

With apologies to Goose Gossage, the day has finally arrived that nerds and their brethren have united to ruin AST. Specifically, Kyle Moran and I have put together some advanced statistics (with the blessing of Shane and webmaster William) to describe the previously unquantifiable. What follows is a primer on how to read/understand these stats, as well as details on how we made these calculations.

Briefly, in every AST matchup, there is one thing over which you have 100% control: how many answers you get right. Your result on that day, however, depends on at least three other factors--two of which you have absolutely no say in. Those are (key terms in bold):

- **Matchup Luck:** How good/bad your opponent is compared to their average performance.
  - A positive number means your opponent has *underachieved* relative to their average; a negative means they've overachieved.
- **Defensive Luck:** How well your opponent limits the value of your correct answers.
  - A positive number means they've defended you *poorly*; a negative means they've been stingy.
- And the one you can somewhat influence, **Defensive Skill:** How well you limit the value of your opponent's correct answers
  - A positive number means you've defended *well*; a negative means you've given up some soft goals.

Each of these reflects variance in one's **score differential** (think plus/minus). We'll use **points** to refer to one's performance in the league standings.

If that's good enough for you, you can click the link [here](#) to see if you've been the recipient of good or bad fortune this past fortnight; note that **Matchup Luck** and **Defensive Luck** have been combined to give a figure called **Net Luck**. If you want to nerd out and see how the statistical sausage is made, read on.

---

Okay, now that we've weeded out the Great Unwashed, we can really dig into how these stats are calculated. We'll start with **Defensive Skill** and **Defensive Luck**, two sides of the same coin.

Defense is the secret sauce that takes trivia from good to great. While it doesn't matter when facing perfect or, um, struggling opponents, for the vast majority of opposition the difference between good defense and bad represents a four-point swing: 1 correct answer can earn 0, 2, or 4 points. But you know this already.

What makes defense so frustrating--and occasionally rewarding--is that it changes from one day to the next. Some days, it is obvious where to assign the K and the HR; others become tactical nightmares.

We accounted for this by calculating--for each day of the fortnight--the expected value of a 1-answer, 2-answer, 3-answer, and 4-answer day. (Again, a 0-answer day will always net 0, and a 5-answer day 10, regardless of defense.) Then, based on how many answers you and your opponent got correct on that day--and the value of those answers--we can see how you performed relative to the baseline.

A quick example: say you tie your opponent 4(2) - 4(3), and across all competitors on the day, the average 3-answer performance earned a score of 5, and the average 2-answer performance earned a score of 3.5.

- Your **Defensive Skill** is your opponent's expected score (5) minus their actual score (4), or 1.
  - This is the same value for your opponent's **Defensive Luck**, but with the sign reversed (i.e. their defensive luck is -1).
- Your **Defensive Luck** is your actual score (4) minus your expected score (3.5), or 0.5.
  - Likewise, your opponent's **Defensive Skill** is the negation of this (-0.5).

Add these stats across the nine matchups of the fortnight, and you have two of the three key metrics.

**Matchup Luck**, the final piece of this triumvirate, accounts for the fact that while Player A may know more questions than Player B over the fortnight, on a given day the five questions Shane has selected may favor Player B. In theory, Player B may have the fewest correct answers in their division, but if Player B faced every other player on *their* worst day of the fortnight, Player B could coast to the title.

To determine how lucky a player is in any given matchup, after each fortnight we divide their correct answers by 9 (Day 10 stats are not included in the standings, and so are irrelevant for this and all future calculations) to find their average expected correct answers. To find a player's Matchup Luck over the fortnight, on each day we compare their daily performance (net of their average) to their opponents' performance (net of their average). The sum of these values across the nine days will reflect *the degree to which a player's opponents outperformed/underperformed their expectation, in terms of correct answers*. This is the **Raw Matchup Luck** figure.

The last thing to do, then, is convert this value into a measure of score differential. To do that, we took the total number of points scored in the fortnight and divided it by the total number of questions answered. We then multiplied that value by the **Raw Matchup Luck** figure for each player to generate the final **Matchup Luck** figure.

So, to recap before you return to the chart [here](#):

- **Defensive Skill** is your opponent's expected score minus their actual score based on their number of correct answers that day, summed across all nine days.

- Your **Defensive Luck** is your actual score minus your expected score based on your number of correct answers that day, summed across all nine days.
- Your **Matchup Luck** is a measure of how many more/fewer questions your opponents have answered correctly relative to their averages, converted into a value that shows how this affected your score differential.
- The sum of the latter two is **Net Luck**.
- The sum of all three is the **Total Score**, which is key for the following section.

Luck matters more for players in the middle of the distribution. In other words, turning thin negative margins into thin positive margins helps a lot more than, say, running up the score vs. *really* running up the score. A dominant player will remain dominant, but for the mediocre, luck is as big a factor as any.