Strategic Modeling Solution

Puzzle: Yannick Yao
Solution: Yannick Yao
Answer: ADMINISTRATOR

This puzzle is a reference to the well-known incremental game *Universal Paperclips*, in which there is one module called Strategic Modeling where multiple strategies play against each other in repeated Prisoner's Dilemma style games, and one need to bet on which strategy will do the best based on the payoff grid.

There are 16 players in this puzzle, each of which has a "real name", which hints at the strategy they will be using, and a "alias" represented by a letter from K to Z. In the six practice rounds, the payoff grid is defined (by the value of α, β, γ, δ), and the full results of all 16^2=256 match-ups are given. From inspecting the results, we can see that each match-up contains exactly 10 rounds, and the payoff corresponds to the row player.

The next big step is to figure out what the 16 players' strategies are, and which player has which alias. With the help of the real names and the large amount of available data, we could determine the player strategies as follows:

[For players with the same adjectives, their initials (A or B) indicates what their first move is, since their strategy depends on their own or their opponent's previous move.]

- Copycat Alice/Betsy [T/Q]: Copies the opponent's previous move. Also known as "tit for tat".
- Envious Eric [U]: Make sure that his payoff is no less than his opponent's payoff.
- Generous Grace [O]: Pick the option that gives the opponent the largest potential payoff.
- Greedy George [K]: Pick the option that gives himself the largest potential payoff.
- Grim Amanda/Brenda [R/L]: Stick with the option that they started with until their opponent picks the other option, and then permanently switches to the other option. Also known as "grim trigger".
- Mean Michael [S]: Pick the option that gives the opponent the least potential payoff.
- Naïve Nancy [Y]: Assume that the opponent pick the same option as her and maximize payoff based on that assumption.

- Prudent Paul [W]: Pick the option that gives himself the largest payoff in the worst case.
- Retrospective Anna/Bela [Z/N]: Pick the option that performs the best against opponent's previous move.
- Simple Aaron/Bobby [P/X]: Always stick with their initial option.
- Wavering Allen/Brian [M/V]: Always pick the opposite option they picked in the previous round.

With these strategies and correspondence, we can now tackle the second part of this puzzle, which is to figure out the payoff tables based on the results of the actual tournament. There are two twists in this part: only the total score is given, and the two worst-performing player are eliminated each round. However, there are some strategies that can help with the deduction. For example, given a payoff table, Eric, Grace, George, Michael, Nancy, Paul will play the same way as either Aaron or Bobby. Based on the total scores, one can figure out which strategies each of them followed, and use this to figure out the size ordering of α, β, γ, δ. Possibly with the help of coding, we can determine the payoff tables as following:

Round	α	β	γ	δ	
1	20	8	5	1	THEA
2	14	19	23	5	NSWE
3	18	9	19	1	RISA
4	4	13	9	14	DMIN
5	9	19	20	18	ISTR
6	1	20	15	18	ATOR

If we convert these values to letters, we get the message THE ANSWER IS **ADMINISTRATOR**.

Author's note: Having two copies of the same strategies was not intended initially, but since there is not good way to canonically define the first move (in the actual game, it actually seems to use the last move of the previous player in their spot as the "previous move", which does not make sense), so I have to include two copies to make sure the distinction gets through. The side effect of this is that the puzzle becomes much larger and time-consuming than I first envisioned, but I still hope the puzzle is somewhat fun.