



Peer Review Statistical Analysis

Author: Lyrie Adler



Overview

This methodology outlines the statistical approach used to analyze entitlement assignments and generate peer-based certification recommendations. Designed to streamline User Access Reviews (UAR) and validate Access Requests, this framework assists in these processes by automatically flagging high-risk outliers for reviewers to investigate while identifying safe, peer-validated requests that are suitable for auto-approval. The primary objective is to determine whether entitlements are distributed randomly throughout the organization or if they exhibit significant department-based patterns. By identifying these patterns, the system creates a scoring model for every user-entitlement assignment.

Part 1: Probability Calculations and Visualization

The foundation of this analysis rests on core probability metrics, starting with the **Overall Entitlement Probability**, or **P(General)**. This metric defines the probability that any eligible user in the organization holds a specific entitlement. While the basic calculation involves dividing the number of users holding the entitlement by the total eligible users, a critical nuance is the **Vendor-Based User Pool**. Since not all entitlements are universally applicable—for instance, a Salesforce role requires a Salesforce account—the eligible user pool is restricted to users belonging to the distinct vendors associated with the entitlement.

Building on the general probability, the analysis calculates **P2**, or the **Department Probability**. This measures the likelihood of a user possessing an entitlement given their membership in a specific department. Like the general metric, the department-level calculation respects the vendor-based user pool, counting only those department members who belong to the relevant vendors.

To ensure these calculations are robust, the system addresses data gaps through a specific fallback mechanism. Since identity providers like Okta or Azure AD do not always contain department data, the system primarily looks at the user account's direct department field. If this data is missing, the system falls back to the linked Person record to fill in the missing department context.

This data is visualized using department distribution charts. For every entitlement, a bar chart compares the P(General) baseline against the P2 for each department. These visuals allow for immediate interpretation: if department bars (P2) align with the baseline (P(General)), the entitlement is independent; if specific department bars spike significantly above the baseline, it indicates elevated access and potential concentration.

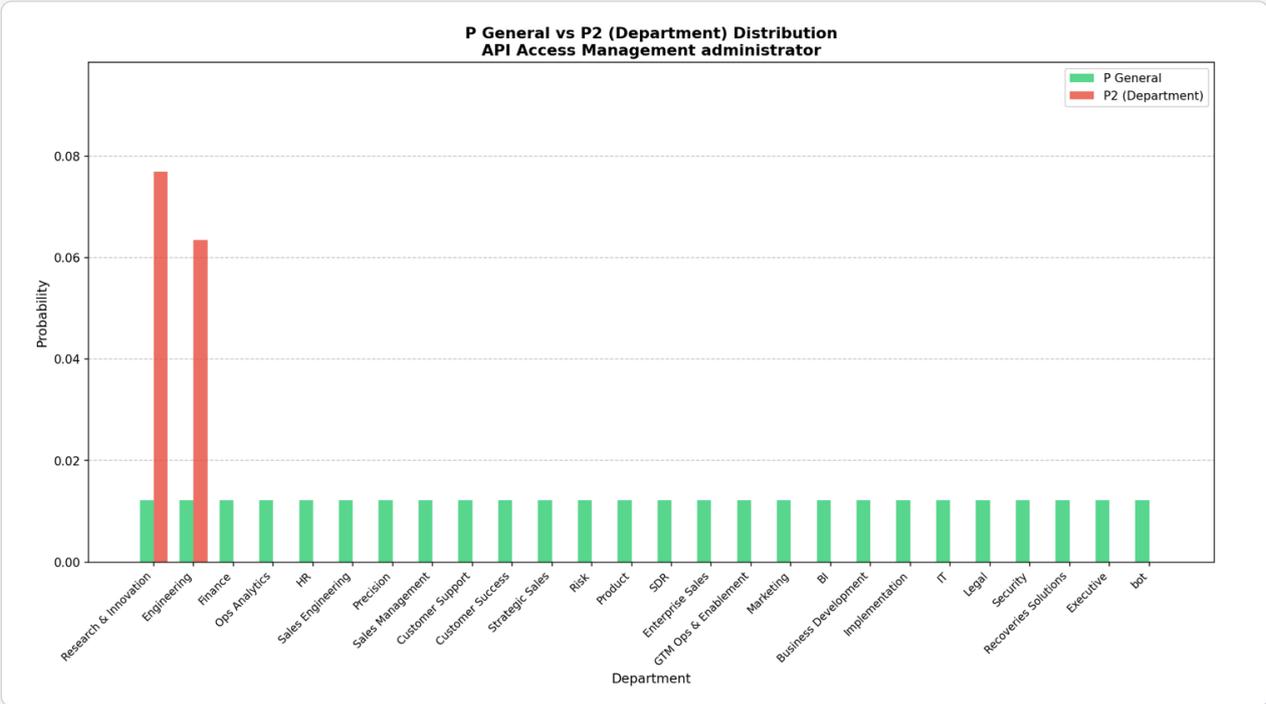


Figure 1 Strong dependency

Note how the red bars (Department Probability) for Engineering and research&innovation spike significantly above the green baseline (General Probability), while other departments have near-zero access. This indicates that access is highly concentrated and dependent on specific job functions.

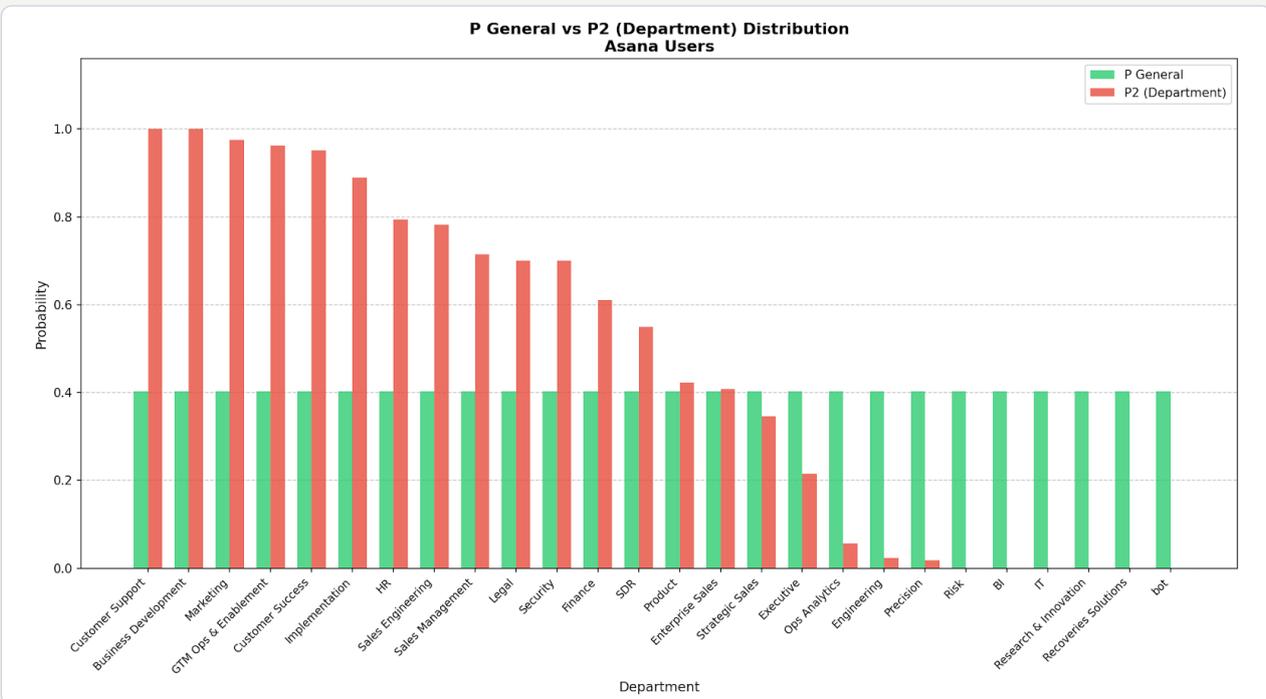


Figure 2 Moderate dependency

While access is widespread, distinct patterns emerge where certain departments (customer support, business development) have notably higher probabilities than others (engineering, precision). The red bars vary in height relative to the green baseline, showing a pattern that is neither purely random nor strictly isolated.

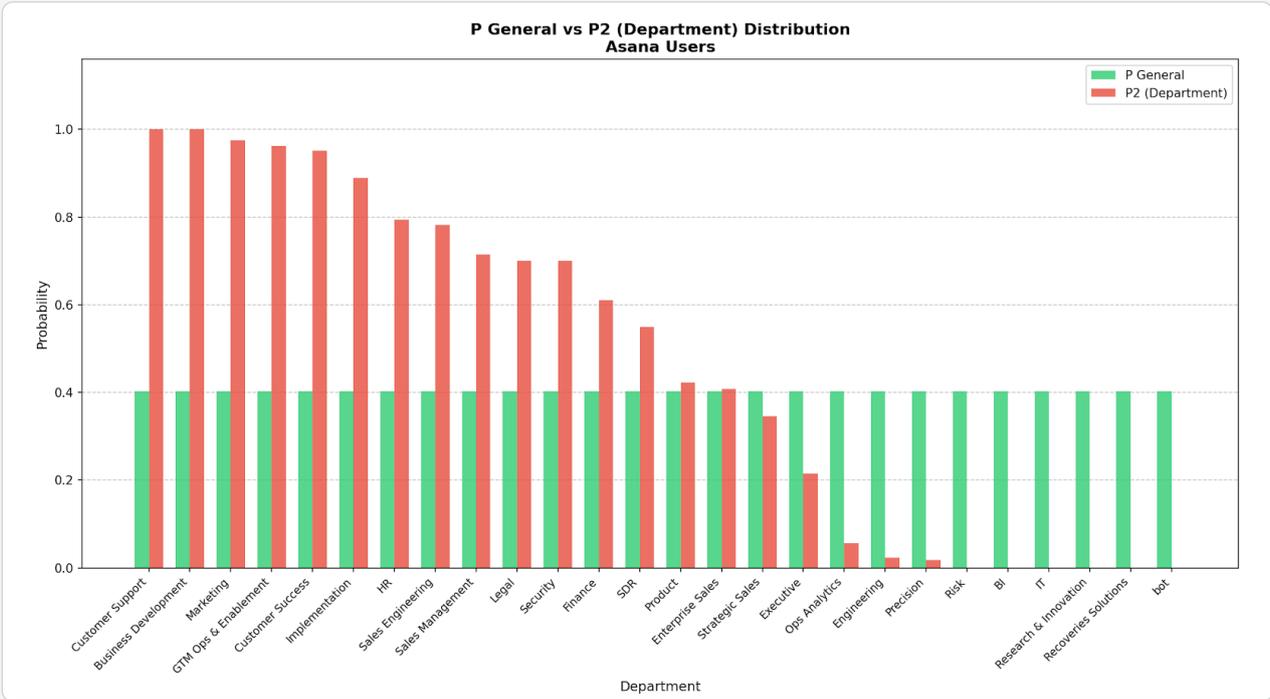


Figure 3 Independent

Independent Distribution. The red bars (Department Probability) are nearly identical in height to the green bars (General Probability) and to each other across all departments. This "flat" distribution indicates that a user's department has no impact on their likelihood of holding this entitlement.

Part 2: Department Dependency Analysis

The dependency analysis classifies each entitlement into one of four categories based on whether its distribution shows statistically significant department patterns. This process relies on a combination of standard statistical tests and practical usage indicators.

2.1 STATISTICAL PARAMETERS

Chi-Square Test for Independence

The primary statistical test used is the Chi-Square Test, which measures whether the distribution of an entitlement across departments is statistically different from what would be expected by random chance. Mathematically, it compares the observed frequencies against expected frequencies using the following formula:

$$\chi^2 = \sum \frac{(Observed - Expected)^2}{Expected}$$

Intuitively, if every user had an equal probability $P(\text{General})$ of receiving an entitlement regardless of their department, we would expect each department to hold a number of entitlements proportional to its size. The Chi-Square test quantifies how far reality deviates from this expectation. The analysis utilizes a standard threshold of $p < 0.05$.

Small Grant Exception: For entitlements with fewer than 20 users, Chi-Square tests often lack the statistical power to identify patterns, even when obvious dependencies exist (such as a 6x lift in a specific team). Therefore, entitlements with fewer than 20 users bypass the Chi-Square requirement and are evaluated solely on practical indicators.

Variance of P2 Across Departments

To measure the practical "spread" of access rates, the analysis calculates the variance of the department probabilities (P2).

$$\text{Variance} = \frac{\sum(P2_{dept} - \text{Mean_P2})^2}{N_{departments}}$$

High variance indicates that an entitlement is common in some departments and rare in others—a clear sign of department-based access. The threshold for high variance is set at $\text{Variance} \geq 0.0369$, which represents the 75th percentile of observed variances. This ensures that only the top 25% most variable entitlements—those with truly distinctive patterns rather than minor fluctuations—are flagged.

Concentration Ratio

While variance measures spread, the Concentration Ratio determines if ownership is dominated by a small number of groups. It is calculated by determining the percentage of entitlement holders located within the top three departments:

$$\text{Concentration} = \frac{\text{Users with Entitlement in Top 3 Departments}}{\text{Total Users with Entitlement}}$$

A threshold of **Concentration $\geq 60\%$** is used. Given that most organizations have 15 to 25 or more departments, finding more than half of all entitlement holders concentrated in just three departments indicates a meaningful organizational association.

Lift (Max P2 / P General)

The analysis includes a "Lift" metric to identify how much more likely the top department is to have the entitlement compared to the organization average.

$$\text{Lift} = \frac{\max(P2 \text{ across all departments})}{P(\text{General})}$$

For example, if the general probability is 2% but the Engineering department has a 12% probability, the Lift is 6x. The analysis uses a threshold of **Lift $\geq 3.0x$** . This parameter is critical because it captures cases that variance and concentration might miss; an entitlement might have low overall variance, but if a single department is significantly more likely to hold it than the average employee, a relationship clearly exists. Furthermore, for the "Small Grant" entitlements that bypass Chi-Square, a high lift serves as a strong primary indicator of dependency.

2.2 CLASSIFICATION LOGIC

The final classification of an entitlement involves a multi-step logic flow that weighs statistical evidence against the practical indicators (Variance, Concentration, and Lift).

Step 1: Significance Check

First, the system determines if the entitlement passes the "significance" check. This occurs if the Chi-Square p-value is less than 0.05 OR if the entitlement falls under the Small Grant Exception (fewer than 20 users).

Step 2: Practical Indicator Count

Next, the system counts how many practical thresholds are met: High Variance (≥ 0.0369), High Concentration ($\geq 60\%$), or High Lift ($\geq 3.0x$).

Step 3: Final Classification

Based on the previous steps, the entitlement is assigned one of four statuses:

- **Strong Dependency:** The entitlement passes the significance check and meets two or more practical indicators. This represents a clear department pattern, such as "GitHub Developer Access," which might show both high lift and high concentration in Engineering.
- **Moderate Dependency:** The entitlement passes the significance check but meets exactly one practical indicator. An example might be "API Admin Access," which shows a high lift but lacks broad concentration.
- **Weak Dependency:** The entitlement is statistically significant ($p < 0.05$) but fails to meet any of the practical indicators (Variance, Concentration, or Lift). These are often broad tools like VPNs with only minor usage fluctuations.
- **Independent:** The entitlement fails the significance check ($p \geq 0.05$ and users ≥ 20). These distributions are effectively random, such as a universal "Everyone" group.

Part 3:

Entitlement Recommendation Scoring

For entitlements classified as having Strong or Moderate Dependency, the system calculates scores to recommend approval or revocation.

A crucial step in scoring is establishing a baseline. Rather than using the organization-wide P(General) or the maximum P2—which can be skewed by unrelated departmental changes or small team sizes—the system uses the **Top 3 Departments Average**. This baseline is the average P2 of the three departments with the highest number of users holding the entitlement, ensuring the comparison is made against departments that genuinely use the entitlement at scale.

The **Approve Score** (0–100) measures how appropriate an assignment is. It combines two components: the **P2 Component**, which assesses the absolute likelihood of the entitlement within the user's department, and the **Ratio Component**, which compares the department's usage to the baseline. Scores above 75 indicate strong appropriateness, while scores below 50 warrant review.

Conversely, the **Revoke Score** (0–100) identifies anomalies. This score is designed conservatively; it only accumulates points when metrics fall below a 5% threshold. If a user's department has less than a 5% probability of holding the entitlement (P2) and falls into the bottom 5% relative to typical users (Ratio), the Revoke Score increases. A score of 80–100 indicates an extremely anomalous assignment that is a strong candidate for revocation.

These scores function together to guide reviewers. An assignment with a high Approve Score and low Revoke Score is a candidate for auto-approval, whereas a low Approve Score paired with a high Revoke Score suggests the access should be revoked. Assignments falling in the middle range are flagged for manual review.

Summary

This peer review analysis provides a comprehensive, data-driven framework for entitlement certification. By calculating department-specific probabilities, classifying entitlements based on statistical evidence, and employing a dual-score system for individual assignments, the methodology enables reviewers to confidently approve appropriate access while focusing their attention on genuine anomalies.