Asymptotic properties of the K-in-a-row design

Nancy Flournoy[1], José Moler[2], Maider Sada[2]

[1] University of Missouri. Columbia, USA. flournoyn@missouri.edu
[2] Universidad Pública de Navarra, Pamplona, Spain. jmoler@unavarra.es, maider.sada@unavarra.es

Abstract

Searching for the level of a stimulus that has a prescribed percentage of success is a common goal for experiments in many fields. We consider experiments in which only a few values of the stimulus can be applied to the experimental units. We assume that the success response probability increases with the stimulus. For instance, in the development of a new drug, the dose which has a toxicity Γ is of interest. When a design with this objective adapts to concentrate experimental units around the unknown target percentile it is called a dose-finding design.

The K-in-a-Row design (KR design) is a rule that sequentially allocates experimental units to one of the permissible stimuli. As we will prove, the KR design tends to concentrate the allocations around the targeted stimulous. In the context of drug development, assume that M doses are permissible and the target toxicity is $\Gamma < 0.5$. Given the last subject experiences the toxicity at dose d_i, the next subject receives the next lower dose. Otherwise, dose d_i is repeated until K consecutive non-toxic responses are observed; in which case the next subject receives d_{i+1}. Even though this kind of experiment is applied to a small number of subjects, the asymptotic properties of the allocation rule are of interest to characterize the allocations sequences, e.g., to explain its bias and suggest improvements.

In this work we prove that the KR design is a Semi-Markov process in discrete time. A new proof, with fewer conditions, of the strong unimodality of the asymptotic allocation proportions is provided. Moreover, although it was previously known that the allocation mode is adjacent to the target dose, we explicitly identify it.

Our proofs do not require embedding the discrete process into a continuous one. The only assumption we make is that the M probabilities of toxicity increase with dose. This more precisely describes the statistical framework as non-parametric. Our statistical model accommodates contexts in which the response function is unknown, and even when the experimenter is suspicious about irregularities in the response function (non-continuity, non-derivability or, even, sudden changes in the slope of the response function).