For the distribution of many resources, competitions are used to decide which *k* of *n* potential competitors are winners (*k*<*n*), typically after making rankings on a criterion. Whereas competitions can be characterized by costs and benefits, we explore their ecology by modeling the probability that a potential competitor will be a winner, taking account of numbers of winners (*k*), potential competitors (*n*), and relative ability level of the individual competitor (*q*). Using simulations, we explore how agents with different profiles of skill levels assess the chances of winning different types of competitions. These lead to an astonishingly simple pattern of results: across different situations and profiles, probabilities cluster at the poles, i.e., close to 0 and 1. In contrast, when we model humans (assuming linear information aggregation) probabilities cluster away from the poles. We therefore use the multiple-cue probability learning paradigm to test whether humans can learn to make the appropriate probability judgments. We find that humans make systematic errors and their estimates do not cluster at 0 and 1. Since the appropriate calculations require considerable computational resources, we explore the possibility of using a simple decision rule. We identify one such heuristic and demonstrate its efficacy. Finally, we discuss changes and uncertainties in the parameters of our model and how these illuminate both the ecology and psychology of competitions. We also suggest extensions to more complex forms of competition.