Web? A comparison of nine search engine transaction logs. *Information Processing & Management, 42*(1), 248–263. https://doi.org/10.1016/j.ipm.2004.10.007

Jansen, B. J., & Schuster, S. (2011). Bidding on the buying funnel for sponsored search and keyword advertising. *Journal of Electronic Commerce Research, 12*(1), 1–18.

Joachims, T., Swaminathan, A., & Schnabel, T. (2017). Unbiased learning-to-rank with biased feedback. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining* (pp. 781–789). https://doi.org/10.1145/3018661.3018699

Li, L., Chu, W., Langford, J., & Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web* (pp. 661–670). https://doi.org/10.1145/1772690.1772758

Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press.

Ricci, F., Rokach, L., & Shapira, B. (2015). *Recommender systems handbook* (2nd ed.). Springer.

Shokouhi, M. (2013). Segmentation of search queries. In *Proceedings of the 22nd International World Wide Web Conference* (pp. 123–124).

APPENDIX A. DATA Sample

query_text	searches	purchases	add_to_cart	purchase_rate	add_to_cart_rate	query_length	query_length_words	query_length_chars	query_length_group
58241	12	0	12	0.0	1.0	1	1	5	1
Catrice	20	3	11	0.15	0.55	1	1	7	1
Celimax	15	0	9	0.0	0.6	1	1	7	1
CeraVe	12	0	4	0.0	0.3333333333333333	1	1	6	1
Dior	33	2	14	0.06060606060606061	0.4242424242424242425	1	1	4	1
Hair Trend	10	0	0	0.0	0.0	2	2	10	2
Huggies	13	1	9	0.07692307692307693	0.6923076923076923	1	1	7	1
Lamel	49	8	27	0.16326530612244897	0.5510204081632653	1	1	5	1
Maybelline	52	1	29	0.019230769230769232	0.5576923076923077	1	1	10	1
Miso	12	12	12	1.0	1.0	1	1	4	1
Re.Care	37	2	25	0.05405405405405406	0.6756756756756757	1	1	7	1
Skin1004 probiocica cream	15	2	10	0.13333333333333333	0.6666666666666666	3	3	25	3
about	91	9	39	0.09890109890109989	0.4285714285714285 5	1	1	5	1
about	59	4	26	0.06779661016949153	0.4406779661016949	1	1	5	1
cat	25	0	10	0.0	0.4	1	1	3	1
catrice	161	20	88	0.12422360248447205	0.5465838509316 77	1	1	7	1
celimax	16	3	13	0.1875	0.8125	1	1	7	1
cha	20	1	7	0.05	0.35	1	1	3	1
dove	119	19	66	0.15966386554621848	0.5546218487394958	1	1	4	1

dove	25	1	12	0.04	0.48	1	1	4	1
esense	14	0	11	0.0	0.7857142857142857	1	1	6	1
essence прайм ер	29	0	16	0.0	0.5517241379310345	2	2	15	2
eveline	248	30	156	0.12096774193548387	0.6290322580645161	1	1	7	1
farm	16	3	12	0.1875	0.75	1	1	4	1
garni	17	4	8	0.23529411764705882	0.47058823529411764	1	1	5	1

APPENDIX B. CODE (GOOGLE COLAB)

This appendix provides access to the Python code used for data processing and analysis.

The full codebase, including data loading, preprocessing, aggregation, and visualization, is available in the following Google Colab notebook:

https://colab.research.google.com/drive/1kKWO07PdFCSds_mcNXoM65ZZ0Il3gD-J#scrollTo=lO2ck8y_IqJM

```
import pandas as pd

query_stats = {}

for file in files:

    print("Processing:", file)

    for chunk in pd.read_csv(
        file,
        chunksize=chunk_size,
        low_memory=False
    ):

        chunk = chunk.dropna(subset=['query_text'])

        chunk['add_to_cart_flag'] = (chunk['total_add_to_cart_count'] >
0).astype(int)

        agg = chunk.groupby('query_text').agg(
            searches=('query_text', 'count'),
            purchases=('had_purchase', 'sum'),
            add_to_cart=('add_to_cart_flag', 'sum')
```

```
        )

        for query, row in agg.iterrows():

            if query not in query_stats:
                query_stats[query] = [0,0,0]

            query_stats[query][0] += row['searches']
            query_stats[query][1] += row['purchases']
            query_stats[query][2] += row['add_to_cart']
```