Analyze the Spire: An analytical approach to improving the Design of a Sequel

Authors

Alex Ma, Brandon Lyman, Sebastian Salas, Akshar Vandara

1. Executive Summary

This research aims to identify player personas in Slay the Spire based on behavioral patterns observed in gameplay data. The game's design allows for free-form deckbuilding alongside random events, offering insight into how players manage uncertainty, resources, and strategic choices. Understanding these patterns can inform the design of Slay the Spire 2, helping developers tailor features to different types of players.

We analyze actions such as card selection, combat goals, risk-taking, and resource use. From this, we identify four distinct player personas: **Strategists**, who focus on long-term planning and efficient resource management, **Bursters**, who prioritize aggressive, high-damage strategies, **Quick Experimenters**, who enjoy trying new mechanics and experiments card combinations and **Gamblers**, who favor high-risk, high-reward choices. These personas can serve as a framework for designing cards, relics, and mechanics that support a wider range of playstyles. The results provide clear guidelines for integrating persona-driven design into future development.

2. Introduction

Slay the Spire is a roguelike deck builder where players can choose 1 of 4 classes and ascend a spire. The player collects various resources throughout a single run and builds their deck over time. Each class provides a different card pool which will be important for later analyses. The game is separated into 3 main acts, each with a boss at the end of the act. A player traverses through various nodes with a wide variety of encounters, from enemies to elites to random events.



Battle Phase from Slay the Spire.

3. Data Sources

3.1 Player Run Data

We acquired internal metrics data directly from the development team of Slay the Spire, which is publicly available via Jake Rabinowitz's Google Drive. These data consist of run history for players of the game between 2018 and 2020. Each record in the dataset provides comprehensive metrics of the player's actions throughout one play session of STS. This will be the foundation of our analysis of player personas.

The data was provided in a .json format, which was unwrapped and loaded into a remote relational MySQL database for accessibility by the entire team. As this dataset is massive (77 million runs), we will begin our analysis with 1269 runs that all occurred on November 12, 2020. We will build our analysis pipelines using this representative dataset, and acquire more data from the collection into our ecosystem should we determine more records are required to meaningfully answer the research questions.

The link to the Data can be found **here**.

As for what each data represents, we provide the following Data Dictionary with each item and its definition as seen in the below table.

Item	Definition
RunID	Unique key to identify each run.
Character	The class selected for the run: Ironclad, Silent, Defect, Watcher.
Win/Loss	Binary outcome — either the player won (finished Act 3 or 4) or lost.
Player Experience	A proxy for player skill or account experience, often captured by previous games or win rate.
Floor Reached	The floor where the player lost or won.
Killed By	Enemy, trap, or condition that led to game over (e.g., "Act 2 Champ", "Hexaghost").
Ascension Level	Challenge rating chosen (from 0 to 20).
Play Time	Total time spent in the run.
Neow's Blessing/Cost	Starting bonus (or trade-off) offered at the beginning of the run.
Relics Obtained	All relics picked up during the run — tracked over time or by rarity.
Relics Picked/Not Picked	Relics rewarded or offered, alongside skip decisions.
Campfire Choices	Actions like Rest, Smith, Remove, Dig, Prayer Wheel done at rest sites.
Card Rewards Picked	Cards that were selected from reward screens.
Card Rewards Not Picked	Cards that were offered but skipped.
Current HP per Floor	Snapshot of health at the end of each floor.

Max HP per Floor	Tracks changes in max HP due to relics, events, or rest sites.
Damage Taken per Floor	Total damage received (can be zero) per floor.
Gold per Floor	Economic value recorded continuously.
Event Choices	Logs player choices on event maps (“? rooms”).
Items Purchased/Purged	Shop activity: what was bought, what was removed
Final Deck	The end-of-run card list — used to analyze synergy, rarity, build type.
Path Chosen per Floor	Map path for each floor.
Potions Spawned/Used	Detailed tracking of potion lifecycle.

4. Methods & Results

- *What analysis did you perform, and what did you learn? This is the bulk of your paper.*
    - *Visualizations*
    - *Tables*
    - *Text explanations and interpretations.*
- *You should demonstrate that you have thought critically about how the data relates to your question of interest.*
- *Did you meet your proposed research objective?*

4.1 Deck Composition and Player Motivation (RQ1)

The purpose of this analysis was to understand the composition of player decks and gather insight about their motivation to play a run of Slay the Spire.

4.1.1 Feature Engineering

We began by selecting and calculating the features we believe would contribute to our understanding of player motivation using domain expertise. This process yielded a wide range of features that we intended to reduce using techniques explained in the next subsection. A list of the 15 features included for the initial analysis, alongside the calculation method and justification for their inclusion is presented in **Table 1**.

4.1.2 Feature Reduction with Unsupervised + Supervised Learning

We conducted an iterative loop of principal component analysis (unsupervised learning), k-means clustering (unsupervised learning), and random forest regression (supervised learning) to reduce the number of features to something we could interpret with visuals and additional statistical methods. A flow chart depicting our process is presented in **Figure 1**. The entire process was conducted using R.

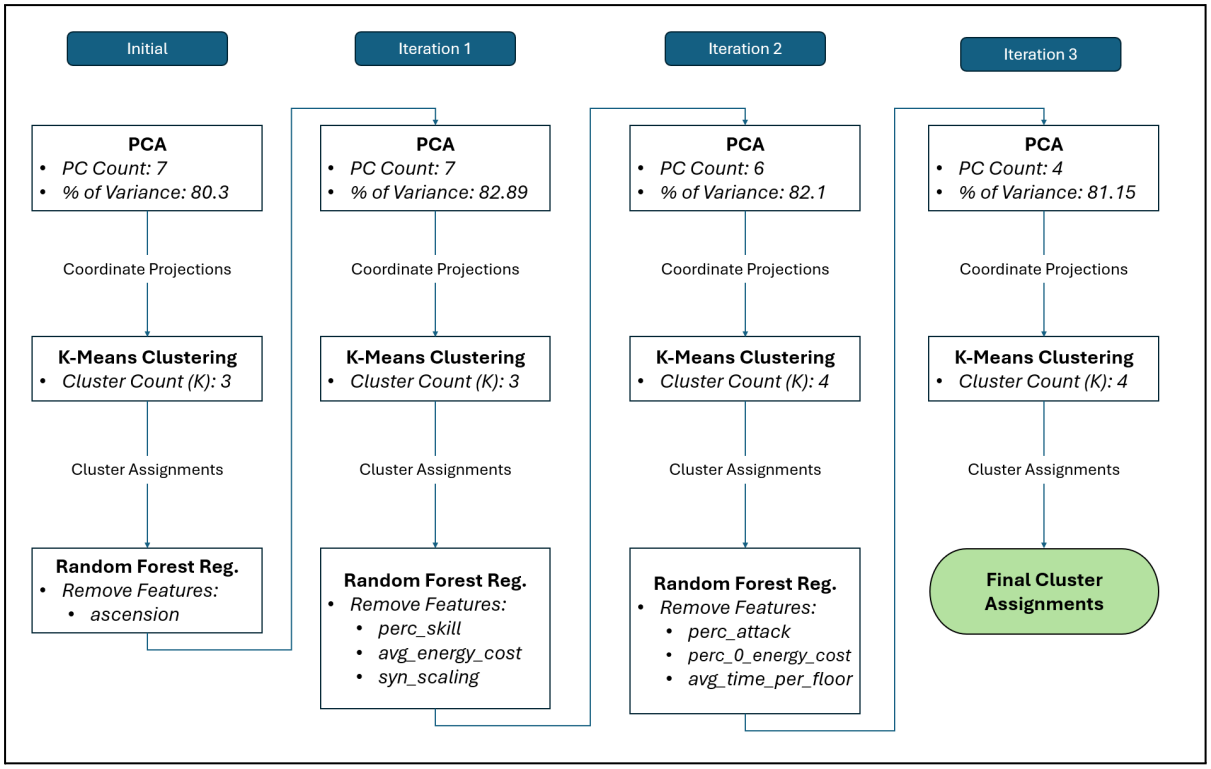


Figure 1. Unsupervised + Supervised Learning Pipeline for Feature Reduction

The contribution of each of these sub processes will be described in the following subsections.

Feature	Description	Semantic Calculation	Justification
ascension	Lore-specific difficulty, selected by the player	Given	Difficulty selection can indicate how much effort the player is looking to put forth on a given run.
avg_time_per_floor	How much time (in seconds) the player spent on each floor (level)	Total Play Time / Floor Reached	Players who spend more time on floors may be thinking more critically about their strategy.
deck_count	How many cards the player had in their deck at the end of a run	Given	Different deck strategies require different sized decks, could indicate a playstyle the player is motivated to execute.
perc_attack	Percentage of the final deck that consisted of attack cards	Count(attack_cards) / deck_count	More attacks could indicate more aggressive play styles.
perc_skill	Percentage of the final deck that consisted of skills	Count(skill_cards) / deck_count	More skills could indicate a more strategic or defensive play style.
perc_uncommon	Percentage of the final deck that consisted of uncommon cards	Count(uncommon_cards) / deck_count	More uncommon cards typically indicate that the player is specializing towards a particular strategy.
perc_rare	Percentage of the final deck that consisted of rare cards	Count(rare_cards) / deck_count	More rare cards typically indicate that the player is specializing towards a particular strategy.
avg_energy_cost	Average energy cost of all cards in the deck.	Sum(energy_cost_per_card) / deck_count	High energy cards may indicate that a player is going for more burst or big plays.
perc_0_cost_cards	Percentage of the final deck that consisted of 0 cost cards.	Count(0_cost_cards) / deck_count	0 cost cards is a common synergy among all characters that may indicate the player is trying to play as many cards as possible in a given turn.
perc_upgraded	Percentage of the final deck that consisted of upgraded cards	Count(upgraded_cards) / deck_count	Upgrading cards could indicate success and snowballing in a run, as the player typically can choose this over healing themselves.
perc_card_draw	Percentage of the final deck that consisted of generation (card draw) cards	Count(card_draw_cards) / deck_count	Card generation is another common synergy among all characters that may indicate the player is trying to find the card that activates a synergy in any given fight - indicating a more strategic approach
syn_offensive	Count of cards that are part of an offensive synergy	Count(vulnerable, poison, shiv, claw, lightning, mark)	Could indicate motivation for high damage output
syn_defensive	Count of cards that are part of a defensive synergy	Count(weak, intangible, artifact, heal, frost)	Could indicate motivation for “turtle” or long-term, patient strategies.
syn_scaling	Count of cards that are part of a scaling (stronger as the battles are longer) synergy	Count(strength, dexterity, focus, mantra)	Could indicate motivation for long-term strategies and big plays
syn_utility	Count of cards that are part of a utility (special effects and resource manipulation) synergy	Count(all_enemies, exhaust, card_draw, energy, x_cost, discard, retain, dark, scry, miracle, stance)	Could indicate motivation for short-term or long-term strategic thinking

Table 1. Features included for initial cluster analysis (See Section 4.1.1).

4.1.2.1 Principal Component Analysis (PCA)

We used PCA to reduce the dimensions of the feature space for visualization purposes in k-means clustering. For this sub process, we used R's *prcomp* method on a dataframe containing scaled values for each of the 15 features presented in Table 1. By generating a summary of the PCA analysis, we were able to see how many principal components were required to capture a cumulative proportion of the variance in the data. We chose 80% as a threshold based on Jolliffe's "Principal Component Analysis" (Jolliffe, 2011). We chose the minimum amount of principal components that met this threshold for each iteration of PCA. The output from the initial PC is provided in **Figure 2**.

```
Importance of components:
                          PC1    PC2     PC3     PC4     PC5     PC6     PC7     PC8
PC9    PC10    PC11
Standard deviation     2.2111 1.3883 1.19658 1.02930 1.01941 0.95404 0.88784 0.79099
0.75691 0.72732 0.65432
Proportion of Variance 0.3259 0.1285 0.09545 0.07063 0.06928 0.06068 0.05255 0.04171
0.03819 0.03527 0.02854
Cumulative Proportion  0.3259 0.4544 0.54987 0.62050 0.68978 0.75046 0.80301 0.84472
0.88292 0.91818 0.94673
                          PC12    PC13   PC14    PC15
Standard deviation     0.59577 0.50005 0.3310 0.29078
Proportion of Variance 0.02366 0.01667 0.0073 0.00564
Cumulative Proportion  0.97039 0.98706 0.9944 1.00000
```

Figure 2. Example PCA output from R (initial iteration). Here, 7 principal components were selected as this was the first number that's cumulative proportion passed the 80% threshold (80.3%).

Next, we plotted the original data, but translated into the new PCA coordinate space, and we examined the first two principal components in a 2D scatter plot. Additionally, we superimposed the loadings from all 15 original features to understand their influence on the projection. The results of this process from the initial iteration are presented in **Figure 3**.

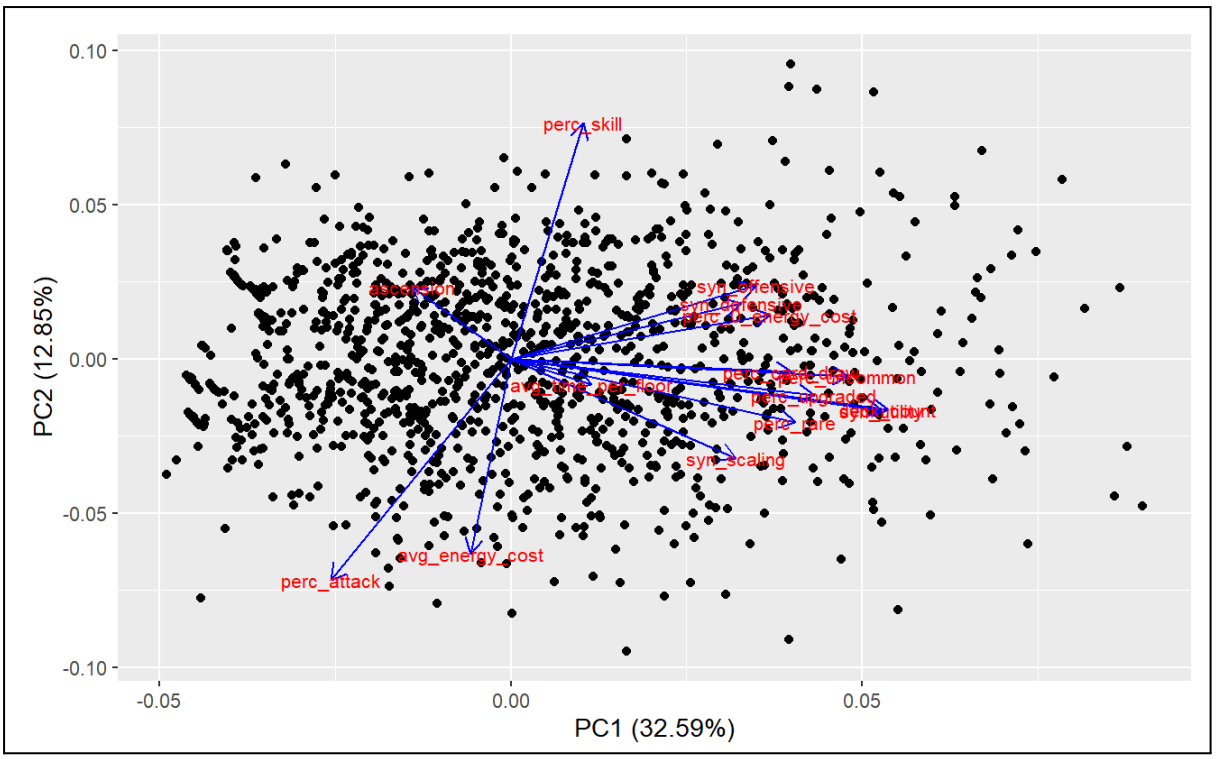


Figure 3. Example PCA projection plot with loadings from original 15 features.

4.1.2.2 K-Means Clustering

Following each iteration of PCA, we conducted k-means clustering analysis to assign a cluster group to each run in the data set. To determine the value of K, we used the elbow method: testing k values between 1 and 10, plotting the within-cluster sum of squares per k value, and looking for the "elbow" or where the difference in sum of squares per k value seems to drop off. In addition, we looked over a colored version of the condensed PCA space to validate the k value suggested by the elbow analysis. An example of the within sum of square plot and the colored version of the PCA space for the first iteration of the k-means clustering sub-process are presented in **Figure 4** and **Figure 5**, respectively.

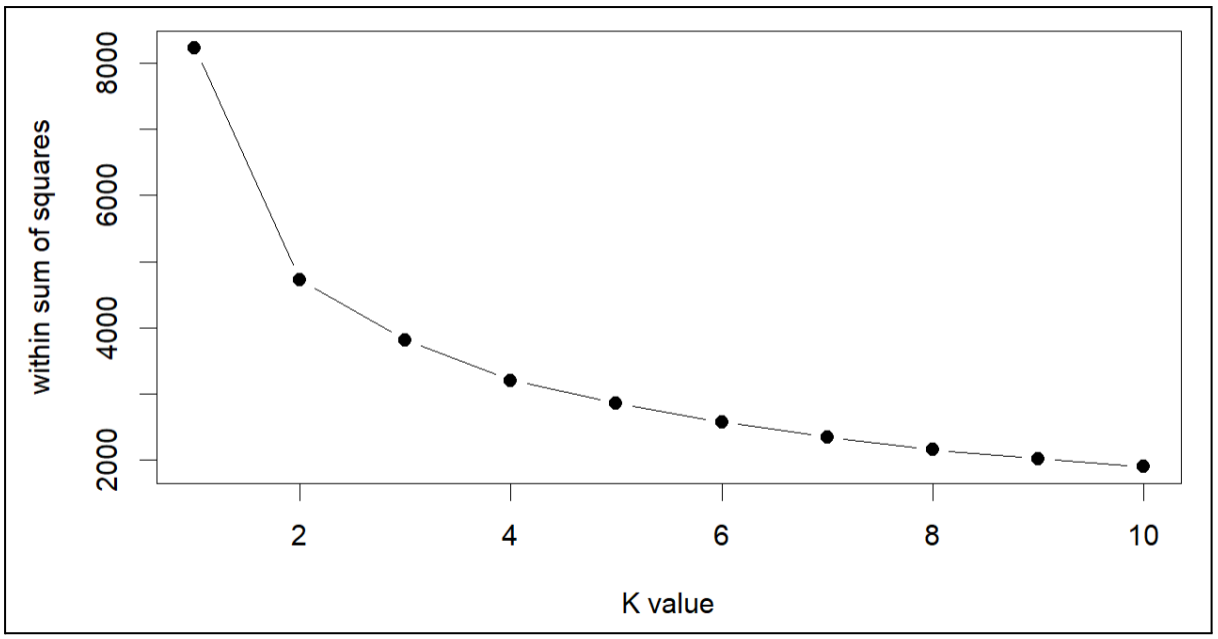


Figure 4. An example of the within sum of square plot for each k value 1:10 used for elbow method analysis for the initial iteration. Here we see the elbow is roughly at k = 3.

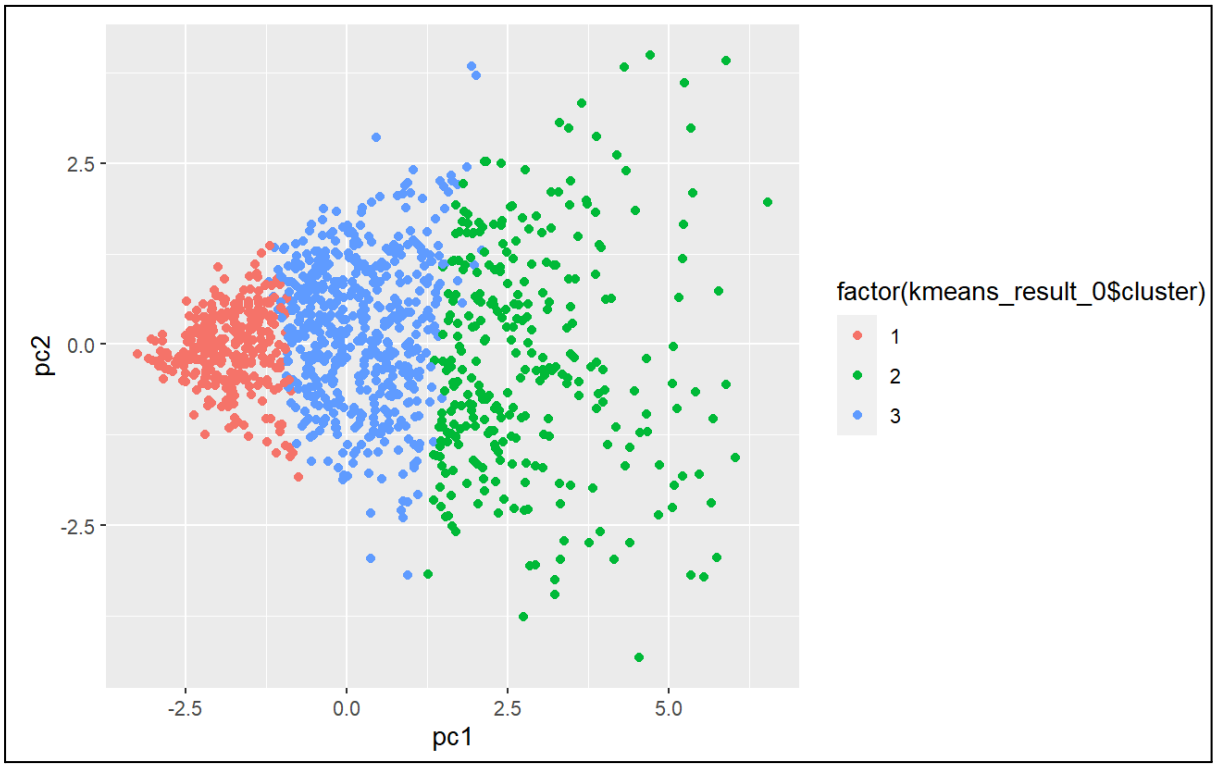


Figure 5. An example of the runs from the data set classified into the initial set of k = 3 clusters, and colored to represent their assigned cluster.

4.1.2.3 Random Forest Regression

The final subprocess that comprised our feature reduction strategy is the supervised learning method *random forest regression.* Using PCA and K-means clustering, we artificially created a target value for supervised learning. We used random forest regression to better understand which of the remaining features in each iteration contributed most to the classification of each datapoint from the original set into the k-means clusters. Resulting from this analysis are two feature importance plots: Mean Decrease in Accuracy (MDA) and Mean Decrease in the Gini Coefficient (MDG). According to Martinez-Tobada and Redondo (2020a, 2020b), the former is the amount of classification accuracy lost when the feature is removed, and the latter is a measure of homogeneity the feature contributes to the random forest. Both have a direct relationship with the importance of the feature in classifying a run into the predetermined cluster.

For each iteration of the random forest regression subprocess, we reviewed the variable importance plots and chose the obvious, least-important features to remove. An example of these plots from the initial iteration are presented in **Figure 6.**

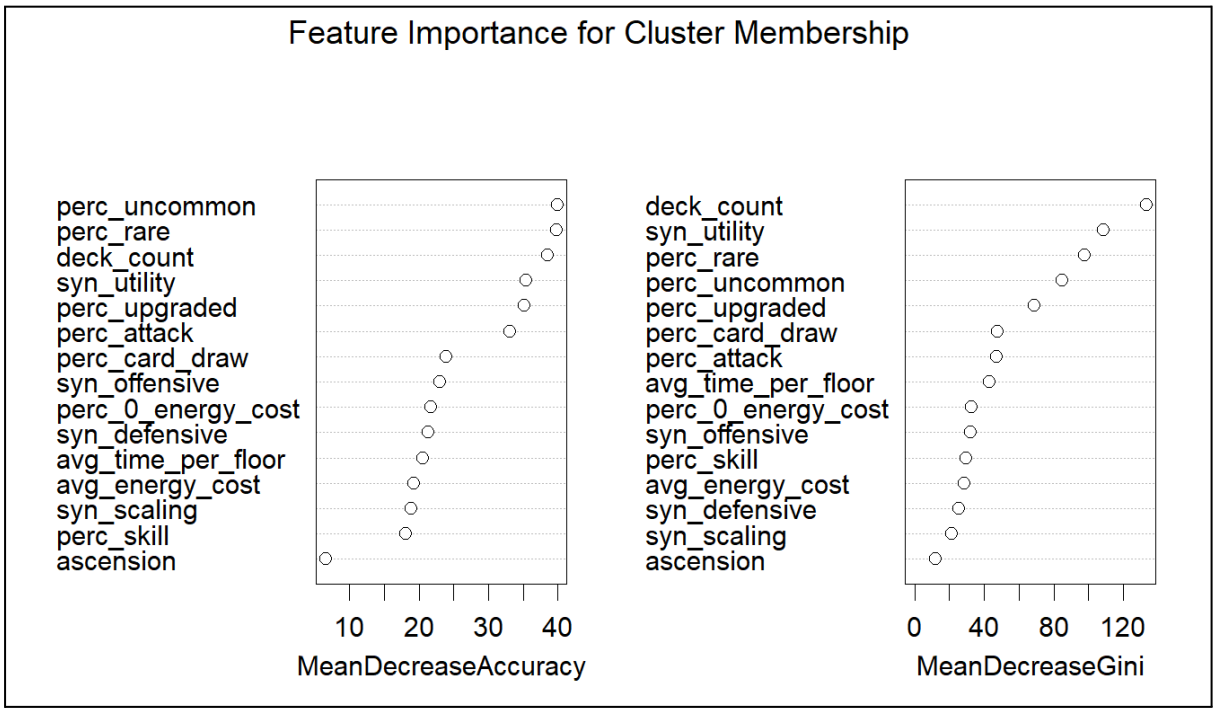


Figure 6. Example feature importance plots (MDA on the left, MDG on the right). In this example, ascension was the only feature that was chosen to be removed for the next iteration.

Resulting from each iteration of the random forest regression analysis was a new set of features that was smaller than the last, and that set was fed back into the PCA sub-process to further reduce the feature count.

4.1.2.4 Iteration

Referring back to **Figure 1**, we repeated this pipeline until we felt meaningful clusters had emerged to give us insight into the motivations or play styles represented by each run in the game. The final clusters, and a deeper exploration of them, is presented in the next section.

4.1.3 Cluster Exploration

The result of the iteration process was a final set of cluster assignments for the original run data set. Two visuals were required to capture the final clusters: a visual representation of the PC space with loadings from the final 8 features, and a colored clustered chart in PC space that denoted the cluster assignments. They are presented in **Figure 7** and **Figure 8**, respectively.

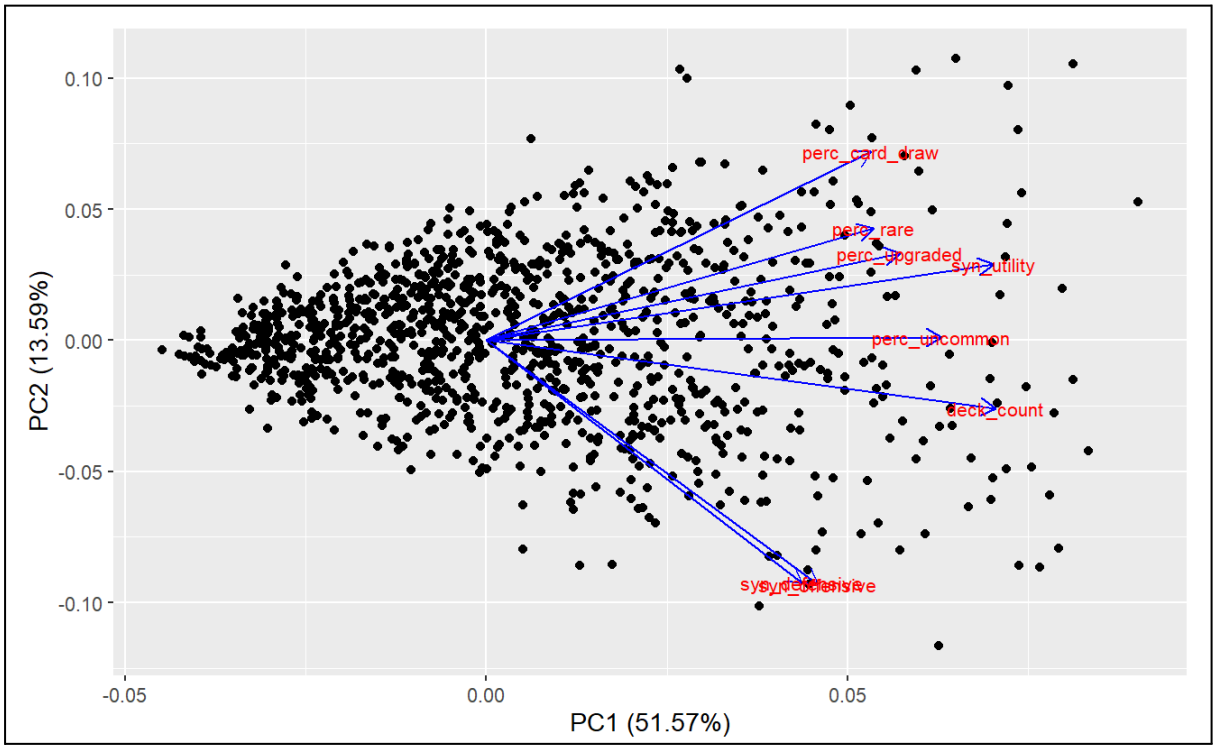


Figure 7. Final PCA projection plot with loadings from the reduced set of 8 features (perc_card_draw, perc_rare, perc_upgraded, syn_utility, perc_uncommon, deck_count, syn_offensive, syn_defensive).

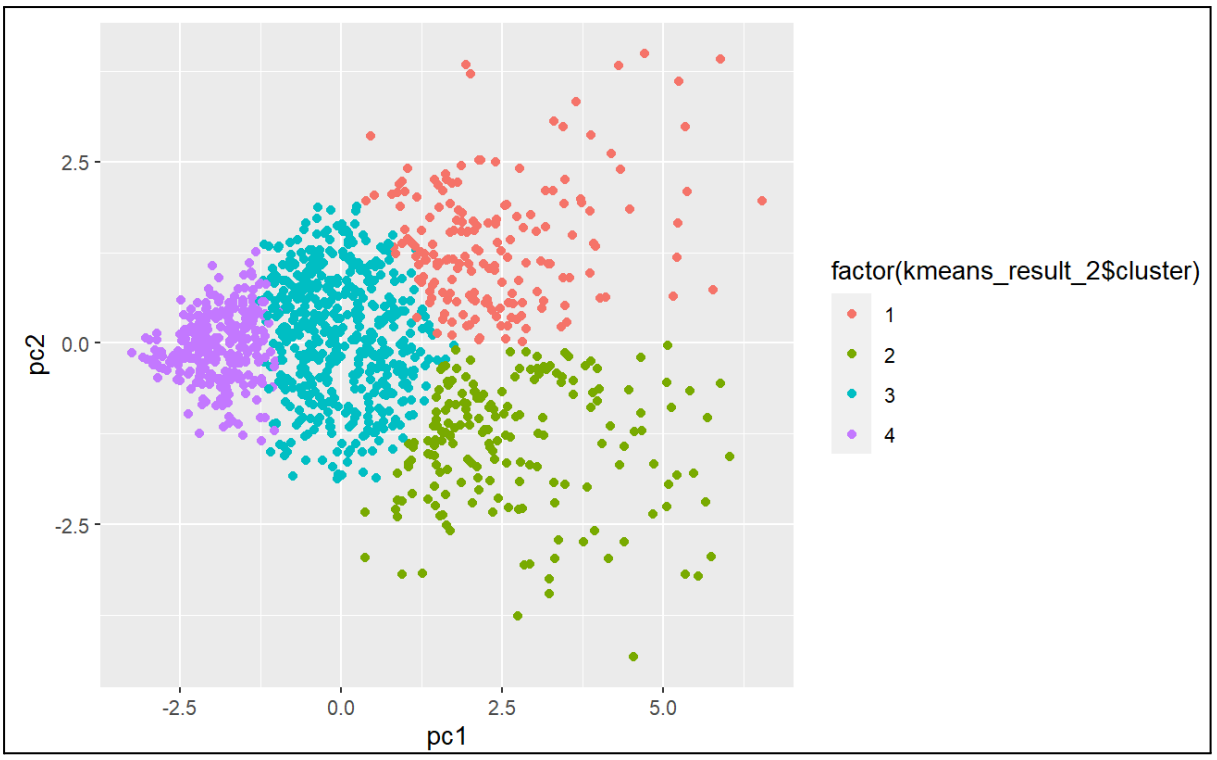


Figure 8. Final visual of the runs from the data set classified into the final set of k = 4 clusters, and colored to represent their assigned cluster.

Starting with these plots, we begin our characterization of each cluster based on these loadings and a profile consisting of: (1) the favorite player character, (2) popular cards, (3) popular relics, (4) average time per floor, (5) win percentage.

4.1.4: Player Personas

In this section we present our player personas that we derived from the analysis. The insight from this section is based on interpretations of visuals presented in the Appendix, Section 7.1. We will put the figure that corresponds to the interpretation in parentheses after each observation.

4.1.4.1 Strategists

Strategists are the players that try to uncover the deeper levels of strategic gameplay to secure wins. They are motivated by discovery and weaving together intricate and deadly synergies. They tend to build their decks around rare cards, value card generation, tend to upgrade important cards when they can, and favor utility synergies (**Figure 7**). They tend to favor the Ironclad and the Watcher, both of which tend to have more difficult-to-uncover synergies (**Figure 31**). Their favorite cards are the likes of *Shrug it off, Flurry of Blows, Prostrate, Empty Mind,* and *Follow-up* **(Figure**

32). These cards typically indicate the player is going for a more nuanced strategy, as these cards are generally dependent on a synergized deck to be effective. Favorite relics like *Molten Egg*, *Juzu Bracelet*, and *Red Mask* indicate they value free upgrades to their synergized cards, they don't want to risk being caught off guard by a hallway fight, and value early enemy damage suppression while they enact their plan (**Figure 33**)**.** Alongside *Bursters*, they have the highest average time per floor, indicating they are thinking carefully about how to play the cards they draw (**Figure 34**). All of this careful planning pays off to the highest win rate of all the personas at 28.33% (**Figure 35**).

4.1.4.2 Bursters

Bursters are all about collecting a lot of high impact cards and crushing opponents with more obvious offensive and defensive synergies (**Figure 7**). They are motivated by asserting their will over the denizens of the Spire. They are Silents and Defects at heart, two spire slayers who are renowned for their high damage potential (**Figure 31**). Unsurprisingly, they prefer cards that make those offensive synergies pop, including *Deadly Poison*, *Noxious fumes*, and *Ball Lightning* **(Figure 32**). Their relic choice points even more in the direction of burst - including *Varja* for a base damage increase, *Bag of Marbles* for the vulnerability to make their hits hurt even more, and *Orichalcum* to protect them when they inevitably forget to play a defensive card (**Figure 33**)**.** As with *Strategists*, they spend a lot of time on each floor in the act of crushing their enemies (**Figure 34**), and deck bloat (**Figure 7**) may be the reason they can't clear out floors more quickly. But, come on, you can't pass up a good card! With a win rate of 23.91%, the second highest of the four clusters, it's hard to argue with all-out assault as a primary method of slaying the spire (**Figure 35**).

4.1.4.3 Quick Experimenters

Quick Experimenters may not necessarily be playing to win. They seem to be motivated by exploring all the cards the Spire has to offer them, and getting quick feedback on what works and what does not. They tend not to specialize too far into any one synergy, offensive, defensive, nor utility (**Figure 7**). They don't have a ton of upgraded cards nor rare or uncommon cards (**Figure 7**). They typically play the Ironclad and the Defect, two spire slayers whose synergies are not obvious, so experimentation is crucial to mastering them (**Figure 31**). Their experimental motivation is also supported by their relatively simple favorite card choices - Sure, *Flex* and *Heavy Blade* go well together, but *Shrug it off* and *Armaments* are generally cards that can be obtained early and rarely support a deeper strategy (**Figure 32**). The relic choice skews towards resource generation - gold from Golden Idol,

Maw Bank, and health from Regal Pillow (**Figure 33**), probably so they can afford more cards to experiment with**.** Time is of the essence with a relatively low average time per floor compared to *Strategists* or *Bursters* (**Figure 34=**)**.** These players just seem to be quickly testing things out and seeing what sticks. About 1.39% of the time, it does, and this quick dabbling turns into a win (**Figure 35**). Perhaps a quick thrill is what they are after, or maybe they are preparing for a more meticulously planned run in the future.

4.1.4.4 Gamblers

Finally, we have the *Gamblers*, or those motivated to playing quickly and seeing if luck is in their favor. They hardly obtain rare or uncommon cards, achieve card generation, or upgrade their cards (**Figure 7**). This player type *did not win* in our data set (**Figure 35**), and they also spent the *least amount of time per floor* (**Figure 34**). They play the Ironclad and the Silent primarily, though, they play all characters more evenly than any other cluster (**Figure 31**). Based on their favorite card choices, they seem to want to get the first synergy that falls into their lap and run with it, evidenced by their affinity for *Flex, Heavy Blade, Sword Boomerang,* and *twin Strike*, all of which are basic cards that support a strength synergy (**Figure 32**). With 4 out of their 5 favorite relics being non-combat relics, they rely on their wit and deck building skills instead of relics that make things easier in combat (**Figure 33)**. The Spire can certainly be a slot-machine if one chooses to play this way.

4.2 Common Card Combinations (RQ2)

This analysis aimed to identify which cards are frequently selected together and to explore what these combinations reveal about the player's strategic goals and priorities during combat encounters.

As part of our data collection process, we compiled a Master Deck dataset that captured the final state of each player’s card deck at the end of a run. This dataset includes key information such as the RunID, CardID, the order in which each card was acquired (DeckCardIndex), the name of the card (CardName), and whether the card was upgraded (Upgraded). An example of this dataset is shown in **Figure 9**.

4.2.1 Data Process for Decks

4.2.1.1 Transformation of the Data

	RunID	CardID	DeckCardIndex	Card	Upgraded
1	1	361	1	AscendersBane	0
2	1	227	2	Strike_P	0
3	1	227	3	Strike_P	0
4	1	227	4	Strike_P	0
5	1	227	5	Strike_P	0
6	1	225	6	Defend_P	0
7	1	225	7	Defend_P	0
8	1	225	8	Defend_P	1
9	1	225	9	Defend_P	1
10	1	226	10	Eruption	1
11	1	228	11	Vigilance	1
12	1	353	12	Doubt	0
13	2	361	1	AscendersBane	0

Figure 9. Database of MasterDeck.

We transformed the data, so each card becomes a column and each row a run. Each cell has a binary value of 0 or 1 if that run has that card in its deck by the end of it. As seen in **Figure 10**.

RunID	AscendersBane	Strike_P	Defend_P	Eruption	Vigilance	Doubt	Strike_R	Defend_R	Bash	WildStrike	Armaments	InfernalBla	Rampage	Carnage	FlameBarrier	Metallicize	Reaper	FireBreathing
1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0
2	1	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1
3	0	0	0	0	0	0	1	1	1	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	1	1	1	0	0	0	0	1	0	1	0	0
6	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
10	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0
11	1	0	0	0	0	0	1	1	1	0	0	0	0	0	0	0	0	0
12	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0
13	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
14	0	0	0	0	0	0	1	1	1	0	0	0	0	0	0	0	0	0
15	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0
16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
17	0	0	0	0	0	0	1	1	1	0	0	1	0	0	0	0	1	0
18	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
19	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0
20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 10. Transformation of Data for processing step.

To prepare the data for analysis, we transformed the dataset so that each card became a column and each row represented a single run. Each cell contained a binary value, 1 if the card was present in the final deck for that run, and 0 otherwise.

4.2.1.2 Supervised Learning and Association Matrix

Following this transformation, we implemented a procedure to explore associations between cards. For each card, we constructed a linear regression model in which the selected card served as the independent variable, while the presence of all other cards acted as dependent variables. This allowed us to observe how the presence of one card related to the presence of others across different runs. We stored the resulting

regression coefficients in a matrix, where each cell captured the strength and direction of association between two cards.

We then visualized this matrix to highlight card associations, generating two versions of the visualization: one showing the top 30 associations and another showing the top 50 (**Figure 11 and 12**). The coefficients in these matrices ranged from -0.3 to 0.3. Positive coefficients indicate that two cards tended to appear together in the same deck, while negative coefficients suggest that the cards rarely co-occurred. A limitation of this approach is that the linear model relied solely on binary values to indicate card presence, which may restrict the nuance of the associations captured.

In the Appendix there is also **Figure 36** that is the same as **Figure 12** but on a higher scale.

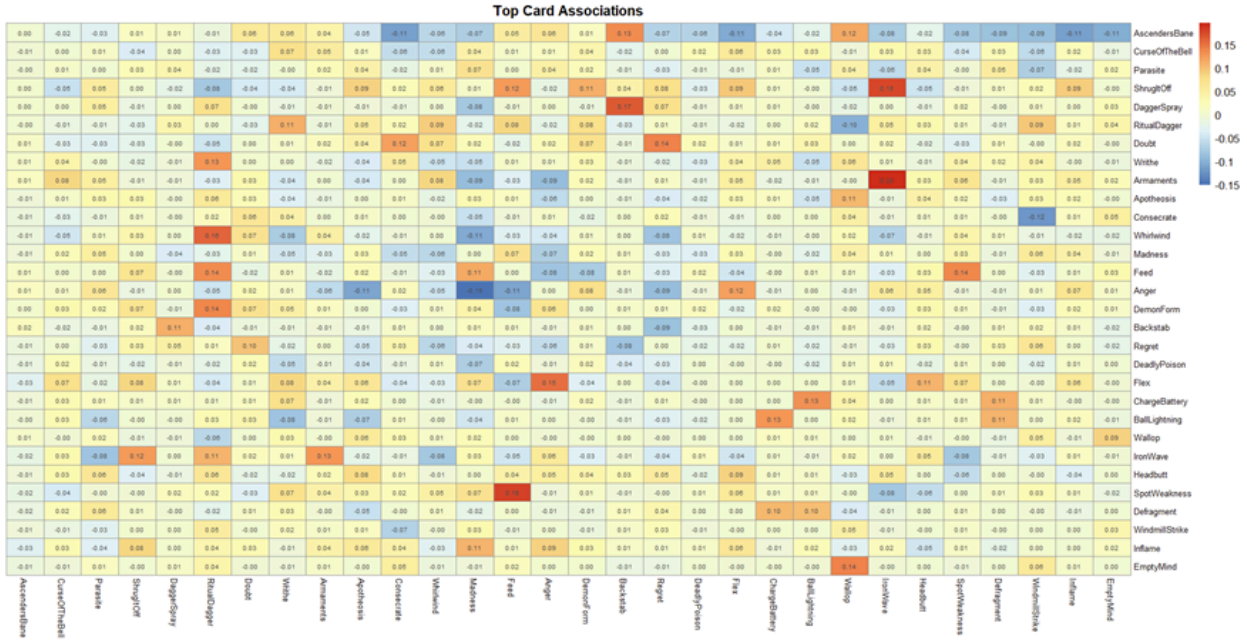


Figure 11. Association Matrix when n=30

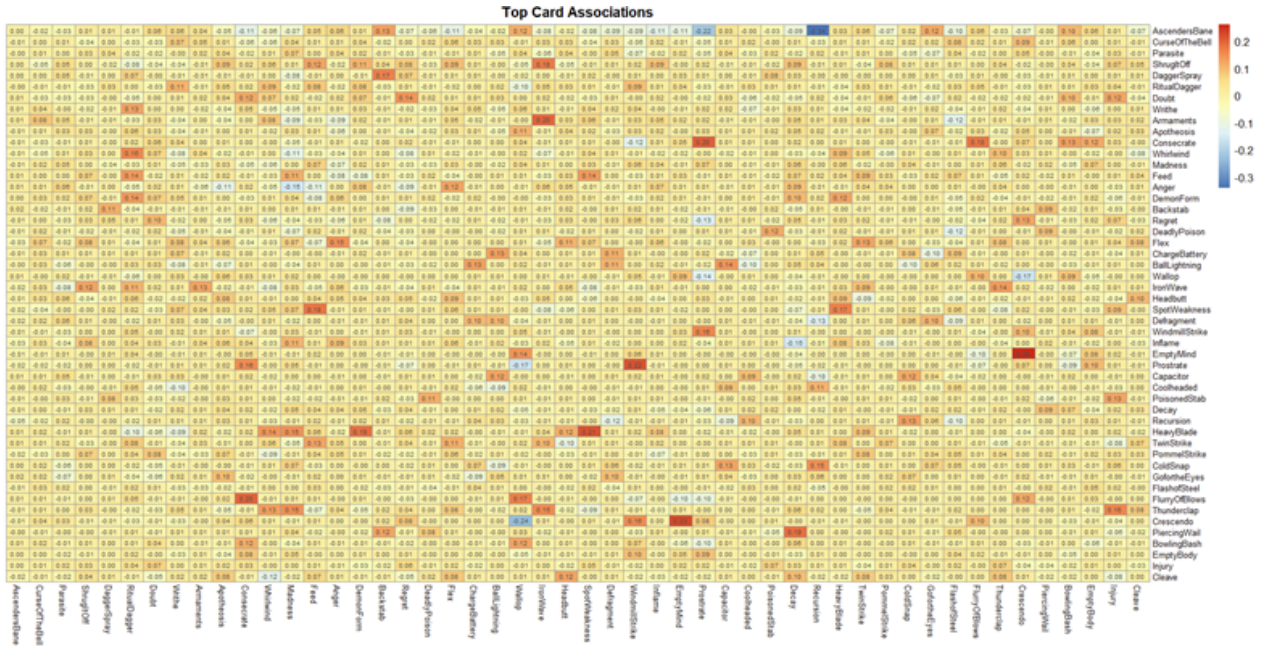


Figure 12. Association Matrix when n=50

Using the association matrix, we identified pairs of cards with strong correlations, which provided insight into patterns of co-selection and potential synergies between cards. To further investigate synergy, we selected two target cards known to work well together and applied logistic regression to model their relationship with other cards in the deck. In this model, the presence of the two selected cards served as predictors, and the outcome variable indicated whether the target card appeared in the deck. The resulting coefficient reflects the likelihood that the target card is included when both selected cards are present. This allows us to predict the probability of specific cards appearing in future decks based on known combinations.

4.2.2 Deck Synergies and their Objective

4.2.2.1 The Ironclad (Red Cards)

For the Ironclad, we predicted the cards of the deck using the most correlated cards. In these cases would be **IronWave + Armaments (0.20), SpotWeakness + Feed (0.18)** and **Anger + Flex (0.15)**. With these cards, we predicted the decks and which cards would appear in it. We graph these decks in Figure n to show them all at once. As seen in **Figure 13**, these are the cards that most appeared for the Ironclad.

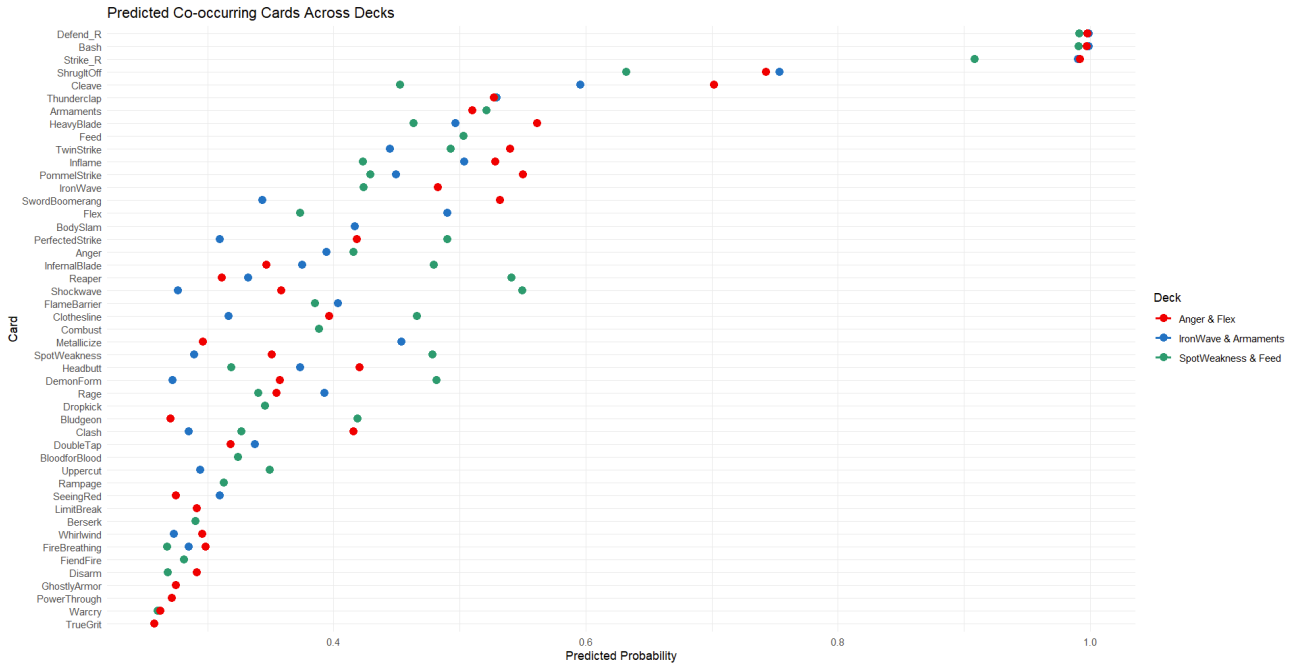


Figure 13. Ironclad decks predicted with the Correlated Cards

The red dots, Anger & Flex focuses on rapid, strength-based damage output through zero-cost cycling cards. *Anger* (0 cost, 6 damage) and *Flex* (temporary Strength boost) form the core of the deck, enabling quick and repeated attacks. Supporting cards such as *Cleave* (area-of-effect damage), *Sword Boomerang* (multiple random hits), and *Heavy Blade* (scales damage with Strength, especially effective when upgraded) amplify this synergy. Due to its straightforward and aggressive playstyle, this deck aligns with the **Quick Experimenter** persona, favoring fast, easy-to-understand strategies.

The blue dots, Iron Wave & Armaments emphasizes defense and sustainability. Core cards include *Iron Wave* (provides both attack and block), *Armaments* (upgrades cards and grants block), and others like *Shrug It Off* (block and card draw), *Metallicize* (gain passive block each turn), *Body Slam* (converts block to damage), and *Flame Barrier* (damage to attackers and block gain). The consistent block generation and damage mitigation reflect a methodical approach, which also fits within the **Quick Experimenter** persona, though with a more defense-oriented twist.

The green dots, the Spot Weakness and Feed follows a more complex strategic path. *Spot Weakness* provides Strength when the enemy intends to attack, and *Feed* increases the player's maximum HP upon defeating an enemy with it. It is supported by high-impact cards such as *Shockwave* (applies Weak and Vulnerable to all enemies), *Reaper* (deals damage and heals), *Demon Form* (grants Strength every turn), and *Bludgeon* (high

single-target damage). The combination of buffing the player and weakening enemies points to a long-term strategy. This playstyle is more aligned with the **Strategist** persona, reflecting deeper planning and optimization.

4.2.2.2 The Silent (Green Cards)

For The Silent, we predicted the cards of the deck using the most correlated cards. In these cases would be **Backstab + DaggerSpray (0.17), Piercing Wail + Decay (Curse) (0.19)** and **Poisoned Stab + Deadly Poison (0.13)**. Using these, we generated predicted decks and identified which additional cards were most likely to appear alongside them. These predictions are visualized in **Figure N**, where the decks are shown collectively to provide a comparative view of their structure and strategy. Two cards, *Neutralize* and *Survivor*, are consistently present across all decks. This is expected, as they are included in the starting Silent deck, and players tend to remove Strike cards before defensive options during runs.

The blue dots, **Backstab & Dagger Spray**, focus on fast-paced, low-cost burst damage. Backstab (0 energy, 6 damage) provides immediate value on turn 1, while Dagger Spray hits all enemies twice, allowing for quick clearing of early encounters. Supporting cards such as *Blade Dance* (multiple Shivs for free attacks), *Cloak and Dagger* (block plus Shiv), and *Backflip* (block and draw) enhance both offensive pressure and efficiency. *Poisoned Stab* appears frequently, offering fallback scaling through light poison damage. Due to its quick, aggressive nature and simple execution, this archetype aligns with the **Quick Experimenter persona**, appealing to players who favor intuitive, high-tempo strategies. As seen in **Figure 14**, these are the cards that most appeared for the Silent.

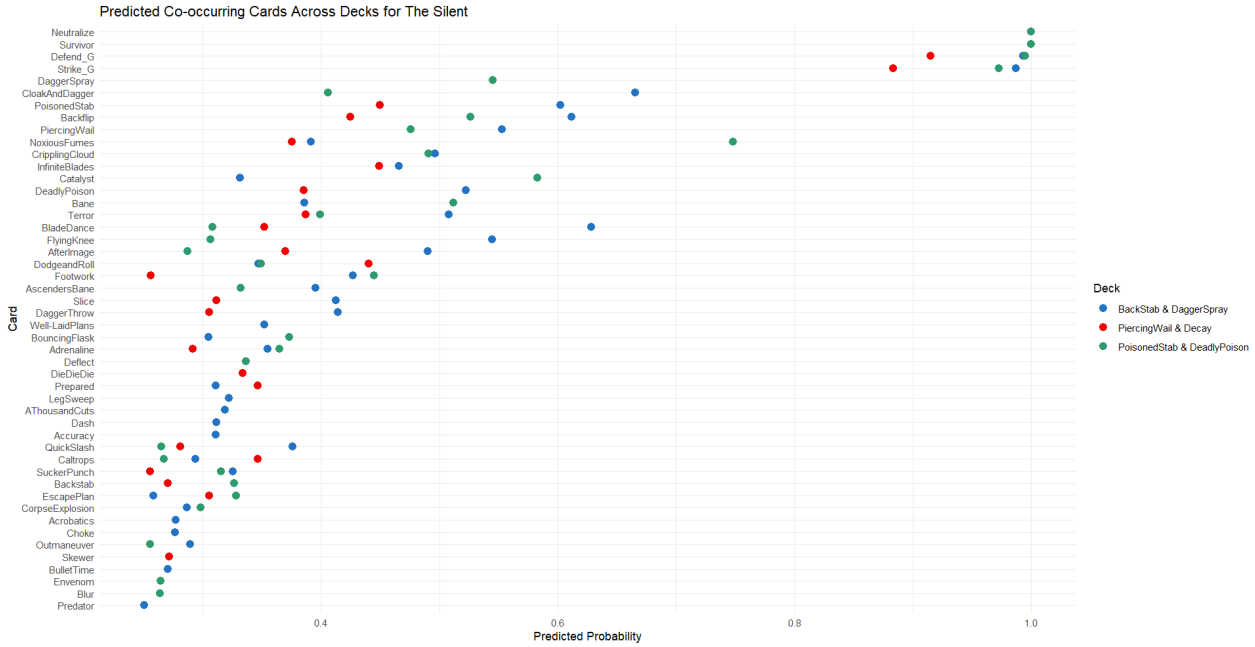


Figure 14. Silent decks predicted with the Correlated Cards

The red dots, **PiercingWail & Decay (Curse)** is an outlier here. This correlation represents the players who prefer to play with The Silent are more risk tolerant. Piercing Wail applies AoE Weak and reduces enemy strength, providing control during high-pressure turns. Decay, a curse that deals HP loss over time, is balanced out by cards like *Prepared* and *Backflip*, which help discard or draw through the deck efficiently. Additional defensive tools such as *Dodge and Roll* and *Cloak and Dagger* reduce incoming damage while enabling small attacks like Shivs. This deck reflects the **Gambler** persona, appealing to players who accept short-term risk in exchange for long-term advantage through careful planning and resource management.

The green dots, **Poisoned Stab & Deadly Poison**, form the foundation of a poison-based, damage-over-time deck. These cards inflict poison early, enabling scaling damage while the player focuses on defense and setup. Key synergies appear with *Noxious Fumes* (passively poisons all enemies), *Catalyst* (amplifies poison), and *Bane* (bonus damage against poisoned foes), creating explosive finishes after initial build-up. *Crippling Cloud*, which poisons and weakens enemies, further supports this debuff-based approach. With a clear strategy and strong synergy, this archetype aligns with the **Strategist** persona, appealing to players who prefer refined, long-term planning and setup for efficient enemy takedowns.

4.2.2.3 The Defect (Blue Cards)

For **The Defect**, we predicted full deck compositions using the most correlated pairs of cards. The three core card pairings were **Ball Lightning + Capacitor (0.14), Ball Lightning + Charge Battery (0.13) and Cold Snap + Recursion (0.13)** which stood out statistically and were used to model the rest of each corresponding deck. The cards *Dualcast* and *Zap* are consistently present across all decks. This is expected, as they are included in the starting Defect deck, and players tend to remove Strike cards before defensive options during runs. The Defect deck revolves around using orbs as a main output of damage, and these cards create this synergy from the beginning of the run. As seen in **Figure 15**, these are the cards that most appeared for the Defect.

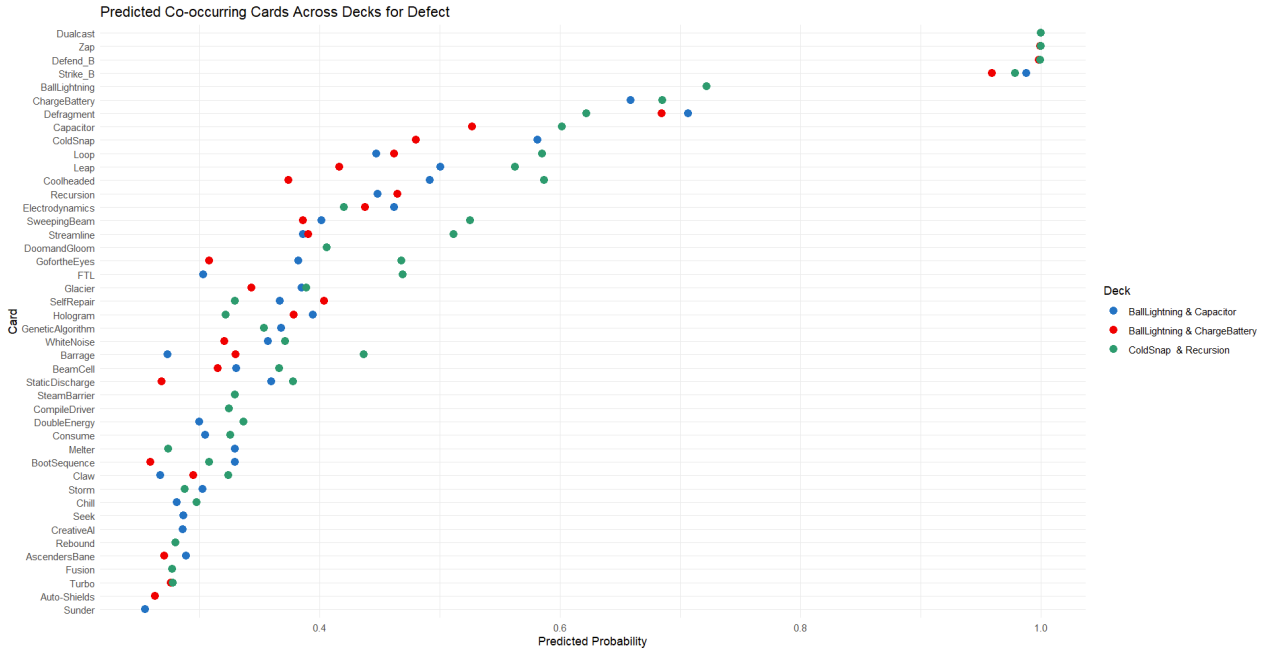


Figure 15. The Defect decks predicted with the Correlated Cards

The blue dots, associated with **Ball Lightning & Capacitor**, represent a deck focused on building and maintaining a strong orb engine. Ball Lightning enables early Lightning orb channels for passive chip damage, while Capacitor increases the number of available orb slots. This allows the player to build a sustained system of orbs, especially when supported by cards like *Hologram* (re-use card in Discard Pile), *Electrodynamics* (add lightning orbs and hit all enemies), *Defragment* (increase effect from orbs), and *ColdSnap* (deals damage and adds a defensirve orb). The deck typically relies on steady, repeatable turn value and passive damage scaling, aligning closely with the **Burster** persona, who favors slow-building synergy and long-term payoff over rapid or flashy damage.

The red dots, tied to **Ball Lightning & Charge Battery**, lean toward a more balanced archetype. Ball Lightning maintains early board control, while Charge Battery provides energy and reliable defense, creating smoother and safer turns. This deck often includes cards like *Loop, Defragment, Streamline Recursion,* and *Self Repair*, which help the deck deal with threats immediately while still building toward orb-based advantages. This configuration represents the **Quick Experimenter** persona, a player who values consistency, efficient turn execution, and responsive deckbuilding.

The green dots, representing **Cold Snap & Recursion**, highlight a defensive Combo-oriented archetype that revolves around Frost orb generation and cycling. Cold Snap offers a stable mix of defense and damage, while Recursion enables looping of powerful orbs. Decks in this group commonly

include *Glacier, Loop, Leap, Cool Headed* and *Static Discharge*, supporting heavier Frost synergy with delayed but reliable power curves. These decks align with the **Strategist** persona, which favors turn planning, synergies that pay off over time, and accurate resource control via orb cycling and defensive setup.

4.2.2.4 The Watcher (Purple Cards)

For **The Watcher**, the most correlated cards were **Empty Mind + Crescendo (0.26), Windmill Strike + Prostrate (0.22)** and **Flurry of Blows + Consecrate (0.20)**. As seen in **Figure 16**, these are the cards that most appeared for the Watcher.

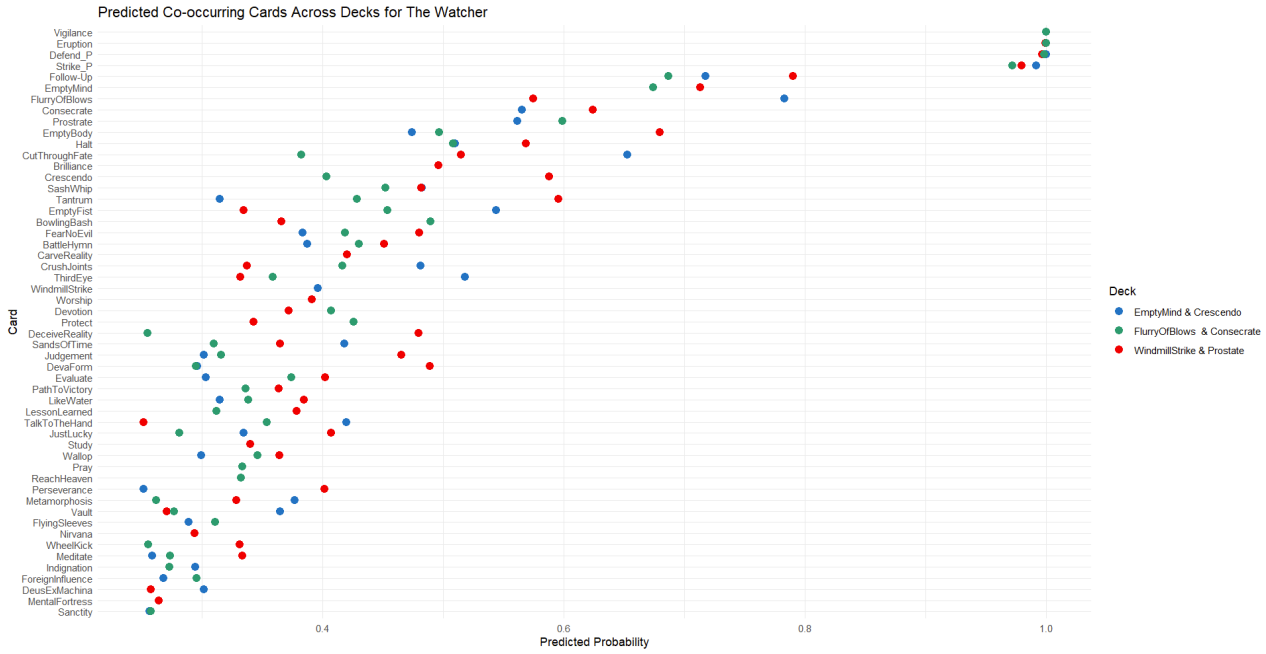


Figure 16. The Watcher decks predicted with the Correlated Cards

The blue dots deck of **Empty Mind + Crescendo** is an archetype focused on high card draw, fluid stance switching, and burst damage turns. *Empty Mind* (card draw + Calm exit) and *Crescendo* (enter Wrath) give the player explosive potential by transitioning into Wrath for a high-damage turn, then resetting into Calm to recover. *Cut Through Fate* (deal damage and draw cards), *Empty Fist* (deals damage and exit Stance), *Crush Joints* (deal damage and apply vulnerable) and *Third Eye* (block and draw cards) let the deck have a consistent damage output while pulling more cards and letting the player transition between Stances even more. This deck fits the **Burster** persona while taking advantage of the Wrath stance to maximise damage output done.

The green dots, **Flurry of Blows + Consecrate** deck is a low cost combo deck while changing stances and having a consistent output of damage. Cards like *Devotion* and *Pray* (gain Mantra), *Protect* (Blocks) and *Bowling Bash* (damage to all enemies) give this deck a way to deal damage, block receiving damage and let the player use the stances as better see fit and gives the deck versatility for any type of situation. This deck would fit the **Quick Experimenter** or **Strategist** persona with their all around cards and versatility but having a strategy for every turn.

The red dots deck, the **Windmill Strike** and **Prostate** gives the player a way to store damage and to gain Mantra, to reach the Stance of Divinity, which is a Stance that triple the damage done. Cards like *Judgement* (defeat enemy with HP below 30), *Deva Form* (gain Enery), *Deceive Reality* (gain Blocks), *Worship* (gains mantra) which gives the player versatility while also waiting for the right opportunity to act. This archetype aligns with the **Strategist** persona, a player who builds structured, delayed value and ultimately looks to execute one or two overwhelming turns. It rewards patience and foresight, waiting for the moment to unleash out.

4.2.3 Characters Chosen and their Cluster

In Figure x we can observe the clusters by the Character chosen by the player. Each character shows how its playstyle attracts different player **personas**, based on our clustering of behavior and preferred card synergy. As seen in **Figure 17**.

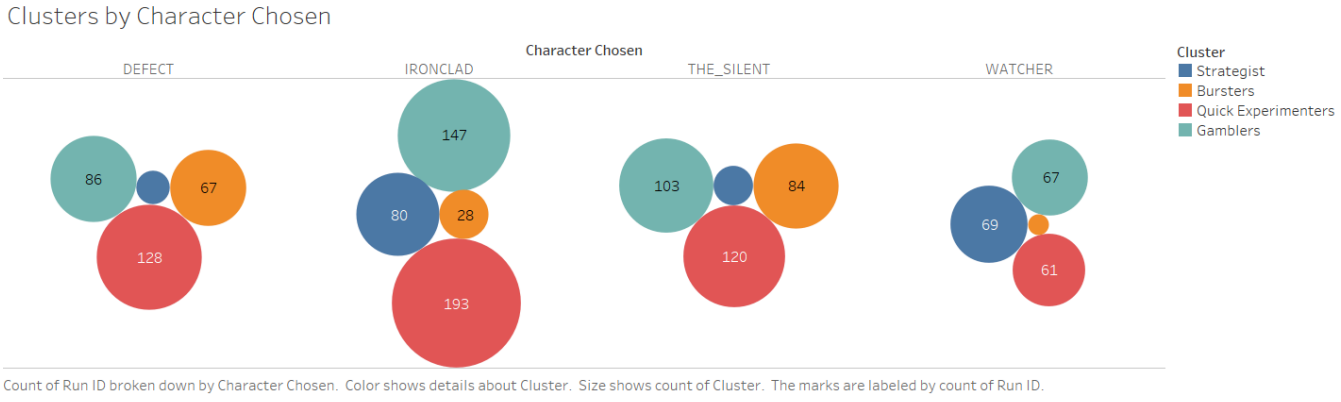


Figure 17. Clusters by Character Chosen

The Ironclad players are predominantly classified as **Quick Experimenters** and **Gamblers**. This aligns with the character's flexibility, offering multiple viable build paths such as strength scaling, block stacking, self-damage, and exhaust-based decks. The high-risk nature of some red cards, particularly those that inflict self-damage or rely on exhausting critical cards, may appeal to **Gambler**-type players, since the red cards can inflict

self-damage or exhaust cards there's a moderate risk of receiving too much damage or not having a card when the player needs.

The Silent players focus on mechanical experimentation and sneaky combo discovery while trying new combinations of synergies for the archetypes of **Quick Experimenter** or **Gambler**, or maximizing damage output as a **Burster.** Silent encourages players to explore dynamic combinations that reward flexible, reactive play. In this context, Quick Experimenters find an ideal environment to test synergy combinations, while Bursters may favor high-output cards.

The Defect players' appeals are around combinations but they are not deep to long-term planners. Likely, the **Quick Experimenters**, has a variety of cards ready due to its mechanics, like orb-cycling and power stacking, needing adaptation over execution. While the **Gamblers** and **Bursters** are a way to experiment with more damage oriented decks.

The Watcher is the most **Strategist** oriented. Unlike other characters, Watcher introduces unique mechanics, such as stances (*Calm, Wrath, Divinity*) and *Mantra*, that require precise turn planning and state management, the player has to focus more on strategies and thinking ahead to get the most of this character. This is reflected in the comparatively high presence of Strategists among character runs.

4.3 Risk (RQ3)

With this analysis we are also interested in determining the risk taking nature of players. In *Slay the Spire*, taking risks is inevitable given the game's mechanics of map traversal, however, for this question we are specifically looking at risk taking within random events, as our other analyses cover more data about the combat portion of the game. Understanding risk taking behavior is also important for improving the design for *Slay the Spire 2* as it can give insights into dominant strategies or what can be fun for certain players, i.e. playstyles.

4.3.1 Risk Classification and Context

In order to begin to answer this question, we first must define what exactly we mean when we say risk in terms of *Slay the Spire*. We classify a risk as an action that may give a reward to the player for potentially losing some resource or incurring a detrimental effect, whether random or not. Risk is also determined by how easily it is to mitigate the risk. Because of these definitions, generally purchasing something using gold is not considered a risk, but trading health for something is considered a risk, as health is a more precious resource and harder to restore. Additionally, for an event to

classify as a risk, the player must also have the choice not to incur the risk, so the player has to actively decide to incur said risks.

We also delineate each risky event into 1 of 3 different severities, which correlate to the amount of risk taken. Generally, the more risk taken the higher the potential rewards as well. The lowest risk events are labeled as Minor Risk, which usually do not include hidden information and have risks that are relatively easy to mitigate. Next are Standard Risk events, which are more risky as they usually include some element of random number generation (RNG) or hidden information. Many of these events add curses to the player's deck, which negatively impact their decks composition. Finally, we have Major Risk events, which have severe downsides such as losing Max Health.

4.3.1.1 Events

Slay the Spire events are classified into 4 different categories. Common, which can appear in every Act, as well as 3 Act specific categories, that can only appear in Acts 1, 2, or 3 respectively. In total, there are 52 unique events in *Slay the Spire*, with 16 being Common, 12 in Act 1, 16 in Act 2, and 8 in Act 3. However, only 31 events have been determined to carry some sort of risk with them.

With our risk classification model, there are 6 Minor Risk events, 16 Standard Risk events, and 9 Major Risk Events

Finally, because events are separated by Act, due to the nature of the game the earlier a risk is incurred the more chances it has to become impactful to the outcome of a run. This is offset by the design of these events, with earlier Acts generally having lighter costs associated with the risky options and later Act events having steep costs. Additionally, curses have a sort of opposite effect where if taken early, the player has more opportunities to remove said curse and mitigate the downsides, depending on the severity of the curse itself.

While still arbitrary, these considerations were also taken when classifying events, but these classifications still serve to better understand the data in a meaningful way.

4.3.2 Visualizations

In order to fully capture the risk taking behavior of players, we have created several different visualizations taking a look at the data from different angles as well as a case study on a specific event. The visualizations are also split up by Act, as later acts have less data due to the nature of the game and less runs making it all the way to the end of the game. We will

also take a look at specific events that are of particular interest or have a large impact on a run.

4.3.2.1 Overall Risk Taking

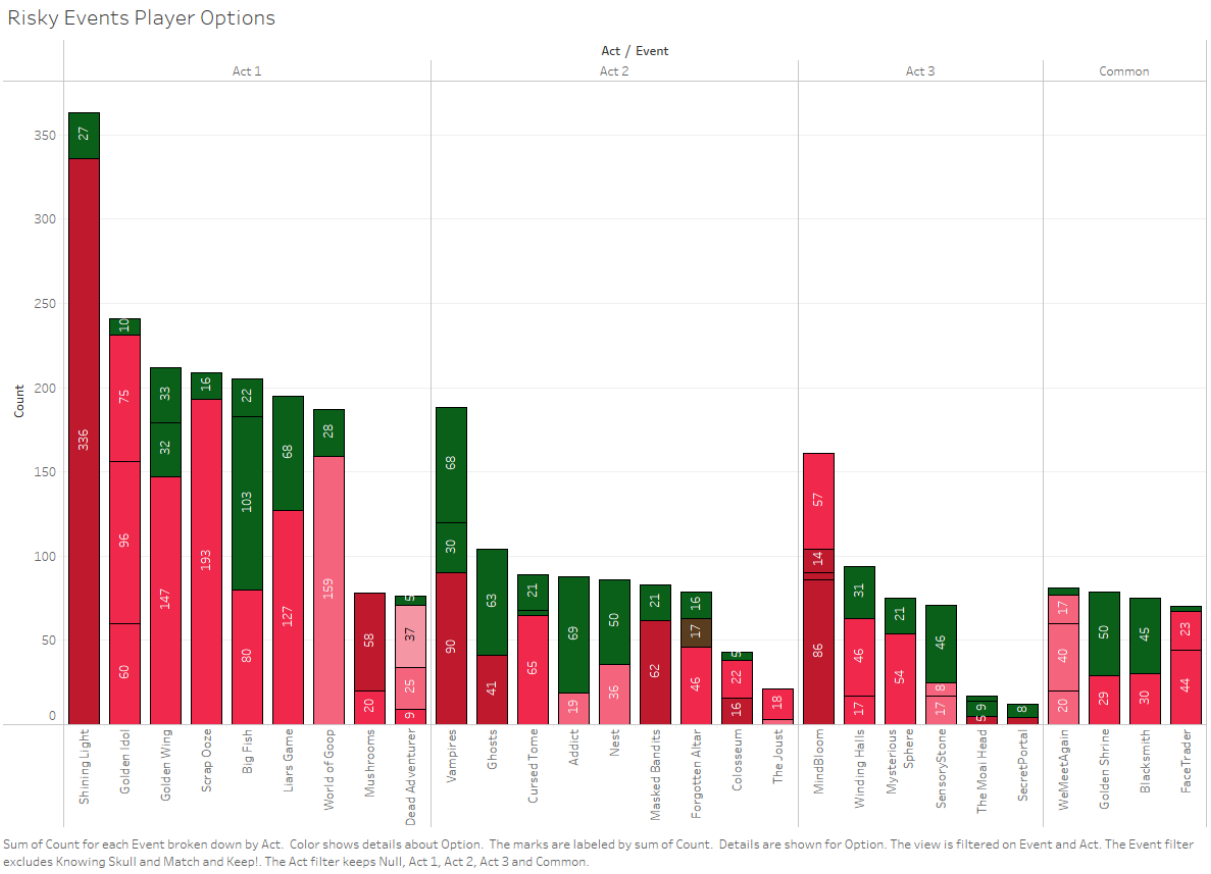


Figure 18 All Risky Events in *Slay The Spire* with different options delineated by Color. Green represents safe options, while Red represents risky options. The darker the shade, the more severe the risk. Ordered by Act and Count

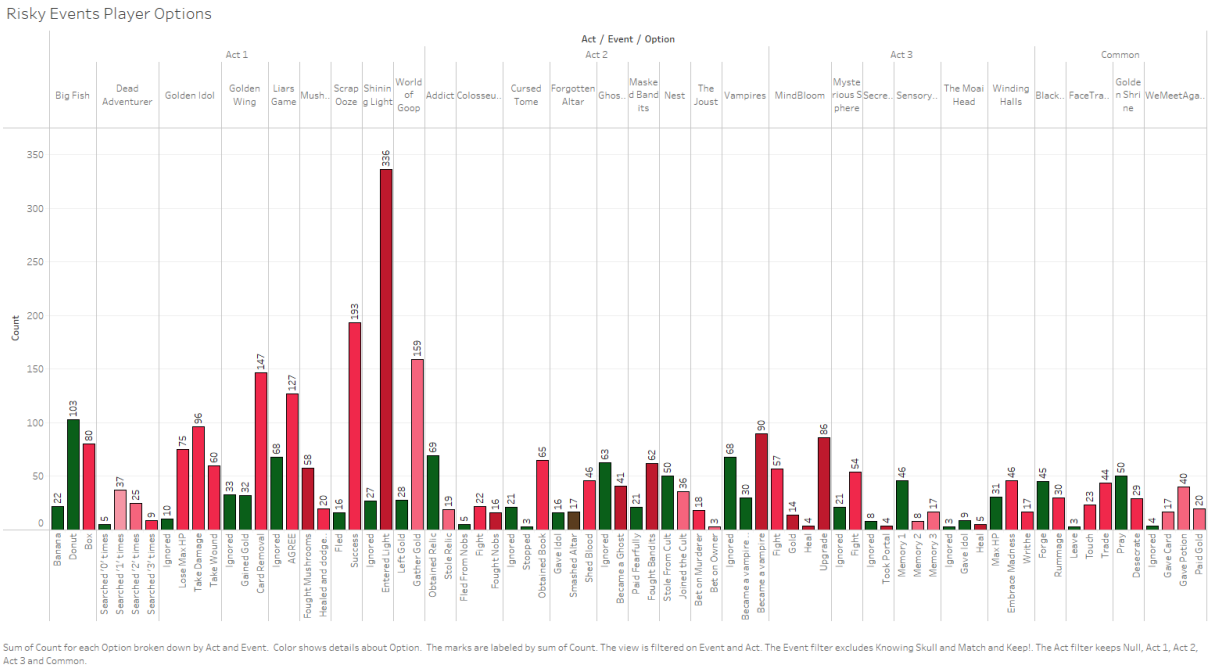


Figure 19 Another version of **Figure 18**, but each event has the individual options separated for more clarity. Some low data counts omitted.

Figure 18 gives an overview of the risky options present in the game. While this is not a great visualization for a deep analysis, it does have some key findings that will help for the rest of the analysis. For starters, it makes clear that there is significantly more data for earlier acts, as all runs will provide data for at least Act 1, but for later acts less runs make it that far. As such it is important for more focused visualizations focusing on specific acts. Second, it makes clear that most of the time, players will choose the riskier option over the safe option. This has some variance due to the exact details of each option, but overall it trends towards a more risk taking playstyle across the board. This makes sense in the context of domain expertise, as many of the additional benefits of the riskier options greatly can help a run if used correctly. In addition, these kinds of choices add to the “fun” of the game, introducing more variety and nuance to a run. In order to gain further insights, the next visualization will be a mapping of risks taken vs not taken based off of the floors traveled.

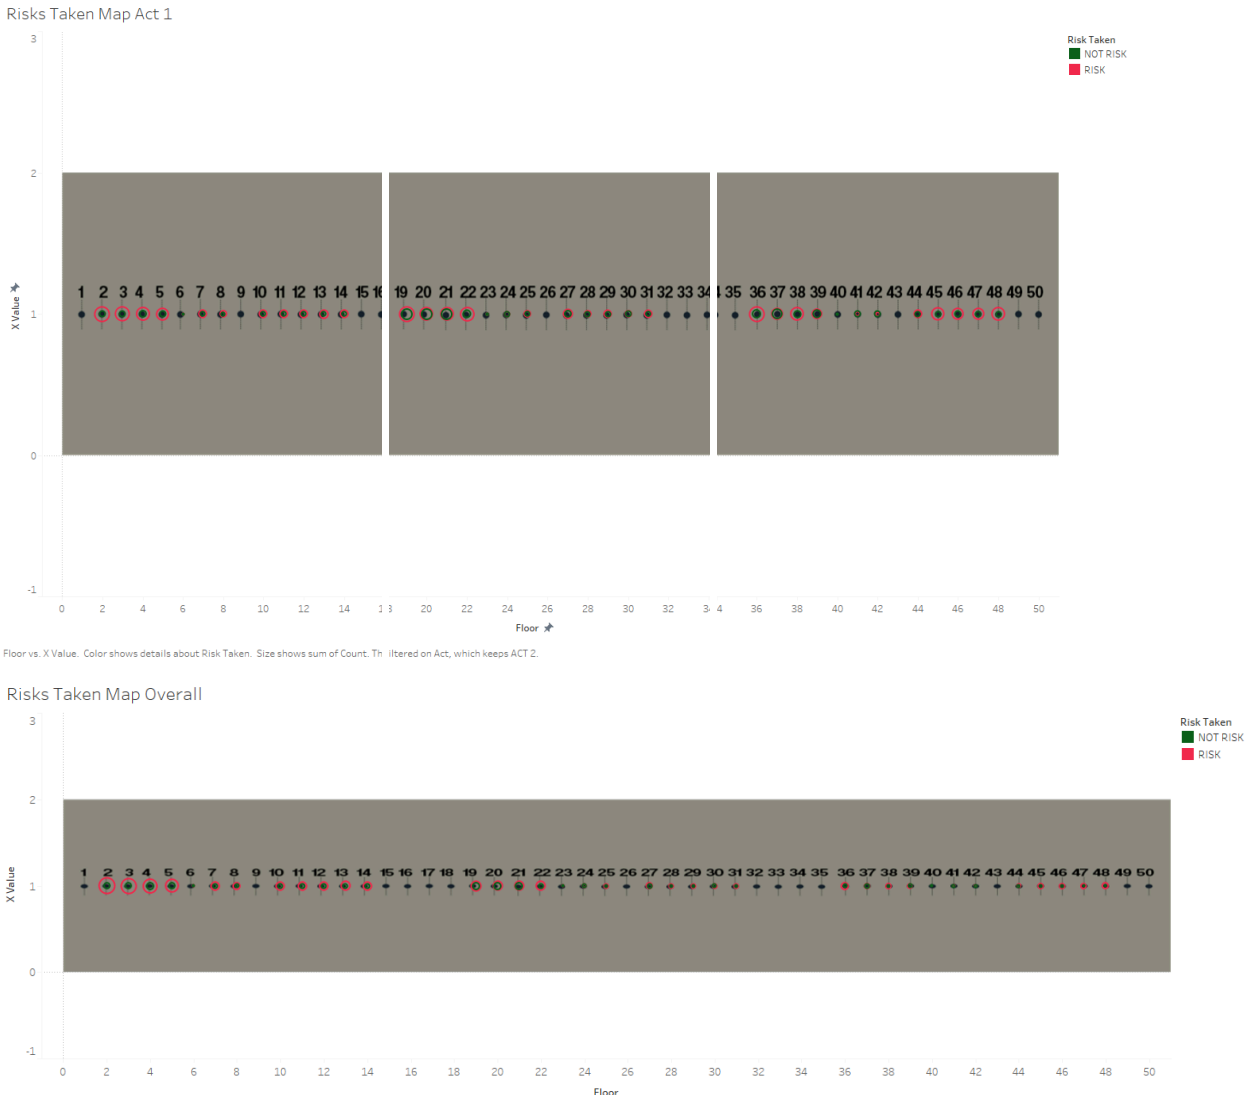


Figure 20 and 21 Floor map of all risks taken. with an overall map as well as a map for each Act. Colors indicate whether a risky option was taken or not, and circle size indicates number of options taken.

Another perspective to take a look at is the overall risk taking nature of players. **Figures 20 and 21** outline this by showing the total risks taken for each floor. This uses a cumulative total of all risky events for a given floor. Any gaps in floors signifies a scripted encounter, using a resting place, treasure chest, or boss. All other floors have the potential to be a random event. **Figure 20** shows the data on a per act basis, which scales the sizes of the circles more appropriately due to there being less data for later acts. As shown here, there are generally more risk taking behaviors over the safe options, even between different acts. From our domain knowledge, this makes sense as sometimes the rewards from one event can counteract the risks from another event. Stronger runs also generally tend to take more

risks, and if these runs do not make it far enough then they do not leave data for the later acts.

Figure 21 also highlights how risk taking takes a slight backseat to the safer options in Act 2 specifically. This could be due to several different reasons, such as the design of events specifically in Act 2 being less favorable, to the overall flow of the game, with Act 2 being a critical juncture in the success of a run.

4.3.2.2 Unique Cases

There are several events included in the above visualizations that have strange quirks about them, be it how the data is present in the database or how the event is designed.

1. Knowing Skull
    a. This event lets the player choose how many times they pay life to get rewards, and its final output is recorded as the sequence of choices the player makes. This makes each choice have a very low sample size, as there are essentially infinitely many choice combinations a player could make (given infinite resources)
2. Cursed Tome
    a. Similar to Knowing Skull, however it has a set end point. The player can choose to stop midway, but they will not receive any reward by doing so. Data is recorded as when the player left the event, but the one piece of data missing is how much HP a player lost if they stopped early. However, this does not impact the data much as very few players stop early.
3. Dead Adventurer
    a. This event has a random chance to spawn a difficult fight, which is the entire risk of the event. Data is recorded as the number of times a player searched, with the odds of the bad outcome happening increasing per search.
4. Match and Keep
    a. A random minigame where the player matches facedown cards to add to their collection. If played correctly, the player can essentially choose which cards to add, but there is a very small chance they obtain a curse without being able to control the outcome. Data is recorded as the number of cards matched, which does not necessarily correlate to any risk.
5. Mind Bloom

    a. This event features all options as being risks, but each option has a different tradeoff.
6. Scrap Ooze
    a. Another event where the player can do it multiple times, although this time it is random when the final reward is given. The data simply shows whether or not the final reward was gotten.
7. Neow
    a. This event is unique in that it only appears at the very start of a run and has different options depending on how far your last run got. In particular, if given the additional options, 2 of these represent significant risk that greatly change how you play your run, and it being at the very start of the run compounds this risk. However, we do not have event data corresponding to this event.

4.3.2.3 The Joust

The Joust is a specific event that features a strict risk reward option to the player. How it works is that there are two options where the player bets the same amount of gold (50). There is a 70% chance to win 100 gold with the 1st option or a 30% chance to win 250 gold with the 2nd option. These values are not shown to the player explicitly, but are told implicitly through context clues. From these values, the expected value decision would be to always choose option 2, however, the player only has 1 chance to bet.

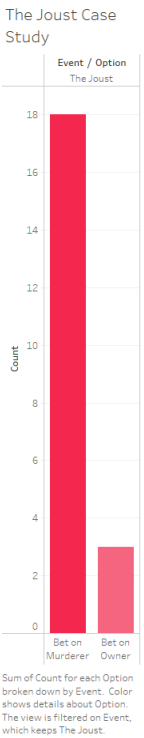


Figure 22 The Joust Case Study. Number of times each option has been chosen in the dataset.

As we can see in **Figure 22**, the majority of players choose the first option. This is interesting because of the odds presented, but in reality the player does not know the actual odds. Interestingly, the safe option is chosen most of the time, rather than the risky option which would be expected from both the overall risk taking nature of players as well as the expected value of the second option being higher. Another factor could be that because the options are opaque, players tend to default to the 1st options rather than the second option.

4.3.2.4 Personas

We can also analyze risk taking using the previously defined Personas from Section 4.1.4. Each run that is recorded has their associated persona, and we can correlate this to the risks they've taken throughout a run.

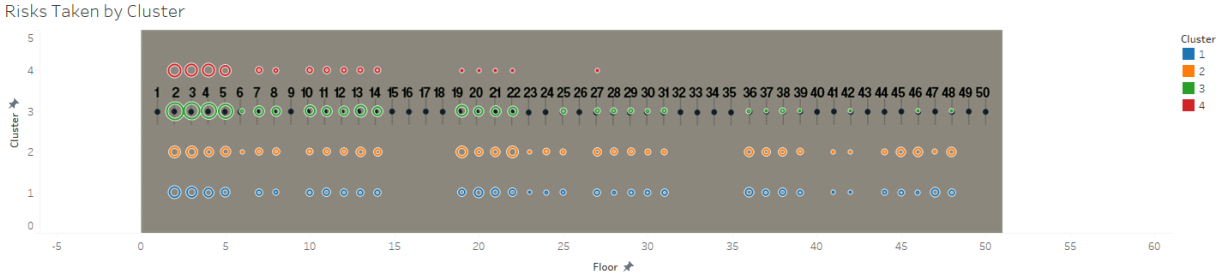


Figure 24 Shows all instances of Risk taking in the dataset, separated by Persona

Figure 24, shows how Personas differ in their risk taking habits. A couple interesting details jump out immediately. One is the **Gambler** persona. From even the name it would be expected that the **Gambler** would be more risk prone, but in reality these builds usually do not last long and are seemling also less popular. As such, there are very few data points and the data points start to dwindle the higher floor you go. Another less extreme example of this is the **Quick Experimenter.** While they take a lot of risk in Act 1, more so than any other cluster, they seemingly take less risks in Act 3 comparatively as well. Finally, our last 2 personas **Busters** and **Strategists** have very similar risk taking natures.

4.4 Resource Behavior (RQ4)

Another subject for us while making analyses was monitoring players' usage patterns for resources and resource handling - such as potions, health,

gold, and items, as victory is also determined by effective resource management. Here we analyze players' resource behavior via data-driven insight and describe results in terms of four player types defined earlier (Strategists, Bursters, Quick Experimenters, Gamblers). We focused our attention on methods (data aggregation, PCA, clustering) and drew conclusions on how different playstyles consume or conserve resources.

4.4.1 Data Aggregation

To get there, we aggregated a set of metrics for each run capturing how the player managed key resources. These metrics include things like consumable usage and investment decisions made during the run, for example:

- **Potions**: the number of potions obtained during the run (potions "spawned") and the number of potions actually used. The difference indicates whether players tend to hoard potions or use them.
- **Gold**: the total amount of gold gained (and indirectly how much was spent). High gold gained combined with low spending would imply resource hoarding, whereas spending gold at shops reflects investment in cards or relics.
- **Items**: the number of items purchased (e.g. cards, relics, or potions bought from merchants) and items purged (cards removed from the deck at cost). These indicate how actively players manage the resources to strengthen or streamline their deck. We also consider the net items gained (purchases minus purges) as a measure of how much a deck grew via shop transactions.
- **Health Management**: various measures of health, including the average percentage of max HP remaining after combats, the minimum HP % reached (a proxy for how close to death the player came at worst), total HP lost over the run, and the variance in HP loss between encounters. These features reflect how safely or aggressively players navigate fights. For e.g. a low minimum HP and high total HP loss suggest risking health as a resource to progress, whereas consistently high HP implies a cautious approach.

Together, these features capture the nuances of a particular run's resource profile. We also recorded character class being used in all the runs, in order to check if certain resource patterns were character-specific.

4.4.2 Principal Component Analysis (PCA)

Considering that some of the resource-related variables could be correlated, we applied PCA to reduce dimensionality and reveal the key

dimensions of resource behavior (Drachen et al., 2019). We decided to check the PCA scores for 16 features initially. After analysing the PCA scores and the graph, we decided to go ahead with the first 6 PCs that covered a substantial portion of variance (80%) (Jolliffe, 2011).

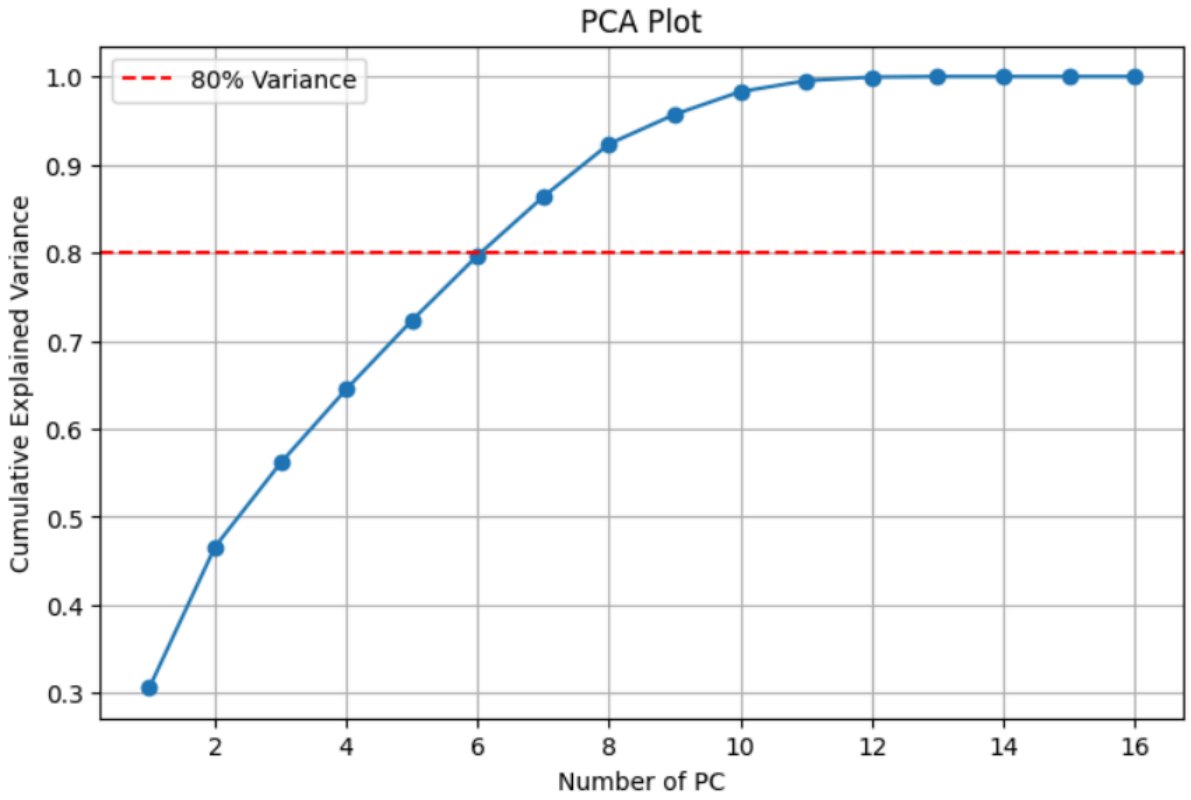


Figure 25 PCA Plot

	PC1	PC2	PC3	PC4	PC5	PC6
Potions Spawned	0.367871	0.275622	0.027304	-0.015949	0.028146	0.049272
Potions Used	0.350065	0.282407	0.021697	0.113760	-0.204994	0.230284
Potion Diff	-0.033254	-0.066508	0.006388	-0.265030	0.476463	-0.383476
Gold	0.042884	0.046355	-0.032725	-0.085753	0.043695	-0.385638
Items Purchased Count	0.327369	0.336100	-0.005203	-0.061435	0.100053	-0.230619
Items Purged Count_x	0.217424	0.185555	0.037700	0.102587	-0.091550	0.341889
Items Diff	0.270385	0.292298	-0.020361	-0.106708	0.144306	-0.383701
Avg_HP_Percentage	-0.273086	0.439393	0.167901	-0.025449	-0.033351	-0.017640
Min_HP_Percentage	-0.357841	0.294418	0.009190	-0.053439	-0.027935	0.004295
Avg_HP_Loss	0.294831	-0.437707	-0.070567	-0.010245	0.017462	0.008872
Total_HP_Loss	0.406956	0.015733	-0.029441	-0.035990	0.029669	0.045678
HP_Loss_Variance	0.235187	-0.354655	0.180214	-0.006562	0.007772	-0.053053
Class_DEFECT	0.012156	-0.019657	-0.316048	-0.619773	-0.503209	-0.046336
Class_IRONCLAD	0.018654	-0.052538	0.783509	0.091408	-0.069622	-0.111911
Class_THE_SILENT	-0.023345	0.025425	-0.446695	0.667571	0.014531	-0.250644
Class_WATCHER	-0.010535	0.060968	-0.126101	-0.201130	0.653932	0.498667

Table 2 Loadings Table

The first two principal components capture a large part of player behavior, mainly highlighting patterns of overall resource consumption, like potions used, gold accumulated, and damage taken (HP loss). Players seem to naturally separate along these axes into groups who aggressively spend and use their resources versus those who cautiously preserve and manage their health, gold, and potions.

Interestingly, class-based features like Ironclad, Silent, Defect, and Watcher begin to appear prominently in later components (PC3 to PC6). This suggests that after the general resource strategies are accounted for, there's still significant variation tied specifically to the character chosen. Silent and Watcher in particular stand out distinctly in PC4 and PC6, reflecting unique play styles such as deck optimization and potion efficiency.

Overall, the **six principal components** effectively summarize about 80% of the variation in resource management behaviors (GeeksforGeeks, n.d.). The scree plot supports this, showing that after around six components, each additional component explains significantly less variation. This implies that the core differences between how players handle their resources are largely captured in these first few components.

This PCA strongly suggests that while resource management in Slay the Spire is nuanced, a handful of behavioral dimensions largely govern player strategy. Potions, gold usage, deck building choices, and class-specific traits are central to understanding player differences. This foundational insight provides a strong basis for further clustering and detailed comparisons across player personas and character classes.

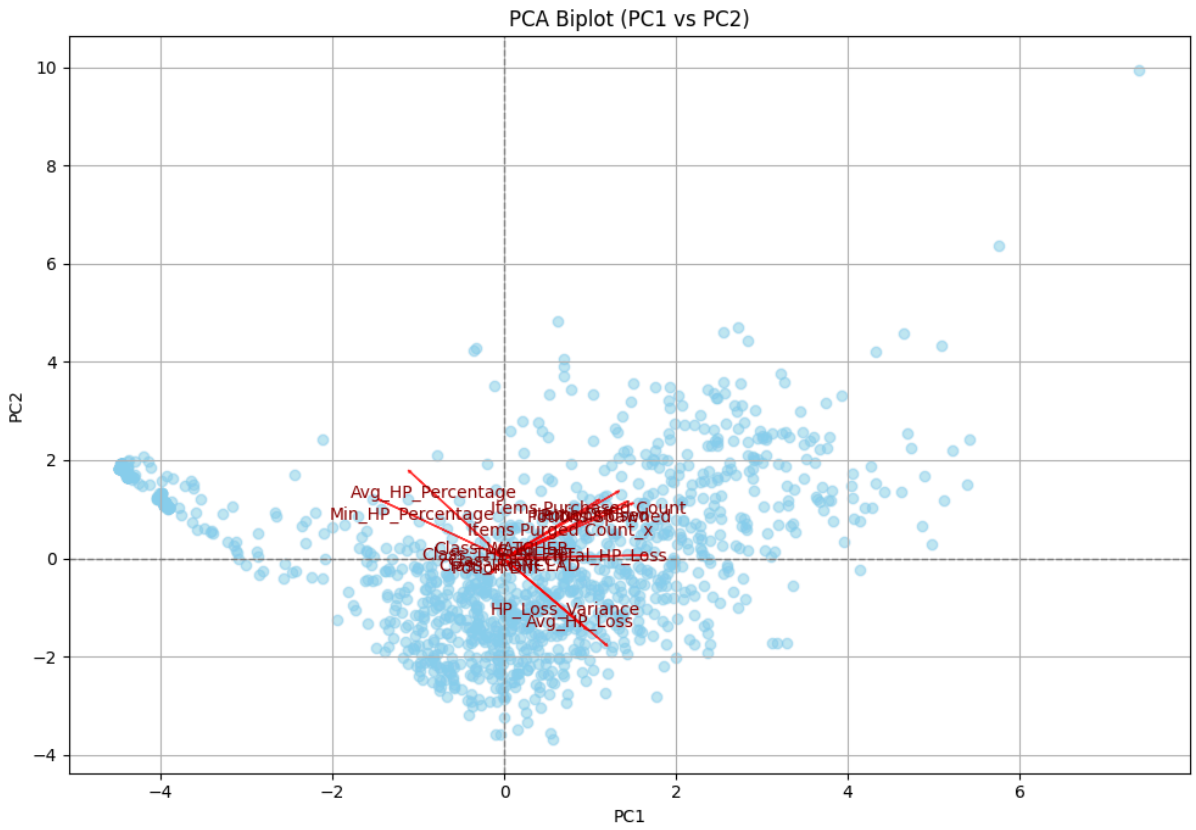


Figure 26 PC1 vs PC2 Biplot

4.4.3 Clustering Resource Profiles

Using the transformed and standardized PCA data, we performed unsupervised clustering to group runs with similar resource behavior. We chose the K-means clustering algorithm, which is commonly used in game analytics to segment player behavior dataarxiv.org. K-means partitions runs into clusters such that each cluster has a characteristic centroid (mean profile) representing a particular style of resource usage.

To determine an appropriate number of clusters k, we employed the elbow method (Figure 27), examining the within-cluster variance for cluster solutions. The elbow plot suggested diminishing returns beyond about 3–4 clusters – there was a sharp drop in variance going from 2 to 3 clusters, and a more gradual improvement thereafter. We ultimately opted for a 3-cluster solution as the partition between them seemed more natural and the least noisy.

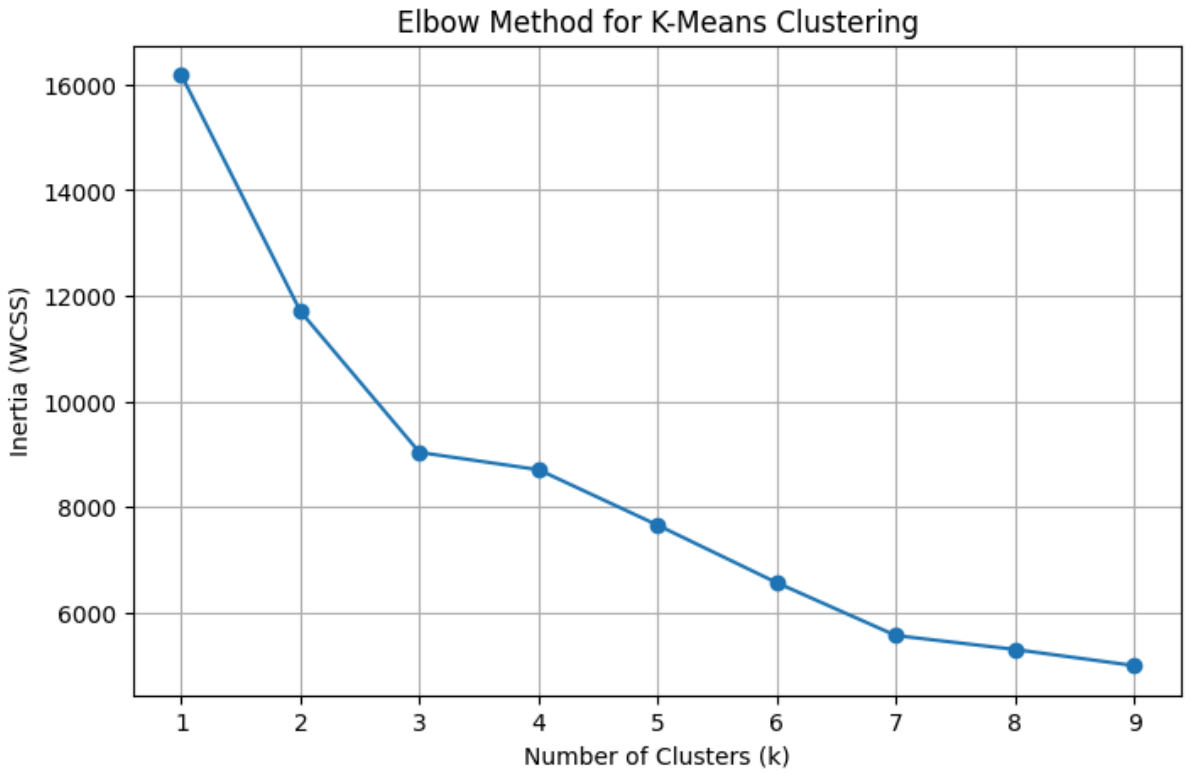


Figure 27 Elbow Method K-Means

Based on our clustering results using K-Means with three clusters, we see clear, intuitive patterns emerging in how players manage resources in Slay the Spire. The clusters effectively break down into distinct playstyles, each characterized by its own unique approach to managing health, gold, potions, and deck composition.

Cluster 0 represents players who engage heavily with potions and items, as indicated by their high average of nearly 40 potions spawned and used, and around 5-6 items purchased per run. They tend to maintain relatively stable health but incur significant cumulative damage overall, suggesting a more aggressive style that balances risks with proactive resource usage. This cluster evenly spans across character classes, though slightly favors Ironclad and Defect, showing that aggressive resource management is a viable strategy regardless of class.

Cluster 1 is strikingly different, highlighting cautious and minimalistic players. These players spawn and use very few potions, rarely purchase or purge items, and maintain an exceptionally high average health percentage throughout their runs (about 92%). Such minimal resource interaction, coupled with low average damage taken, indicates a playstyle focused strongly on defensive strategies and careful resource preservation. This

approach is again quite evenly distributed across all four classes, suggesting that conservative play is universally effective.

Cluster 2 captures an intermediate group. These players spawn and use potions moderately (around 15 per run), purchase few items, and generally experience significant health fluctuations, as reflected by their lower average HP percentage and relatively high average damage taken per floor. This suggests a middle-ground approach, combining aspects of both resource spending and preservation. Again, the class distribution is fairly balanced but slightly favors Ironclad and Defect, possibly indicating these classes naturally support a mixed resource-management style.

Based on these patterns, we can tentatively label the clusters for easier reference. **Cluster 0** could be seen as the **"Active Spenders"**, players who fully engage with the game's resource systems, leveraging potions and shop items to push through challenges while accepting moderate HP loss. **Cluster 1** fits the profile of **"Efficient Conservers"**, who manage to preserve health and minimize spending, relying instead on core deck strength and careful decision-making. Lastly, **Cluster 2** could be called **"Balanced Adapters"**, blending resource use with risk, showing flexibility in strategy that adapts based on moment-to-moment demands. These names help clarify the practical differences between the clusters and offer a foundation for future behavioral tagging or player-type personalization.

Overall, we can observe that the patterns are not strictly dictated by class choice alone, indicating that individual playstyle choices significantly impact resource usage strategies. These resource-based archetypes aligns with the type of player modeling approaches discussed in prior game analytics literature (Drachen et al., 2019)

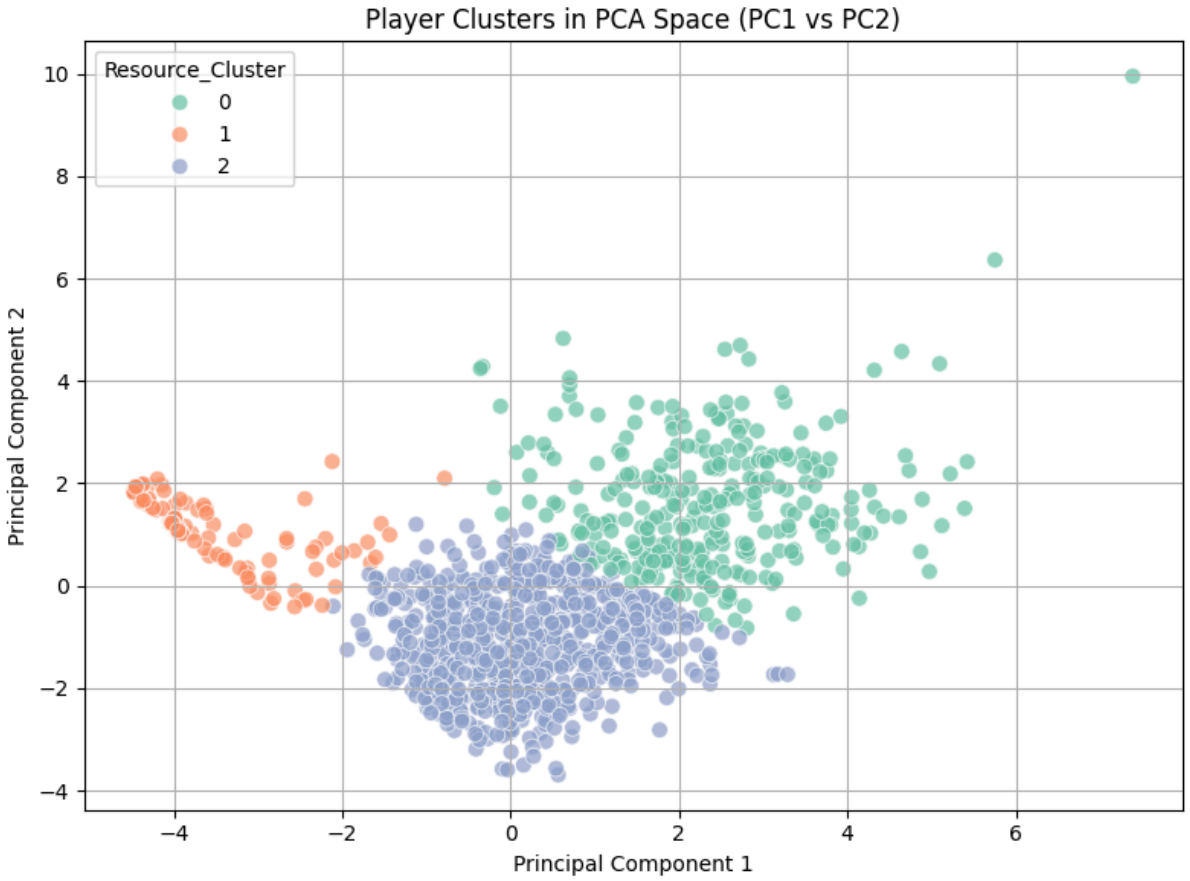


Figure 28 PC1 vs PC2 K=3 Clusters

4.4.4 Cluster Profiles and Archetypes

4.4.4.1 Resource Clusters vs Character Classes

Character Resource_Cluster	DEFECT	IRONCLAD	THE_SILENT	WATCHER
0	0.23	0.36	0.23	0.17
1	0.22	0.30	0.29	0.19
2	0.23	0.37	0.26	0.14

Table 3 Confusion Matrix - Resource Clusters vs Character Classes

Looking at the class breakdown across the three resource management clusters, one thing that stands out immediately is how evenly distributed the classes are. No single class overwhelmingly dominates any one cluster, but there are subtle differences worth noting. **Cluster 2 ('Balanced Adapters')** has the highest proportion of Ironclad players (37%), which makes sense given Ironclad's built-in healing ability that allows for a bit more flexible decision-making with potions and damage intake. On the other hand, the Watcher has the lowest presence across all clusters, particularly in Cluster 2 (just 14%), possibly reflecting the more volatile or specialized nature of Watcher runs, which might not align well with balanced or mixed-resource styles.

The PCA scatter plot supports this interpretation. Players are spread across the PC1–PC2 space in a way that doesn't show hard separations between classes, but there's a definite gradient in how the resource clusters are shaped. You can see clear layering in the markers, particularly with **Cluster 1 ("Efficient Conservers")** concentrated more in the lower-right quadrant, suggesting these players cluster around low resource usage and stable HP trends. **Cluster 0 ("Active Spenders")** tends to occupy more scattered regions, especially toward higher PC1 values, which reflects their high interaction with items and potions, as observed in the cluster profiles.

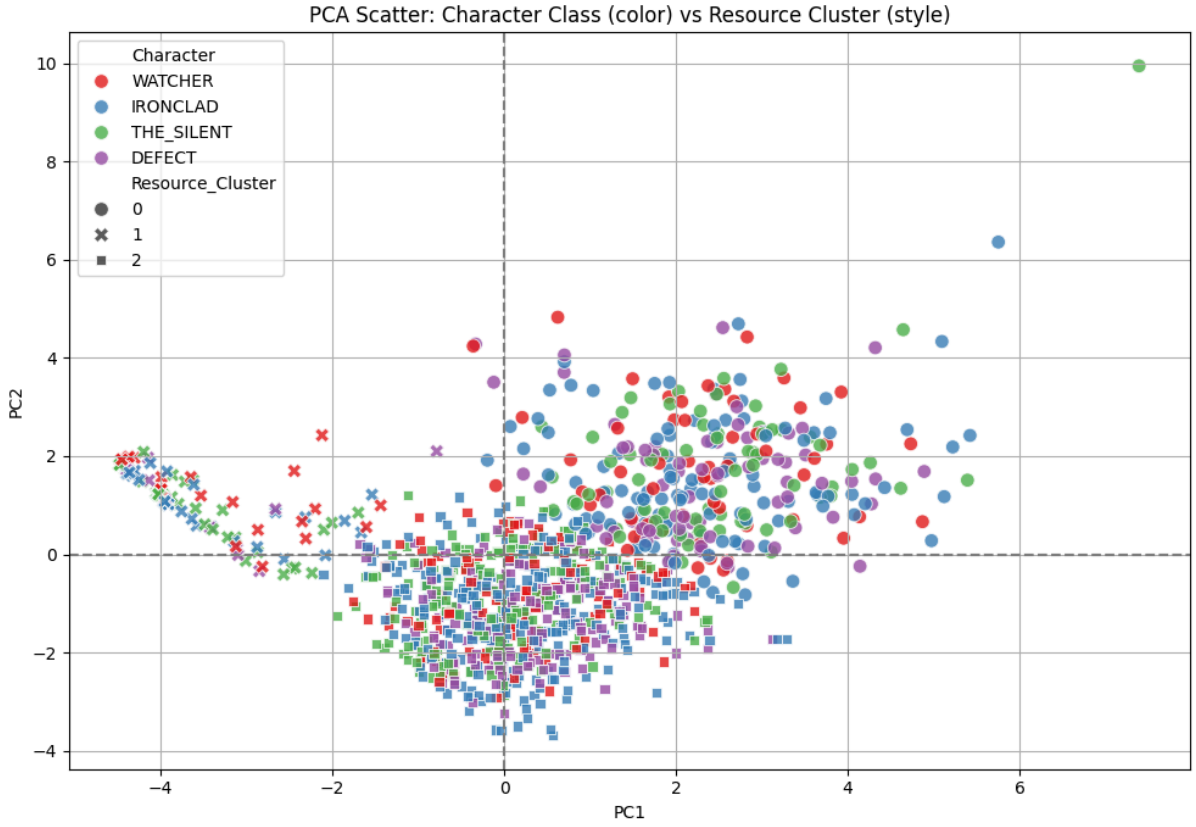


Figure 29 PCA Scatter Plot - Resource Clusters vs Character Classes

Taken together, this suggests that the resource management style in Slay the Spire is not rigidly tied to class, players across classes adopt a variety of strategies. However, some class mechanics do slightly nudge players toward certain tendencies. Ironclads show up more in mixed-style and active-use clusters, while Silents and Defects are more spread out, adapting to either cautious or aggressive strategies. This supports the idea that while class matters, the player's approach to risk, economy, and resource investment plays an equally big role in defining how a run unfolds.

4.4.4.2 Resource Clusters vs Persona Clusters

Persona_Cluster Resource_Cluster	1	2	3	4
0	0.40	0.41	0.19	0.00
1	0.00	0.00	0.06	0.93
2	0.07	0.07	0.59	0.27

Table 4 Confusion Matrix - Resource Clusters vs Player Personas

The relationship between player personas and how they manage their resources reveals a surprisingly tight fit, especially when you consider the motivations and behaviors defined in the original persona analysis. **Gamblers**, for instance, are overwhelmingly represented in Resource **Cluster 1 ("Efficient Conservers")**. At first glance, it seems like a mismatch, since Gamblers are described as impulsive, fast-moving, and luck-driven. But in the context of resource management, it actually makes sense: they're not hoarding potions or spending gold at the shop, they're speeding through runs, grabbing whatever synergy comes first, and seeing if it sticks. The low resource usage isn't because they're careful but it's because they don't bother engaging deeply with those systems at all.

On the opposite end, **Strategists and Bursters**, the thoughtful planners and the synergy-obsessed powerhouses are solidly parked in Resource **Cluster 0 ("Active Spenders")**. These are the players who are deep in the realm of optimization. They tend to slowly build toward a deadly combo or stacking high-impact relics and they're interacting with potions, shops, and their decks much more often. Strategists lean into synergy and utility, often spending carefully but consistently, while Bursters are all about grabbing

power when it shows up, no matter the cost. The scatter plot shows them spread confidently through the more active regions of the PCA space, reflecting their deeper engagement with in-game systems.

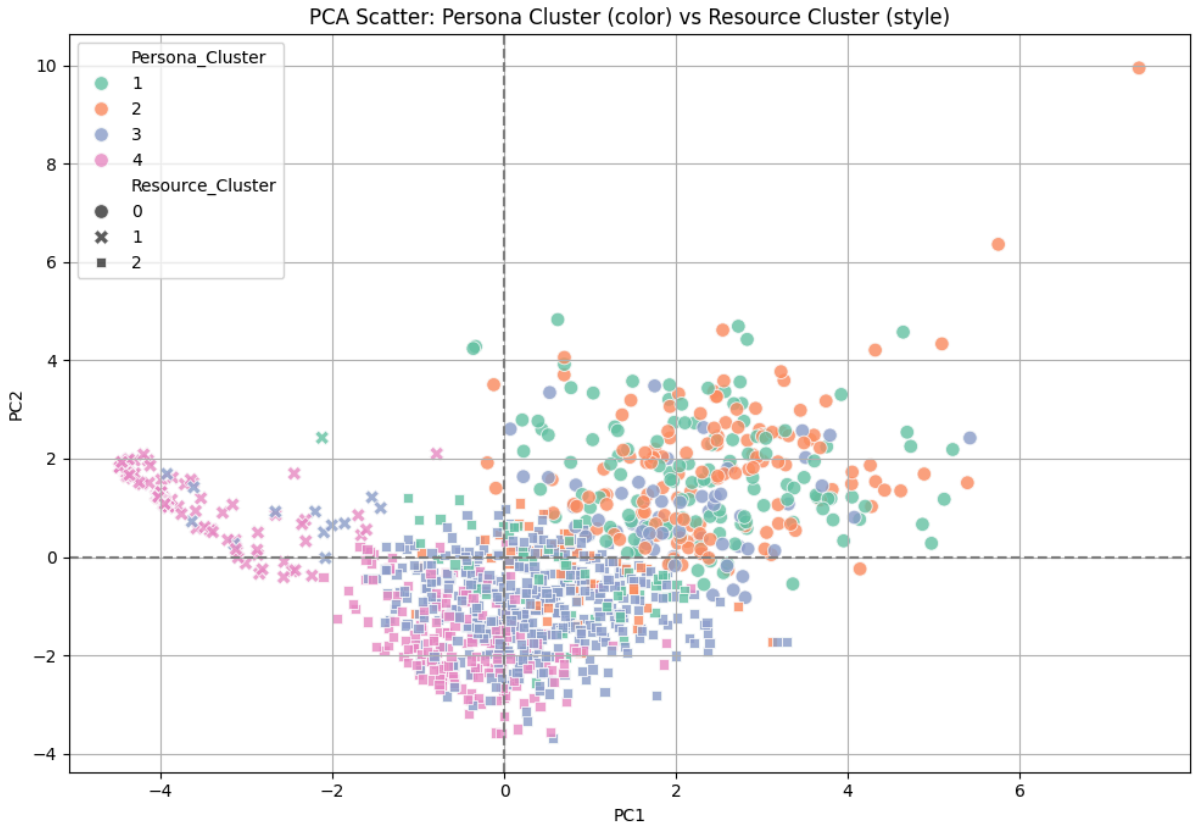


Figure 30 PCA Scatter Plot - Resource Clusters vs Player Personas

Quick Experimenters show a more diffuse spread, but primarily land in Resource **Cluster 2 (“Balanced Adapters”)**. This makes sense. They’re not committing to one synergy or strategy; they’re feeling things out. Their deck-building is broad, their time-per-floor is low, and they make just enough use of resources to test mechanics, but not enough to optimize them. Their resource behavior reflects their mindset: experiment now, maybe commit later. What’s compelling is that these patterns aren’t forced, they emerge naturally from the data, validating the idea that player motivation meaningfully shapes in-game behavior.

5. Future Work

With the data we have collected it would be simple to expand the scope of this project, there are still questions we could ask about the mental model of the player and the choices they made throught their runs. Some questions that were still unanswered were “What are the choices that players prefer over others? And what do they skip to obtain them?”, with

this question we mean to explore what relics and cards does the player obtain while not choosing others and how that relates to their strategy and synergy of their run. Other questions would be in what floor does the player prefer to smith and upgrade cards and with what range of HP they prefer that over healing themselves. The deck analysis and synergy can be expanded with the upgraded cards.

The Risk behavior analysis could be heavily expanded upon with more data points. Because we only took a snapshot of the entire dataset to keep it consistent with updates, and the fact that there are 52 unique events, each event only had so many data points associated with them. This is compounded by the fact that later acts had less data by their nature, and some events are also rarer to encounter, such as The Joust (Section 4.3.2.3). There could be potentially additional findings or insights gained from analyzing a larger dataset.

In resource behavior aspects, it could be worthy to explore resource management strategies within individual character classes to uncover more class-specific behaviors that may be obscured in the combined analysis. Also could explore incorporating elements like event choices or map pathing which could yield deeper insights into how players adapt their resource decisions throughout a run, distinguishing between consistent playstyles and situational responses.

6. Bibliography / Works Cited

DeLise, T. (2020). *Data Segmentation via t-SNE, DBSCAN, and Random Forest* (No. arXiv:2010.13682; Version 1). arXiv. https://doi.org/10.48550/arXiv.2010.13682

Jolliffe, I. (2011). Principal Component Analysis. In *International Encyclopedia of Statistical Science* (pp. 1094–1096). Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-04898-2_455

Martinez-Taboada, F., & Redondo, J. I. (2020a). The SIESTA (SEAAV Integrated evaluation sedation tool for anaesthesia) project: Initial development of a multifactorial sedation assessment tool for dogs. *PLOS ONE*, *15*(4), e0230799. https://doi.org/10.1371/journal.pone.0230799

Martinez-Taboada, F., & Redondo, J. I. (2020b, April 1). *Variable importance plot (mean decrease accuracy and mean decrease Gini).* [Figure]. Figshare; PLOS ONE. https://doi.org/10.1371/journal.pone.0230799.g002

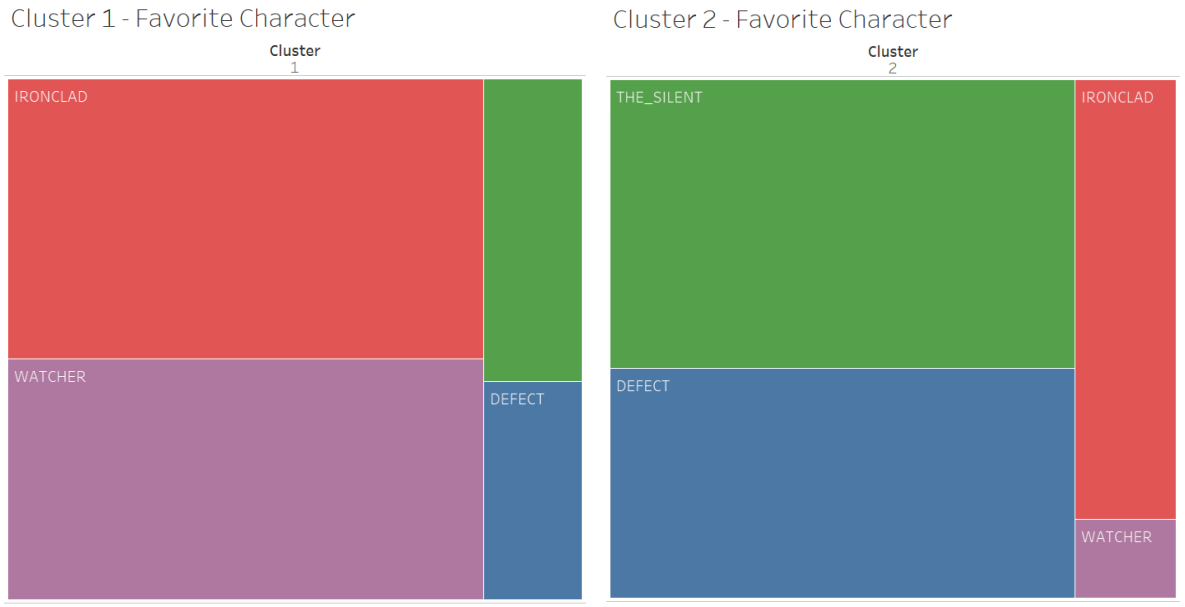
Principal Component Analysis(PCA). (18:43:59+00:00). GeeksforGeeks. https://www.geeksforgeeks.org/principal-component-analysis-pca/

Drachen, A., Thurau, C., Sifa, R., & Bauckhage, C. (2014). A comparison of methods for player clustering via behavioral telemetry. *ArXiv, abs/1407.3950*. https://doi.org/10.48550/arXiv.1407.3950

7. Appendix

7.1 Cluster Exploration

7.1.1 Favorite Player Character by Cluster



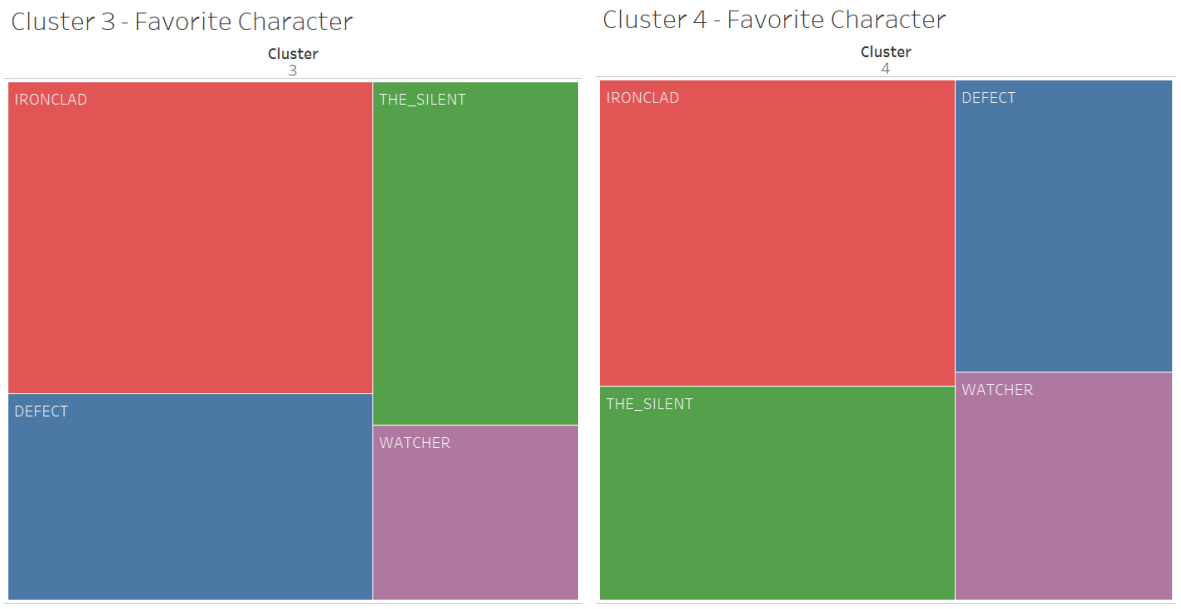


Figure 31. Tree Maps of the favorite character breakdown by cluster. Note that the colors are selected to represent the characters for ease of interpretation.

7.1.2 Popular Cards by Cluster

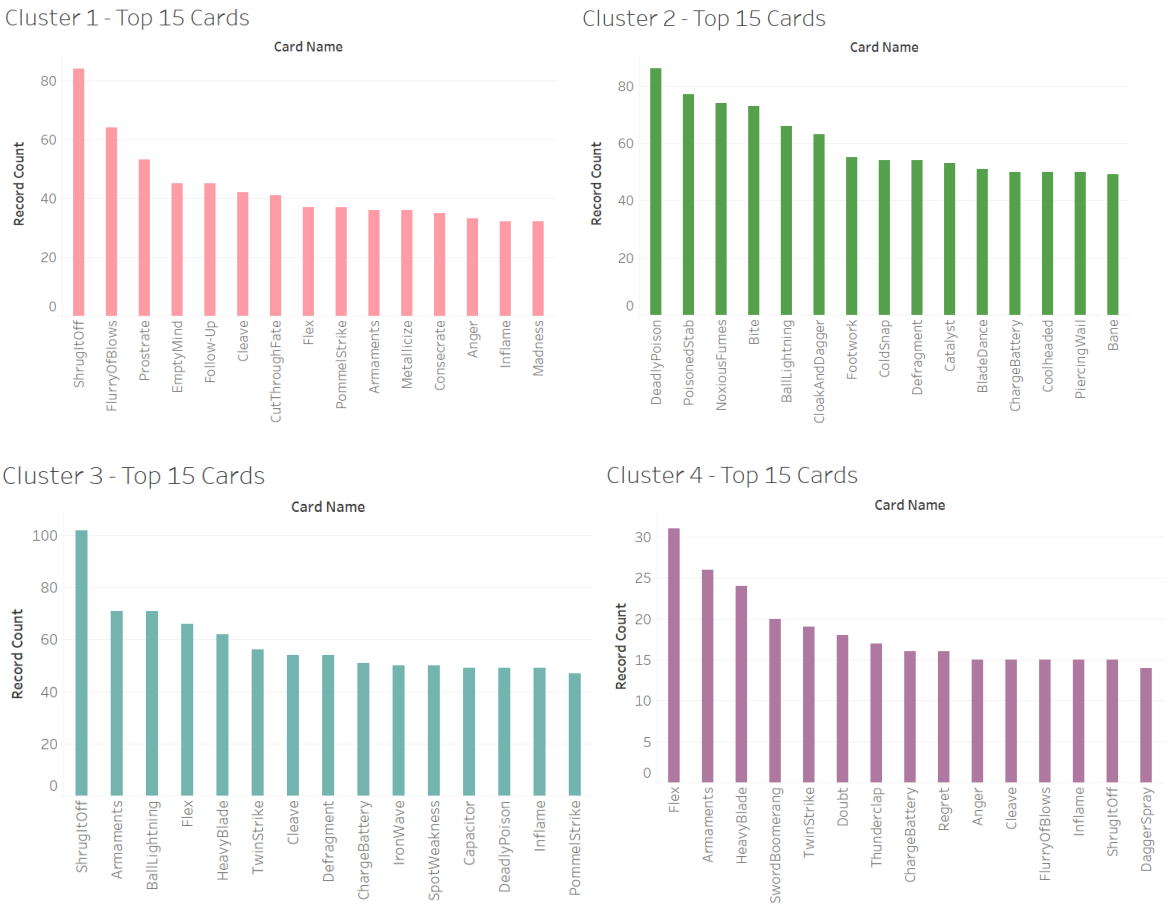


Figure 32. Bar graphs of the popular cards breakdown by cluster. Note that the colors correspond to those assigned by R during the final k-means cluster analysis for ease of interpretation.

7.1.3 Popular Relics by Cluster



Figure 33. Bar graphs of the popular relic breakdown by cluster.

7.1.4 Average Time Per Floor by Cluster

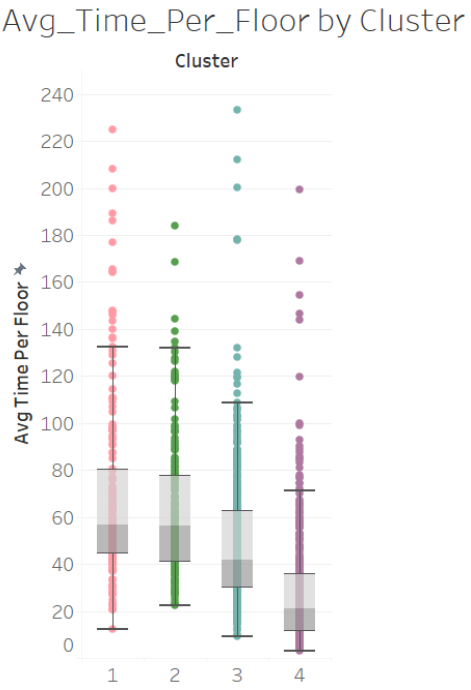


Figure 34. Box plots representing the average time per floor by cluster.

7.1.5 Win Rate by Cluster

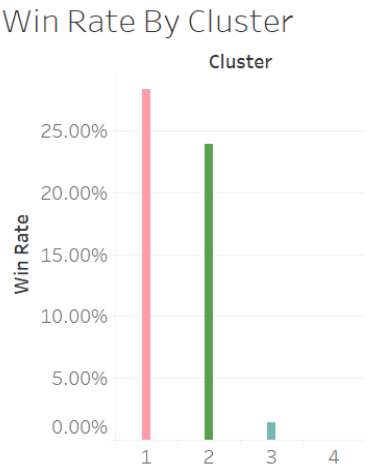


Figure 35. A simple bar graph depicting the win rate by cluster.

Figure 36. Association Matrix with n=50

